

# Jennifer Wortman Vaughan

Microsoft Research, New York City  
300 Lafayette Street  
New York, NY 10012

jenn@microsoft.com  
jenn@jennwv.com  
<http://jennwv.com>

Updated May 2024

## Research Interests

---

I am interested in the interaction between people and AI systems. In recent years, I have focused on this interaction in the context of responsible AI—specifically transparency, intelligibility, and fairness in machine learning—as part of Microsoft’s FATE group and Co-chair of Microsoft’s Aether Working Group on Transparency. Previously, I often studied this interaction in the context of prediction markets and other crowdsourcing systems. My research background is in machine learning and algorithmic economics. However, because of the central role that people play in the machine learning lifecycle, I now weave human-subject experiments and even qualitative methods into my research to better understand human behavior in sociotechnical systems.

## Education

---

### University of Pennsylvania, Philadelphia, PA

Ph.D., Computer and Information Science, August 2009

Dissertation: *Learning from Collective Preferences, Behavior, and Beliefs*

Advisor: Michael Kearns

M.S.E., Computer and Information Science, August 2006

### Stanford University, Stanford, CA

M.S., Computer Science (Specialization in AI), September 2004

### Boston University, Boston, MA

B.A., Computer Science (Minor in Philosophy), Magna Cum Laude, May 2002

## Employment

---

### Microsoft Research, New York, NY

Senior Principal Researcher, September 2019–Present

Principal Researcher\*, September 2014–September 2019

Senior Researcher\*, October 2012–September 2014

\* The Microsoft Research title-to-level map was updated in 2019. Present-day titles are used here.

### University of California, Los Angeles, Los Angeles, CA

Adjunct Assistant Professor, September 2012–June 2016

Assistant Professor, Computer Science Department, September 2010–September 2012

**Harvard University**, Cambridge, MA

Computing Innovation Fellow (Postdoc), September 2009–August 2010

**Google**, New York, NY

Intern, Research Group, June 2008–September 2008

**Yahoo!**, New York, NY

Intern, Yahoo! Research, May 2007–August 2007

**SiteAdvisor**, Boston, MA

Consultant, May 2005–November 2005

**Stanford University**, Stanford, CA

Research Assistant, March 2003–August 2004

**Intuit**, Waltham, MA

Developer (part time/summer position in college), May 2000–August 2002

## **Teaching Experience**

---

### **Tutorial Presentations**

#### **Transparency and Intelligence Throughout the Machine Learning Lifecycle**

Microsoft Research Webinar, 2020

#### **Challenges of Incorporating Algorithmic Fairness into Industry Practice** (with Henriette Cramer, Kenneth Holstein, Hal Daumé III, Miroslav Dudík, Hanna Wallach, Jean Garcia-Gathright, and Sravana Reddy)

Tutorial at FAccT 2019

#### **Machine Learning & Fairness** (with Hanna Wallach)

Microsoft Research Webinar, 2019

#### **Making Better Use of the Crowd**

Tutorial at KDD 2017

Tutorial at ACL 2017

Tutorial at NeurIPS 2016 (as “Crowdsourcing: Beyond Label Generation”)

#### **Prediction, Belief, and Markets** (with Jake Abernethy)

Tutorial at AAAI 2013

Tutorial at KDD 2012

Santa Cruz Machine Learning Summer School, 2012 (as “Learning and Markets”)

Tutorial at ICML 2012

### **Courses Taught**

**CS 112: Modeling Uncertainty in Information Systems**, UCLA, Spring 2011 and 2012  
Upper division undergraduate course covering the basics of probability, Markov chains, and statistical inference. Previously titled *Computer System Modeling Fundamentals*.

**CS 260: Machine Learning Theory**, UCLA, Fall 2010 and 2011

New graduate course that I designed in Fall 2010. Covers the theoretical foundations of machine learning, including the PAC model, online learning, SVMs, and boosting.

**CS 269: Mathematical Frameworks for Social Computing**, UCLA, Winter 2012  
Seminar-style course exploring theoretical models and algorithms for social computing.

### Courses TAed

**CSE 112: Networked Life**, University of Pennsylvania, Spring 2006

**CIS 520: AI and Machine Learning**, University of Pennsylvania, Fall 2005

### Additional Instruction/Mentorship

**Data Science Summer School**, Microsoft Research, New York City, Summer 2015  
Intensive eight-week hands-on introduction to data science for college students in the New York City area aimed at increasing diversity in computer science.

**Additional Guest Lectures**, Microsoft (many), Boston University (2020), New York University (2013), Columbia University (2013)

### Book Chapters

---

**A Human-Centered Agenda for Intelligible Machine Learning**. Jennifer Wortman Vaughan and Hanna Wallach. In *Machines We Trust: Perspectives on Dependable AI*, edited by Marcello Pelillo and Teresa Scantamburlo, MIT Press, 2021.

### Journal Articles

---

**AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap**. Q. Vera Liao and Jennifer Wortman Vaughan. *Harvard Data Science Review*, 2024 (to appear).

**An Equivalence Between Fair Division and Wagering Mechanisms**. Rupert Freeman, Jens Witkowski, Jennifer Wortman Vaughan, and David M. Pennock. *Management Science*, 2024 (to appear).

**Greedy Algorithm almost Dominates in Smoothed Contextual Bandits**. Manish Raghavan, Aleksandrs Slivkins, Jennifer Wortman Vaughan, and Zhiwei Steven Wu. *SIAM Journal of Computing*, 52(2): 487–524, 2023.

**Incentive-Compatible Forecasting Competitions**. Jens Witkowski, Rupert Freeman, Jennifer Wortman Vaughan, David M. Pennock, Andreas Krause. *Management Science*, 69(3): 1354–1374, 2023.

**Truthful Aggregation of Budget Proposals**. Rupert Freeman, David Pennock, Dominik Peters, and Jennifer Wortman Vaughan. *Journal of Economic Theory*, 193, 2021.

**Oracle-Efficient Online Learning and Auction Design**. Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. *Journal of the ACM*, 67(5), 2020.

**The Possibilities and Limitations of Private Prediction Markets**. Rachel Cummings, David M. Pennock, and Jennifer Wortman Vaughan. *ACM Transactions on Economics and Computation*, 8(3), 2020.

