

Manish Raghavan

mragh@mit.edu

<https://mraghavan.github.io/>

EMPLOYMENT

Drew Houston (2005) Career Development Professor August 2022–Present
Massachusetts Institute of Technology, Cambridge, MA
Sloan School of Management and Department of Electrical Engineering and Computer Science

Postdoctoral Fellow August 2021–August 2022
Harvard Center for Research on Computation and Society, Cambridge, MA

Visiting Researcher September 2019–August 2021
Facebook, New York, NY

EDUCATION

Ph.D., Computer Science
Cornell University

Thesis: *The Societal Impacts of Algorithmic Decision-Making*. Advised by Jon Kleinberg.

M.S., Computer Science
Cornell University

B.S., Electrical Engineering and Computer Science
University of California, Berkeley

PUBLICATIONS

Conferences

- N. Dalvi, M. Olteanu, M. Raghavan, and P. Bohannon. *Deduplicating a Places Database*. In *Proceedings of the 23rd International World Wide Web Conference*. April 2014
- J. Kleinberg, S. Oren, and M. Raghavan. *Planning Problems for Sophisticated Agents with Present Bias*. In *Proceedings of the 17th ACM Conference on Economics and Computation*. July 2016
- J. Kleinberg, S. Mullainathan, and M. Raghavan. *Inherent Trade-offs in the Fair Determination of Risk Scores*. In *The 8th Innovations in Theoretical Computer Science Conference*. January 2017
- J. Kleinberg, S. Oren, and M. Raghavan. *Planning with Multiple Biases*. In *Proceedings of the 18th ACM Conference on Economics and Computation*. June 2017
- G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, K. Weinberger. *On Fairness and Calibration*. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems*. December 2017
- J. Kleinberg, M. Raghavan. *Selection Problems in the Presence of Implicit Bias*. In *The 9th Innovations in Theoretical Computer Science Conference*. January 2018
- M. Raghavan, A. Anderson, J. Kleinberg. *Mapping the Invocation Structure of Online Political Interaction*. In *Proceedings of the 27th International World Wide Web Conference*. April 2018
- M. Raghavan, A. Slivkins, J. W. Vaughan, Z. S. Wu. *The Externalities of Exploration and How Data Diversity Helps Exploitation*. In *Conference on Learning Theory*. July 2018
- M. Raghavan, M. Purohit, S. Gollapudi. *Hiring Under Uncertainty*. In *International Conference on Machine Learning*. June 2019
- R. Abebe, S. Barocas, J. Kleinberg, K. Levy, M. Raghavan, D. G. Robinson. *Roles for Computing in Social Change*. In *Proceedings of the Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020

- S. Barocas, A. D. Selbst, M. Raghavan. *The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons*. In *Proceedings of the Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020
- M. Raghavan, S. Barocas, J. Kleinberg, K. Levy. *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*. In *Proceedings of the Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020
- J. Finocchiaro, R. Maio, F. Monachou, G. K. Patro, M. Raghavan, A.-A. Stoica, and S. Tsirtis. *Bridging Machine Learning and Mechanism Design towards Algorithmic Fairness*. In *Proceedings of the Fourth ACM Conference on Fairness, Accountability, and Transparency*. January 2021
- J. Kleinberg, S. Oren, M. Raghavan, and N. Sklar. *Stochastic Model for Sunk Cost Bias*. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence*. July 2021
- E. Black, M. Raghavan, and S. Barocas. *Model Multiplicity: Opportunities, Concerns, and Solutions*. In *Proceedings of the Fifth ACM Conference on Fairness, Accountability, and Transparency*. June 2022
- C. Hays, Z. Schutzman, M. Raghavan, E. Walk, P. Zimmer. *Simplistic Collection and Labeling Practices Limit the Utility of Benchmark Datasets for Twitter Bot Detection*. In *The Web Conference*. April 2023
- R. Alur, L. Laine, D. Li, M. Raghavan, D. Shah, D. Shung. *Auditing for Human Expertise*. In *Proceedings of the 37th Annual Conference on Neural Information Processing Systems*. December 2023
- C. Dwork, C. Hays, J. Kleinberg, M. Raghavan. *Content Moderation and the Formation of Online Communities: A Theoretical Framework*. In *The Web Conference*. May 2024
- K. Peng, M. Raghavan, E. Pierson, J. Kleinberg, N. Garg. *Reconciling the Accuracy-Diversity Trade-Off in Recommendations*. In *The Web Conference*. May 2024
- C. Dwork, C. Hays, J. Kleinberg, M. Raghavan. *Equilibria, Efficiency, and Inequality in Network Formation for Hiring and Opportunity*. In *Proceedings of the 25th ACM Conference on Economics and Computation*. July 2024
- R. Alur, M. Raghavan, D. Shah. *Human Expertise in Algorithmic Prediction*. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems*. December 2024
- B. Laufer, M. Raghavan, S. Barocas. *What Constitutes a Less Discriminatory Algorithm?*. In *ACM Symposium on Computer Science and Law*. March 2025
- Maria Ballesteros, Cynthia Dwork, Gary King, Conlan Olson, Manish Raghavan. *Evaluating the Impacts of Swapping on the US Decennial Census*. In *ACM Symposium on Computer Science and Law*. March 2025
- C. Hays, M. Raghavan. *Double Machine Learning for Causal Inference under Shared-State Interference*. In *International Conference on Machine Learning*. July 2025
- N. Jo, K. Creel, A. Wilson, M. Raghavan. *Homogeneous Algorithms Can Reduce Competition in Personalized Pricing*. In *Neural Information Processing Systems*. December 2025
- D. Shanmugam, S. Sadhuka, M. Raghavan, J. Gutttag, B. Berger, E. Pierson. *Evaluating multiple models using labeled and unlabeled data*. In *Neural Information Processing Systems*. December 2025
- C. Hays, B. Laufer, S. Barocas, M. Raghavan. *Statistical Guarantees in the Search for Less Discriminatory Algorithms*. In *International Conference on Learning Representations*. April 2026
- C. Park, K. Donahue, M. Raghavan. *When to Ask a Question: Understanding Communication Strategies in Generative AI Tools*. In *Foundations of Responsible Computing*. June 2026
- C. Hays, R. Li, B. Flanigan, M. Raghavan. *Strategic Candidacy in Generative AI Arenas*. In *International Conference on Machine Learning*. July 2026

Journals

- J. Kleinberg, M. Raghavan. *How Do Classifiers Induce Agents To Invest Effort Strategically?*. In *ACM Transactions on Economics and Computation*. October 2020
- Also appeared as:
- *How Do Classifiers Induce Agents To Invest Effort Strategically?*. In *Proceedings of the 20th ACM Conference on Economics and Computation*. June 2019
 - *Designing Evaluation Rules that are Robust to Strategic Behavior*. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, Sister Conference Track*. April 2020
 - *Algorithmic Classification and Strategic Effort*. In *SIGecom Exchanges*. November 2020
 - J. Kleinberg and M. Raghavan. *Algorithmic Monoculture and Social Welfare*. In *Proceedings of the*

National Academy of Sciences 118(22). June 2021

- M. Raghavan, A. Slivkins, J. W. Vaughan, Z. S. Wu. *Greedy Algorithm Almost Dominates in Smoothed Contextual Bandits*. In *SIAM Journal on Computing*. April 2023
- J. Kleinberg, S. Mullainathan, and M. Raghavan. *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*. In *Management Science*. November 2023

Also appeared as:

The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. June 2022

- C. K. Morewedge, S. Mullainathan, H. F. Naushan, C. R. Sunstein, J. Kleinberg, M. Raghavan, and J. O. Ludwig. *Human Bias in Algorithm Design*. In *Nature Human Behavior*. November 2023
- J. Kleinberg, J. Ludwig, S. Mullainathan, and M. Raghavan. *The Inversion Problem: Why Algorithms Should Infer Mental State and Not Just Predict Behavior*. In *Perspectives on Psychological Science*. December 2023
- M. Raghavan and P. T. Kim. *Limitations of the “Four-Fifths Rule” and Statistical Parity Tests for Measuring Fairness*. In *Georgetown Law Technology Review*. January 2024
- M. Raghavan. *What Should We Do when Our Ideas of Fairness Conflict?*. In *Communications of the ACM*. January 2024
- C. Dwork, K. Greenewald, M. Raghavan. *Synthetic Census Data Generation via Multidimensional Multiset Sum*. In *Journal of Privacy and Confidentiality*. August 2025

Also appeared as:

Synthetic Census Data Generation via Multidimensional Multiset Sum. In *Foundations of Responsible Computing*. June 2024

Other

- M. Olteanu, N. Dalvi, and M. Raghavan. *Identifying descriptive terms associated with a physical location from a location store*. In *U.S. Patent No. 9613054*. April 2017
- M. Raghavan, S. Barocas. *Challenges for mitigating bias in algorithmic hiring*. In *The Brookings Institution*. December 2019

TEACHING

6.3950/6.3952: Artificial Intelligence, Decision-Making, and Society. *MIT*

Fall 2022 (with Aleksander Mądry and Ashia Wilson); Fall 2023 (with Ashia Wilson)

15.563: Artificial Intelligence for Business. *MIT*

Spring 2023; Spring 2024; Spring 2026

24/6.C40: Ethics of Computing. *MIT*

Fall 2025 (with Brad Skow)

AWARDS AND HONORS

Cornell University Fellowship	2016–2017
NSF GRFP Fellowship	2017–2021
Microsoft Research PhD Fellowship	2018–2020
SIGecom Doctoral Dissertation Award Honorable Mention	2021
ACM Doctoral Dissertation Award	2021
Cornell University Department of Computer Science PhD Dissertation Award	2022
ACM EC Exemplary Applied Modeling Track Paper Award	2022
The Web Conference Best Paper Award	2023

INVITED TALKS

Planning Problems for Sophisticated Agents with Present Bias

<i>Social Impact through Network Science.</i> Venice, Italy	June 8–10, 2016
<i>Young Researcher Workshop on Economics and Computation.</i> Tel Aviv, Israel	January 1–5, 2017
Algorithmic Fairness and Bias	
<i>Deloitte Data Scientist Speaker Series.</i> Virtual	September 7, 2018
Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices	
<i>Fairness and Equity in Hiring Selection Algorithms.</i> Upturn, Washington D.C.	October 4, 2019
<i>ASSA Session on Auditing and Regulating AI Systems.</i> Virtual	January 7, 2022
<i>5th IDSC of IZA Workshop: Matching Workers and Jobs Online (Keynote).</i> Virtual	September 16, 2022
Algorithmic Fairness in Practice	
<i>Walmart Community of Practice.</i> Virtual	August 26, 2020
The Societal Impacts of Algorithmic Decision-Making	
<i>Netflix Research Seminar.</i>	May 21, 2021
<i>Sai University Frontiers Seminar.</i> Virtual	November 23, 2021
The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization	
<i>NYU Data Science Seminar.</i> Virtual	April 6, 2022
<i>Harvard EconCS Seminar.</i> Virtual	April 8, 2022
<i>Facebook Core Data Science Seminar.</i> Virtual	May 6, 2022
<i>MIT Initiative on the Digital Economy seminar.</i> Virtual	October 4, 2022
<i>INFORMS Annual Meeting.</i> Indianapolis, IN	October 18, 2022
<i>Opportunity Insights.</i> Cambridge, MA	October 26, 2022
<i>JHU Behavioral Science Forum on Technology and Social Change.</i> Baltimore, MD	October 28, 2022
<i>UIUC Responsible Data Science and AI seminar.</i> Virtual	November 4, 2022
<i>INFORMS Annual Meeting.</i> Phoenix, AZ	October 16, 2023
Measuring Fairness	
<i>FinRegLab Advisory Board.</i> Virtual	July 14, 2021
Understanding Societal Impacts through Machine Learning and Mechanism Design: Automated Hiring as a Case Study	
<i>Israel Algorithmic Game Theory Seminar.</i> Virtual	October 19, 2021
The Future of AI in Financial Services (Panel)	
<i>Artificial Intelligence and the Economy.</i> U.S. Dept. of Commerce, Washington D.C.	April 27, 2022
Reflections on Incentivizing Behavior through Evaluation	
<i>Algorithmic Contract Design: Present and Future (EC '22 Workshop).</i> Boulder, CO	July 15, 2022
Responsible AI and the Role of Measurement	
<i>Network Architecture for Geeks.</i> Virtual	September 21, 2022
<i>Sai University Frontiers Seminar.</i> Virtual	November 17, 2022
<i>Wells Fargo Data Ethics.</i> Virtual	April 26, 2023
<i>Fujitsu x MIT Knowledge Session.</i> Virtual	May 16, 2023
Is AI the New HR? Protecting Civil Rights at Work (Panel)	
<i>South by Southwest.</i> Austin, TX	March 11, 2023
Generative AI and Employment	

