Journal of the American Medical Informatics Association, 29(12), 2022, 2178–2181
https://doi.org/10.1093/jamia/ocac156
Advance Access Publication Date: 1 September 2022
Perspective

Perspective

Picture a data scientist: a call to action for increasing diversity, equity, and inclusion in the age of AI

Anne A.H. de Hond[1,2], Marieke M. van Buchem[1,2], and Tina Hernandez-Boussard[2,3,4]

[1]Clinical AI Implementation and Research Lab, Leiden University Medical Center, Leiden, The Netherlands, [2]Department of Medicine (Biomedical Informatics), Stanford University, Stanford, California, USA, [3]Department of Biomedical Data Science, Stanford University, Stanford, California, USA and [4]Department of Epidemiology & Population Health (By Courtesy), Stanford University, Stanford, California, USA

Corresponding Author: Anne A.H. de Hond, Clinical AI Implementation and Research Lab, Leiden University Medical Center, Albinusdreef 2, Leiden 2300 RC, The Netherlands; a.a.h.de_hond@lumc.nl

Received 9 June 2022; Revised 18 August 2022; Editorial Decision 19 August 2022; Accepted 26 August 2022

ABSTRACT

The lack of diversity, equity, and inclusion continues to hamper the artificial intelligence (AI) field and is especially problematic for healthcare applications. In this article, we expand on the need for diversity, equity, and inclusion, specifically focusing on the composition of AI teams. We call to action leaders at all levels to make team inclusivity and diversity the centerpieces of AI development, not the afterthought. These recommendations take into consideration mitigation at several levels, including outreach programs at the local level, diversity statements at the academic level, and regulatory steps at the federal level.

INTRODUCTION

When we ask children to draw a scientist, less than 28% draw a woman[1]. Where this picture may be outdated for professions such as medicine and biology, it remains woefully accurate for data scientists working on artificial intelligence (AI). Globally, women hold less than one-third of data science jobs and this share has seen a small decline since 2018.[2] Data scientists from ethnic and racial minority backgrounds are especially underrepresented. Moreover, authorship for scientific publications within the AI field is unbalanced across gender[3,4] and if current trends hold, gender parity for data science authors will not be reached for another century.[5] In addition, grant funding is negatively biased against consortia with a higher share of female[6–8] and Black principal investigators.[9]

Previous work asserts that a diverse informatics workforce broadens the research agenda and facilitates the development of equity-centered technologies.[10] Here, we build on this argument by focusing on recent research pertaining to diversity, equity, and inclusion in the data science field. This perspective demonstrates how a lack of diversity, equity, and inclusion in the data science profession might hamper the ethical and responsible development of AI and is especially problematic for its application in healthcare. We outline potential solutions and urge leadership to provide consistent and coordinated guidance to increase team diversity.

FAIRNESS IN AI ALGORITHMS

The promise of AI for healthcare is undeniable, but critics of the technology have emphasized the risks posed by biased or unfair AI algorithms exacerbating societal inequalities.[11] Bias in AI algorithms is common and can for example be caused by measurement error, missing data, or underrepresentation.[12,13] Here, we focus on the type of algorithmic bias that can systematically and harmfully disadvantage a particular group by producing discriminatory predictions for gender, race, or other protected identities.[14,15]

Healthcare algorithms are rife with algorithmic bias. Recently, an algorithm widely used by health systems for resource allocation was found to display racial bias.[16] According to this algorithm, Black patients were only provided with similar levels of care when they were sicker than White patients. Many other examples exist, such as an AI underdiagnosing under-served populations such as female patients in chest X-ray pathology classification.[17] Also, models predicting in-hospital mortality for intensive care unit (ICU) patients showed poor model calibration for Black, Hispanic, and Asian patients with respect to the White majority group.[18]

Moreover, the development of modern-day health technologies often lacks the inclusivity needed to service a wide range of people. Take Apple's 2014 Health application: the application was marketed to provide a comprehensive health check, but its team overlooked the charting of women's menstrual cycles.[19] Forgetting one of the oldest forms of self-tracking, the application illustrates a blind spot in women's needs. As another example, inclusive treatment may be lacking for certain medical procedures, such as pulse oximetry and mechanical ventilation. Pulse oximetry measurements were found to systematically overestimate oxygen levels for darker-skinned patients[20,21] and Black ICU patients were found to be less likely to receive ventilation treatment and more likely to receive a shorter treatment duration.[22]

