Hoda Heidari

CONTACT INFORMATION
Gates Hillman Center, Office #8229
Pittsburgh, PA 15213

E-mail: hheidari@cmu.edu
Phone: 412-268-6869

RESEARCH INTERESTS

- Responsible AI, AI Ethics
- AI Accountability and Governance
- Fairness and Bias through ML
- Algorithmic Economics

POSITIONS

K&L Gates Career Development Assistant Professor in Ethics & Computational Technologies
Carnegie Mellon University, Pittsburgh, PA Jul 2023–present

Assistant Professor, Machine Learning Department (100% appointment, tenure home)
School of Computer Science, **Carnegie Mellon University** Jul 2020–present

- *Co-founder and co-lead of the Responsible AI Initiative at CMU*
- *CMU Affiliations:* Institute for Software, Systems, and Society (0% appointment), Human -Computer Interaction Institute (HCII), CyLab Security and Privacy Institute, Block Center for Technology and Society, and Heinz College of Information Systems and Public Policy

Post-doctoral Associate, **Cornell University**, Ithaca, NY Sep 2019–Jul 2020

- *Supervisor:* Professors Jon Kleinberg, Karen Levy, and Solon Barocas
- *Affiliation:* AI, Policy, and Practice initiative, Computer & Information Science Department

Post-doctoral Fellow, **ETH Zürich**, Switzerland Aug 2017–Aug 2019

- *Supervisor:* Professor Andreas Krause
- *Affiliation:* Learning and Adaptive Systems Group at the Machine Learning Institute

EDUCATION

University of Pennsylvania, Philadelphia, PA 2017
Ph.D. in Computer and Information Science

- *Thesis*: "Essays in Algorithmic Market Design under Social Constraints"
- *Advisors*: Professors Michael Kearns and Ali Jadbabaie
- *Committee*: Rakesh Vohra, Aaron Roth, Shivani Agrawal, Vahab S. Mirrokni
- *Elective Coursework*: Microeconomics I (Decision Theory and General Equilibrium Theory), Microeconomics II (Game Theory and Mechanism Design), Market Design, Discrete Convexity & Submodularity, Convex Optimization, Computational Learning Theory

Wharton School of Business, Philadelphia, PA 2017
M.Sc. in Statistics

- *Thesis:* "What Can Machine Learning Teach Econometrics?"
- *Supervisor*: Prof. Dylan Small
- *Coursework*: Applied Econometrics I, Mathematical Statistics, Probability, Stochastic Processes, Statistical Learning Theory, Data Analysis, Statistical Computing

Sharif University of Technology, Tehran, Iran 2011
B.Sc. in Computer Engineering

- *Thesis:* "Toward Optimal Vaccination Strategies for Probabilistic Models"
  Joint work with Z. Abbassi, supervised by Prof. Mohammad Ghodsi.
- *Elective Coursework*: Linear Algebra, Algorithmic Game Theory (graduate course)

INDUSTRY EXPERIENCE

Summer Intern, **Google Research**, New York 2015
- *Mentors:* Hossein Azari, Mohammad Mahdian, Umar Syed, Sergei Vassilvitskii

Summer Intern, **Microsoft Research**, New York 2014
- *Mentors:* Sebastien Lahaie, David Pennock, Jenn Wortman Vaughan

PEER-REVIEWED PUBLICATIONS

[1] *A Taxonomy of Human and ML Strengths in Decision-Making to Investigate Human-ML Complementarity*. Charvi Rastogi*, Leqi Liu*, Ken Holstein, Hoda Heidari. The AAAI Conference on Human Computation and Crowdsourcing (HCOMP), 2023.

[2] *The AI Incident Database as an Educational Tool to Raise Awareness of Harms: A Classroom Exploration of Efficacy, Limitations, & Future Design Improvements*. Michael Feffer, Nikolas Martelaro, and Hoda Heidari. ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), 2023.

