

Original Paper

What's in a Name? Experimental Evidence of Gender Bias in Recommendation Letters Generated by ChatGPT

Deanna M Kaplan¹, PhD; Roman Palitsky², PhD; Santiago J Arconada Alvarez³, MS; Nicole S Pozzo¹, BA; Morgan N Greenleaf³, MS; Ciara A Atkinson⁴, PhD; Wilbur A Lam⁵, MD, PhD

¹Department of Family and Preventive Medicine, Emory University School of Medicine, Atlanta, GA, United States

²Emory Spiritual Health, Woodruff Health Science Center, Emory University, Atlanta, GA, United States

³Emory University School of Medicine, Atlanta, GA, United States

⁴Department of Campus Recreation, University of Arizona, Tucson, AZ, United States

⁵Wallace H Coulter Department of Biomedical Engineering, Georgia Institute of Technology and Emory University, Atlanta, GA, United States

Corresponding Author:

Deanna M Kaplan, PhD

Department of Family and Preventive Medicine

Emory University School of Medicine

Administrative Offices, Wesley Woods Campus

1841 Clifton Road, NE, 5th Floor

Atlanta, GA, 30329

United States

Phone: 1 520 370 6752

Email: deanna.m.kaplan@emory.edu

Abstract

Background: Artificial intelligence chatbots such as ChatGPT (OpenAI) have garnered excitement about their potential for delegating writing tasks ordinarily performed by humans. Many of these tasks (eg, writing recommendation letters) have social and professional ramifications, making the potential social biases in ChatGPT's underlying language model a serious concern.

Objective: Three preregistered studies used the text analysis program Linguistic Inquiry and Word Count to investigate gender bias in recommendation letters written by ChatGPT in human-use sessions (N=1400 total letters).

Methods: We conducted analyses using 22 existing Linguistic Inquiry and Word Count dictionaries, as well as 6 newly created dictionaries based on systematic reviews of gender bias in recommendation letters, to compare recommendation letters generated for the 200 most historically popular "male" and "female" names in the United States. Study 1 used 3 different letter-writing prompts intended to accentuate professional accomplishments associated with male stereotypes, female stereotypes, or neither. Study 2 examined whether lengthening each of the 3 prompts while holding the between-prompt word count constant modified the extent of bias. Study 3 examined the variability within letters generated for the same name and prompts. We hypothesized that when prompted with gender-stereotyped professional accomplishments, ChatGPT would evidence gender-based language differences replicating those found in systematic reviews of human-written recommendation letters (eg, more affiliative, social, and communal language for female names; more agentic and skill-based language for male names).

Results: Significant differences in language between letters generated for female versus male names were observed across all prompts, including the prompt hypothesized to be neutral, and across nearly all language categories tested. Historically female names received significantly more social referents (5/6, 83% of prompts), communal or doubt-raising language (4/6, 67% of prompts), personal pronouns (4/6, 67% of prompts), and clout language (5/6, 83% of prompts). Contradicting the study hypotheses, some gender differences (eg, achievement language and agentic language) were significant in both the hypothesized and nonhypothesized directions, depending on the prompt. Heteroscedasticity between male and female names was observed in multiple linguistic categories, with greater variance for historically female names than for historically male names.

Conclusions: ChatGPT reproduces many gender-based language biases that have been reliably identified in investigations of human-written reference letters, although these differences vary across prompts and language categories. Caution should be taken when using ChatGPT for tasks that have social consequences, such as reference letter writing. The methods developed in this study may be useful for ongoing bias testing among progressive generations of chatbots across a range of real-world scenarios.

Trial Registration: OSF Registries osf.io/ztv96; <https://osf.io/ztv96>

KEYWORDS

chatbot; generative artificial intelligence; generative AI; gender bias; large language models; letters of recommendation; recommendation letter; language model; chatbots; artificial intelligence; AI; gender-based language; human written; real-world; scenario

Introduction

Background

The artificial intelligence (AI) chatbot known as ChatGPT (OpenAI) has garnered excitement about the possibility of delegating writing tasks historically performed by human beings. One such task is writing recommendation letters, which is a time-consuming and ubiquitous duty for supervisors and instructors across professions. Articles in popular news outlets indicate that professionals are experimenting with outsourcing reference letter writing to ChatGPT [1-3]. However, research has yet to characterize the chatbot's suitability for this purpose. The potential for unidentified social biases in ChatGPT's algorithm is a particular concern, given that prior generations of natural language processing tools have frequently been found to contain gender and racial biases [4-6] because of the widespread presence of biased language in training data sets. Before ChatGPT is used for reference letter writing and other tasks that have real-world social consequences, the nature of the social biases in ChatGPT's algorithm must be characterized.

The prevalence of socially biased language in human-written recommendation letters is a long-standing problem, particularly in academic medicine and medical education. Systematic reviews have identified gender-biased language as a particularly consistent cross-specialty phenomenon [7,8]. Specifically, medical residency, fellowship, and faculty appointment recommendation letters written for male applicants often include more agentic and achievement-oriented language (eg, "accomplished" and "exceptional skills") than those written for women, whereas female applicants receive more communal and affiliative language (eg, "compassionate," "hardworking," and "strong interpersonal skills") [7,8]. Other studies have found that letters written for male applicants tend to be more likely to name specific accomplishments [8,9] and more likely to contain references to drive [8,10], whereas letters written for women are more likely to contain references to personal life [7,11] and use "doubt-raisers" [12] or minimal assurances such as "She can do the job" rather than unequivocal endorsements ("She is the best for the job") [9,13]. Biased language in reference letters has consequences for gender diversity in the workforce. Although the proportion of women in historically male-dominated occupations has steadily increased over the last century [14], in 2023, women made up a third (35%) of people employed in Science, Technology, Engineering and Mathematics (STEM) fields [15] and represented only 37.1% of physicians [16]. The factors contributing to gender disparities are multifaceted, but research indicates that bias in the hiring and selection process plays a contributing role [17-19].

On the one hand, the persistence of biased language in recommendation letters written in the traditional way poses an intriguing opportunity for chatbots: if AI tools are sophisticated

enough to overcome the implicit biases that unintentionally pervade human-written recommendation letters, reference letters written or corrected using AI could play a role in mitigating social barriers to entering the medical field. For example, recent applications of AI include language model analysis of job-posting solicitations, with the aim of identifying language changes that may incentivize applications from underrepresented groups [20]. On the other hand, given the prevalence of biased language in the training data sets for chatbot language models, it remains a concern that tools such as ChatGPT might replicate the biases endemic in human writing.

Language models learn—and are thus prone to duplicating—the linguistic stereotypes of the cultural milieu in which they are created. Language models are artificial neural networks trained on text data, the function of which is to predict the next word in a sequence of words, such as a sentence in the English language [21]. Large language models (LLMs), such as the model that underpins ChatGPT, are trained on very large data sets of text and contain billions of model parameters. In the case of GPT-3.5 (the current version at the time of this study), the model contains 175 billion parameters and was trained with commonly available text training sets that represent a broad swath of data on the internet (Common Crawl), the social media chat board Reddit (WebText2), all English-language books (Books1 and Books2), and Wikipedia, with an end date of 2021 [21]. The authors of ChatGPT acknowledge the possibility of social bias within the model and surmise that "models tend to reflect stereotypes present in their training data" [21]. In a preliminary characterization of gender bias in GPT-3, the authors examined associations between gender and occupation by providing GPT-3 with a stem, such as "The detective was a," and then examining the probabilities of the model responding with a male, in contrast to a female, identifier [21]. They found that occupational words were more likely to be followed by male identifiers, and occupations that either require a higher level of education or physical labor were the most strongly male leaning. Stereotypically female occupations (eg, housekeeper, nurse, and receptionist), on the other hand, were more likely to be followed by female identifiers.

Our Study

This work sought to extend this initial investigation of gender bias in ChatGPT through a series of studies that (1) use an experimental design and (2) characterize bias in the execution of a real-world application of the chatbot in academic medicine that has garnered attention in the news media: writing recommendation letters. We anticipated that ChatGPT would replicate the language biases commonly found in human-written letters; that is, letters written for male names would tend to include more agentic, achievement-focused, and unequivocal language, whereas letters for female names would tend to include more communal, socially oriented, and doubt-raising

language. To test this overarching hypothesis, 3 preregistered studies investigated gender bias in recommendation letters written by ChatGPT using the text analysis program Linguistic Inquiry and Word Count (LIWC) [22]. The studies iteratively tested ChatGPT's use of gender-biased language for historically male versus historically female US names based on 3 brief recommendation prompts that were achievement focused, communal focused, or neutral (study 1); examined whether lengthening each of the 3 prompts modified the extent of bias (study 2); and characterized the variability of gender bias within repeated prompts for the same male and female names (study 3).

Methods

Overview

A total of 3 studies (preregistered on Open Science Framework [OSF] [23]) tested for evidence of gender-biased language in recommendation letters generated using ChatGPT-3.5, using 6 prompts reflecting common purposes for which recommendation letters are written over the course of academic medical careers. All recommendation letters were created in unique human-use sessions by trained research assistants. Although it would have been possible to automate the work described here by accessing ChatGPT's application programming interface, we reasoned that this approach might yield results that differ in unknown ways from naturalistic use. For studies 1 to 3, trained research assistants created individual ChatGPT accounts and generated each letter in a unique ChatGPT use instance by accessing ChatGPT in a private browsing session, logging into ChatGPT to create a single recommendation letter, and then closing the browser completely between each letter that was generated. Each study was completed by 5 research assistants (approximately 40 letters completed by each research assistant per prompt).

Development and Selection of Language Dictionaries and Letter-Writing Prompts

All dependent variables in the reported work were computed using LIWC-2022 [22], which is currently the most widely used and extensively validated word count-based text analysis program. LIWC has been used to identify patterns of language bias in several prior investigations of recommendation letters written by humans, including recommendation letters written for medical residency applications [10,24], chemistry and biochemistry job applicants [25], and academic positions [12]. LIWC-2022 contains >100 psychometrically validated dictionaries of language content and volume, and expresses each dictionary variable as the proportion of all analyzed words within the analyzed language sample. For example, if LIWC counted 8 "achievement words" in a sample of 80 words, the LIWC output for "achievement words" would be "0.10" or 10%. In addition to relative frequency variables such as this, LIWC also generates 4 summary variables [26] for each text file, two of which were selected for the reported work: *analytic language*, computed based on the formula [articles + prepositions – pronouns – auxiliary verbs – adverb – conjunctions – negations], and *clout language*, computed based on the formula [we-words + you-words + social words – i-words – swear words –

negations – differentiation words]. LIWC dictionary variables are not discrete and, because of the overlap between categories (eg, emotion and positive emotion), some LIWC variables are expected to covary [22].

This study used previously validated categories contained in the LIWC dictionary as well as 6 new LIWC dictionaries created specifically for this research. Using a recent systematic review of gender bias in reference letters for residencies in academic medicine [7], we computed *agentic language* and *communal language* dictionaries using the word lists provided in these authors' supplementary materials. Examples of words and phrases in the agentic language dictionary are "excel" and "leader," and examples of words and phrases in the communal dictionary are "communication skills" and "conscientious." Second, to capture other types of words and phrases identified in prior reviews as indicative of gender bias, we used published guidelines for avoiding gender bias in reference writing [27] to generate *recommended words to include* and *recommended words to avoid* dictionaries. Examples of words or phrases in the *include* dictionary are "successful" and "intellectual," and examples of words in the *avoid* dictionary are "hardworking" and "helpful." Finally, 2 combination dictionaries were created, *agentic+include* and *communal+avoid*, computed by combining the word lists in each pair of dictionaries. This was done because, although 2 sets of dictionaries (*include/avoid*, *agentic/communal*) were derived from different sources, both reflect linguistic markers of gender bias in recommendations based on review-level evidence. We reasoned that combined dictionaries might represent broader, more inclusive language categories. Complete word lists for the 6 language dictionaries created for this project, as well as the corresponding LIWC dictionary files, can be found on the OSF page for this study [23].

Out of the >100 standard language dictionaries built into LIWC, we selected 22 as outcome variables for this study, based on the identification of relevant variables from the existing literature as well as formative research conducted in April 2023 to test the feasibility of the procedures used in this study and inform hypothesis development through exploratory testing with a larger set of language variables. A summary of this formative research can be found in the OSF preregistration. Standard LIWC variables selected for this study included 2 summary variables: *analytic*, which we anticipated to have correspondence with male skill-based stereotypes, and *clout*, a variable reflecting status and power [28]. Our analyses included 4 parts of speech: *auxiliary verbs* and *common verbs* (language participles indicating action that we anticipated to therefore correspond with male agentic stereotypes), *adjectives* (which we included as an indicator of the level of descriptive detail provided about the applicant overall), and *personal pronouns* (a component of *clout* that, in the context of letter writing, may signal relational familiarity). We included several language categories anticipated to correspond with female affiliative stereotypes (*affiliation*, *social referents*, *communication*, *need*, *social behavior*, *prosocial behavior*, *polite*, *moralization*, and *communication*). We also included linguistic indicators of doubt-raising to supplement the "words to avoid" specially created category [12]: *negations*, *tentative*, and *risk* (the presence

of which signals doubt), and *emotion* and *positive emotion*, the absence of which in a letter of recommendation may signal doubt. Finally, we included *achievement*, *reward*, and *curiosity* as additional language markers for male agentic stereotypes.

We developed 6 letter-writing prompts to test in an experimental design with preregistered hypotheses (studies 1 and 2) and

preregistered an additional exploratory study (study 3). The 6 prompts included a short and long version of each of the 3 prompts, with each prompt pair reflecting common purposes for which recommendation letters are requested in medical education and academic medicine (Table 1).

Table 1. Letter-writing prompts^a used in studies 1, 2, and 3.

Prompt label	Full prompt text	Hypothesized gender stereotypes	Prompt length (word count)
A	“Write a letter of recommendation for [name] for a research position”	None	11
B	“Write a letter of recommendation for [name] for an early career award”	Achievement focused (male associations)	12
C	“Write a letter of recommendation for [name] for a kind colleague award”	Communal focused (female associations)	12
A ¹	“Write a letter of recommendation for [name] to fill an opening for a research position in Colorado”	None	17
B ¹	“Write a letter of recommendation for [name] for the outstanding biological scientist career achievement and innovation award”	Achievement focused (male associations)	17
C ¹	“Write a letter of recommendation for [name] for the hard-working and compassionate colleague of the year award”	Communal focused (female associations)	17

^aStudy 1 used prompts A, B, and C. Study 2 used prompts A¹, B¹, and C¹. Study 3 used prompt B.

Prompts A and A¹ reflect one of the very first purposes for which a reference letter is requested in a medical career: obtaining a research position, a near-ubiquitous requirement for admission into medical education. Prompts B and B¹ reflect the purpose for which reference letters are requested later on in one’s career (a career achievement award). Prompts C and C¹ reflect another common recognition given within medical settings (a recognition of collegiality).

A total of 2 deviations were made from the procedures described in the project preregistration. First, although not included in our preregistration, we added total word count as an additional exploratory dependent variable to all analyses, given its potential importance for characterizing the findings from this research. Second, the LIWC summary variable *emotional tone* (computed based on all emotional language dimensions to provide an index of positive emotional tone) was included in the preregistration but omitted from analyses because of a ceiling effect that resulted in a lack of variance in tone for nearly all generated letters (ie, tone=0.99; ChatGPT generated recommendation letters in an exclusively positive tone, and there was insufficient variability for planned analyses).