DIVERSITY, EQUITY, AND INCLUSION FOR AI DEVELOPMENT

Clearly algorithmic bias, non-inclusive design, and biased medical procedures can eventually lead to inequitable outcomes across subgroups and societal harm. Despite its tremendous potential for all populations, AI applications can and will exacerbate inequalities when algorithmic bias is left unchecked. Increasing the diversity of AI development teams is one mitigation step that could help to counteract algorithmic bias. Some preliminary work examines how team diversity may contribute to mitigating bias throughout the AI lifecycle. First, including a diverse perspective from designers, coders, health practitioners, and end-users may lead to products that better serve the needs of their respective communities.[23–25] For example, diverse team members may aid in better anticipating the likely impacts of certain model choices on different subgroups and modes of failure.[25] Second, a recent study found that AI developers from the same demographic group were more likely to develop models with the same prediction errors compared to AI developers outside of that group.[26] Hence, composing a diverse development team from the get-go may assist in addressing algorithmic bias by averaging out these prediction errors across developer subgroups.[24] Finally, team diversity can aid in broadening the actual questions being addressed by AI, for example developing a model predicting appointment no-shows versus a model predicting barriers to appointments, such as timing of appointments. This is not to say that team diversity is a panacea nor that women can only design products for women, Black people for Black people, etc. On the contrary: team diversity is an important option in the assortment of mitigation strategies available, such as improving the representativeness of the AI training data[25] and educating developers on fair AI practices.[26] A combination of these strategies is needed to effectively combat algorithmic bias.

WHAT DRIVES THE LACK OF DIVERSITY, EQUITY, AND INCLUSION?

Several causes may contribute to the lack of diversity in AI development teams, such as the lack of role models and the "leaky pipeline," a metaphor describing how marginalized groups progressively leave science, technology, engineering, and mathematics (STEM) subjects in the period between preschool and college. Work environments that lack a diversity focus may be another contributor.[27] The latter may be especially problematic when good professional traits recognized in the majority group are valued differently in minority groups. For example, ambition is valued more negatively in female academics than in male academics.[28] Moreover, professionals with minority backgrounds may be further disadvantaged by what is referred to as a "minority tax": overtime spent on diversity initiatives that come at the expense of other activities more directly beneficial to one's career.[29] In academia, this may consist of time spent on committees to meet diversity quotas rather than on promoting their career through research and teaching.

A CALL TO ACTION

We encourage leadership at all levels to provide consistent and coordinated guidance to increase team diversity. Various mitigation steps exist in the literature. At the community and local governance level, encouraging involvement in outreach programs could inspire members from minority backgrounds to pursue STEM subjects. For example, the American Medical Informatics Association's (AMIA) First Look program introduces women to informatics and provides mentoring and career resources.[10] For the academic sector, incentive structures promoting diversity could be tied to the hiring and retention of diverse faculty, journal editors, and senior leadership, especially where AI research is conducted.[30–32] A great example is the Diversity, Equity, and Inclusion (DEI) Task Force that advises AMIA on these matters.[33] When evaluating and allocating grants, committees can require diversity of the investigative team, including diversity of senior investigators and key personnel.[9] Academic success and impact could be valued beyond citations to also encompass mentoring, teaching, and well-being.[34] In industry, raising awareness about biases via seminars or programs, helping employees to better understand their biases, and targeted hiring practices to increase diversity may aid to address the problem.[26,31,35,36] At the federal level, regulatory oversight could be expanded to include evaluations on the performance of AI systems across populations, ensuring the reliability of AI tools across underrepresented populations.[37,38] In tandem with regulatory oversight, monitoring, and auditing are needed at all levels to secure the fair and inclusive use of AI. At the level of the manufacturer, surveillance systems and vigilance in terms of incident reports and safety warnings that include societal or population harm should be employed to provide post-market monitoring of algorithmic bias.[39,40] Moreover, the care organization where the AI is to be implemented is advised to draft a monitoring plan, ensuring that their target population is represented in the training data. In this plan aspects such as the monitoring of the expected and unexpected effects of the AI on clinical practice may be described.[40]

CONCLUSION

The problems associated with biased or unfair AI go beyond mere technological challenges. It is a problem ingrained in the inequities of our societies and requires structural change at the organizational level to affect change at the technological level. Diverse and inclusive AI teams form an important mitigation strategy toward achieving equitable and fair AI development. Team inclusivity and diversity should therefore become the centerpieces of AI development, not the afterthought. This will benefit not only the AI technology itself, but the entire society that may one day be reliant on it.