[3] *Toward Operationalizing Pipeline-aware ML Fairness: A Research Agenda for Developing Practical Guidelines and Tools.* Emily Black, Rakshit Naidu, Rayid Ghani, Kit T. Rodolfa, Daniel E. Ho, and Hoda Heidari. ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), 2023.

[4] *Cognitive Diversity and Affinity Bias in Team Formation Dynamics*. Solon Barocas, Karen Levy, Jon Kleinberg, and Hoda Heidari. ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), 2023.

[5] *Perspectives on Incorporating Expert Feedback into Model Updates*. Valerie Chen, Umang Bhatt, Hoda Heidari, Adrian Weller, and Ameet Talwalkar. Pattern (Volume 4, Issue 7), 2023.

[6] *From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research*. Michael Feffer, Michael Skirpan, Zachary Chase Lipton, and Hoda Heidari. AAAI /ACM Conference on Artificial Intelligence, Ethics, and Society (AIES), 2023.

[7] *A Validity Perspective on Evaluating the Justified Use of Data-driven Decision-making Algorithms*. Amanda Coston, Anna Kawakami, Haiyi Zhu, Ken Holstein and Hoda Heidari. The IEEE Conference on Secure & Trustworthy Machine Learning (SAT-ML), 2023. Best paper award.

[8] *Local Justice & ML: Modeling and Inferring Dynamic Ethical Judgments Around High-stakes Allocations*. Violet Chen, Joshua Williams, Derek Leben, and Hoda Heidari. The AAAI Conference on Artificial Intelligence (AAAI) , 2023.

[9] *Moral Machine or Tyranny of the Majority?*. Michael Feffer, Hoda Heidari, Zachary Chase Lipton. AAAI, 2023.

[10] *Bayesian Persuasion for Algorithmic Recourse*. Keegan Harris, Valerie Chen, Joon Sik Kim, Ameet Talwalkar, Hoda Heidari, and Steven Wu. NeurIPS, 2022.

[11] *Four Years of FAccT: A Reflexive, Mixed-Methods Analysis of Research Contributions, Shortcomings, and Future Prospects*. Ben Laufer, Sameer Jain*, A. Feder Cooper*, Jon Kleinberg, and Hoda Heidari. ACM conference on Fairness, Accountability, and Transparency (FAccT), 2022.

[12] *Strategic Instrumental Variable Regression: Recovering Causal Relationships From Strategic Responses*. K. Harris, D. Ngo, L. Stapleton, H. Heidari, and S. Wu. The International Conference on Machine Learning (ICML), 2022.

[13] *Stateful Strategic Regression*. K. Harris, H. Heidari, and S. Wu. Neural and Information Processing Systems (NeurIPS), 2021.

[14] *On Modeling Human Perceptions of Allocation Policies with Uncertain Outcomes*. H. Heidari, S. Barocas, J. Kleinberg, and K. Levy. The ACM Conference on Economics and Computation (EC), 2021. Exemplary track award.

[15] *Addressing the Long-term Impact of ML Decisions via Policy Regret*. D. Lindner, H. Heidari, and A. Krause. The International Joint Conference on Artificial Intelligence (IJCAI), 2021.

[16] *Fair equality of chances: fairness for statistical prediction-based decision-making*. M. Loi, A. Herlitz, and H. Heidari. AIES, 2021.

[17] *A Human-in-the-loop Framework to Construct Context-aware Mathematical Notions of Outcome Fairness*. M. Yaghini, A. Krause, and H. Heidari. AIES, 2021.

[18] *Allocating Opportunities in a Dynamic Model of Intergenerational Mobility*. Hoda Heidari and Jon Kleinberg. ACM conference on Fairness, Accountability, and Transparency (FAccT), 2021. Best paper award.

[19] *On the Desiderata for Online Altruism: Nudging for Equitable Donations*. Nuno Mota, Abhijnan Chakraborty, Asia J. Biega, Krishna P. Gummadi, Hoda Heidari: Proceedings of the ACM on Human-Computer Interaction (CSCW), Volume 4, 2020.

[20] *Algorithmic Notions vs. Human Perceptions of Fairness: A Descriptive Approach to Selecting a Suitable Fairness Metric*. Megha Srivastava, Hoda Heidari, and Andreas Krause. Proceedings of the SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2019.

[21] *On the Long-term Impact of Algorithmic Decision Policies: Effort Unfairness and Feature Segregation through Social Learning*. Hoda Heidari, Vedant Nanda, and Krishna P. Gummadi. The International Conference on Machine Learning (ICML), 2019.

[22] *On the Impact of Choice Architectures on Inequality in Online Donation Platforms*. Abhijnan Chakraborty, Nuno Mota, Asia J. Biega, Krishna P. Gummadi and Hoda Heidari. The Web Conference (WWW), 2019.

[23] *A Moral Framework for Understanding of Fair ML through Economic Models of Equality of Opportunity*. Hoda Heidari, Michele Loi, Krishna P. Gummadi, and Andreas Krause. FAccT, 2019.

[24] *Fairness Behind a Veil of Ignorance: A Welfare Analysis for Automated Decision Making*. Hoda Heidari, Claudio Ferrari, Krishna P. Gummadi, and Andreas Krause. NeurIPS, 2018.

[25] *A Unified Approach to Quantifying Algorithmic Unfairness: Measuring Individual and Group Unfairness via Inequality Indices.* Till Speicher*, Hoda Heidari*, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, Muhammad Bilal Zafar. KDD, 2018.

[26] *Preventing Disparate Treatment in Sequential Decision Making.* Hoda Heidari and Andreas Krause. IJCAI, 2018.

[27] *Fairness in Criminal Justice Risk Assessments: The State of the Art.* Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Sociological Methods and Research, 2018.

[28] *Pricing a Low-regret Seller*. Hoda Heidari, Mohammad Mahdian, Umar Syed, Sergei Vassilvistskii, and Sadra Yazdanbod. ICML, 2016.

[29] *Tight Policy Regret Bounds for Improving and Decaying Bandits*. Hoda Heidari, Michael Kearns and Aaron Roth. IJCAI, 2016.

[30] *Integrating Market Makers, Limit Orders, and Continuous Trade in Prediction Markets*. Hoda Heidari, Sebastien Lahaie, David Pennock, and Jenn Wortman Vaughan. EC, 2015.
Full version was invited for publication, and appeared in the ACM Transactions on Economics and Computation (TEAC).

[31] *Learning from Contagion (Without Timestamps)*. Kareem Amin, Hoda Heidari, and Michael Kearns. ICML, 2014.

[32] *New Models for Competitive Contagion*. With Moez Draief and Michael Kearns. AAAI, 2014.

[33] *Competitive Contagion in Networks*. Sanjeev Goyal, Hoda Heidari, and Michael Kearns. Games and Economic Behavior Journal (GEB), 2014.

[34] *Depth-Workload Tradeoffs for Workforce Organization*. Hoda Heidari and Michael Kearns. The Conference on Human Computation & Crowdsourcing (HCOMP), 2013.