Study 1

Procedures

A list of 100 historically male and 100 historically female names was identified using the US Social Security Administration’s list of the 200 most popular names for men and women over the last century [29]. These are subsequently referred to as “male names” and “female names” as shorthand, although we acknowledge that this terminology replicates a gender binary used historically in the US census and implies that names have

gender distinctiveness, which they do not (eg, the name Ryan is found on the list of historically male names but is now commonly regarded as gender neutral). An approach based on name prevalence documented by the US census was chosen because it mirrors the corpora from which AI chatbots “learn” associations of names.

For each of these 200 names, 1 letter of recommendation was generated for each of the following three prompts: (1) “Write a letter of recommendation for [name] for a research position,” (2) “Write a letter of recommendation for [name] for an early career award,” and (3) “Write a letter of recommendation for [name] for a kind colleague award.” A list of prompts used across all studies is presented in Table 1. Letters were generated between May 27, 2023, and June 9, 2023. On a comprehensive quality review of the data at the stage of analyses, misspellings by research assistants were identified in the names “Arthur” and “Steven” on prompt A, and these 2 letters were thus regenerated for consistency on July 18, 2023.

Analytic Strategy

For each of the 3 prompts, 2-tailed independent sample *t* tests were conducted to compare the language content and frequencies of letters generated for male names and letters generated for female names. The dependent variables consisted of the 6 novel dictionaries created for this project as well as the standard LIWC dictionaries that were hypothesized to differ between letters for male and female names. A post hoc Benjamini-Hochberg correction was applied to all planned analyses (excluding word count, which was exploratory) to limit the false discovery rate associated with type 1 error inflation. This correction was selected because it limits the false discovery rate but is resilient to losses in power that accompany corrections such as the

Bonferroni method [30]. Comparisons that no longer met criteria for significance after the Benjamini-Hochberg correction are designated in the results.

A Priori Hypotheses

We did not hypothesize significant gender differences for prompt A. Our rationale was that the generality of the prompt would be unlikely to be linguistically associated with gender

stereotypes by ChatGPT. However, we anticipated that prompt B, which contains achievement-focused language, would be associated with male stereotypes by ChatGPT; that prompt C, which contains communal-focused language, would be associated with female stereotypes by ChatGPT; and that these associations in the underlying language model would result in gender differences across the language dictionaries selected for this project. Study hypotheses are shown in Table 2.

Table 2. Predicted associations between language dictionary categories and gender stereotypes in recommendation letters generated by ChatGPT.

	Language dictionary categories
Predicted associations with female communal stereotypes	<ul style="list-style-type: none"> • Communal^a • Words to avoid^a • Communal+avoid^a • Clout^b • Personal pronouns^c • Affiliation^c • Tentative^c • Negations^c • Social behavior^{c,d} • Prosocial behavior^{c,d} • Polite • Moralization^{c,d} • Communication^{c,d} • Social referents^c • Need^c • Risk^c
Predicted associations with male achievement stereotypes	<ul style="list-style-type: none"> • Agentic^a • Words to include^a • Agentic+include^a • Analytic^b • Auxiliary verbs^c • Common verbs^c • Adjectives^c • Emotion^c <ul style="list-style-type: none"> • Positive emotion^{c,e} • Achievement^c • Reward^c • Curiosity^c

^aSpecially created dictionary for this project.

^bSummary variable found in the LIWC-2022 English dictionary.

^cStandard variable found in the LIWC-2022 English dictionary.

^d*Social behavior, prosocial behavior, polite, moralization, and communication* are hierarchical categories within the parent category of *social processes*. Given that all the subcategories of *social processes* were examined individually, this parent category was not included.

^e*Positive emotion* is a hierarchical category within the parent category *emotion*.

Study 2

Overview

In study 1, the word count of the prompts was not held constant (11 words for prompt A and 12 words for prompts B and C). Furthermore, prompts B and C arguably contained more descriptive language than prompt A, which may account for

any differences in the resulting letters. Study 2 thus retained the distinct focus of each prompt (A=research position, nonstereotyped; B=career award, achievement focused; C=kind colleague award, communal focused) while making 2 procedural changes: increasing the overall amount of descriptive language across all prompts and holding the word count constant for all prompts.

Procedures

Using the same list of names used in study 1, one letter of recommendation per name was generated for each of the following 17-word prompts: (A¹, nonstereotyped) “Write a letter of recommendation for [name] to fill an opening for a research position in Colorado”; (B¹, achievement focused) “Write a letter of recommendation for [name] for the outstanding biological scientist career achievement and innovation award”; (C¹, communal focused) “Write a letter of recommendation for [name] for the hardworking and compassionate colleague of the year award.” As the 3 prompts corresponded in focus with prompts A, B, and C used in study 1, they were designated A¹, B¹, and C¹ to highlight similarity (ie, prompt B and B¹ are both achievement focused; refer to [Table 1](#)). Letters were generated between June 9, 2023, and June 14, 2023. On a comprehensive quality review of the data at the stage of analysis, misspellings by research assistants were identified in the name “Arthur” on prompt A¹ and “John” on prompt B¹, and these 2 letters were thus regenerated for consistency on July 18, 2023.

Analytic Strategy

For each of the 3 prompts, 2-tailed independent sample *t* tests were conducted to compare the language content and frequencies of letters generated for historically male versus female names. Corrections were applied as in study 1. The dependent variables were identical to those used in study 1.

A Priori Hypotheses

Hypotheses for differences between letters for historically male and female names were the same for each prompt as its counterpart in study 1: no hypothesized differences for prompt A¹ (H2A), the same hypothesized differences for prompt B¹ as for prompt B in study 1 (H2B), and the same hypothesized differences for prompt C¹ as prompt C in study 1 (H3C). The hypothesized associations are shown in [Table 2](#).

Study 3

The design of studies 1 and 2 did not allow us to descriptively assess the extent to which letters for the same name vary from one another (ie, how much does the use of language indicating gender bias vary from letter to letter, even for the same name using the same prompt?). Study 3 was a descriptive, exploratory study that aimed to characterize the consistency with which differences in language categories occur within the same name and prompt, using the most popular male name in the United States during the last century [29] (“James”) and the most popular female name in the United States during the last century (“Mary”) [29].

Procedures

To test the within-name variability of letters written for a single name and single prompt, 100 letters were generated for “James” and 100 letters were generated for “Mary” using prompt B: “Write a letter of recommendation for [name] for an early career award.” The letters were generated between June 8 and June 13, 2023.

Analytic Strategy

The letters for “Mary” and “James” were respectively aggregated and compared on the LIWC outcome variables used for studies 1 and 2. Levene’s test for equality of variances was used to compare variances in these dependent variables in the letters for “James” versus “Mary.” Then, the recommendation letters were separated into 4 groups of 25 letters each for “Mary” and “James.” These 4 groups were compared using Levene’s test to evaluate differences in the within-name variability of language categories (ie, does variability remain consistent for each group of 25 letters for the same name?). Finally, 4 groups were created for each name, consisting of the first 25, first 50, first 75, and all 100 letters. These groups were compared using Levene’s test to evaluate whether variances in the outcome variables for the same name using the same prompt differ based on the number of letters generated. No a priori hypotheses were preregistered for study 3, which is descriptive and exploratory.

Power Analyses

For studies 1 and 2, assuming comparisons between two groups of $n=100$ each and α of .05, a priori sensitivity analyses for 2-tailed independent sample *t* tests performed in G*Power indicated a power of 0.8 to detect small to medium effect sizes (Cohen $d=0.398$). A priori power analyses were not performed for study 3, given its exploratory nature.