FUNDING

Research reported in this publication was supported by the National Library of Medicine of the National Institutes of Health under Award Number R01LM013362. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

AUTHOR CONTRIBUTIONS

AAHdH, MMvB, and TH-B conceived the idea, wrote the initial draft, edited, and approved the final manuscript.

ACKNOWLEDGMENTS

We would like to thank Dr. Fatima Rodriguez and Dr. Ilse Kant for their valuable input on this manuscript.

CONFLICT OF INTEREST STATEMENT

None declared.

DATA AVAILABILITY

No new data were generated or analyzed in support of this research.

REFERENCES

1. Miller DI, Nolla KM, Eagly AH, Uttal DH. The development of children's gender-science stereotypes: a meta-analysis of 5 decades of U.S. draw-a-scientist studies. *Child Dev* 2018; 89 (6): 1943–55.
2. World Economic Forum. Global Gender Gap Report 2021. 2021. https://www3.weforum.org/docs/WEF_GGGR_2021.pdf. Accessed July 15, 2022.
3. Celi LA, Cellini J, Charpignon M-L, *et al.*; for MIT Critical Data. Sources of bias in artificial intelligence that perpetuate healthcare disparities—a global review. *PLoS Digit Health* 2022; 1 (3): e0000022.
4. Holman L, Stuart-Fox D, Hauser CE. The gender gap in science: how long until women are equally represented? *PLoS Biol* 2018; 16 (4): e2004956.
5. Wang LL, Stanovsky G, Weihs L, Etzioni O. Gender trends in computer science authorship. *Commun ACM* 2021; 64 (3): 78–84.
6. Bianchini S, Llerena P, Öcalan-Özel S, Özel E. Gender diversity of research consortia contributes to funding decisions in a multi-stage grant peer-review process. *Humanit Soc Sci Commun* 2022; 9 (1): 195.
7. Safdar B, Naveed S, Chaudhary AMD, Saboor S, Zeshan M, Khosa F. Gender disparity in grants and awards at the National Institute of Health. *Cureus* 2021; 13 (4): e14644.
8. Witteman HO, Hendricks M, Straus S, Tannenbaum C. Female grant applicants are equally successful when peer reviewers assess the science, but not when they assess the scientist. *bioRxiv* 2018: 232868. doi: 10.1101/232868.
9. Taffe MA, Gilpin NW. Racial inequity in grant funding from the US National Institutes of Health. *eLife* 2021; 10: e65697.
10. Bright TJ, Williams KS, Rajamani S, *et al.* Making the case for workforce diversity in biomedical informatics to help achieve equity-centered care: a look at the AMIA First Look Program. *J Am Med Inform Assoc* 2021; 29 (1): 171–5.
11. Zou J, Schiebinger L. AI can be sexist and racist—it's time to make it fair. *Nature* 2018; 559 (7714): 324–6.
12. Roselli D, Matthews, J, Talagala, N. Managing bias in AI. In: Companion Proceedings of The 2019 World Wide Web Conference. 2019. doi: 10.1145/3308560.3317590.
13. Parikh RB, Teeple S, Navathe AS. Addressing bias in artificial intelligence in health care. *JAMA* 2019; 322 (24): 2377–8.
14. Pfohl SR, Foryciarz A, Shah NH. An empirical characterization of fair machine learning for clinical risk prediction. *J Biomed Inform* 2021; 113: 103621.
15. McCradden MD, Joshi S, Mazwi M, Anderson JA. Ethical limitations of algorithmic fairness solutions in health care machine learning. *Lancet Digit Health* 2020; 2 (5): e221–3.
16. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 2019; 366 (6464): 447–53.
17. 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): 2176–82.
18. Sarkar R, Martin C, Mattie H, Gichoya JW, Stone DJ, Celi LA. Performance of intensive care unit severity scoring systems across different ethnicities in the USA: a retrospective observational study. *Lancet Digit Health* 2021; 3 (4): e241–9.
19. Eveleth R. How self-tracking apps exclude women. *The Atlantic*. 2014.
20. Sjoding MW, Dickson RP, Iwashyna TJ, Gay SE, Valley TS. Racial bias in pulse oximetry measurement. *N Engl J Med* 2020; 383 (25): 2477–8.
21. Wong A-KI, Charpignon M, Kim H, *et al.* Analysis of discrepancies between pulse oximetry and arterial oxygen saturation measurements by race and ethnicity and association with organ dysfunction and mortality. *JAMA Netw Open* 2021; 4 (11): e2131674.
22. Meng C, Trinh L, Xu N, Enouen J, Liu Y. Interpretability and fairness evaluation of deep learning models on MIMIC-IV dataset. *Sci Rep* 2022; 12 (1): 7166.
23. de Hond AAH, Leeuwenberg AM, Hooft L, *et al.* Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: a scoping review. *NPJ Digit Med* 2022; 5 (1): 2.
24. Ng M, Kapur, K.D. Blizinsky S, Hernandez-Boussard T. The AI life cycle: a holistic approach to creating ethical AI for health decisions. *Nat Med*. in press.
25. Fazelpour S, De-Arteaga M. Diversity in sociotechnical machine learning systems. *Big Data Soci* 2022; 9 (1): 205395172210820.
26. Cowgill B, Dell'Acqua F, Deng S, Hsu D, Verma N, Chaintreau A. Biased programmers? Or biased data? A field experiment in operationalizing AI ethics. In: Proceedings of the 21st ACM Conference on Economics and Computation. 2020. doi: 10.2139/ssrn.3615404.
27. The Lancet Digital H. All things being equal: diversity in STEM. *Lancet Digit Health* 2020; 2 (4): e149.
28. Chubb J, Derrick GE. The impact a-gender: gendered orientations towards research Impact and its evaluation. *Palgrave Commun* 2020; 6 (1): 72.
29. Williamson T, Goodwin CR, Ubel PA. Minority tax reform—avoiding overtaxing minorities when we need them most. *N Engl J Med* 2021; 384 (20): 1877–9.
30. Editorial. Nature's under-representation of women. *Nature* 2018; 558 (7710): 344.
31. West SM, Whittaker, M, Crawford, K. *Discriminating Systems: Gender, Race, and Power in AI*. AI Now Institute; 2019. https://ainowinstitute.org/discriminatingsystems.pdf. Accessed July 15, 2022.
32. Greider Carol W, Sheltzer Jason M, Cantalupo Nancy C, *et al.* Increasing gender diversity in the STEM research workforce. *Science* 2019; 366 (6466): 692–5.
33. Bakken S. Toward diversity, equity, and inclusion in informatics, health care, and society. *J Am Med Inform Assoc* 2020; 27 (11): 1639–40.
34. Davies SW, Putnam HM, Ainsworth T, *et al.* Promoting inclusive metrics of success and impact to dismantle a discriminatory reward system in science. *PLoS Biol* 2021; 19 (6): e3001282.