PEER-REVIEWED NON-ARCHIVAL PUBLICATIONS

- *Assessing AI Impact Assessments: A Classroom Study*. Nari Johnson and Hoda Heidari. The Regulatable ML Workshop at NeurIPS, 2023
- *Allocating Bonus Points in Sequential Matchings with Preference Dynamics*. Meirav Segal, Leqi Liu, Anne-Marie George, Christos Dimitrakakis1, Hoda Heidari. Workshop on Algorithmic Fairness through the Lens of Time at NeurIPS 2023
- *Recentering Validity Considerations through Early-Stage Deliberations Around AI and Policy Design*. Anna Kawakami, Amanda Costosn, Haiyi Zhu, Hoda Heidari, Kenneth Holstein. Workshop on Designing Policy and Technology Simultaneously at the ACM Conference on Human Factors in Computing Systems (CHI) 2023.
- *A Unifying Framework for Combining Complementary Strengths of Humans and ML toward Better Predictive Decision-Making*. Charvi Rastogi, Leqi Liu, Ken Holstein, and Hoda Heidari. Poster at EAAMO, 2022.
- *A Sandbox Tool to (Bias)Stress-Test Fairness Algorithms*. Nil-Jana Akpinar, Manish Nagireddy, Logan Stapleton, Hao-Fei Cheng, Haiyi Zhu, Steven Wu, and Hoda Heidari. Poster at EAAMO, 2022.

WORKING PAPERS

- *The Situate AI Guidebook: Co-Designing a Toolkit to Support Multi-Stakeholder Early-stage Deliberations Around Public Sector AI Proposals*. Anna Kawakami, Amanda Coston, Hoda Heidari*, Kenneth Holstein*, Haiyi Zhu*. 2023
- *Studying Up Public Sector AI: How Networks of Power Relations Shape Agency Decisions Around AI Design and Use*. Anna Kawakami, Amanda Coston, Hoda Heidari*, Kenneth Holstein*, Haiyi Zhu*. 2023
- *Stagnated XAI? A Mixed-method Account of Persisting Predicaments and Path to Progress*. Rebecca Yu and Hoda Heidari. 2023
- *RELand: Risk Estimation of Landmines via Interpretable Invariant Risk Minimization*. Mateo Dulce Rubio, Siqi Zeng, Qi Wang, Didier Alvarado, Francisco Moreno, Hoda Heidari, Fei Fang. 2023
- *Beneficent Intelligence: A Capability Approach to Modeling Benefit, Assistance, and Associated Moral Failures through AI Systems*. Alex J. London* and Hoda Heidari*. 2023.
- *Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models*. Ben Laufer, Jon Kleinberg*, and Hoda Heidari*. 2023.
- *Toward Generating Actionable Counterfactual Explanations via Posterior Inference*. Joshua Williams, Anurag Katakkar, Hoda Heidari, and Zico Kolter. 2022.

INVITED ARTICLES

- Invited Book Chapter on *"Fairness in AI"* 2023
  To be published by ACM as part of a textbook on "Artificial Intelligence and Society"
- *On Modeling Human Perceptions of Allocation Policies with Uncertain Outcomes.* 2023
  Invited letter at *SIGEcom Exchanges*
- *Allocating Opportunities in a Dynamic Model of Intergenerational Mobility.* 2022
  Invited for publication and presentation at the Sister Conference Best Paper Track of IJCAI

AWARDS

- Winner of a **Best Paper Award** at The IEEE Conference on Secure and Trustworthy Machine Learning (SAT-ML) 2023
- Winner of an **Exemplary Track Award** at The ACM Conference on Economics and Computation (EC) 2021
- Winner of a **Best Paper Award** at the ACM Conference on Fairness, Accountability, and Transparency (FAccT) 2021
- Selected to deliver a lecture through the *AAAI New Faculty Highlights Program* 2021
- **Facebook Research Award (100K)** to build *"A Tool to Study the Efficacy of Fairness Algorithms on Specific Bias Types"*. With co-PIs S. Wu and H. Zhu 2021
- **J. P. Morgan and Chase Individual Faculty Award** (10K) 2021
- Selected to participate in the *Rising Stars in EECS* workshop 2019
- *Travel Award*, Women in Machine Learning (WiML) workshop 2017
- *Travel Award*, Women in Theory (WiT) workshop 2016
- *Silver Medal* in the National Mathematical Olympiad, Iran 2006