- Making Better Use of the Crowd.** Jennifer Wortman Vaughan. *Journal of Machine Learning Research*, 18(193):1–46, 2018.
- Integrating Market Makers, Limit Orders, and Continuous Trade in Prediction Markets.** Hoda Heidari, Sébastien Lahaie, David M. Pennock, and Jennifer Wortman Vaughan. *ACM Transactions on Economics and Computation (Special Issue on EC 2015)*, 6(3–4):Article 15, 2018.
- Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems.** Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. *Journal of Artificial Intelligence Research*, 55:317–359, 2016.
- Belief Aggregation with Automated Market Makers.** Rajiv Sethi and Jennifer Wortman Vaughan. *Computational Economics*, 48(1):155–178, 2016.
- An Axiomatic Characterization of Wagering Mechanisms.** Nicolas S. Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. *Journal of Economic Theory*, 156:389–416, 2015.
- Computational Social Science and Social Computing.** (Guest editorial) Winter Mason, Jennifer Wortman Vaughan, and Hanna Wallach. *Machine Learning Journal (Special Issue on Computational Social Science and Social Computing)*, 95(3):257–260, 2014.
- Efficient Market Making via Convex Optimization, and a Connection to Online Learning.** Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. *ACM Transactions on Economics and Computation*, 1(2):Article 12, 2013.
- A Theory of Learning from Different Domains.** Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. *Machine Learning Journal (Special Issue on Learning from Multiple Sources)*, 79(1–2):151–175, 2010.
- The True Sample Complexity of Active Learning.** Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman Vaughan. *Machine Learning Journal (Special Issue on COLT 2008)*, 80(2–3):111–139, 2010.
- Maintaining Equilibria During Exploration in Sponsored Search Auctions.** John Langford, Lihong Li, Yevgeniy Vorobeychik, and Jennifer Wortman. *Algorithmica (Special Issue on Internet Markets)*, 58(4): 990–1021, 2010.
- Behavioral Experiments on Biased Voting in Networks.** Michael Kearns, Stephen Judd, Jinsong Tan, and Jennifer Wortman. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.
- Learning from Multiple Sources.** Koby Crammer, Michael Kearns, and Jennifer Wortman. *Journal of Machine Learning Research*, 9:1757–1774, 2008.
- Regret to the Best Vs. Regret to the Average.** Eyal Even-Dar, Michael Kearns, Yishay Mansour, and Jennifer Wortman. *Machine Learning Journal (Special Issue on COLT 2007)*, 72(1–2):21–37, 2008.

## Conference Publications

---

- “I’m Not Sure, But...”: Examining the Impact of Large Language Models’ Uncertainty Expression on User Reliance and Trust.** Sunnie S. Y. Kim, Q. Vera Liao, Mihaela Vorvoreanu, Stephanie Ballard, and Jennifer Wortman Vaughan. *7th ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2024.
- Tinker, Tailor, Configure, Customize: The Articulation Work of Contextualizing AI Fairness Checklists.** Michael Madaio, Jingya Chen, Hanna Wallach, and Jennifer Wortman Vaughan. *27th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2024.
- Understanding the Role of Human Intuition on Reliance in Human-AI Decision-Making with Explanations.** Valerie Chen, Q. Vera Liao, Jennifer Wortman Vaughan, and Gagan Bansal. *26th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2023.
- Designerly Understanding: Information Needs for Model Transparency to Support Design Ideation for AI-Powered User Experience.** Q. Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. *2023 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2023.
- GAM Coach: Towards Interactive and User-Centered Algorithmic Recourse.** Zijie J. Wang, Jennifer Wortman Vaughan, Rich Caruana, and Duen Horng Chau. *2023 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2023.
- REAL ML: Recognizing, Exploring, and Articulating Limitations in Machine Learning Research.** Jessie J. Smith, Saleema Amershi, Solon Barocas, Hanna Wallach, and Jennifer Wortman Vaughan. *5th ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2022.
- Interpretability, Then What? Editing Machine Learning Models to Reflect Human Knowledge and Values.** Zijie J. Wang, Alex Kale, Harsha Nori, Peter Stella, Mark Nunnally, Duen Horng Chau, Mihaela Vorvoreanu, Jennifer Wortman Vaughan, and Rich Caruana. *28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2022.
- Understanding Machine Learning Practitioners’ Data Documentation Perceptions, Needs, Challenges, and Desiderata.** Amy K. Heger, Liz B. Marquis, Mihaela Vorvoreanu, Hanna Wallach, and Jennifer Wortman Vaughan. *25th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2022.
- Assessing the Fairness of AI Systems: AI Practitioners’ Processes, Challenges, and Needs for Support.** Michael Madaio, Lisa Egede, Hariharan Subramonyam, Jennifer Wortman Vaughan, and Hanna Wallach. *25th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2022.
- From Human Explanation to Model Interpretability: A Framework Based on Weight of Evidence.** David Alvarez-Melis, Harmanpreet Kaur, Hal Daumé III, Hanna Wallach, and Jennifer Wortman Vaughan. *9th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2021.

- Designing Disaggregated Evaluations of AI Systems: Choices, Considerations, and Tradeoffs.** Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Kronen, Meredith Ringel Morris, Jennifer Wortman Vaughan, Duncan Wadsworth, and Hanna Wallach. *4th AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 2021.
- Manipulating and Measuring Model Interpretability.** Forough Poursabzi-Sangdeh, Daniel G. Goldstein, Jake M. Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. *2021 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2021.
- No-Regret and Incentive-Compatible Online Learning.** Rupert Freeman, David M. Pennock, Chara Podimata, and Jennifer Wortman Vaughan. *Thirty-seventh International Conference on Machine Learning (ICML)*, 2020.
- Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning.** Harmanpreet Kaur, Harsha Nori, Samuel Jenkins, Rich Caruana, Hanna Wallach, and Jennifer Wortman Vaughan. *2020 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2020.
- Honorable mention.**
- Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI.** Michael Madaio, Luke Stark, Hanna Wallach, and Jennifer Wortman Vaughan. *2020 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2020.
- Best paper award.**
- Truthful Aggregation of Budget Proposals.** Rupert Freeman, David Pennock, Dominik Peters, and Jennifer Wortman Vaughan. *Twentieth ACM Conference on Economics and Computation (EC)*, 2019.
- Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?** Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miroslav Dudík, and Hanna Wallach. *2019 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2019.
- Understanding the Effect of Accuracy on Trust in Machine Learning Models.** Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. *2019 ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2019.
- Honorable mention.**
- Using Search Queries to Understand Health Information Needs in Africa.** Rediet Abebe, Shawndra Hill, Jennifer Wortman Vaughan, Peter M. Small, and H. Andrew Schwartz. *Thirteenth International AAAI Conference on Web and Social Media (ICWSM)*, 2019.
- Group Fairness for Indivisible Goods Allocation.** Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. *Thirty-third AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- An Equivalence Between Wagering and Fair-Division Mechanisms.** Rupert Freeman, David Pennock, and Jennifer Wortman Vaughan. *Thirty-third AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

**The Disparate Effects of Strategic Manipulation.** Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. *Second ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2019.