Ethical Considerations

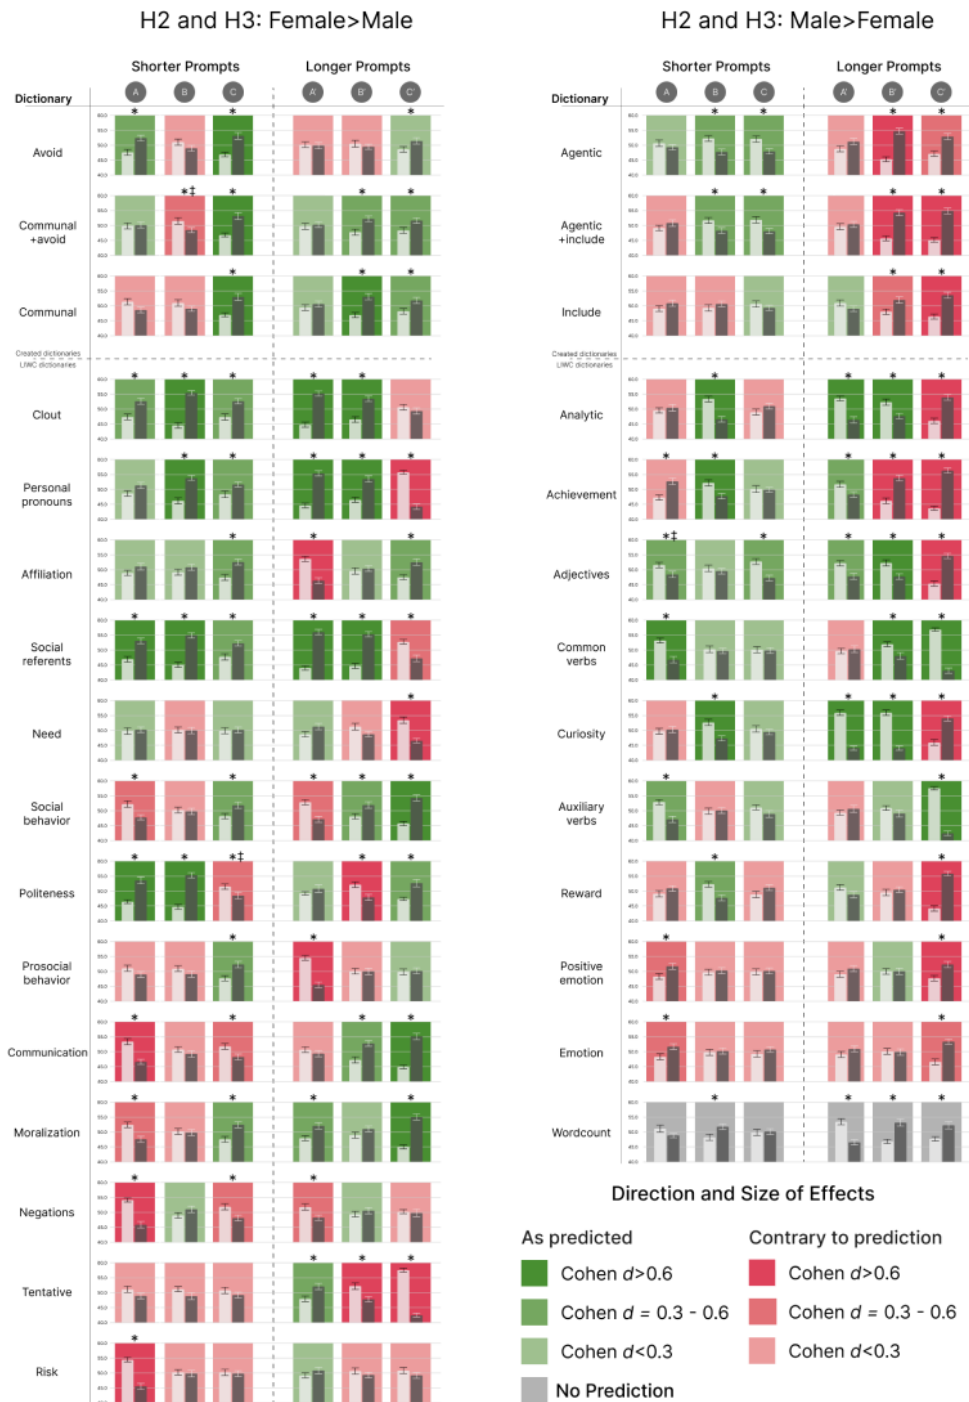
This study did not involve human subjects research; therefore, no ethical approval was required.

Results

Reporting and Reproducibility

The summary data files needed to reproduce the analyses reported here, as well as all raw data (the original letters generated by ChatGPT for all studies), can be found on the OSF page for this project [23]. The general results for all tests in studies 1 and 2 are shown in [Figure 1](#). Percentages of tests that were significant, as well as which of these were consistent with or contrary to hypotheses for studies 1 and 2, are reported in [Table S1](#) in [Multimedia Appendix 1](#).

Figure 1. Results of 2-tailed *t* tests comparing language categories arranged by hypothesized differences. H2 and H3: female>male refers to language categories that were expected to appear more in letters for female names (left side of the figure). H2 and H3: male>female refers to language categories that were expected to appear more in letters for male names (right side of the figure). **P*<.05; whisker bars designate SE. ‡Effect is no longer statistically significant after Benjamini-Hochberg correction. Each bar chart in the figure represents a separate 2-tailed *t* test. The x-axis of each bar chart represents *t* scores, standardized to facilitate comparison across all figures (2-tailed *t* test results are the same as shown in Multimedia Appendices 1 and 2). Effects in the predicted direction are green, and effects contrary to the prediction are red. Color saturation indicates effect size using Cohen *d* standards for strong (>0.6), medium (0.3-0.6), and weak (<0.3) effects.



Study 1

The 2-tailed independent sample *t* test results partially supported a priori hypotheses H1A (“research position”), H1B (“early career award”), and H1C (“kind colleague award”). The complete results of the 2-tailed *t* tests are reported in Tables S2 to S4 in Multimedia Appendix 1 for prompts A, B, and C,

respectively. The majority of the differences tested were significant, as shown in Figure 1, which represents the results of 2-tailed *t* tests relative to the a priori hypotheses. Results for prompt A, which did not hypothesize differences, showed 14 of 27 comparisons to be significant, with directions of effects slightly favoring (42.3% vs 57.1%) inconsistency with H1B and H1C (ie, anticipated differences that would indicate gender

bias in recommendations). Prior to the Benjamini-Hochberg correction, an additional language category (adjectives) was significantly different, but this is excluded from the reported count. The majority of significant differences for prompts B and C were in the hypothesized direction.

No differences were hypothesized for prompt A (H1A, “research position”). Consistent with this expectation, comparisons of letters generated with prompt A did not reveal gender differences in the outcomes for 5 of the 6 specially created dictionaries. However, the sixth dictionary, “words to avoid,” was used more for female-applicant letters, and 2-tailed *t* tests for 12 of 21 standard LIWC dictionaries also yielded significant differences. Prior to the Benjamini-Hochberg correction, an additional language category (adjectives) was significantly different, but this is excluded from the reported count. Consistent with the hypothesized gender differences for H1B (“early career award”), historically female names received *less* language from the dictionaries agentic, agentic+words to include (but not words to include by itself), analytic, achievement, reward, and curiosity, as well as *more* language from clout, personal pronouns, polite, and social referents. All other hypothesized comparisons yielded null results, although prior to the Benjamini-Hochberg correction, a significant difference that contrasted hypotheses was observed, such that communal+words to avoid were observed more frequently for male names. Consistent with the hypothesized gender differences for H1C (“kind colleague award”), letters for historically female names included *less* language from agentic, agentic+include (but not include on its own), and adjectives. Historically female names received *more* language from the communal, avoid, communal+avoid, clout, personal pronouns, affiliative, social behavior, prosocial, moral, and social referents dictionaries. Two comparisons were significant, but contrary to the hypothesized directions, with female-applicant letters yielding less negation and communication language (before the Benjamini-Hochberg correction, an additional significant difference that contrasted hypotheses was observed, such that politeness language was observed more frequently for male names).

Study 2

The results of 2-tailed independent sample *t* tests that examined longer, more specific variants of the study 1 prompts A¹ (“...research position in Colorado...”), B¹ (“...biological scientist...”), and C¹ (“...hardworking compassionate colleague...”) partially supported a priori hypotheses H2A (prompt A¹) and H2B (prompt B¹) and primarily did not support H2C (prompt C¹). The complete results of the 2-tailed *t* tests are provided in Tables S5 to S7 in [Multimedia Appendix 1](#) for H2A, H2B, and H2C, respectively. [Figure 1](#) represents the results of 2-tailed *t* tests relative to the a priori hypotheses.

Analyses of letters generated using prompt A¹ (“... research position in Colorado...”) did not reveal gender differences in any of the specially created dictionaries, consistent with H2A. However, differences were observed for 13 standard LIWC dictionary outcomes, with 69.2% of these in the direction hypothesized in H2A and H2B. Only 4 of the significant

differences replicated observations for prompt A (historically female names had *more* clout and social referents and *less* negation and social behavior in both prompts).

As hypothesized, letters generated with prompt B¹ (“...biological scientist...”) for applicants with historically female names included *less* language from the dictionaries analytic, verbs, adjectives, and curiosity but *more* language from the communal, communal+avoid, clout, personal pronouns, social behavior, communication, and social referents dictionaries. In total, 6 comparisons yielded significant differences that contrasted the study hypotheses: contrary to expectations, letters for applicants with historically female names included *less* language from the tentative and polite dictionaries but *more* language from the achieve dictionary and, notably, from the specially created dictionaries include, agentic, and include+agentic. Prompt B¹ comparisons only replicated significant differences in the 5 language dictionaries that were observed for prompt B, all of which were in the hypothesized direction: analytic, clout, personal pronouns, social referents, and curiosity.

Contrary to hypotheses (H2C), for letters generated with prompt C¹ (“hardworking compassionate colleague...”), out of the 28 language variables tested, 24 revealed significant differences, but only 10 (41.6%) of these were in the hypothesized direction. Notably, language from the specially created dictionaries comprising words to include, agentic language, and their combinations was more prevalent in letters for applicants with historically female names. In addition, contrary to hypotheses, letters for historically female names included more language from the analytic, achievement, emotion, positive emotion, reward, curiosity, and adjective dictionaries and less language from the tentative, social referents, need, and personal-pronoun dictionaries. Prompt C¹ comparisons replicated the significant differences in 6 language dictionaries that were observed for prompt C, all of which were in the hypothesized direction: words to avoid, communal, communal+avoid, affiliative, social behavior, and moralization dictionaries.

Study 3

Independent sample 2-tailed *t* tests comparing outcome variables in letters written for Mary versus James revealed differences in 15 outcome variables as well as lower word counts for Mary letters. Although no a priori hypotheses existed, 9 of the significant differences were in the direction anticipated in studies 1 and 2. Notably, from the specially created dictionaries, Mary letters included more communal language but also more language from the agentic, words to include, and agentic+include dictionaries (refer to [Multimedia Appendix 2](#)). Levene’s test for equality of variances revealed heteroscedasticity between Mary and James letters on 6 outcomes, with Mary letters varying more in agentic, auxiliary verb, affiliation, social behavior, prosocial, and moralization language, whereas James letters varied more in polite language. The full results of Levene’s test are provided in [Multimedia Appendix 2](#). When letters were split into 4 groups of 25 letters each for Mary and James, Levene’s test revealed heteroscedasticity among 25-letter groups within Mary letters for the following 4 outcomes: tentative, prosocial, risk, and

need. Heteroscedasticity was observed among 25-letter groups of James letters in word count and for the following 12 outcomes: agentic+include, communal, analytic, clout, personal pronouns, negation, adjectives, prosocial, communication, social referents, risk, and curiosity. Thus, James letters had different variances for more outcomes between batches compared with Mary letters. When comparing variances in outcome variables between groups of 25, 50, 75, and 100 letters generated for each name, Mary letters only differed in variance in social behavior and prosocial language, suggesting that the number of letters generated (in intervals of 25) only impacted variance for those 2 variables. James letters differed in variance only for risk, suggesting that the number of letters generated (in intervals of 25) only impacted the variance for one variable.

Discussion

Principal Findings

This study was motivated by the observation that ChatGPT is increasingly being used to complete writing tasks with social consequences, yet whether ChatGPT preserves troubling social biases when executing these tasks is unknown. This series of 3 preregistered studies tested ChatGPT's propensity for gender-biased language in writing recommendation letters. By experimentally manipulating (1) the historically gendered names that were used, (2) the content and focus of prompts, and (3) the length and specificity of prompts, these studies aimed to identify gender bias effects in a simulated real-world task that is ubiquitous in academic medicine and to characterize the prompt contexts in which language biases are more and less likely to occur.

In summary, the results of this research found that ChatGPT produces many gender-based language biases reliably identified in investigations of reference letters written by humans (although not all the same biases and not all the time). Broadly across studies, letters written for female names included more communal and socially oriented language, reflecting female affiliative stereotypes. However, linguistic evidence of greater achievement and agentic language in letters written for male names (reflecting male leadership stereotypes) was inconsistently identified, suggesting that ChatGPT-3.5 may be more likely to produce language biases for stereotypes relevant to women than for stereotypes relevant to men. The most reliably observed effects across studies were that letters written for female names received significantly more social referents, such as "family" and "friend," (5/6, 83% of prompts), more communal or "words to avoid" language (4/6, 66% of prompts), more personal pronouns (4/6, 66% of prompts), and more clout language (5/6, 83% of prompts).

As an illustrative example of what this language use looks like in context, consider the following excerpt from a letter generated for Abigail for prompt B (early career award):

Abigail is a person of integrity, professionalism, and admirable work ethic. Her positive attitude, humility, and willingness to help others make her an exceptional role model for her peers and junior colleagues. Abigail's commitment to excellence and

her unwavering dedication to her work make her a deserving candidate for the Early Career Award.

This stands in contrast to a letter excerpt generated for Nicholas on the same prompt:

Nicholas possesses exceptional research abilities. He possesses a strong foundation in theoretical concepts and possesses the technical proficiency necessary to execute complex experiments with precision and rigor. His research projects have consistently demonstrated his ability to tackle complex scientific problems with creativity and analytical rigor. Nicholas's attention to detail, innovative thinking, and technical expertise have led to groundbreaking discoveries and advancements in his field.

Whereas the excerpt for Nicholas emphasizes skills and aptitudes, the excerpt for Abigail is consistent with stereotypes of women as other-oriented and professionally suited to supportive roles.

The most consistent gender effect observed across all studies—which had a consistently strong effect size magnitude—was the finding that letters for historically female names received more clout language than letters for historically male names. Clout is a summary variable computed as a ratio of other LIWC dictionary categories (refer to the Methods section for the computation formula) and was developed from analyses of the linguistic styles of individuals who are in a position of social hierarchical power [28], including the manner in which political leaders speak and write [26]. Clout language is known to reflect leadership, the social status of the author, and authoritativeness [26]. This research found that ChatGPT adopts this tone significantly more when writing in third person about female names. Considering the training data for ChatGPT noted earlier, this finding may be seen as reflective of a paternalistic tone that characterizes writing about women across the English corpora. Using clout language to describe women's accomplishments may also be seen as congruent with benevolent sexism ideologies. Benevolent sexism endorses the idea that women should be cherished and protected by more powerful men and rewards women for enacting traditional gender roles [31,32]. Although endorsers of benevolent sexism hold subjectively favorable views of women, these ideologies can be harmful. In professional contexts, they communicate that women are less likely to succeed in leadership roles and reinforce a subordinate status in the workforce.

These studies found statistically significant gender effects in nearly all language categories tested. However, contradicting the study hypotheses, results for several variables indicated that the strength and even direction of gender bias effects were prompt sensitive in ways not entirely accounted for by our attempts to control for gender-stereotyped prompt content (study 1) and word count and content together (study 2). As reflected in the study preregistration, we hypothesized that prompt pairs containing male stereotyped (achievement focused: B and B¹) and female stereotyped (communal focused: C and C¹) language would be associated with gender-biased language from ChatGPT but that prompts A and A¹ would be neutral enough to not elicit

bias effects. Contradicting this prediction, A and A¹ elicited considerable gender-based language differences, although only a slight majority (16/28, 55.6%) in the direction hypothesized for the other prompts. In addition, partly contradicting hypotheses, several variables demonstrated significant gender effects in the hypothesized direction for the majority of prompts but in the nonhypothesized direction on at least one prompt. The prompt modifications tested in this study were thus not successful at eliminating statistically significant language differences between genders (an absence of bias) but instead sometimes reversed the direction of significant effects.

In some cases, it appears that extending the length of the prompts may have impacted the direction of the gender effect: for agentic language, female names received significantly *less* agentic language on prompts B and C but significantly *more* agentic language on B¹ and C¹ (the longer versions of these prompts). However, in other cases, language content but not length seemed to matter. Prompt C¹—the prompt containing the most communal-focused language—had the effect of reversing a statistically significant bias, such that it was significant in the *opposite* direction on 14 of 28 language categories. Intriguingly, this included a reversal of the observed male bias for agentic, “words to include,” and achievement language categories such that, for prompt C¹, female names received significantly more language from those categories. Further research is needed to understand the prompt contexts that drive gender associations in ChatGPT.

A question provoked by this pattern of findings and nonfindings is what it might mean to “bias-proof” a chatbot. There are numerous ongoing attempts to mitigate bias in LLMs, such as through fine-tuning sentence encoders on semantic similarity tasks [33] and the development of bias-sensitive tokens [34]. The success of these tools in mitigating bias is commonly assessed through word vector associations tests that measure how closely associated specific words or phrases are with respect to sets of attribute words such as “male” versus “female” [35,36] (although other measures of association exist as well [37-39]). However, the foundational challenge to mitigating bias in chatbots may be as conceptual as it is computational; association tests rely on human categories of bias, which means that bias testing may only ever be as good as the a priori assumptions that inform the test. Given that LLM-based chatbots are increasingly being used by the public for the execution of everyday tasks, we suggest that, in addition to bias tests based on measures of association, tests of how biases manifest in real-world use cases (such as those used in this study) should be an additional step that informs progressive generations of AI chatbots. Tests based on simulated real-world tasks do 2 things that association measures cannot: (1) identify how biases manifest in tasks for which chatbots are actually being used and (2) test for bias according to criteria that make a difference in people’s lives, which may not always be anticipated by association tests alone. We hope that this study offers an experimental paradigm that can be adapted and extended by others for further investigations of social bias in chatbots designed for public use.

Limitations

This research has several limitations. First, our studies exclusively used names that have historically been popular in the United States. We relied on lists of names drawn from US census data, which means that many names that are highly popular worldwide cannot be found on this list. This study is also unable to address the question of how ChatGPT responds to either less common names or names that are commonly perceived to be gender neutral. Relatedly, this research used a binary operationalization of gender for its analyses and is unable to address the question of how the results would be impacted by the inclusion of additional information about gender, such as nonbinary gender identity or the use of specific pronouns. Characterizing ChatGPT’s capacity for inclusivity—for example, specifying a reference letter written using they/them pronouns—is an important direction for further research. In addition to examining more gender-inclusive prompts, further research should also investigate other potential forms of name-based social biases that our study was not designed to investigate (eg, racial, ethnic, and nationality biases that could potentially be associated with some names).

Although systematic reviews of biased language in letter writing have predominantly focused on gender differences, it is important to note that studies have also identified biased language in letters written for Black, Indigenous, and people of color when compared with those written for White applicants [40,41]. Specific names themselves may also evoke stereotypes in ways that are distinct from gender effects alone. For example, Zwebner et al [42] recently offered experimental evidence that appearance stereotypes associated with names are powerful enough to lead people to make choices about their facial appearances that correspond to social stereotypes of their names. Further research should investigate name-based stereotypes that may exist in ChatGPT, ideally using international (rather than specific to the United States) lists of names.

The reported research also evaluated only 3 categories of letter-writing prompts. These categories do not represent natural kinds but were created based on stereotypes and the prototypicality of common purposes for which reference letters are requested over the course of academic medical careers. It is unknown whether these results would generalize to other purposes for which reference letters are requested in medicine, such as for specific clinical positions or for recognitions specific to the provision of certain types of clinical care (eg, “cardiologist of the year”). It is also unknown whether these results would generalize to fields outside of academic medicine, which also commonly require recommendation letters (eg, arts foundation fellowships or office promotions), or to other types of content that ChatGPT is used to generate. Furthermore, an important limitation of LIWC, the software used in this study for language analyses, is that it provides count-based results. LIWC is therefore able to account for the frequencies of word choice but not the contexts in which word choices are made. Future research can use alternate approaches to language analyses (eg, topic modeling) [43] to add increased context to the effects observed here.

It is also important to acknowledge that these results reflect ChatGPT's performance during the summer of 2023 using ChatGPT-3.5. The nature and magnitude of social biases within ChatGPT's language model may change with future updates to the language model. Future research should investigate change in biased language effects over time.

Conclusions

This research found evidence of gender bias in recommendation letters written for applicants with historically female versus male names by ChatGPT. Although the nature and direction of biased language depended on the prompt, the overall results point to a practical recommendation: gender-based differences in language are prevalent in ChatGPT-generated reference letters in ways that may negatively impact applicants with stereotypically female names and caution should be taken when using ChatGPT for assistance with this task (and perhaps others that have social consequences). Given the influential role that recommendation letters play in the selection and hiring process

[44], gender bias in letters may negatively impact women's advancement in medicine and other fields that emphasize agentic and leadership qualities, even if they possess these qualities to a similar degree as their male counterparts. This has the potential to contribute to unfairness in the selection and hiring process in these professions, ultimately reinforcing the existing gender disparities in the workforce. As noted earlier, despite a steady increase in the proportion of women in historically male-dominated professions [14], women remain underrepresented in top executive and leadership roles, particularly in medicine and other STEM fields. There are published checklists [27,45] that aim to help letter writers identify and correct for unintended gender-biased language in recommendation letters. At least for the time being, these results suggest that using resources such as these remains an important step in letter writing with the assistance of ChatGPT and any other task in medicine that relies on providing ChatGPT with a person's name.

Acknowledgments

The authors would like to thank Esha Babu, Haden Fulkerson, Sam Chao, Shiyoon Kim, and Tra My Duong for their assistance in conducting the studies reported in this manuscript. The authors also thank Dr Mary-Frances O'Connor and Dr Matthias Mehl, who consulted about the conceptualization of this project.

Data Availability

The data necessary to reproduce the results reported in this manuscript, all study materials, language dictionaries, and the study preregistration can be found on the Open Science Framework [23].

Authors' Contributions

Conceptualization was completed by DMK, RP, SJAA, MNG, NSP, and WAL. Data collection was completed by DMK and NSP. Data management was completed by NSP. Data analysis was completed by RP and DMK. Data visualization was completed by SJAA and RP. Writing of the original draft was completed by DMK, RP, SJAA, MNG, NSP, and CAA. Writing and editing was completed by DMK, RP, SJAA, MNG, NSP, CAA, and WAL. Funding acquisition was completed by WAL.

Conflicts of Interest

None declared.