GRANTS

- Co-PI, **PwC Research Grant** (500K) for *"Developing an AI Assurance Process"*. With PI R. Ghani and Co-PI A. J. London 2023
- Co-PI, **PIT-UN grant** (90K) on *Developing a Community Education Toolkit for Public Social Service AI: Training Next-Generation Multidisciplinary Social Workers.* With PI. H. Shen and others 2023
- Senior Personnel, the **NSF AI Institute for Societal Decision-Making (AI-SDM)** (20M).

with PI. Aarti Singh and others 2023
- Invited research member at the **Simons Laufer Mathematical Sciences Institute**. 2023
- **PwC Research Grant** (300K) for *"Robust and Fair AI Systems in Dynamic Environments"*. With PI Z. Lipton and Co-PIs A. Risteski, and D. Childers 2022
- PI, **NSF FAI grant (600K)** on *"Fair AI in Public Policy: Achieving Fair Societal Outcomes in ML Applications to Education, Criminal Justice, and Health and Human Services"*. With co-PIs Chouldechova, Ghani, Lipton and Rodolfa. 2021
- PI, **J. P. Morgan and Chase Individual Faculty Award** (10K) 2021
- PI, **CMU CyLab grant (50K)** *"On the Impact of Algorithmic Fairness Metrics and Methods on Trust in Machine Learning Systems"*. 2021

ADVISING

CMU Students

- Nari Johnson, doctoral student at MLD. Summer 2023–present
- Anusha Sinha, doctoral student at Societal Computing. Co-advised with Steven Wu. Spring 2023–present
- Rebecca Yu, doctoral student at MLD. Fall 2022–present
- Michael Feffer, doctoral student at Societal Computing. Co-advised with Z. Lipton. Fall 2021–present
- Keegan Harris, doctoral student at MLD. Co-advised with S. Wu. Fall 2020–Fall 2022
- Atishay Jain, master student at MLD. Fall 2022–Spring 2023
- Eric Liang, master student at MLD. Co-advised with R. Ghani. Spring 2022–Fall 2022
- Siqi Zeng, master student at MLD. Co-advised with F. Fang. Fall 2022–Spring 2023
- Qiqi Xu, master student at MLD. Co-advised with N. Shah. Fall 2021–Spring 2022
- Sameer Jain, master student at LTI. Spring–Fall 2021
- Anurag Katakkar, master student at LTI. Spring 2021
- Manish Nagireddy, bachelor student at SCS. Summer–Fall 2021

Graduate Thesis Committee

- Matteo Dulce Rubio, Doctoral student at Heinz, CMU.
- Nil-Jana Akpinar, Doctoral student at MLD and Statistics, CMU.
- Emily Black, Doctoral student at CSD, CMU.
- Anthonia Carter, Doctoral student at Information Science, Cornell University.
- Violet Chen, Doctoral student at Tepper, CMU.
- Amanda Coston, Doctoral student at MLD and public policy, CMU.
- Leqi Liu, Doctoral student at MLD, CMU.
- Charvi Rastogi, Doctoral student at MLD, CMU.

TEACHING

Instructor, CMU

- **Responsible AI** (10-735/80-831) Spring'24
  - Co-taught with Alex J. London
  - Cross-listed between the School of Computer Science (Machine Learning) and Dietrich College of Humanities and Social Sciences (Department of Philosophy)
- **Introduction to Machine Learning** (10-601) Spring'24
  - Co-taught with Henry Chai and Matthew Gormley
- **Machine Learning, Ethics, and Society** (10-613/713) Fall'21, Spring'23
- **Mathematical/Computational Foundations of ML** (10-606/7) Fall'22
- **Probabilistic Graphical Models** (10-708) Spring'22
  - Co-taught with Andrej Risteski

Instructor, ETH Zürich and CMU

- Designed and taught a new course on **Fairness, Explainability, and Accountability for Machine Learning** Spring 2019, Fall 2020