**The Externalities of Exploration and How Data Diversity Helps Exploitation.** Manish Raghavan, Aleksandrs Slivkins, Jennifer Wortman Vaughan, and Zhiwei Steven Wu. *Thirty-first Annual Conference on Learning Theory (COLT)*, 2018.

**Incentive-Compatible Forecasting Competitions.** Jens Witkowski, Rupert Freeman, Jennifer Wortman Vaughan, David Pennock, and Andreas Krause. *Thirty-second AAAI Conference on Artificial Intelligence (AAAI)*, 2018.

**A Decomposition of Forecast Error in Prediction Markets.** Miroslav Dudík, Sébastien Lahaie, Ryan Rogers, and Jennifer Wortman Vaughan. *Advances in Neural Information Processing Systems 30 (NeurIPS)*, 2017.

**Oracle-Efficient Learning and Auction Design.** Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. *Fifty-eighth Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2017.

**The Double Clinching Auction for Wagering.** Rupert Freeman, David M. Pennock, and Jennifer Wortman Vaughan. *Eighteenth ACM Conference on Economics and Computation (EC)*, 2017.

**Bounded Rationality in Wagering Mechanisms.** David M. Pennock, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. *Thirty-second Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.

**The Possibilities and Limitations of Private Prediction Markets.** Rachel Cummings, David M. Pennock, and Jennifer Wortman Vaughan. *Seventeenth ACM Conference on Economics and Computation (EC)*, 2016.

**The Communication Network Within the Crowd.** Ming Yin, Mary Gray, Siddharth Suri, and Jennifer Wortman Vaughan. *Twenty-fifth International World Wide Web Conference (WWW)*, 2016.

**Integrating Market Makers, Limit Orders, and Continuous Trade in Prediction Markets.** Hoda Heidari, Sébastien Lahaie, David Pennock, and Jennifer Wortman Vaughan. *Sixteenth ACM Conference on Economics and Computation (EC)*, 2015.

**Incentivizing High Quality Crowdsourcing.** Chien-Ju Ho, Aleksandrs Slivkins, Siddharth Suri, and Jennifer Wortman Vaughan. *Twenty-fourth International World Wide Web Conference (WWW)*, 2015.

**Nominee for the Best Paper Award (8 papers out of 929 submissions).**

**Market Makers with Decreasing Utility for Information.** Miroslav Dudík, Rafael Frongillo, and Jennifer Wortman Vaughan. *Thirtieth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.

**A General Volume-Parameterized Market Making Framework.** Jacob Abernethy, Rafael Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. *Fifteenth ACM Conference on Economics and Computation (EC)*, 2014.

- Removing Arbitrage from Wagering Mechanisms.** Yiling Chen, Nikhil R. Devanur, David Pennock, and Jennifer Wortman Vaughan. *Fifteenth ACM Conference on Economics and Computation (EC)*, 2014.
- Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems.** Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. *Fifteenth ACM Conference on Economics and Computation (EC)*, 2014.
- An Axiomatic Characterization of Adaptive-Liquidity Market Makers.** Xiaolong Li and Jennifer Wortman Vaughan. *Fourteenth ACM Conference on Electronic Commerce (EC)*, 2013.
- Cost Function Market Makers for Measurable Spaces.** Yiling Chen, Michael Ruberry, and Jennifer Wortman Vaughan. *Fourteenth ACM Conference on Electronic Commerce (EC)*, 2013.
- Adaptive Task Assignment for Crowdsourced Classification.** Chien-Ju Ho, Shahin Jabbari, and Jennifer Wortman Vaughan. *Thirtieth International Conference on Machine Learning (ICML)*, 2013.
- Designing Informative Securities.** Yiling Chen, Mike Ruberry, and Jennifer Wortman Vaughan. *Twenty-eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.
- Online Task Assignment in Crowdsourcing Markets.** Chien-Ju Ho and Jennifer Wortman Vaughan. *Twenty-sixth AAAI Conference on Artificial Intelligence (AAAI)*, 2012.
- An Optimization-Based Framework for Automated Market-Making.** Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. *Twelfth ACM Conference on Electronic Commerce (EC)*, 2011.
- Evolution with Drifting Targets.** Varun Kanade, Leslie G. Valiant, and Jennifer Wortman Vaughan. *Twenty-third Annual Conference on Learning Theory (COLT)*, 2010.
- Regret Minimization with Concept Drift.** Koby Crammer, Eyal Even-Dar, Yishay Mansour, and Jennifer Wortman Vaughan. *Twenty-third Annual Conference on Learning Theory (COLT)*, 2010.
- A New Understanding of Prediction Markets Via No-Regret Learning.** Yiling Chen and Jennifer Wortman Vaughan. *Eleventh ACM Conference on Electronic Commerce (EC)*, 2010.
- Censored Exploration and the Dark Pool Problem.** Kuzman Ganchev, Michael Kearns, Yuriy Nevmyvaka, and Jennifer Wortman Vaughan. *Twenty-fifth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009.
- Winner of a Best Student Paper Award.**
- Complexity of Combinatorial Market Makers.** Yiling Chen, Lance Fortnow, Nicolas Lambert, David Pennock, and Jennifer Wortman. *Ninth ACM Conference on Electronic Commerce (EC)*, 2008.



**Self-Financed Wagering Mechanisms for Forecasting.** Nicolas Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David Pennock. *Ninth ACM Conference on Electronic Commerce (EC)*, 2008.

**Winner of an Outstanding Paper Award.**

**Learning from Collective Behavior.** Michael Kearns and Jennifer Wortman. *Twenty-first Annual Conference on Learning Theory (COLT)*, 2008.

**The True Sample Complexity of Active Learning.** Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman. *Twenty-first Annual Conference on Learning Theory (COLT)*, 2008.

**Winner of a Best Student Paper Award.**

**Exploration Scavenging.** John Langford, Alexander Strehl, and Jennifer Wortman. *Twenty-fifth International Conference on Machine Learning (ICML)*, 2008.

**Privacy-Preserving Belief Propagation and Sampling.** Michael Kearns, Jinsong Tan, and Jennifer Wortman. *Advances in Neural Information Processing Systems 20 (NeurIPS)*, 2007.

**Learning Bounds for Domain Adaptation.** John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. *Advances in Neural Information Processing Systems 20 (NeurIPS)*, 2007.

**Maintaining Equilibria During Exploration in Sponsored Search Auctions.** Jennifer Wortman, Yevgeniy Vorobeychik, Lihong Li, and John Langford. *Third International Workshop on Internet and Network Economics (WINE)*, 2007.

**Sponsored Search with Contexts.** Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. *Third International Workshop on Internet and Network Economics (WINE)*, 2007.

**Regret to the Best Vs. Regret to the Average.** Eyal Even-Dar, Michael Kearns, Yishay Mansour, and Jennifer Wortman. *Twentieth Annual Conference on Learning Theory (COLT)*, 2007.

**Winner of a Best Student Paper Award.**

**Learning from Multiple Sources.** Koby Crammer, Michael Kearns, and Jennifer Wortman. *Advances in Neural Information Processing Systems 19 (NeurIPS)*, 2006.

**Risk-Sensitive Online Learning.** Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. *Seventeenth International Conference on Algorithmic Learning Theory (ALT)*, 2006.

**Learning from Data of Variable Quality.** Koby Crammer, Michael Kearns, and Jennifer Wortman. *Advances in Neural Information Processing Systems 18 (NeurIPS)*, 2005.

**Run the GAMUT: A Comprehensive Approach to Evaluating Game-Theoretic Algorithms.** Eugene Nudelman, Jennifer Wortman, Yoav Shoham, and Kevin Leyton-Brown. *Third International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2004.

## Other Published Articles

---

- The CLeAR Documentation Framework for AI Transparency: Recommendations for Practitioners & Context for Policymakers.** Kasia Chmielinski, Sarah Newman, Chris N. Kranzinger, Michael Hind, Jennifer Wortman Vaughan, Margaret Mitchell, Julia Stoyanovich, Angelina McMillan-Major, Emily McReynolds, Kathleen Esfahany, Mary L. Gray, Audrey Chang, and Maui Hudson. *Harvard Kennedy School Shorenstein Center Discussion Paper*, April 2024.
- Summarize with Caution: Comparing Global Feature Attributions.** Alex Okeson, Rich Caruana, Nick Craswell, Kori Inkpen, Scott M. Lundberg, Harsha Nori, Hanna Wallach, and Jennifer Wortman Vaughan. *IEEE Data Engineering Bulletin*, December 2021.
- Datasheets for Datasets.** Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. *Communications of the ACM*, 64(12):86–92, December 2021.
- Responsible Computing During COVID-19 and Beyond.** Solon Barocas, Asia J. Biega, Margarita Boyarskaya, Kate Crawford, Hal Daumé III, Miroslav Dudík, Benjamin Fish, Mary L. Gray, Brent Hecht, Alexandra Olteanu, Forough Poursabzi-Sangdeh, Luke Stark, Jennifer Wortman Vaughan, Hanna Wallach, Marion Zepf. *Communications of the ACM*, 64(7): 30–32, July 2021.
- Toward Fairness in AI for People with Disabilities: A Research Roadmap.** Anhong Guo, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach and Meredith Ringel Morris. *ACM SIGACCESS Newsletter*, 125, October 2019.
- Incentives and the Crowd.** Jennifer Wortman Vaughan. *ACM XRDS (Crossroads)*, 24(1):42–46, September 2017.
- Mathematical Foundations of Social Computing.** Yiling Chen, Arpita Ghosh, Michael Kearns, Tim Roughgarden, and Jennifer Wortman Vaughan. *Communications of the ACM*, 59(12):102–108, December 2016.
- Incentivizing High Quality Crowdwork.** Chien-Ju Ho, Aleksandrs Slivkins, Siddharth Suri, and Jennifer Wortman Vaughan. *ACM SIGecom Exchanges*, 14(2), December 2015. (Based on the WWW 2015 paper.)
- Online Decision Making in Crowdsourcing Markets: Theoretical Challenges.** Aleksandrs Slivkins and Jennifer Wortman Vaughan. *ACM SIGecom Exchanges*, 12(2), December 2013.
- Connections Between Markets and Learning.** Yiling Chen and Jennifer Wortman Vaughan. *ACM SIGecom Exchanges*, 9(1), June 2010. (Based on the EC 2010 paper.)
- Censored Exploration and the Dark Pool Problem.** Kuzman Ganchev, Michael Kearns, Yuriy Nevmyvaka, and Jennifer Wortman Vaughan. *Communications of the ACM*, Research Highlights, May 2010. (Invited, based on the UAI 2009 paper.)

## Selected Workshop Papers

---

My research has appeared in many lightly peer-reviewed, non-archival workshops including:

NeurIPS Workshop on Human-Centered AI, December 2022  
ICML DataPerf Workshop on Benchmarking Data for Data-Centric AI, July 2022  
NeurIPS Workshop on Human-Centered AI, December 2021  
NeurIPS Workshop on Bridging the Gap: From ML Research to Clinical Practice, Dec 2021  
3rd KDD Workshop on Data Science with Humans in the Loop, August 2021  
CHI 2021 Human-Centered XAI Workshop, May 2021  
CSCW 2021 Workshop on Beyond Checklist Approaches to Ethics in Design, October 2020  
2nd AAMAS Games, Agents, and Incentives Workshop, May 2020  
CHI 2020 Fair & Responsible AI Workshop, April 2020  
FAccT 2020 CRAFT on Bridging the Gap from AI Ethics Research to Practice, January 2020  
NeurIPS 2019 Workshop on Human-Centric Machine Learning, December 2019  
NeurIPS 2019 Workshop on Machine Learning with Guarantees, December 2019  
ASSETS 2019 Workshop on AI Fairness for People with Disabilities, October 2019  
5th International Conference on Computational Social Science (IC2S2), July 2019  
3rd EC Workshop on Mechanism Design for Social Good (MD4SG), June 2019  
AAMAS 2019 Workshop on Fair Allocation in Multiagent Systems, May 2019  
NeurIPS 2018 Workshop on Critiquing and Correcting Trends in ML, December 2018  
ICML/IJCAI Workshop on Human Interpretability in Machine Learning, July 2018  
Fairness, Accountability, and Transparency in Machine Learning (FATML), July 2018  
EC Workshop on Opinion Aggregation, Dynamics, and Elicitation, June 2018  
NeurIPS Workshop on Machine Learning for the Developing World, December 2017  
NeurIPS Workshop on Transparent & Interpretable ML in Safety Critical Envs., Dec. 2017  
NeurIPS Workshop on Prioritising Online Content, December 2017  
3rd EC Workshop on Algorithmic Game Theory and Data Science, June 2017  
CHI Workshop on Designing for Uncertainty in HCI, May 2017  
Conference on Digital Experimentation @ MIT, October 2016  
2nd Annual International Conference on Computational Social Science (IC2S2), June 2016  
ICML Workshop on Theory and Practice of Differential Privacy, June 2016  
Conference on Digital Experimentation @ MIT, October 2015  
5th EC Workshop on Social Computing and User Generated Content, June 2015  
NeurIPS Workshop on Crowdsourcing: Theory, Algorithms, and Applications, Dec. 2013  
NSF/NBER/CEME Seminar in Math. Economics and General Equilibrium, October 2013  
3rd EC Workshop on Social Computing and User Generated Content, June 2013  
AAAI Fourth Human Computation Workshop (HCOMP), July 2012  
ICML Workshop on Markets, Mechanisms, and Multi-agent Models, June 2012  
NeurIPS Workshop on Domain Adaptation, December 2011  
NeurIPS Workshop on Comp. Social Science and the Wisdom of Crowds, December 2010  
NeurIPS Workshop on Advances in Ranking, December 2009  
NeurIPS Workshop on Principles of Learning Problem Design, December 2007  
DIMACS Workshop on the Boundary between Econ. Theory and CS, October 2007  
Third Workshop on Sponsored Search Auctions at WWW, May 2007  
NeurIPS Workshop on On-line Trading of Exploration and Exploitation, December 2006  
15th International Conference on Game Theory at Stony Brook, July 2004  
Many iterations of the New York Computer Science and Economics (NYCE) Day  
Many iterations of the New York Academy of Sciences Machine Learning Symposium