Multimedia Appendix 1

Additional reported results for studies 1 and 2.
[\[DOCX File, 54 KB-Multimedia Appendix 1\]](#)

Multimedia Appendix 2

Additional reported results for study 3.
[\[DOCX File, 27 KB-Multimedia Appendix 2\]](#)

References

1. Bogost I. The end of recommendation letters. The Atlantic. Apr 20, 2023. URL: <https://www.theatlantic.com/technology/archive/2023/04/chatgpt-ai-college-professors/673796/> [accessed 2023-07-21]
2. I wrote an MBA recommendation letter for my employee using ChatGPT — and they got in!. Medium. Jan 16, 2023. URL: <https://generativeai.pub/i-wrote-an-mba-recommendation-letter-for-my-employee-using-chatgpt-and-they-got-in-2a6a9b948e37> [accessed 2023-07-21]
3. Chaturvedi O. A ChatGPT generated recommendation letter helps a student to get a prestigious scholarship. Techstory. Apr 27, 2023. URL: <https://techstory.in/a-chatgpt-generated-recommendation-letter-helps-a-student-to-get-a-prestigious-scholarship/> [accessed 2023-07-21]

4. Bolukbasi T, Chang KW, Zou J, Saligrama V, Kalai A. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. arXiv. Preprint posted online July 21, 2016. [FREE Full text]
5. Caliskan A, Bryson JJ, Narayanan A. Semantics derived automatically from language corpora contain human-like biases. *Science*. Apr 14, 2017;356(6334):183-186. [FREE Full text] [doi: [10.1126/science.aal4230](https://doi.org/10.1126/science.aal4230)] [Medline: [28408601](https://pubmed.ncbi.nlm.nih.gov/28408601/)]
6. Sap M, Card D, Gabriel S, Choi Y, Smith NA. The risk of racial bias in hate speech detection. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. 2019. Presented at: 57th Annual Meeting of the Association for Computational Linguistics; July 28-August 2, 2019, 2019; Florence, Italy. [doi: [10.18653/v1/p19-1163](https://doi.org/10.18653/v1/p19-1163)]
7. Khan S, Kirubarajan A, Shamsheri T, Clayton A, Mehta G. Gender bias in reference letters for residency and academic medicine: a systematic review. *Postgrad Med J*. May 22, 2023;99(1170):272-278. [doi: [10.1136/postgradmedj-2021-140045](https://doi.org/10.1136/postgradmedj-2021-140045)] [Medline: [37222712](https://pubmed.ncbi.nlm.nih.gov/37222712/)]
8. Trix F, Psenka C. Exploring the color of glass: letters of recommendation for female and male medical faculty. *Discourse Soc*. Jul 25, 2016;14(2):191-220. [doi: [10.1177/0957926503014002277](https://doi.org/10.1177/0957926503014002277)]
9. Turrentine FE, Dreisbach CN, St Ivany AR, Hanks JB, Schroen AT. Influence of gender on surgical residency applicants' recommendation letters. *J Am Coll Surg*. Apr 2019;228(4):356-65.e3. [doi: [10.1016/j.jamcollsurg.2018.12.020](https://doi.org/10.1016/j.jamcollsurg.2018.12.020)] [Medline: [30630084](https://pubmed.ncbi.nlm.nih.gov/30630084/)]
10. Filippou P, Mahajan S, Deal A, Wallen EM, Tan HJ, Pruthi RS, et al. The presence of gender bias in letters of recommendations written for urology residency applicants. *Urology*. Dec 2019;134:56-61. [doi: [10.1016/j.urology.2019.05.065](https://doi.org/10.1016/j.urology.2019.05.065)] [Medline: [31491451](https://pubmed.ncbi.nlm.nih.gov/31491451/)]
11. Hoffman A, Ghoubril R, McCormick M, Matemavi P, Cusick R. Exploring the gender gap: letters of recommendation to pediatric surgery fellowship. *Am J Surg*. Jun 2020;219(6):932-936. [doi: [10.1016/j.amjsurg.2019.08.005](https://doi.org/10.1016/j.amjsurg.2019.08.005)] [Medline: [31416590](https://pubmed.ncbi.nlm.nih.gov/31416590/)]
12. Madera JM, Hebl MR, Dial H, Martin R, Valian V. Raising doubt in letters of recommendation for academia: gender differences and their impact. *J Bus Psychol*. Apr 26, 2018;34(3):287-303. [doi: [10.1007/s10869-018-9541-1](https://doi.org/10.1007/s10869-018-9541-1)]
13. Dutt K, Pfaff DL, Bernstein AF, Dillard JS, Block CJ. Gender differences in recommendation letters for postdoctoral fellowships in geoscience. *Nature Geosci*. Oct 3, 2016;9(11):805-808. [doi: [10.1038/ngeo2819](https://doi.org/10.1038/ngeo2819)]
14. England P, Levine A, Mishel E. Progress toward gender equality in the United States has slowed or stalled. *Proc Natl Acad Sci U S A*. Mar 31, 2020;117(13):6990-6997. [FREE Full text] [doi: [10.1073/pnas.1918891117](https://doi.org/10.1073/pnas.1918891117)] [Medline: [32229559](https://pubmed.ncbi.nlm.nih.gov/32229559/)]
15. Diversity and STEM: women, minorities, and persons with disabilities. National Center for Science and Engineering Statistics (NCSES). 2023. URL: <https://nces.nsf.gov/pubs/nsf23315/> [accessed 2024-01-19]
16. Active physicians by sex and specialty, 2021. Association of American Medical Colleges. URL: <https://www.aamc.org/data-reports/workforce/data/active-physicians-sex-specialty-2021> [accessed 2023-08-13]
17. Eaton AA, Saunders JF, Jacobson RK, West K. How gender and race stereotypes impact the advancement of scholars in STEM: professors' biased evaluations of physics and biology post-doctoral candidates. *Sex Roles*. Jun 3, 2019;82(3-4):127-141. [doi: [10.1007/s11199-019-01052-w](https://doi.org/10.1007/s11199-019-01052-w)]
18. Koch AJ, D'Mello SD, Sackett PR. A meta-analysis of gender stereotypes and bias in experimental simulations of employment decision making. *J Appl Psychol*. Jan 2015;100(1):128-161. [doi: [10.1037/a0036734](https://doi.org/10.1037/a0036734)] [Medline: [24865576](https://pubmed.ncbi.nlm.nih.gov/24865576/)]
19. Moss-Racusin CA, Dovidio JF, Brescoll VL, Graham MJ, Handelsman J. Science faculty's subtle gender biases favor male students. *Proc Natl Acad Sci U S A*. Oct 09, 2012;109(41):16474-16479. [FREE Full text] [doi: [10.1073/pnas.1211286109](https://doi.org/10.1073/pnas.1211286109)] [Medline: [22988126](https://pubmed.ncbi.nlm.nih.gov/22988126/)]
20. Hire and retain a diverse team. Textio. URL: <https://textio.com/> [accessed 2023-08-13]
21. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. Language models are few-shot learners. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020 Presented at: NIPS '20; December 6-12, 2020, 2020;1877-1901; Vancouver, BC. URL: <https://dl.acm.org/doi/abs/10.5555/3495724.3495883>
22. Boyd RL, Ashokkumar A, Seraj S, Pennebaker JW. The development and psychometric properties of LIWC-22. University of Texas at Austin. 2022. URL: <https://www.liwc.app/static/documents/LIWC-22%20Manual%20-%20Development%20and%20Psychometrics.pdf> [accessed 2024-01-19]
23. Kaplan DM, Palitsky R. Investigating evidence of gender bias in letters of recommendation generated by ChatGPT. *Open Science Framework*. Jun 12, 2023. URL: <https://osf.io/g6b8y/> [accessed 2024-01-22]
24. Girgis MY, Qazi S, Patel A, Yu D, Lu X, Sowards J. Gender and racial bias in letters of recommendation for orthopedic surgery residency positions. *J Surg Educ*. Jan 2023;80(1):127-134. [doi: [10.1016/j.jsurg.2022.08.021](https://doi.org/10.1016/j.jsurg.2022.08.021)] [Medline: [36151044](https://pubmed.ncbi.nlm.nih.gov/36151044/)]
25. Schmader T, Whitehead J, Wysocki VH. A linguistic comparison of letters of recommendation for male and female chemistry and biochemistry job applicants. *Sex Roles*. Aug 17, 2007;57(7-8):509-514. [FREE Full text] [doi: [10.1007/s11199-007-9291-4](https://doi.org/10.1007/s11199-007-9291-4)] [Medline: [18953419](https://pubmed.ncbi.nlm.nih.gov/18953419/)]
26. Jordan KN, Sterling J, Pennebaker JW, Boyd RL. Examining long-term trends in politics and culture through language of political leaders and cultural institutions. *Proc Natl Acad Sci U S A*. Feb 26, 2019;116(9):3476-3481. [FREE Full text] [doi: [10.1073/pnas.1811987116](https://doi.org/10.1073/pnas.1811987116)] [Medline: [30808741](https://pubmed.ncbi.nlm.nih.gov/30808741/)]
27. Avoiding gender bias in reference writing. University of Arizona Commission on the Status of Women. URL: https://csw.arizona.edu/sites/default/files/avoiding_gender_bias_in_letter_of_reference_writing.pdf [accessed 2024-01-19]