Guest Lecturer, CMU

- *"Fairness, Transparency, and Accountability in ML"* at 10-701. Spring'21, Spring'23
- *"Fairness, Transparency, and Accountability in ML"* at 10-315. Spring'22
- *"Formalizations of Fairness in ML"* at 17-200. Fall'21
- *"Societal Computing Practicuum"* (17-994). Fall'20, Fall'21, Fall'22

Teaching Assistant, ETH Zürich

- *Introduction to Machine Learning* Prof. A. Krause Spring 2018
- *Probabilistic Artificial Intelligence*. Prof. A. Krause Fall 2017

SERVICE

Organizer

- Co-organizer and speaker at the NeurIPS Tutorial on *"Governance & Accountability for Machine Learning: Existing Tools, Ongoing Efforts, & Future Directions"* 2023
- Chair and organizer of the 2022 PI Meeting for the NSF Program on Fairness in AI in Collaboration with Amazon 2022
- Tutorial Co-chair for the ACM FAccT conference 2022
- NeurIPS Workshop on *"Learning and Decision-Making with Strategic Feedback"* 2022
- ICLR Workshop on *"Responsible AI"* 2021
- Workshop on *"Intersectionality in Fair Machine Learning: Where Are We and Where Should We Go from Here?"* 2020
  At MOSAIC 2020: An Annual Conference on Intersectionality.
- NeurIPS Workshop on *"Human-centric Machine Learning"* 2019
- Tutorial on *"Economic Theories of Distributive Justice for Fair ML"* 2019
  Presented at the 30th Web Conference
- Weekly reading group on the *"Societal Aspects of AI"*, ETHZ 2017, 2018

University Service

- Executive leadership of the *Responsible AI initiative* 2021-present
- Faculty hiring committee member for a joint SCS-Dietrich search. Fall'22-present
- *SCS council member* discussing school-wide topics with the dean, School of Computer Science at CMU 2021-present
- Admissions committee member for MLD 2020
- Admissions committee member for Societal Computing 2020, 2021, 2022

Senior Program Committee

- The Conference on Economics and Computation (EC) 2022
- ACM Conference on Fairness, Accountability, and Transparency (FAccT) 2022
- The AAAI Conference on Artificial Intelligence (AAAI) 2022
- The Conference on Neural Information Processing Systems (NeurIPS) 2021
- The International Conference in Machine Learning (ICML) 2021

Program Committee

- ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*) 2019-21
- The AAAI Conference on Human Computation and Crowdsourcing (HCOMP) 2019
- The International Conference in Machine Learning (ICML) 2016-2020
- The Conference on Economics and Computation (EC) 2015, 2018
- The Conference on Artificial Intelligence (AAAI) 2018, 2020

Ad Hoc Reviewer

- Serving on an *NSF Panel* to evaluate and discuss research grant proposals April 2021
- Journal of Machine Learning Research (JMLR)
- The Conference on Neural Information Processing Systems (NeurIPS)
- The Operations Research Journal (OR)
- Mathematics of Operations Research (MOR)
- The Network Science Journal
- Transactions on Economics and Computation (TEAC)
- Transactions on Signal Processing (TSP)
- Information Processing Letters (IPL)

External Reviewer

- Conference on Learning Theory (COLT)
- The Conference on Web and Internet Economics (WINE)
- The Symposium on Discrete Algorithms (SODA) (2013, 2017)
- The Symposium on Foundations of Computer Science (FOCS)
- The International Joint Conference on Artificial Intelligence (IJCAI)
- Conference on Decision and Control (CDC)

Outreach & Pipeline Building Activities

- Speaker at the Pittsburgh Women in Mathematics and Computing Symposium (WMCS) hosted at CMU Feb 2023
- Team mentor, OurCS conference, CMU Oct 2022
- Group leader at the Undergraduate Research Mixer, SCS CMU 2021
- Member of the Wellness Network at MLD, SCS CMU 2020-2021
- Panelist at the MLD Admissions Workshop for prospective students, SCS CMU 2020
- Panelist at the Academic Job Market discussion, SCS CMU 2020
- Interviewee, Women@SCS Faculty Lunch Interview Series, SCS CMU 2020
- Scientific Body Representative at a *Faculty Recruiting Committee*. Department of Computer Science, ETHZ Mar 2019
- Expert Feedback at *ETH Week: Manufacturing the Future*, ETHZ Sep 2017

