

ALEXANDER WILLIAMS TOLBERT

Assistant Professor of Data and Decision Sciences, Emory University

*Secondary Appointments: Philosophy, Computer Science, African American Studies
Faculty Fellow, Emory Center for Ethics; Affiliate, Center for Mind, Brain, and Culture*

alexander.tolbert@emory.edu | +1 (470) 547-2185

EDUCATION

- **PhD in Philosophy**, University of Pennsylvania, 2019–2023
Advisors: Michael Kearns, Quayshawn Spencer, Anita Allen
Committee: Samuel Freeman, Aaron Roth, Scott Weinstein
- **M.A. in Statistics**, Wharton School, University of Pennsylvania, 2019–2023
- **M.S. in Biochemistry**, Virginia Tech, 2017–2019
- **M.A. in Philosophy**, Virginia Tech, 2017–2019
- **B.S. in Biology**, University of Mobile, 2009–2013

ACADEMIC APPOINTMENTS

- **Assistant Professor** (2024–present), Department of Data and Decision Sciences, Emory University
- **Postdoctoral Researcher** (2023–2024), Emory University
- **Visiting Graduate Scholar** (2020), CUNY Graduate Center (Supervisor: Charles Mills)

AREAS OF SPECIALIZATION

Philosophy of Science, Machine Learning, Causal Inference, Philosophy, Politics, and Economics (PPE)

BOOKS

- Alexander W. Tolbert (Ed.). *The Breaking Point: Philosophy, Politics, and Economics in an Age of Power, Peril, and Possibility*. New York: Springer, under contract (forthcoming 2027). Edited with prelude, introduction, six thematic unit essays, conclusion, and epilogue by Alexander W. Tolbert, with additional submission of chapter(s).

PUBLICATIONS

PEER-REVIEWED ARTICLES

Tolbert, A. W. (forthcoming). Why AI ethics is not about AI. *Philosophical Studies*.

- Blum, A., Diana, E., Ravichandran, K., & Tolbert, A. W. (2026). A theoretical model for grit in pursuing ambitious ends [Forthcoming]. *Proceedings of the AAAI Conference on Artificial Intelligence*. <https://doi.org/https://doi.org/10.48550/arXiv.2503.02952>
- Tolbert, A. W. (2025). The racial algorithm: Reassessing scientific legitimacy in a value-laden world [Forthcoming]. *Critical Philosophy of Race*. <https://philpapers.org/rec/TOLTRA-2>
- Tolbert, A. W., & Smith, B. (2025). The problem of generics in LLM training. *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 1275–1280. <https://doi.org/10.1145/3715275.3732085>
- Behzad, T., Casacuberta, S., Diana, E. R., & Tolbert, A. W. (2025). Reconciling predictive multiplicity in practice. *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 3350–3369. <https://doi.org/10.1145/3715275.3732216>
- Shukla, P., Chinchure, A., Diana, E., Tolbert, A. W., Hosanagar, K., Balasubramanian, V. N., Sigal, L., & Turk, M. A. (2025). Mitigate one, skew another? tackling intersectional biases in text-to-image models [Forthcoming, accepted to Findings of ACL at EMNLP 2025]. *Findings of the Association for Computational Linguistics: EMNLP 2025*. <https://doi.org/10.48550/arXiv.2505.17280>
- Tolbert, A. W. (2025). An epistemic argument for reparations: A solution to the problem of social stratification in causal modeling [Published online 18 April 2025]. *Synthese*, 205(179). <https://doi.org/10.1007/s11229-024-04901-8>
- Tolbert, A. W. (2025). Restricted racial realism: Heterogeneous effects and the instability of race in social science. *Philosophy of the Social Sciences*, 55(2), 146–164. <https://doi.org/10.1177/00483931241299884>
- Tolbert, A. W. (2025). Why causal inference is necessary for algorithmic fairness [Published online 10 September 2025]. *Synthese*, 206(162). <https://doi.org/10.1007/s11229-025-05044-0>
- Roth, A., & Tolbert, A. W. (2025). Resolving the reference class problem at scale. *Philosophy of Science*, 1–15. <https://doi.org/10.1017/psa.2025.17>
- Tolbert, A. W. (2024). Causal agnosticism about race: Variable selection problems in causal inference. *Philosophy of Science*, 91(5), 1098–1108. <https://doi.org/10.1017/psa.2023.166>
- Tolbert, A. W. (2025). Philosophy of race and causation ii: Potential replies [Forthcoming]. *Philosophy Compass*.
- Tolbert, A. W. (2025). Philosophy of race and causation i: Challenges to race as a causal variable [Forthcoming]. *Philosophy Compass*.
- Diana, E., Tolbert, A. W., Ravichandran, K., & Blum, A. (2025). Pessimism traps and algorithmic interventions. *6th Symposium on Foundations of Responsible Computing (FORC 2025)*, 329, 5:1–5:19. <https://doi.org/10.4230/LIPICs.FORC.2025.5>
- Tolbert, A. W. (2025). Correcting underrepresentation and intersectional bias for classification [Forthcoming]. *American Philosophical Quarterly*. <https://philpapers.org/rec/TOLCUA>
- Roth, A., Tolbert, A. W., & Weinstein, S. (2023). Reconciling individual probability forecasts. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 101–110. <https://doi.org/10.1145/3593013.3593980>

BOOK CHAPTERS

- Tolbert, A. W. (2027). Race as a concept and as a cause [Invited contribution, Forthcoming]. In S.-O. Hansson & T. Grüne-Yanoff (Eds.), *Comprehensive philosophy of science*. Elsevier.
- Tolbert, A. W. (2026). Contemporary issues in the philosophy of universal basic income: From AI to reparations and beyond [Forthcoming]. In L. Liu & R. K. Caputo (Eds.), *The politics of basic income in the public and private sector: International experiences and initiatives*. Palgrave Macmillan.
- Tolbert, A. W. (2026). Algorithmic abolitionism and the racial algorithm [Invited contribution, Forthcoming]. In K. C. Elliott & T. Richards (Eds.), *The routledge handbook of values and science*. Routledge. <https://philpapers.org/rec/TOLAAA-6>

REVISE AND RESUBMIT

- Tolbert, A. W. (2025). Race as a statistical regime [Revise and resubmit]. *Journal of Causal Inference*.
- Tolbert, A. W. (2025). A capabilities approach for causal fairness [Revise and resubmit]. *Observational Studies*.

PREPRINTS

- Bouyamoun, A., & Tolbert, A. W. (2025). Escaping the subprime trap in algorithmic lending [Preprint]. *arXiv preprint arXiv:2502.17816*. <https://doi.org/https://doi.org/10.48550/arXiv.2502.17816>
- Huang, J., Lu, J., & Tolbert, A. W. (2025). Causal feature learning in the social sciences [Preprint]. *arXiv preprint arXiv:2503.12784*. <https://doi.org/https://doi.org/10.48550/arXiv.2503.12784>

SELECTED TALKS AND PRESENTATIONS

CONFERENCES

- *Pacific APA, 2025* – Pessimism Traps and a Theoretical Model of Grit
- *SWEET Conference, Florida, 2025* – Theoretical Model of Grit
- *PPE Colloquium, Buffalo, 2024* – Causal Agnosticism About Race
- *PhilML'24, Tübingen, 2024* – Causal Agnosticism About Race
- *PSA Biennial Meeting, 2022* – Causal Race Agnosticism
- *PSA Biennial Meeting, 2021* – Against Predictive Invariance
- *APA Pacific Division, 2021* – Causation and Algorithmic Fairness
- *APA Eastern Division, 2021* – Causation and Algorithmic Fairness
- *BIAS Conference, 2021* – Against Predictive Invariance

COLLOQUIA AND INVITED TALKS

- *AI Symposium, Morgan State, 2025* – Pessimism Traps and a Theoretical Model of Grit
- *Tepper OR Colloquium, CMU, 2024* – Causal Agnosticism About Race
- *First Paris Conference on Frontiers of Economics and Philosophy, 2024* – Algorithms, Justice, and the Urban Ghetto
- *TTIC Colloquium, 2024* – Causal Agnosticism About Race
- *Philosophy of AI (PhAI), 2023* – Algorithms, Justice, and the Urban Ghetto
- *TOC4Fairness Seminar, Simons Foundation, 2022* – Reference Class Problems, Model Multiplicity
- *Penn Law Symposium on Innovation and Criminal Justice, 2022* – Algorithms in Criminal Justice (Panel)
- *Ernest E. Just Biomedical Society, Penn, 2020* – Causation and Algorithmic Fairness
- *Philosophy and Physical Computing, 2019* – Dynamics and Biochemical Kinds
- *Virginia Tech Biochemistry Speakers Series, 2019* – Inositol Kinase Function
- *Philosophy and Physical Computing, 2018* – Finding Natural Kinds

PROFESSIONAL SERVICE

- Guest Editor, *American Philosophical Quarterly*, Special Issue on AI (2025)
- Co-Organizer, AI, Systems & Society Conference, Emory University (2025)
- Founder, Black Philosophy of AI Roundtable, Emory University (2024)

UNIVERSITY AND GRADUATE SERVICE

- Organizer, AI Systems and Society Reading Group, Emory QTM (2023–2024)
- Founder, W. E. B. Du Bois & Ida B. Wells Symposium on Socially Aware Data Science (2022–)
- Founder, Penn Philosophy of Computation and Data Workshop (2021–)
- Organizer, Penn-Rutgers Philosophy of Race Reading Group (2020–)
- Africana Summer Institute Fellow, University of Pennsylvania (2020–2021)
- Co-President, Minorities and Philosophy (MAP), Virginia Tech (2018–2019)
- Editorial Assistant, *HOPOS* (2018–2019)

PROFESSIONAL EXPERIENCE

- **Amazon Web Services (Remote)** – Research Scientist Intern Summer 2021
Project: Algorithmic Framework for Bias
Supervisor: Peter Hallinan | Language: Python
- **Amazon Web Services (Remote)** – Research Scientist II Intern Summer–Fall 2022
Project: Algorithmic Harm and Bias Database
Supervisor: Peter Hallinan | Language: Python
- **Wharton Sports Analytics and Business Initiative, UPenn** – Instructor Summer 2022
Moneyball Academy: teaching data science and coding (R) to underrepresented students

AWARDS AND GRANTS

- Fontaine Fellow, University of Pennsylvania (2019–2023)
- Graduate Scholar, MAOP Fellowship, Virginia Tech (2017–2019)

TECHNICAL SKILLS

Proficient: R, Python

Familiar: SQL, Java

Tools: LaTeX, Git

REFERENCES

Available upon request.