28. Kacewicz E, Pennebaker JW, Davis M, Jeon M, Graesser AC. Pronoun use reflects standings in social hierarchies. *J Lang Soc Psychol*. Sep 19, 2013;33(2):125-143. [doi: [10.1177/0261927x13502654](https://doi.org/10.1177/0261927x13502654)]
29. Top names over the last 100 years. Social Security Administration. URL: <https://www.ssa.gov/oact/babynames/decades/century.html> [accessed 2024-01-19]
30. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B Stat Methodol*. Dec 05, 2018;57(1):289-300. [doi: [10.1111/j.2517-6161.1995.tb02031.x](https://doi.org/10.1111/j.2517-6161.1995.tb02031.x)]
31. Connor RA, Glick P, Fiske ST. Ambivalent sexism in the twenty-first century. In: Sibley SG, Barlow FK, editors. *The Cambridge Handbook of the Psychology of Prejudice*. Cambridge, MA. Cambridge University Press; 2016.
32. Glick P, Fiske ST. The ambivalent sexism inventory: differentiating hostile and benevolent sexism. *J Pers Soc Psychol*. Mar 1996;70(3):491-512. [doi: [10.1037/0022-3514.70.3.491](https://doi.org/10.1037/0022-3514.70.3.491)]
33. Dolci T. Fine-tuning language models to mitigate gender bias in sentence encoders. 2022. Presented at: IEEE Eighth International Conference on Big Data Computing Service and Applications (BigDataService); August 15-18, 2022, 2022; Newark, CA. [doi: [10.1109/bigdataservice55688.2022.00036](https://doi.org/10.1109/bigdataservice55688.2022.00036)]
34. Liang PP, Wu C, Morency LP, Salakhutdinov R. Towards understanding and mitigating social biases in language models. arXiv. Preprint posted online June 24, 2021.
35. Chaloner K, Maldonado A. Measuring gender bias in word embeddings across domains and discovering new gender bias word categories. 2019. Presented at: First Workshop on Gender Bias in Natural Language Processing; August 2, 2019, 2019; Florence, Italy. [doi: [10.18653/v1/w19-3804](https://doi.org/10.18653/v1/w19-3804)]
36. Cer D, Yang Y, Kong SY, Hua N, Limtiaco N, St. John R, et al. Universal sentence encoder. arXiv. Preprint posted online March 29, 2018.
37. Brunet ME, Alkalay-Houlihan C, Anderson A, Zemel R. Understanding the origins of bias in word embeddings. arXiv. Preprint posted online October 8, 2018. [FREE Full text]
38. Wolfe R, Caliskan A. VAST: the valence-assessing semantics test for contextualizing language models. arXiv. Preprint posted online March 14, 2022. [doi: [10.1609/aaai.v36i10.21400](https://doi.org/10.1609/aaai.v36i10.21400)]
39. Ethayarajh K, Duvenaud D, Hirst G. Understanding undesirable word embedding associations. arXiv. Preprint posted online August 18, 2019. [doi: [10.18653/v1/p19-1166](https://doi.org/10.18653/v1/p19-1166)]
40. Akos P, Kretchmar J. Gender and ethnic bias in letters of recommendation: considerations for school counselors. *Prof School Counsel*. Feb 15, 2018;20(1). [doi: [10.5330/1096-2409-20.1.102](https://doi.org/10.5330/1096-2409-20.1.102)]
41. Grimm LJ, Redmond RA, Campbell JC, Rosette AS. Gender and racial bias in radiology residency letters of recommendation. *J Am Coll Radiol*. Jan 2020;17(1 Pt A):64-71. [doi: [10.1016/j.jacr.2019.08.008](https://doi.org/10.1016/j.jacr.2019.08.008)] [Medline: [31494103](https://pubmed.ncbi.nlm.nih.gov/31494103/)]
42. Zwebner Y, Sellier AL, Rosenfeld N, Goldenberg J, Mayo R. We look like our names: the manifestation of name stereotypes in facial appearance. *J Pers Soc Psychol*. Apr 2017;112(4):527-554. [doi: [10.1037/pspa0000076](https://doi.org/10.1037/pspa0000076)] [Medline: [28240942](https://pubmed.ncbi.nlm.nih.gov/28240942/)]
43. Vayansky I, Kumar SA. A review of topic modeling methods. *Inf Syst*. Dec 2020;94:101582. [doi: [10.1016/j.is.2020.101582](https://doi.org/10.1016/j.is.2020.101582)]
44. Marwan Y, Waly F, Algarni N, Addar A, Saran N, Snell L. The role of letters of recommendation in the selection process of surgical residents in Canada: a national survey of program directors. *J Surg Educ*. Jul 2017;74(4):762-767. [doi: [10.1016/j.jsurg.2017.01.006](https://doi.org/10.1016/j.jsurg.2017.01.006)] [Medline: [28126378](https://pubmed.ncbi.nlm.nih.gov/28126378/)]
45. Avoiding unintended gender bias in letters of recommendation. National Center for Women & Information Technology. 2010. URL: <https://tinyurl.com/2c8eh8v7> [accessed 2024-01-11]

Abbreviations

LIWC: Linguistic Inquiry and Word Count

LLM: large language model

OSF: Open Science Framework

STEM: Science, Technology, Engineering and Mathematics

Edited by T de Azevedo Cardoso, G Eysenbach; submitted 14.08.23; peer-reviewed by A Mihalache, K Wolsiefer; comments to author 02.10.23; revised version received 12.11.23; accepted 27.11.23; published 05.03.24

Please cite as:

Kaplan DM, Palitsky R, Arconada Alvarez SJ, Pozzo NS, Greenleaf MN, Atkinson CA, Lam WA

What's in a Name? Experimental Evidence of Gender Bias in Recommendation Letters Generated by ChatGPT

J Med Internet Res 2024;26:e51837

URL: <https://www.jmir.org/2024/1/e51837>

doi: [10.2196/51837](https://doi.org/10.2196/51837)

PMID: [38441945](https://pubmed.ncbi.nlm.nih.gov/38441945/)

©Deanna M Kaplan, Roman Palitsky, Santiago J Arconada Alvarez, Nicole S Pozzo, Morgan N Greenleaf, Ciara A Atkinson, Wilbur A Lam. Originally published in the Journal of Medical Internet Research (<https://www.jmir.org>), 05.03.2024. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in the Journal of Medical Internet Research, is properly cited. The complete bibliographic information, a link to the original publication on <https://www.jmir.org/>, as well as this copyright and license information must be included.