ACADEMIC PRESENTATIONS & DISCUSSIONS

- Speaker at the *Algorithmic Fairness Seminar* at Stanford CS Nov'23
- Speaker/panelist at the workshop *"Toward Algorithmic Justice in Precision Medicine"* at UCSF Nov'23
- Q&A moderator for the University Lecture *"ChatGPT and Your Brain"* by Nita Farahany at CMU Nov'23
- Invited speaker at the FDS Workshop *"Theory and Practice of Foundation Models"* hosted by Google and Yale Oct'23
- Invited speaker at the DIMACS Workshop on *"Foundation Models, Large Language Models, and Game Theory"* at Rutgers University Oct'23
- Invited speaker at *"Fair and Socially Aware Practices in Operations Management"* session at INFORMS Oct'23
- Invited speaker at the Connections Workshop: *"Algorithms, Fairness, and Equity"* at the Simons Laufer Mathematical Sciences Institute, Berkeley Aug'23
- Invited panelist at AIES panel on *"AI for society: Developing, Deploying, and Auditing Public-Facing AI"* Aug'23
- Invited speaker at the IDEAS Summer Program *"on human-AI collaboration"* at Northeastern University Aug'23

- Invited discussant at the *the Cowles Foundation Annual Conference on Econometrics*, Yale University Jun 2023
- Invited panelist at TILOS (an NSF AI Institute) *Early Career Development Panel on AI Ethics* May 2023
- Invited keynote speaker at the *"Responsible AI in the Natural Sciences"* workshop hosted at CMU May 2023
- Invited speaker at *"No Joke: we are talking about LLMs"*, Pittsburgh Apr 2023
- Invited speaker at the *Women in Data Science, CMU*, Pittsburgh 2023
- Invited panelist at MIT's *Institute for Data, Systems, and Society (IDSS)* celebration, on Equity, justice, AI and automation Apr 2023
- Invited speaker at the *Human + AI Conference* at University of Chicago Oct 2022
- Invited panelist at the *64th National Association for Business Economics (NABE) Annual Meeting* Oct 2022
- Invited speaker at the *Workshop in AI+Economics*, ETH Zürich 2022
- Invited speaker at the *Workshop on Fairness and Machine Learning in Health*, Millan, Itally 2022
- Invited speaker at the *European Workshop on Algorithmic Fairness*, Zürich 2022
- Invited speaker at the *The Annual ACM SIGecom Winter Meeting* 2022
- Invited speaker at the *CROSSING Conference* by the German Science Foundation 2021
- Invited speaker at the *Trustworthy ML initiative* 2021
- Invited speaker at the *Humanising Machine Intelligence Research Seminar* 2021 organized by Prof. S. Lazar and Dr. C, Adamson at the Australian National University.
- Invited panelist at the *"Roundtable on Data Privacy in Black Communities"* 2021 organized by the Joint Center for Political and Economic Studies
- Invited speaker at the *"Foundations of Algorithmic Fairness"* workshop 2021 organized by European Lab for Learning & Intelligent Systems (ELLIS)
- Invited speaker at the *"Max Planck Symposium on Computing and Society"* 2021 organized by Max Planck Institute
- Invited speaker at the *"Artificial Intelligence and Ethics"* conference 2020 organized by Duke-Kunshan University
- Invited speaker at the workshop on *"Algorithmic Fairness through the Lens of Causality and Interpretability"*, co-located with NeurIPS 2020
- Invited discussant at the *"Philosophical Questions about AI, Law, & Governance"* 2018 Workshop organized by the *Faculty of Law, UZH* (Prof. Christoph Graber), and *Berkman Klein Center of Internet and Society, Harvard University*
- Invited discussant at the *"Governance of Decision-making Algorithms—How to Address Risks?"*. Expert Workshop organized by *Swiss Re* Jul 2018
- Invited discussant at the *"UZH Digital Forum: Can Algorithms be Fair?"* Dec 2018 The *Digital Society Initiative* & the *Swiss Alliance for Data-Intensive Services*
- Invited panel at the *"AI and Intellectual Property"* workshop Jun 2019 Organized by the *the Max Planck Institute for Innovation and Competition*
- Invited speaker at the *"Workshop on Ethical, Social and Governance Issues in AI"* 2018 co-located with NeurIPS