35. Roper Rachel L. Does gender bias still affect women in science? *Microbiol Mol Biol Rev* 2019; 83 (3): e00018–19.
36. Régner I, Thinus-Blanc C, Netter A, Schmader T, Huguet P. Committees with implicit biases promote fewer women when they do not believe gender bias exists. *Nat Hum Behav* 2019; 3 (11): 1171–9.
37. Lee NT, Lai, S. *The U.S. Can Improve Its AI Governance Strategy by Addressing Online Biases*. Brookings; 2022. https://www.brookings.edu/blog/techtank/2022/05/17/the-u-s-can-improve-its-ai-governance-strategy-by-addressing-online-biases/. Accessed July 15, 2022.
38. Lander E, Nelson A. ICYMI: WIRED (Opinion): Americanc Need a Bill of Rights for an AI-powered WorldThe White House. 2021. https://www.whitehouse.gov/ostp/news-updates/2021/10/22/icymi-wired-opinion-americans-need-a-bill-of-rights-for-an-ai-powered-world/. Accessed July 15, 2022.
39. Information Commissioner's Office. Guidance on the AI Auditing Framework: Draft Guidance for Consultation. 2020. https://ico.org.uk/media/about-the-ico/consultations/2617219/guidance-on-the-ai-auditing-framework-draft-for-consultation.pdf. Accessed July 15, 2022.
40. van Smeden M, Moons C, Hooft L, Kant I, van Os H, Chavannes N. Guideline for High-Quality Diagnostic and Prognostic Applications of AI in Healthcare. 2021. https://www.datavoorgezondheid.nl/wegwijzer-ai-in-de-zorg/documenten/publicaties/2021/12/17/guideline-for-high-quality-diagnostic-and-prognostic-applications-of-ai-in-healthcare. Accessed July 15, 2022.