<i>MIT Generative AI week.</i> Cambridge, MA	November 30, 2023
Bias & discrimination in algorithmic decision-making <i>40th Annual NABE Economic Policy Conference.</i> Washington DC	January 16, 2024
Human Expertise in Algorithmic Prediction <i>Bridging Prediction and Intervention Problems in Social Systems.</i> Banff, Canada	June 2, 2024
<i>Humans, Algorithmic Decision-Making and Society: Modeling Interactions and Impact (ICML '24 Workshop).</i> Vienna, Austria	July 27, 2024
<i>TTIC Workshop on Data-Driven Decision Processes.</i> Chicago, IL	August 28, 2024
<i>NSF-SDM Workshop on Human-AI Teaming.</i> Pittsburgh, PA	September 26, 2024
Synthetic Data and Swapping for the US Decennial Census <i>CS+Law Virtual Workshop Series.</i> Virtual	November 15, 2024
The Role of Information in Human-AI Collaboration <i>Tuck Operations and Management Science Seminar.</i> Hanover, NH	December 5, 2024
<i>Simons Collaboration on Theory of Algorithmic Fairness.</i> New York, NY	February 6, 2025
<i>UC Berkeley CLIMB Seminar.</i> Berkeley, CA	March 20, 2025
<i>UC Santa Barbara Technology Management distinguished speaker series.</i> Santa Barbara, CA	May 16, 2025
Competition and Diversity in Generative AI <i>SIGecom Winter Meeting.</i> Virtual	March 6, 2025
<i>ICML Workshop on Reliable and Responsible Foundation Models (invited).</i> Vancouver, Canada	July 19, 2025
<i>MIT Sloan Operations Management Seminar.</i> Cambridge, MA	September 22, 2025
<i>New Directions in Social Algorithms.</i> New Haven, CT	October 17, 2025
<i>MSR New England Econ+Comp/CSS Seminar.</i> Cambridge, MA	December 9, 2025

SERVICE AND ORGANIZATION

International Advisory Board, Sai University (Chennai, India) 2021–Present

Program Committees

- ACM Conference on Fairness, Accountability, and Transparency: 2020, 2021, 2022 (Area Chair), 2023 (Area Chair)
- ACM Conference on Economics and Computation: 2019, 2020, 2021, 2022
- ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization: 2021, 2022, 2023 (Area Chair)
- Conference on Neural Information Processing Systems: 2021
- Conference on Web and Internet Economics: 2021
- AAAI Conference on Artificial Intelligence: 2020, 2021
- International Conference on Learning Representations: 2021
- The Web Conference: 2024 (Senior Area Chair)

Workshops and Tutorials

- Fairness and Discrimination through the Dual Lenses of Mechanism Design and Machine Learning (EC '21)
- Fair and Socially Responsible ML for Recommendations: Challenges and Perspectives (NeurIPS '22)

Journal Reviewing

Journal of Machine Learning Research, Proceedings of the National Academy of Sciences, Artificial Intelligence, Management Science, Patterns, and others

Coaching

Harvard Men's Soccer Club

2021–Present