Portrayal of Medical Students in Artificial Intelligence-Generated Images

Molinna Bui, BS; Paul H. Yi, MD; Ifeanyichukwu Onuh, BS; Ian Kuckelman, MD; Andrew B. Ross, MD, MPH

ABSTRACT

Introduction: The adoption of artificial intelligence (AI) in image generation raises concerns about potential bias, as these technologies may not accurately reflect the populations represented in the images they create. This study examined whether AI-generated images of medical students accurately represent the diversity of the current US medical student population.

Methods: Using the DALL-E (Open AI) image-generation algorithm, we created 300 images with the text prompt "medical student." Two researchers independently analyzed images for demographic indicators, including perceived sex, race/ethnicity, age group, setting, and attire. Descriptive statistics summarized the data, and subgroup analyses assessed differences in portrayals by sex and race/ethnicity. Demographic proportions in the virtual cohort were graphically compared with Association of American Medical Colleges enrollment data.

Results: Of the 300 generated images, 227 (76%) were females and 223 (74%) were White, indicating overrepresentation compared with actual medical school demographics. Black and Latino/Hispanic students were more commonly depicted in scrubs compared to White students, who were often portrayed in white coats or collared shirts ($P = .002$). No images represented Native American/Alaskan Native or Native Hawaiian/Pacific Islander students.

Conclusions: AI-generated images of medical students demonstrated significant demographic disparities, indicating potential bias in these technologies. Such biased portrayals may perpetuate stereotypes and hinder diversity efforts. Future research should identify and address these biases to promote more equitable and inclusive applications of AI tools.

INTRODUCTION

Advancements in artificial intelligence (AI) have revolutionized the generation of realistic images from simple text prompts. These AI applications rely on extensive datasets of preexisting images for training, a practice that could inadvertently propagate bias if the datasets fail to represent the diversity of the general population.[1] AI has had a transformative impact in the health care sector, fostering innovation in diverse fields such as radiology, pathology, gastroenterology, and ophthalmology.[2] However, adoption of AI in these domains has raised concerns about potential bias.[3]

Limited research has examined bias in AI-generated images for specific professions, although general studies found that images of people with lighter skin tones perceived as White are overrepresented.[4,5] A recent study in Australia demonstrated that DALL-E-generated images of medical students favored males with lighter skin tones; however, a more granular analysis of demographics and student depiction has not been conducted.[6]

AI image generation technology is so new that little attention has been given to its applications. With ever-increasing demand for visual content, these algorithms may be used to produce stock images for web content, presentations, and marketing materials, including those in medicine and medical education. Understanding and mitigating biases in AI-generated images is essential to ensure fair and inclusive use across different sectors.[7,8]

In this context, our study aimed to investigate whether AI-generated images of medical students accurately reflect the diversity of the US medical student population. We specifically examined whether

• • •

Author affiliations: University of Wisconsin (UW) School of Medicine and Public Health, Madison, Wisconsin (Bui, Onuh); St. Jude Children's Research Hospital, Department of Radiology Memphis, Tennessee (Yi); Resurrection Medical Center, Transitional Year Residency Program, Chicago, Illinois (Kuckelman); Department of Radiology, UW School of Medicine and Public Health, Madison, Wisconsin (Ross).

Corresponding Author: Andrew Ross; Department of Radiology, University of Wisconsin School of Medicine and Public Health, 600 Highland Ave, Madison, WI 53792; email ARoss@uwhealth.org; ORCID ID 0000-0002-4866-6923

these images displayed evidence of bias towards certain demographic groups. Investigating potential bias in AI-generated images is critical because perceptions of diversity in medicine can significantly influence inclusivity and equity within the profession.[9] Consequently, this study assessed the presence and extent of bias in AI image-generation algorithms by evaluating differences in portrayals of medical students by race/ethnicity and sex.

Figure 1. Representative Images Generated by Dall-E 2 With the Prompt "Medical Student"

Black students were more likely to be portrayed in scrubs than White students who were more frequently portrayed in a white coat ($P < .001$).

METHODS

The study was exempt from institutional review board review as it did not involve human subjects. We generated 300 images of medical students by entering the text prompt, "medical student" into the DALL-E (Open AI) image-generation algorithm between the dates February 1 and March 1, 2023. We selected 300 images to provide reasonable confidence intervals for proportions, ensure adequate cell counts in chi-square subgroup analysis, and maintain feasibility for independent double rating and manual data extraction.