## **Selected Keynotes and Talks**

---

### **Human-Centered Approaches to AI Transparency (And Two Studies on Uncertainty Communication)**

NYU Tandon Center for Responsible AI, April 2024

Microsoft Research AI & Society seminar series, March 2024

### **Some Very Human Challenges in Responsible AI (Or Why My Research Trajectory Took a Surprising Turn)**

University of Washington RAISE Seminar, March 2023

10th AAAI Conf. on Human Computation and Crowdsourcing (Keynote), November 2022

University of Chicago Human + AI Conference, October 2022

### **Intelligibility Throughout the Machine Learning Life Cycle**

Rutgers Workshop on Democratizing AI, April 2022

HRI Workshop on Fairness and Transparency in Human-Robot Interaction, March 2022

AT&T Data Science Summit, November 2021

UbiComp Workshop on Reviewable and Auditable Pervasive Systems, September 2021

Microsoft Data Science & Law Forum 3.0, April 2021

UC Santa Barbara Center for Responsible ML Distinguished Lecture, January 2021

NeurIPS 2020 Workshop on Human and Model in the Loop Evaluation and Training Strategies (HAMLETS), December 2020

Trustworthy ML Initiative Seminar, December 2020

Max Planck Institute for Software Systems Distinguished Lecture, November 2020

University of Michigan Data Science/Comp. Social Science Seminar, November 2020

Samsung AI Forum, November 2020

Boston University Course on Fairness, Accountability, Transparency in AI, October 2020

University of Bath Centre for Doctoral Training in Accountable, Responsible and Transparent AI seminar, October 2020

Open Data Science Conference (ODSC) West, October 2020

IEEE VIS Workshop on Vis Trust and Experience in Visual Analytics, October 2020

CSAIL-MSR Trustworthy and Robust AI (TRAC) Workshop, November 2019

CHI Workshop Where is the Human? Bridging the Gap Between AI and HCI, May 2019

### **AI Fairness in Industry: From Principles to Practice**

The 17th Conference on Web and Internet Economics (WINE) (Keynote), December 2021

Joint Statistical Meetings (JSM) session on Data and Algorithmic Bias, August 2020

### **Why Is Fair Machine Learning Hard and How Can Theory Help?**

30th International Conf. on Algorithmic Learning Theory (ALT) (Keynote), March 2019

UPenn Warren Center Theory of Computer Science Talk Series, March 2019

### **Fairness and Intelligibility in Machine Learning Systems**

FTC Hearing on Algorithms, AI, and Predictive Analytics (Keynote), November 2018

### **Uncovering New Challenges in Fair and Interpretable Machine Learning**

Google Research Seminar, New York, NY, May 2018

### **Manipulating and Measuring Model Interpretability**

NYU Tandon Machine Learning Reading Group, Brooklyn, NY, October 2018

Microsoft Design+AI Seminar Series, Redmond, WA, October 2018

NeurIPS Interpretable Machine Learning Symposium, Long Beach, CA, Dec. 2017

### **Building Fair and Transparent AI (And Why You Should Care)**

Institute for Human Rights and Business briefing event, New York, NY, December 2017

### **Nine Things I Wish I Had Known the First Time I Came to NeurIPS**

12th Annual Women in Machine Learning Workshop (*Opening address*), Long Beach, CA, December 2017

### **The Role of Human Computation in AI, The Human Components of Machine Learning, or other similar variants**

HRL Colloquium, Malibu, CA, January 2018

NYU NLP and Text as Data Speaker Series, New York, NY, November 2007

Spotify Research Seminar, New York, NY, November 2007

Yale Data Science Workshop on Comp. Social Science, New Haven, CT, October 2017

Broadening Participation in Data Mining, Halifax, NS, August 2017

Microsoft Research NYC 5th Anniversary Celebration, New York, NY, May 2017

### **Self-Financed Wagering Mechanisms: What's Been Done and What's to Come**

EC Workshop on Forecasting, Cambridge, MA, June 2017

### **Crowd Behavior and Implications on Research**

CU Boulder Computer Science Colloquium, Boulder, CO, May 2017

### **The Communication Network Within the Crowd**

NeurIPS Workshop on Crowdsourcing and ML, Barcelona, Spain, December 2016

Bloomberg Data Science Seminar, New York, NY, August 2016

### **The Past, Present, and Future of Women in Machine Learning**

10th Annual Women in Machine Learning Workshop (*Opening address*), joint with Hanna Wallach and Amy Greenwald, Montreal, QC, December 2015