Public Policy Contributions & Outreach

- Led a CMU effort to produce **"A Responsible Voter's Guide to Generative AI and Political Campaigns"**. Report to be published in January 2024. Fall 2023
- Participated in Workshop on *"Operationalizing the Measure Function of the NIST AI RMF"* hosted by Northwestern University Fall 2023
- Participated in a **AI expert roundtable discussion on Capitol Hill** with members of the

New Democrat Coalition Artificial Intelligence (AI) Working Group and the Problem Solvers Caucus.
  - CMU News: Carnegie Mellon University's Block Center for Technology and Society Shares AI Expertise with Policymakers in Washington, D.C.
- Co-chair of **MetroLab's Task Force on *GenAI in Local Government***
  - Task: investigate how GenAI can improve city services; develop processes and policies to safeguard and ensure just, equitable and reliable use.
  - Media report: How Can Local Governments Safely Use Generative AI?
- Organizer of the CMU-RAI Workshop in collaboration with NIST on **"Operationalizing the NIST AI Risk Management Framework"** Jul 2023
  - CMU News: CMU and NIST Team To Manage AI Risk
  - Whitepaper from the event: Operationalizing the NIST AI RMF Framework
- Organized a round table to provide expert response to NTIA's Request for Comment On *AI Accountability Policy* Summer 2023
- Invited participant at the *private discussion* with **Assistant Secretary of Commerce for Communications & Information & NTIA Administrator** on AI accountability at University of Pittsburgh Apr 2023
- Invited speaker at the *"Introduction to AI in Municipal Government"*. A partnership of the National League of Cities and the AAAS Center for Scientific Evidence in Public Issues Aug'23
  - Event report: Exploring AI Applications in City Government: The Promise and the Risks by Joshua Pine and Lena Geraghty
- Organized and moderated a policy salon on *"Accountablity for AI"* through the Block center for Technology and Society Spring 2022
  - Policy memo: **"Toward AI Accountability: Policy Ideas for Moving Beyond a Self-Regulatory Approach"**
  - Memo shared with multiple Federal agencies and the offices of several US representatives at their request.

Media Coverage

- "How AI reduces the world to stereotypes" by Victoria Turk at Rest of World Oct'23
- "Weighing the AI Threat By Incident Reports" by Pam Baker at InformationWeek Sep'23
- "CMU's Summer of AI Experts on the Hill" by Peter Kerwin at CMU News Sep'23
- "Our Region's Business - Responsible AI Initiative" by Bill Flanagan, WPXI Jun'23
- "On Artificial Intelligence" with Julie Mason at SiriusXM Feb'23
- "Responsible AI: Heidari Applies Social Awareness to Developing Technologies" by Susie Cribbs, The Link Aug'22
- "Why CMU launched the Responsible AI initiative" by Sophie Burkholder at Technical.ly Apr'22
- "Responsible AI Initiative launches at CMU following panel discussion including government, industry leaders" by Nate Doughty at PittburghInno Apr'22
- *"5 Great Human-Centered AI Papers from 2018"* by Emma Brunskill, Medium 2018
- *"Fair Algorithms"* in the ETH podcast series 2019