Two researchers independently analyzed each image for demographic indicators. Because artificially generated images of virtual people have no inherent biological sex or racial/ethnic identity, demographic indicators were based on researcher perception. Although "perceived sex" or "perceived race" may be more accurate, for brevity the terms are used without qualification throughout this manuscript.

Demographic variables included sex (male, female, nonbinary); race/ethnicity based on Association of American Medical Colleges (AAMC) categories (White, Black, Latino/Hispanic, Asian, Native American/Alaskan Native, and Native Hawaiian/Pacific Islander); and age group (20-25, 26-30, 31-35, 36-40, and 40+). Additional portrayal variables included attire (collared shirt, scrubs, or white coat) and setting (outdoors, hospital, classroom, or plain background).

We summarized data using descriptive statistics and calculated interobserver variability. To explore differences in how students were portrayed by sex and race, we conducted subgroup analyses using 2-way tabulations and chi-square tests to compare attire and setting across demographic groups.

Finally, we graphically compared the proportion of sex and racial/ethnic groups represented in the virtual cohort with those reported in the AAMC medical school enrollment data for academic years 2019-2020 through 2022-2023.[10] Statistical significance was set at $P < .05$, and all analyses were performed in Stata version 17 (StataCorp LLC).

Figure 2. Comparison of the Proportion of Men and Women in AI-generated Images of Medical Students and Actual Medical School Matriculants From AAMC Data for the Years 2019-2023

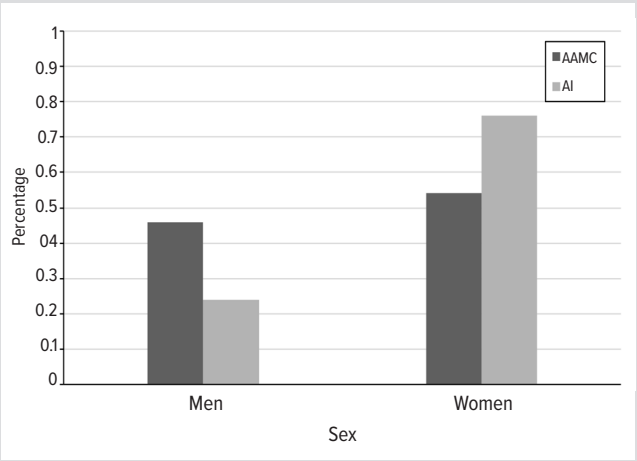

Abbreviations: AAMC, Association of American Medical Colleges; AI, artificial intelligence.

RESULTS

Of the 300 AI-generated images of medical students, 227 (76%) were female and 73 (24%) were male. The most common racial group was White (223 [74%]), followed by Asian (46 [15%]), Black (19 [6%]), and Latino/Hispanic (12 [4%]), There were no images assessed to be representative of Native American/Alaskan Native or Native Hawaiian/Pacific Islander students.

Most students fell into the age range 26-30 years (226 [75%]), followed by age 20-25 (42 [14%]), age 31-35 (28 [9%]), age 36-40 (4 [1%]); none were assessed as older than 40 years.

For setting, most students were depicted against a plain background (163 [54%]), followed by classroom (57 [19%]), hospital (51 [17%]), and outdoors (29 [10%]). Scrubs were the most

common attire (184 [61%]), followed by white coats (105 [35%]) and collared shirts (11 [4%]).

Black and Latino/Hispanic students were more commonly portrayed in scrubs (84% and 83%, respectively) compared with White students (61%; $P = .002$), who were more frequently depicted in white coats or collared shirts (39%) (Figure 1). Outfit and setting did not vary significantly by sex. White students (74%) and females (76%) were overrepresented compared with actual medical school matriculant demographics (53% and 54%, respectively) (Figures 2 and 3).

Interrater agreement was as follows: sex, $\kappa = 0.96$ (excellent); age, $\kappa = 0.28$ (fair); race/ethnicity, $\kappa = 0.83$ (excellent); outfit, $\kappa = 1.0$ (perfect); setting, $\kappa = 0.69$ (substantial).

Figure 3. Comparison of Proportion by Race/Ethnicity in AI-generated Images of Medical Students and Actual Medical School Matriculants From AAMC Data for the Years 2019-2023

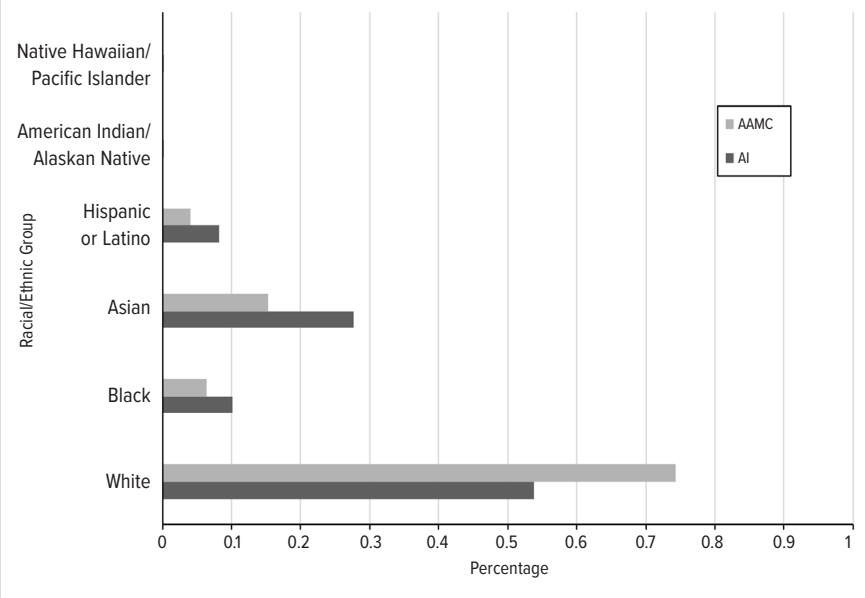

Abbreviations: AAMC, Association of American Medical Colleges; AI, artificial intelligence.

DISCUSSION

Our study revealed significant disparities in the representation of demographic groups in AI-generated images of medical students, highlighting the potential for bias in these technologies. Female and White students were notably overrepresented compared with national medical school enrollment data. The source of this bias likely lies in the diversity – or lack thereof – within the algorithm's training set. As AI-generated images increasingly appear in online content, these discrepancies could perpetuate misrepresentations.

Beyond the demographic imbalances, we observed disparities in portrayal across racial and ethnic groups. White students were more frequently depicted in white coats, whereas Black and Latino students were predominantly shown in scrubs. These differences carry symbolic significance: the formal white coat often represents the physician role and professional authority, while scrubs are more commonly associated with other health care roles. Such portrayals may mirror real-world marginalization experienced by medical students of color, raising concerns about bias influencing perceptions of professional status.[11]

Although these biased portrayals occur in the digital realm, their impact may extend to real world perceptions. AI systems trained on large quantities of online content reflect the bias we encounter daily on the internet. Given the growing demand for online content, AI-generated images are likely to be used in a multitude of applications and platforms, including stock photography, social media, presentations, and marketing materials, including those for medical schools. Unaddressed, these biases could reinforce stereotypes and hinder efforts to promote diversity and inclusivity in medicine.

Our findings align with previous research showing that AI-generated images favor individuals with lighter skin tones perceived as White.[5] This bias may be especially pronounced for high-status roles such as physician or medical student.[12] AI systems trained on existing datasets can perpetuate or even amplify human biases.[13] While our study focused on AI-generated images of medical students, these findings are consistent with a growing body of evidence indicating the potential for bias across multiple AI technologies.[14-16]

Limitations

Our study has several limitations. We used only 1 image-generation application (DALL-E) during a specific time frame, so generalizability of other platforms or later versions is unknown. Given the rapid evolution of AI image generation, these findings are likely to change over time. Future studies should examine newer iterations of DALL-E and compare results to other AI image-generating programs to assess stability of these findings over time.

An additional limitation is that the demographic indicators (sex, race/ethnicity, age) were based on researcher perception. Because virtual humans lack biologic sex or racial identity, this approach was unavoidable. Other investigators have analyzed skin tones or features associated with gender or race quantitatively,[5,12] but it is unclear whether this approach is superior. Ultimately, perception matters because it influences how viewers interpret these images. Interrater agreement was high for most variables, suggesting reliable assessments; agreement was lower for age, although this was not central to our hypothesis.

CONCLUSIONS

This study provides evidence of bias in AI-generated images of medical students primarily related to race. The implications of these findings underscore the importance of addressing bias in AI-generated images used across applications, including medical education. Future research should further investigate the sources of bias to inform strategies that promote equitable and inclusive use of AI tools.

Financial disclosures: None declared.

Funding/support: None declared.

REFERENCES

1. Hartsfield T. How do DALL-E, Midjourney, Stable Diffusion, and other forms of generative AI work? Big Think. September 23, 2022. Accessed October 8, 2022. https://bigthink.com/the-future/dall-e-midjourney-stable-diffusion-models-generative-ai

2. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. *Nat Med*. 2022;28(1):31-38. doi:10.1038/s41591-021-01614-0

3. Belenguer L. AI bias: exploring discriminatory algorithmic decision-making models and the application of possible machine-centric solutions adapted from the pharmaceutical industry. *AI Ethics*. 2022;2(4):771-787. doi:10.1007/s43681-022-00138-8

4. Salminen J, Jung S, Chowdhury S, Jansen BJ. Analyzing demographic bias in artificially generated facial pictures. In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA'20)*. Association for Computing Machinery; 2020:1-8.

5. Luccioni AS, Akiki C, Mitchell M, Jernite Y. Stable bias: analyzing societal representations in diffusion models. arXiv. Preprint posted online March 20, 2023. Last updated November 9, 2023. Accessed February 12, 2025. doi:10.48550/arXiv.2303.11408

6. Currie G, Currie J, Anderson S, Hewis J. Gender bias in generative artificial intelligence text-to-image depiction of medical students. *Health Educ J.* 2024;83(7):732-746. doi:10.1177/00178969241274621

7. Barr A. Google mistakenly tags black people as 'gorillas,' showing limits of algorithms. *Wall Street Journal*. July 1, 2015. Accessed January 15, 2025. https://www.wsj.com/articles/BL-DGB-42522

8. DeCamp M, Lindvall C. Mitigating bias in AI at the point of care. *Science*. 2023;381(6654):150-152. doi:10.1126/science.adh2713

9. Kelly-Blake K, Garrison NA, Fletcher FE, et al. Rationales for expanding minority physician representation in the workforce: a scoping review. *Med Educ.* 2018;52(9):925-935. doi:10.1111/medu.13618

10. American Association of Medical Colleges. Matriculants to U.S. MD-granting medical schools by race, selected combinations of race/ethnicity and gender, 2019-2020 through 2022-2023. Published October 18, 2022. Accessed May 25, 2023. https://www.aamc.org/media/74066/download?attachment

11. Beagan BL. 'Is this worth getting into a big fuss over?' Everyday racism in medical school. *Med Educ*. 2003;37(10):852-860. doi:10.1046/j.1365-2923.2003.01622.x

12. Nicoletti L, Bass D. Humans are biased. Generative AI is even worse. *Bloomberg*. June 9, 2023. Accessed January 15, 2025. https://www.bloomberg.com/graphics/2023-generative-ai-bias/

13. Cho MK. Rising to the challenge of bias in health care AI. *Nat Med.* 2021;27(12):2079-2081. doi:10.1038/s41591-021-01577-2

14. Timmons AC, Duong JB, Simo Fiallo N, et al. A call to action on assessing and mitigating bias in artificial intelligence applications for mental health. *Perspect Psychol Sci.* 2023;18(5):1062-1096. doi:10.1177/17456916221134490

15. Parikh RB, Teeple S, Navathe AS. Addressing bias in artificial intelligence in health care. *JAMA*. 2019;322(24):2377-2378. doi:10.1001/jama.2019.18058

16. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nat Med*. 2021;27(12):2176-2182. doi:10.1038/s41591-021-01595-0

advancing the art & science of medicine in the midwest

WMJ

WMJ (ISSN 2379-3961) is published through a collaboration between The Medical College of Wisconsin and The University of Wisconsin School of Medicine and Public Health. The mission of *WMJ* is to provide an opportunity to publish original research, case reports, review articles, and essays about current medical and public health issues.

Visit www.wmjonline.org to learn more.