

# Manish Raghavan

mraghavan@seas.harvard.edu  
<https://mraghavan.github.io/>

## EDUCATION

---

**Ph.D., Computer Science** 2016–2021  
*Cornell University*  
Advisor: Jon Kleinberg

**M.S., Computer Science** 2018  
*Cornell University*  
GPA: 4.00

**B.S., Electrical Engineering and Computer Science** 2012–2016  
*University of California, Berkeley*  
GPA: 4.00

## EMPLOYMENT

---

**Assistant Professor** Beginning August 2022  
*Massachusetts Institute of Technology, Cambridge, MA*  
*Sloan School of Management and Department of Electrical Engineering and Computer Science*

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

**Research Intern** May–September 2019  
**Visiting Researcher** September 2019–August 2021  
*Facebook, New York, NY*  
Evaluated social impacts of products  
Developed educational materials for issues related to algorithmic fairness

**Software Engineering Intern** May–August 2018  
*Google, Mountain View, CA*  
Analyzed user behavior on social media  
Developed and analyzed algorithms for stochastic probing problems (results published at ICML 2019)

**Research Intern** May–August 2017  
*Microsoft Research, New York, NY*  
Researched algorithmic fairness in contextual bandit settings (results published at COLT 2018)

**Teaching Assistant** January 2014–May 2016  
*UC Berkeley*  
CS 61B: Data Structures and Algorithms; CS 70: Discrete Math and Probability Theory; CS 170: Introduction to CS Theory  
Taught 30–50-student sections, held office hours, and developed course materials (4 semesters total)

**Research Intern** May–August 2015  
*Cornell University*  
Developed a graph-theoretic model for sophisticated present-biased agents (results published at EC 2016)

## PUBLICATIONS

---

1. N. Dalvi, M. Olteanu, M. Raghavan, and P. Bohannon. *Deduplicating a Places Database*. In *Proc. 23rd International World Wide Web Conference*. April 2014
2. J. Kleinberg, S. Oren, and M. Raghavan. *Planning Problems for Sophisticated Agents with Present Bias*. In *Proc. 17th ACM Conference on Economics and Computation*. July 2016
3. 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
4. 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
5. J. Kleinberg, S. Oren, and M. Raghavan. *Planning with Multiple Biases*. In *Proc. 18th ACM Conference on Economics and Computation*. June 2017
6. G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, K. Weinberger. *On Fairness and Calibration*. In *Proc. 31st Annual Conference on Neural Information Processing Systems*. December 2017
7. J. Kleinberg, M. Raghavan. *Selection Problems in the Presence of Implicit Bias*. In *The 9th Innovations in Theoretical Computer Science Conference*. January 2018
8. M. Raghavan, A. Anderson, J. Kleinberg. *Mapping the Invocation Structure of Online Political Interaction*. In *Proc. 27rd International World Wide Web Conference*. April 2018
9. 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
10. M. Raghavan, M. Purohit, S. Gollapudi. *Hiring Under Uncertainty*. In *International Conference on Machine Learning*. June 2019
11. J. Kleinberg, M. Raghavan. *How Do Classifiers Induce Agents To Invest Effort Strategically?*. In *Proc. 20th ACM Conference on Economics and Computation*. June 2019; Also appeared in *ACM Transactions on Economics and Computing*. October 2020, and as *Designing Evaluation Rules that are Robust to Strategic Behavior* in *Proc. 34th AAAI Conference on Artificial Intelligence, Sister Conference Track*, April 2020.
12. M. Raghavan, S. Barocas. *Challenges for mitigating bias in algorithmic hiring*. In *The Brookings Institution*. December 2019
13. R. Abebe, S. Barocas, J. Kleinberg, K. Levy, M. Raghavan, D. G. Robinson. *Roles for Computing in Social Change*. In *Proc. Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020
14. S. Barocas, A. D. Selbst, M. Raghavan. *The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons*. In *Proc. Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020
15. M. Raghavan, S. Barocas, J. Kleinberg, K. Levy. *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*. In *Proc. Third ACM Conference on Fairness, Accountability, and Transparency*. January 2020
16. 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 *Proc. Fourth ACM Conference on Fairness, Accountability, and Transparency*. January 2021
17. J. Kleinberg and M. Raghavan. *Algorithmic Monoculture and Social Welfare*. In *Proc. National Academy of Sciences 118(22)*. June 2021