### **Crowdsourcing Your Data**

Strata session on Hardcore Data Science, New York, NY, September 2015

### **Integrating Market Makers, Limit Orders, and Continuous Trade in Prediction Markets**

International Symposium on Mathematical Programming, Pittsburgh, PA, July 2015

### **Incentivizing High Quality Crowdwork**

NYU Machine Learning Seminar, New York, NY, February 2015

NeurIPS Workshop on Crowdsourcing and ML, Montreal, QC, December 2014

### **Market Making with Decreasing Value of Information**

NeurIPS Workshop on Trans. ML and E-Commerce, Montreal, QC, December 2014

CMU-MSR Mindswap on the Interface Between CS and Econ, New York, NY, April 2014

### **An Optimization-Based Framework for Automated Market-Making, Combinatorial Prediction Markets via Convex Cost Functions, or other variants**

New York Area Theory Day (*Invited*), Columbia, New York, NY, April 2015

Duke Machine Learning Seminar, Durham, NC, October 2014

Yahoo! Research, New York, NY, May 2014

NYU Courant Machine Learning Seminar, New York, NY, November 2012  
LA Machine Learning Meetup, hosted by eHarmony, Santa Monica, CA, July 2012  
Microsoft Research, New York, NY, July 2012  
LogicBlox, Atlanta, GA, June 2012  
Google Research, New York, NY, May 2012  
Symantec Research Labs, Culver City, CA, April 2012  
SoCal Symposium on Network Econ and Game Theory, Pasadena, CA, November 2011  
Microsoft Research Algorithms and Game Theory Seminar, Cambridge, UK, June 2011  
Caltech Social and Information Sciences Lab Seminar, Pasadena, CA, February 2011  
University of Southern California, CS Colloquium, Los Angeles, CA, January 2011  
Caltech Rigorous Systems Research Group Seminar, Pasadena, CA, November 2010

### **Aggregating Human Predictions Via Markets**

AAAI 3rd Human Computation Workshop, San Francisco, CA, August 2011

### **A New Understanding of Prediction Markets Via No-Regret Learning**

Harvard Center for Research on Computation and Society, Cambridge, MA, April 2010  
UC Berkeley, Berkeley, CA, April 2010

### **Censored Exploration and the Dark Pool Problem**

UCLA, Statistics Department Seminar, Los Angeles, CA, February 2011  
Boston University, Computer Science Colloquium, Boston, MA, February 2010  
UMass, Amherst, Machine Learning Lunch, Amherst, MA, December 2009  
Brown University, Machine Learning Reading Group, Providence, RI, November 2009  
UMass, Boston, Computer Science Colloquium, Boston, MA, November 2009  
Microsoft Research New England, Econ Reading Group, Cambridge, MA, October 2009  
Harvard University, EconCS Seminar, Cambridge, MA, October 2009

### **Learning from Collective Preferences, Behavior, and Beliefs**

Harvard University, Theory of Computation/EconCS Seminar, Cambridge, MA, May 2009  
UCLA, Computer Science Colloquium, Los Angeles, CA, April 2009  
University of Michigan, CSE Colloquium, Ann Arbor, MI, April 2009  
New York University, Computer Science Colloquium, New York, NY, April 2009  
Toyota Technological Institute at Chicago, Chicago, IL, March 2009  
Cornell University, Computer Science Colloquium, Ithaca, NY, March 2009  
Yahoo! Research, Santa Clara, CA, March 2009  
UC Berkeley, Center for Intelligence Systems Seminar, Berkeley, CA, February 2009  
Duke University, Computer Science Colloquium, Durham, NC, February 2009  
Microsoft Research New England, Cambridge, MA, January 2009

### **Learning from Collective Behavior**

Rutgers University, Computer Science Colloquium, Piscataway, NJ, March 2008

### **Regret to the Best Vs. Regret to the Average**

Yahoo! Research, New York, NY, January 2007

## **Selected Panels and Discussions**

---

Panelist, NeurIPS 2023 tutorial on Data-Centric AI for Reliable and Responsible AI: From Theory to Practice, December 2023

Panelist, National Academies Meeting of the Committee on Understanding and Addressing Misinformation about Science panel on “Advancements in Information Technology and Implications for Addressing Misinformation,” July 2023

Panelist, NIST Responsible AI Research Round Table, July 2023

Panelist, WeCREATE Inspiration Session on Responsible AI, May 2023

Panelist, NeurIPS 2022 tutorial on Incentive-Aware Machine Learning, December 2022

Panel Moderator, “Is a Human in the Loop the Solution?” NIST Workshop on Building the NIST AI Risk Management Framework, May 2022

Panelist, “Artificial Intelligence and Democratic Values,” White House Office of Science and Technology Policy’s Public Event on the Bill of Rights for Automated Society, November 2021

Panelist, “Practical Challenges to Recent Advances in Machine Learning,” Federal Reserve Conference on Artificial Intelligence in Consumer Finance: Defining and Insuring Fairness, November 2021

Panelist, “Responsible AI Systems and Experiences” Round Table, 47th International Conference on Very Large Data Bases (VLDB), August 2021

Panelist, “Attributes of AI Trustworthiness: Fair, Private, Contestable,” National Academies of Science, Engineering, and Medicine Workshop on Assessing and Improving AI Trustworthiness: Current Contexts, Potential Paths, March 2021

Keynote Discussant, ACM Conference on Fairness, Accountability, and Transparency (FAccT), January 2019

Panelist, “The Ghost in the Smart Machine,” The Economist’s Innovation Summit, March 2018

Panelist, “AI and Machine Learning 101,” AI: The Economic and Policy Implications, Technology Policy Institute, September 2016

## **Awards and Honors**

---

**Presidential Early Career Award for Scientists and Engineers (PECASE)**, “the highest honor bestowed by the United States Government on science and engineering professionals in the early stages of their independent research careers,” 2012

**Symantec Term Chair in Computer Science**, UCLA, 2011–2015

**NSF Faculty Early Career Development (CAREER) Award**, 2011–2014

**NSF Computing Innovation Fellowship**, 2009–2010

**Morris & Dorothy Rubinoff Award (Co-winner)**, presented to “a graduate degree candidate whose dissertation has resulted in or could lead to innovative applications of computer technology,” University of Pennsylvania, 2009

### **Paper Awards**

Several **Best Paper** and **Honorable Mention** awards at CHI

**Best Paper Award**, NeurIPS Workshop on Bridging the Gap: From Machine Learning Research to Clinical Practice, 2021

**Nominee, Best Paper Award** (8 papers out of 929 submissions), 24th International World Wide Web Conference, 2015

**Best Student Paper Award**, 25th Conference on Uncertainty in Artificial Intelligence, 2009

**Student Paper Award, Second Place**, New York Academy of Sciences Third Annual Symposium on Machine Learning, 2008

**Outstanding Paper Award**, Ninth ACM Conference on Electronic Commerce, 2008

**Best Student Paper Award**, 21st Annual Conference on Learning Theory, 2008

**Student Paper Award, First Place**, New York Academy of Sciences Second Annual Symposium on Machine Learning, 2007

**Best Student Paper Award**, 20th Annual Conference on Learning Theory, 2007

**Phi Beta Kappa**, Boston University, 2002

**Academic Achievement Award in Computer Science**, Boston University, 2002

**Boston Scholars Scholarship**, Merit-based full-tuition scholarship, 1998–2002

### **Selected Funding**

---

#### **For Research Projects**

**IARPA Hybrid Forecast Competition (HRL MATRICS team)**, Role: Advisor, 2017–2018

**CAREER: Learning- and Incentives-Based Techniques for Aggregating Community-Generated Data**, National Science Foundation, IIS-1054911, Role: PI, 2011–2015

**CIFellows Project: Crowdsourcing for Science Education**, Computing Research Association, CIF-B-17, Role: PI/Mentor for postdoctoral fellow Ricky Sethi, 2010–2011

**CIFellowship: Foundational Understanding of Learning from the Collective and Problems in Learning and Reasoning**, Computing Research Association, CIF-246, Role: Postdoc fellow with mentors Yiling Chen and Leslie Valiant (Harvard), 2009–2010

#### **For Events**

**Theoretical Foundations of Social Computing**, CCC visioning workshops program, Role: Co-organizer with Yiling Chen, Arpita Ghosh, and Tim Roughgarden, 2015

**Collaborative Research: Workshop for Women in Machine Learning**, National Science Foundation, IIS-1036868, Role: PI with co-PI Hanna Wallach, 2010–2014



## Patents Filed

---

**Computer-Based Data Collection Using a Prediction Market with a Liquidity Reducing Cost Function.** With Miro Dudík and Rafael Frongillo. US patent filed November 2014.

**Computer System for Multiple User, Multiple Event Real-Time Online Wagering.** With Yiling Chen, Nikhil Devanur, and David Pennock. US Patent #9947174, filed October 2014, granted April 2018.

**System and Method for Automated Market Making.** With Jake Abernethy and Yiling Chen. US patent filed May 2012.

## Students and Interns Supervised

---

### Doctoral Dissertations Supervised

Chien-Ju Ho, UCLA, 2010–2015, winner of the Google Outstanding Graduate Research Award at UCLA (next position: postdoc at Cornell; now an Assistant Professor in Computer Science at Washington University in St. Louis)

### Other Students and Postdocs Supervised at UCLA

Garret Buell, MS received 2012 (next position: Software Engineer at Google)  
Xinlei Chen, Cross-disciplinary Scholars in Science and Technology summer intern, 2011 (next position upon graduation: doctoral student at CMU)  
Shahin Jabbari, doctoral student, 2011–2013 (next position: doctoral student at UPenn)  
Xiaolong Li, MS received 2012 (next position: doctoral student at UT Austin)  
Ricky Sethi, Computing Innovation Fellow, 2010–2011 (next position: postdoc at USC ISI)  
Petch Wannissorn, MS received Winter 2012 (next position: MS student in Software Management at CMU)

### Interns Mentored or Co-Mentored at Microsoft Research

Wesley Hanwen Deng, Carnegie Mellon, 2023  
Kevin Feng, University of Washington, 2023  
Sunnie S. Y. Kim, Princeton University, 2023  
Zana Buçinca, Harvard University, 2022  
Valerie Chen, Carnegie Mellon University, 2022  
Priyanka Nanayakkara, Northwestern University, 2022  
Helena Vasconcelos, Stanford University, 2022  
Alex Kale, University of Washington, 2021  
Mahsan Nourani, University of Florida, 2021  
Jessie Smith, University of Colorado, Boulder, 2021  
Zijie Jay Wang, Georgia Tech, 2021  
Andrew Anderson, Oregon State University, 2020  
Lisa Egede, University of Oklahoma, 2020  
Liz Marquis, University of Michigan, 2020  
Alex Okeson, University of Washington, 2020  
Hari Subramonyam, University of Michigan, 2020  
Meg Young, University of Washington, 2020  
Marah Abdin, University of Rochester, 2019  
Anhong Guo, Carnegie Mellon, 2019  
Harmanpreet Kaur, University of Michigan, 2019

Michael Madaio, Carnegie Mellon, 2019  
Samir Passi, Cornell University, 2019  
Chara Podimata, Harvard University, 2019  
Samira Samadi, Georgia Tech, 2019  
Angela Zhou, Cornell Tech, 2019  
David Alvarez Melis, MIT, 2018  
Ken Holstein, Carnegie Mellon, 2018  
Lily Hu, Harvard University, 2018  
Rediet Abebe, Cornell, 2017  
Forough Poursabzi-Sangdeh, CU Boulder, 2017  
Manish Raghavan, Cornell, 2017  
Rupert Freeman, Duke University, 2016 and 2017  
Nika Haghtalab, Carnegie Mellon University, 2016  
Ryan Rogers, University of Pennsylvania, 2016  
Rachel Cummings, California Institute of Technology, 2015  
Ming Yin, Harvard University, 2015  
Hoda Heidari, University of Pennsylvania, 2014  
Chien-Ju Ho, UCLA, 2013 and 2014  
Alice Gao, Harvard University, 2013

#### **Other Ph.D. Committees**

**UCLA Dissertation Committees:** Nick Mastronarde (EE, 2011), Mahsan Rofouei (CS, 2012), Vidyut Samanta (CS, 2012), Michael Shindler (CS, 2011)

**UCLA Qualifying Exam Committees:** Dawn Chen (Psychology, 2012), Jihyoung Kim (CS, 2011), Brent Longstaff (CS, 2012), Bobak Mortazavi (CS, 2012), Roozbeh Mottaghi (CS, 2010), Mahsan Rofouei (CS, 2011), Vidyut Samanta (CS, 2011), Wen-Yun Yang (Bioinformatics, 2011)

**Elsewhere:** Debajyoti Ray (qualifying exam, CalTech, 2011), Michael Ruberry (Harvard, 2013), Chara Podimata (Harvard, 2022), Harmanpreet Kaur (University of Michigan, 2023), Mahsan Nourani (University of Florida, 2023), Nina Grgić-Hlača (Max Planck Institute, in progress), Jessie Smith (CU Boulder, in progress)

## **Other Professional Activities and Academic Service**

---

#### **Program Chair Roles**

**Program Co-chair**, NeurIPS 2021

**Program Co-chair and General Co-chair**, HCOMP 2019

#### **Other Conference Chair Roles**

**Workshops Chair**, NeurIPS 2019

**Press Co-chair**, ICML 2019

**Tutorials Co-chair**, NeurIPS 2017 and 2018 (reappointed for second year)

**Workshops Co-chair**, ACM EC 2017 and 2018 (two-year term)

**Publicity Co-chair**, HCOMP 2017

**Workshop Selection Committee**, NeurIPS 2016, 2020, 2022

**Local Arrangements Co-chair**, ACM EC 2010

## Conference Steering/Executive Committees

**Steering Committee**, FAccT (2017–2023)

**Steering Committee**, HCOMP (2019–2020)

**Steering Committee**, WINE (2018–2019)

**Secretary-Treasurer**, ACM Special Interest Group on Economics and Computation (SIGecom), elected position, 2015–2019

## Conference Program Committees

**Senior/Top Level Program Committee Member/Area Chair or equivalent**, AAAI (2013), ACM EC (2015, 2016), COLT (2010, 2011), FAccT (2021, 2023, 2024), HCOMP (2015, 2016), ICML (2012), IJCAI (2015), NeurIPS (2013, 2014, 2016), WWW (Track Co-chair, 2018), UAI (2012)

**General Program Committee Member/Formal Reviewer or equivalent**, AAAI (2008), ACM EC (2010, 2011, 2013), AISTATS (2009, 2012), CHI (2023), FAccT (2019), FORC (2020), IC2S2 (2017), ICML (2007, 2008, 2009), IJCAI (2005, 2009), NeurIPS (2008, 2009, 2015), UAI (2008), WWW (2016, 2017)

**Ethics Reviewer**, NeurIPS (2022)

## Workshop Organization and Committees

**TTIC Summer Workshop on Learning in the Presence of Strategic Behavior** (2018)  
Co-organized with Nika Haghtalab, Yishay Mansour, Tim Roughgarden, Vasilis Syrgkanis

**NeurIPS Workshop on Learning in the Presence of Strategic Behavior** (2017)  
Co-organized with Nika Haghtalab, Yishay Mansour, Tim Roughgarden, Vasilis Syrgkanis

**HCOMP Workshop on Mathematical Foundations of Human Computation** (2016)  
Co-organized with Shuchi Chawla, Chien-Ju Ho, Michael Kearns, and Santosh Vempala

**CCC Visioning Workshop on Theoretical Foundations of Social Computing** (2015)  
Co-organized with Yiling Chen, Arpita Ghosh, and Tim Roughgarden

**6th Annual New York Computer Science and Economics (NYCE) Day** (2013)  
Co-organized with Jason Hartline and Vahab Mirrokni

**NeurIPS Workshop on Crowdsourcing: Theory, Applications, and Algorithms** (2013)  
Co-organized with Xi Chen, Nikhil Devanur, Alexander Ihler, Qiang Liu, Dengyong Zhou

**ICML Workshop on Markets, Mechanisms, and Multi-Agent Models** (2012)  
Co-organized with Amos Storkey and Jake Abernethy

**2nd NeurIPS Workshop on Comp. Social Science and the Wisdom of Crowds** (2011)  
Co-organized with Winter Mason and Hanna Wallach

**NeurIPS Workshop on Relations Between Machine Learning Problems** (2011)  
Co-organized with Bob Williamson, John Langford, Ulrike von Luxburg, and Mark Reid

**NeurIPS Workshop on Computational Social Science and the Wisdom of Crowds** (2010)  
Co-organized with Hanna Wallach

**First Workshop for Women in Machine Learning (WiML)** (2006)  
Co-organized with Hanna Wallach and Lisa Wainer

**Other Workshop Program Committees:** AAAI Human Computation Workshop (2011, 2012), Black in AI (2018), EC Workshops on Social Computing and User Generated Content (2011, 2012, 2015), EC/WWW Workshop on Crowdsourcing and Online Behavioral Experiments (2015, 2016), Northeast Student Colloquium on AI (2007), SIGAI Career

Network Conference (2015), EC Workshop on Opinion Aggregation, Dynamics, and Elicitation (2018), Women in Machine Learning Workshop (several years), NeurIPS Human-Centered AI Workshop (2022), CHI Workshop on Trust and Reliance in AI-Human Teams/Workflows (2022, 2023, 2024), IJCAI Explainable AI Workshop (2024)

### **Journal Boards and Reviews**

**Associate Editor**, ACM Transactions on Economics and Computation, 2015–2020

**Guest Editor**, Machine Learning Journal Special Issue on Computational Social Science and the Wisdom of Crowds, 2014

**Editorial Board Member**, Machine Learning Journal, 2010–2020

**Editorial Board Member**, Journal of Artificial Intelligence Research, 2010–2013

**External Reviewer**, assorted journals

**Grant Reviewer**, National Science Foundation (2010, 2011, 2013, 2015), U.S.-Israel Binational Science Foundation (2011), Computing Innovation Fellows (2020)

**Steering Committee Member**, Partnership on AI ABOUT ML program, 2019–Present

### **Selected Departmental/Lab Service**

#### **Microsoft Research, New York City**

Organizer/Co-organizer of many internal TechFest workshops and panels

Creator and organizer, MSR-NYC Postdoc Mentoring Program, 2016–Present

Lead organizer, MSR-NYC Core lab offsite, 2024

Organizing committee, MSR New England/New York/Montreal Retreat, 2018

Organizing committee, MSR FATE/AI & Society Summit, 2018

Organizing committee, Microsoft Research NYC 5th Anniversary Celebration, 2017

Organizing committee, MSR New England/New York Retreat, 2014

#### **UCLA Computer Science**

Admission, TA, and Fellowship Committee, 2010–2011

Prospective Student Visit Day Coordinator, 2011, 2012

#### **University of Pennsylvania Computer and Information Science**

Graduate Student Representative, 2006–2007

Organizer, Machine Learning Lunch, 2006–2007

Organizer, CISTers group for female graduate students and faculty in CS, 2005–2008

### **Additional Diversity & Inclusion Efforts**

**Advisory Committee**, Learning Theory Alliance, 2021–Present

**Senior Advisor**, Workshop for Women in Machine Learning (WiML), 2018–2020, 2023–Present

**Executive Board Member**, Workshop for Women in Machine Learning (WiML), 2009–2012 and 2014–2018

**Program Co-Chair**, First Celebration of Women in Computing in Southern California (CWIC-SoCal), Santa Ana, CA, 2012

**Co-founder**, Workshop for Women in Machine Learning (WiML), held annually since 2006