18. J. Kleinberg, S. Oren, M. Raghavan, and N. Sklar. *Stochastic Model for Sunk Cost Bias*. In *Proc. 37th Conference on Uncertainty in Artificial Intelligence*. July 2021
19. E. Black, M. Raghavan, and S. Barocas. *Model Multiplicity: Opportunities, Concerns, and Solutions*. In *Proc. Fourth ACM Conference on Fairness, Accountability, and Transparency*. June 2022
20. J. Kleinberg, S. Mullainathan, and M. Raghavan. *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*. In *Proc. 23rd ACM Conference on Economics and Computation*. July 2022

## AWARDS AND HONORS

---

|  |           |
|--|-----------|
| <b>Regents' and Chancellor's Scholar, UC Berkeley</b>        | 2012–2016 |
| <b>Outstanding Graduate Student Instructor, UC Berkeley</b>  | 2015–2016 |
| <b>Hertz Foundation Fellowship Finalist</b>                  | 2016      |
| <b>Cornell University Fellowship Recipient</b>               | 2016–2017 |
| <b>NSF GRFP Fellowship Recipient</b>                         | 2017–2021 |
| <b>Microsoft Research PhD Fellowship Recipient</b>           | 2018–2020 |
| <b>SIGecom Doctoral Dissertation Award Honorable Mention</b> | 2021      |

## INVITED TALKS AND WORKSHOPS

---

|   |                   |
|---|-------------------|
| <b>Social Impact through Network Science</b>  | June 8–10, 2016   |
| Venice, Italy. <i>Planning Problems for Sophisticated Agents with Present Bias</i>                        |                   |
| <b>Fairness, Accountability, and Transparency in Machine Learning</b>                                     | November 18, 2016 |
| New York, New York. <i>Inherent Trade-Offs in the Fair Determination of Risk Scores</i>                   |                   |
| <b>Young Researcher Workshop on Economics and Computation</b>   | January 1–5, 2017 |
| Tel Aviv, Israel. <i>Planning Problems for Sophisticated Agents with Present Bias</i>                     |                   |
| <b>Workshop on Prioritising Online Content</b>  | December 9, 2017  |
| Long Beach, California. <i>The Externalities of Exploration and How Data Diversity Helps Exploitation</i> |                   |
| <b>Deloitte Data Scientist Speaker Series</b>   | September 7, 2018 |
| Virtual. <i>Algorithmic Fairness and Bias</i>   |                   |
| <b>Workshop on Workshop on Ethical, Social and Governance Issues in AI</b>                                | December 7, 2018  |
| Montreal, Canada. <i>How Do Classifiers Induce Agents To Invest Effort Strategically?</i>                 |                   |
| <b>Privacy Law Scholars Conference</b>  | May 30–31, 2019   |
| Berkeley, California. <i>Formalism, Computing, and Social Change</i>                                      |                   |
| <b>Learning in the Presence of Strategic Behavior</b>   | June 28, 2019     |
| Phoenix, Arizona. <i>How Do Classifiers Induce Agents To Invest Effort Strategically?</i>                 |                   |
| <b>Mechanism Design for Social Good</b>   | June 28, 2019     |
| Phoenix, Arizona. <i>Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices</i>           |                   |
| <b>NeurIPS 2019 Workshop on Robust AI in Financial Services</b>   | December 13, 2019 |
| Vancouver, Canada. <i>The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons</i> |                   |
| <b>Workshop on Human Interpretability in Machine Learning</b>   | July 17, 2020     |
| Virtual. <i>The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons</i>           |                   |
| <b>Workshop on Participatory Approaches to Machine Learning</b>   | July 17, 2020     |
| Virtual. <i>The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons</i>           |                   |
| <b>Workshop on Law &amp; Machine Learning</b>   | July 17, 2020     |
| Virtual. <i>Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices</i>                    |                   |
| <b>AI for Social Good (AI4SG 2020)</b>  | July 20, 2020     |
| Virtual. <i>Fairness and Discrimination in Mechanism Design and Machine Learning</i>                      |                   |
| <b>Walmart Community of Practice</b>  | August 26, 2020   |
| Virtual. <i>Algorithmic Fairness in Practice</i>  |                   |
| <b>Netflix Research Seminar</b>   | May 21, 2021      |

Virtual. *The Societal Impacts of Algorithmic Decision-Making*  
**Israel Algorithmic Game Theory Seminar** October 19, 2021  
Virtual. *Understanding Societal Impacts through Machine Learning and Mechanism Design: Automated Hiring as a Case Study*

**Workshop on Explainable AI in Finance** November 3, 2021  
Virtual. *Explanations in Whose Interests?*

**NYU Data Science Seminar** April 6, 2022  
Virtual. *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*

**Harvard EconCS Seminar** April 8, 2022  
Virtual. *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*

**Facebook Core Data Science Seminar** May 6, 2022  
Virtual. *The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization*