

Kush R. Varshney

IBM Thomas J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598
+1-914-945-1628

krvarshn@us.ibm.com
[@krvarshney](http://krvarshney.github.io)

Education:

Massachusetts Institute of Technology, Cambridge, MA, USA

Doctor of Philosophy in Electrical Engineering and Computer Science, 2010

Thesis: Frugal Hypothesis Testing and Classification

Committee: Alan S. Willsky (advisor), John W. Fisher, III, Polina Golland, and Joshua B. Tenenbaum

Minor: linguistics

Electrical Engineer, 2010

Master of Science in Electrical Engineering and Computer Science, 2006

Thesis: Joint Anisotropy Characterization and Image Formation in Wide-Angle Synthetic Aperture Radar

Cumulative GPA 5.00 (A = 5.00)

Cornell University, Ithaca, NY, USA

Bachelor of Science (magna cum laude) in Electrical and Computer Engineering, 2004

Honors Project: Greedy Postprocessing for Spatial Error Concealment in MPEG Video

Advisor: Sheila S. Hemami

Minor: computer science

Cumulative GPA 3.93 (A = 4.00)

Research and Industry Experience:

IBM Fellow, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, since April 2024.

Distinguished research scientist and senior manager, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, July 2023–April 2024.

Distinguished research staff member and manager, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, April 2020–July 2023.

Principal research staff member and manager, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, March 2018–April 2020. (Visiting IBM Research – Africa, Nairobi, Kenya, August 2019–November 2019.)

Research staff member and manager, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, December 2016–March 2018.

Research staff member, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, November 2010–December 2016.

Post doctoral researcher, *IBM Research*, Thomas J. Watson Research Center, Yorktown Heights, NY, USA, April–November 2010.

Research assistant, Laboratory for Information and Decision Systems, *Massachusetts Institute of Technology*, Cambridge, MA, USA, October 2004–February 2010.

Intern, Systems and Decision Sciences Section, National Security Engineering Division, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, May–August 2009. (Department of Energy P Clearance.)

Visiting researcher, Laboratoire de Mathématiques Appliquées aux Systèmes, *École Centrale Paris*, Châtenay-Malabry, France, May–August 2006.

Co-op intern, Enterprise Server Products Group, *Sun Microsystems*, Burlington, MA, USA, September 2002–January 2003 and May–August 2003.

Intern, Air Traffic Systems Division, *Sensis Corporation*, DeWitt, NY, USA, May–August 2001.

Summer research volunteer, Nuclear Medicine Division, *State University of New York Health Science Center*, Syracuse, NY, USA, June–September 1998.

Honors:

IEEE Signal Processing Society Distinguished Industry Speaker, 2024–2025.

IEEE Signal Processing Society Industrial Innovation Award, 2023.

Innovation Award for AI for Good, VentureBeat, 2022.

- Winner for project with Urban Institute;
- Runner-Up for project with Change Machine.

Data4Good Award, CDO Magazine and EDM Council, 2022.

- For project with Neighborhood Trust Financial Partners;
- For project with Greater DC Diaper Bank.

Winner, Science and Innovation Management Breakthrough, Falling Walls Science Summit, 2022 (for AI Fairness 360).

Tech Spotlight Runner-Up, Belfer Center for Science and International Affairs, Harvard Kennedy School, 2020 (for AI Fairness 360).

IBM Corporate Technical Award

- AI-Powered Employee Journey, 2021.
- Trustworthy AI, 2021.

IBM Research Technical Accomplishments

- Extraordinary Accomplishment
 - Research Contributions to Workforce Innovation & Enterprise Transformation,[†] 2015.
- Outstanding Accomplishment
 - AI Governance Innovations to watsonx.governance and Open Source, 2025;
 - FactSheets, 2023;
 - Trustworthy AI,^{*} 2019;

- Computational Creativity, 2014;
- GMU Proactive Retention, 2013.
- Accomplishment
 - Granite Guardian, 2025;
 - Multidimensional Subset Scanning Applications, 2022;
 - FactSheets, * 2022;
 - Science of Uncertainty Quantification, 2021;
 - Data-Driven Discovery in Global Health, 2021;
 - Dictionary- and Sparsity-Driven Imaging, * 2018;
 - IBM Leadership in Fight Against Ebola, * 2015;
 - WellPoint Health Insurance Exchanges Analytics, 2013;
 - Computational Creativity for Culinary Recipes, ‡ 2013;
 - Analytics-Driven Proactive Retention in the Growth Markets Unit, † 2012;
 - Business Impact of Outsourcing Analytics, † 2011;
 - Software Group Sales Analytics, * 2011.

*Also IBM Outstanding Technical Achievement Award

†Also IBM Research Division Award

‡Also IBM Outstanding Innovation Award

Paper Awards

- Best Paper Addressing Opportunities in AI, Computing Community Consortium / Schmidt Futures Computer Science for Social Good White Paper Competition, 2019.
- Best Research Paper Honorable Mention, SIAM International Conference on Data Mining, 2015.
- Best Social Good Paper Award, ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2014.
- Best Paper Award, IEEE International Conference on Service Operations and Logistics, and Informatics, 2013.
- Best Student Paper Travel Award, International Conference on Information Fusion, 2009.

Contributor to WellPoint Team, Gerstner Award for Client Excellence (most prestigious internal award at IBM), 2013.

IBM Eminence and Excellence Award

- Research Contributions to Trust and Transparency for AI on the IBM Cloud, 2018.
- WellPoint/IBM Healthcare Predictive Analytics Project, 2012.

National Science Foundation Graduate Research Fellow, 2004–2009.

John McMullen Dean's Scholar, 2000–2004.

Eta Kappa Nu (electrical and computer engineering honor society), inducted 2003.

Tau Beta Pi (engineering honor society), inducted 2003.

Lockheed Martin Award for Academic Excellence (awarded to five outstanding sophomores in School of Electrical and Computer Engineering, Cornell University), 2002.

Rensselaer Medal, 1999.

Research Interests:

Signal processing: detection and estimation theory, sparse signal representation, quantization

Machine learning: AI governance, alignment, safety, interpretable models, fairness and privacy, topological data analysis

Image processing: image formation, image segmentation, level set methods, remote sensing

Applications: workforce analytics, public affairs, olfaction, international development, health care

Publications:

Books and Booklets

Trustworthy Machine Learning. Kush R. Varshney. Independently Published, Chappaqua, NY, USA, 2022.

AI Fairness: How to Measure and Reduce Unwanted Bias in Machine Learning. Trisha Mahoney, Kush R. Varshney, and Michael Hind. O'Reilly Media, Sebastopol, CA, USA, 2020.

Book Chapters

Human-Centered Explainable AI (XAI): From Algorithms to User Experiences. Q. Vera Liao and Kush R. Varshney. *Human-Centered Machine Learning*. Rebecca Fiebrink, Marco Gillies, and Gonzalo Ramos, editors. Cambridge University Press, 2026.

Try to See Things from My Point of View: Empathy, AI, and the Right to Be an Exception. William G. Kidder, Jason D'Cruz, and Kush R. Varshney. *Empathy and Artificial Intelligence: Challenges, Advances and Ethical Considerations*. Cambridge University Press, 2026.

Learning Interpretable Classification Rules with Boolean Compressed Sensing. Dmitry M. Malioutov, Kush R. Varshney, Amin Emad, and Sanjeeb Dash. *Transparent Data Mining for Big and Small Data*, p. 95–121. Tania Cerquitelli, Daniele Quercia, and Frank Pasquale, editors. Cham, Switzerland: Springer, 2017.

Legislative Prediction with Political and Social Network Analysis. Jun Wang, Kush R. Varshney, and Aleksandra Mojsilović. *Encyclopedia of Social Network Analysis and Mining*, p. 804–811. Reda S. Alhadj and Jon G. Rokne, editors. Heidelberg, Germany: Springer, 2014.

Automatic Fingerprint Matching Systems. Kush R. Varshney. *Glimpses of Systems Theory and Novel Applications: Felicitation Volume in Honour of Professor Raj Kumar Varshney*, p. 149–164. Harjinder Singh Sekhon et al., editors. Aligarh, India: Navin Press, 2005.

Journal and Magazine Articles

Who's Sorry Now: User Preferences Among Rote, Empathic, and Explanatory Apologies from LLM Chatbots. Zahra Ashktorab, Alessandra Buccella, Jason D'Cruz, Zoë Fowler, Andrew Gill, Kei Yan Leung, P.D. Magnus, John Richards, and Kush R. Varshney. *ACM Transactions on Computer-Human Interaction*, accepted.

When Trust is Zero Sum: Automation's Threat to Epistemic Agency. Emmie Malone, Saleh Afroogh, Jason D'Cruz, and Kush R. Varshney. *Ethics and Information Technology*, vol. 27, p. 29, Jun. 2025.

Rethinking Machine Unlearning for Large Language Models. Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, and Yang Liu. *Nature Machine Intelligence*, vol. 7, p. 181–194, Feb. 2025.

Alignment Studio: Aligning Large Language Models to Particular Contextual Regulations. Swapnaja Achintalwar, Ioana Baldini, Djallel Bouneffouf, Joan Byamugisha, Maria Chang, Pierre Dognin, Eitan Farchi, Ndivhuwo Makondo, Aleksandra Mojsilovic, Manish Nagireddy, Karthikeyan Natesan Ramamurthy, Inkit Padhi, Orna Raz, Jesus Rios, Prasanna Sattigeri, Moninder Singh, Sipiwi Thwala, Rosario A. Uceda-Sosa, and Kush R. Varshney. *IEEE Internet Computing*, vol. 28, no. 5, p. 28–36, September/October 2024.

Künstliche Intelligenz: Benchmarks für Generative Sprachmodelle im Überblick. Andreas Christian and Kush R. Varshney. *iX Magazin für Professionelle Informationstechnik*, May 2024.

A Synergistic Future for AI and Ecology. Barbara A. Han, Kush R. Varshney, Shannon LaDeau, Ajit Subramaniam, Kathleen C. Weathers, and Jacob Zwart. *Proceedings of the National Academy of Sciences of the United States of America*, vol. 120, no. 38, p. e2220283120, September 2023.

Humble AI. Bran Knowles, Jason D’Cruz, John Richards, and Kush R. Varshney. *Communications of the ACM*, vol. 66, no. 9, p. 73–79, September 2023.

Skin Tone Analysis for Representation in Educational Materials (STAR-ED) Using Machine Learning. Girmaw Abebe Tadesse, Celia Cintas, Kush R. Varshney, Peter Staar, Chinyere Agunwa, Skyler Speakman, Justin Jia, Elizabeth Bailey, Ademide Adekun, Jules B. Lipoff, Ginikanwa Onyekaba, Jenna C. Lester, Veronica Rotemberg, James Zou, and Roxana Daneshjou. *npj Digital Medicine*, vol. 6, p. 151, August 2023.

The Incentive Gap in Data Work in the Era of Large Models. Katy Ilonka Gero, Payel Das, Pierre Dognin, Inkit Padhi, Prasanna Sattigeri, and Kush R. Varshney. *Nature Machine Intelligence*, vol. 5, no. 6, p. 565–567, June 2023.

Human-Centered Explainability for Life Sciences, Healthcare and Medical Informatics. Sanjoy Dey, Prithwish Chakraborty, Bum Chul Kwon, Amit Dhurandhar, Mohamed Ghalwash, Fernando J. Suarez Saiz, Kenney Ng, Daby Sow, Kush R. Varshney, and Pablo Meyer. *Patterns*, vol. 3, no. 5, p. 100493, May 2022.

A Human-Centered Methodology for Creating AI FactSheets. John Richards, David Piorkowski, Michael Hind, Stephanie Houde, Aleksandra Mojsilović, and Kush R. Varshney. *Bulletin of the Technical Committee on Data Engineering*, vol. 44, no. 4, p. 47–58, December 2021.

Interventional Fairness with Indirect Knowledge of Unobserved Protected Attributes. Sainyam Galhotra, Karthikeyan Shanmugam, Prasanna Sattigeri, and Kush R. Varshney. *Entropy*, vol. 23, no. 12, p. 1571, November 2021.

Socially Responsible AI Algorithms: Issues, Purposes, and Challenges. Lu Cheng, Kush R. Varshney, and Huan Lu. *Journal of Artificial Intelligence Research*, vol. 71, p. 1137–1181, August 2021.

AI Explainability 360: An Extensible Toolkit for Understanding Data and Machine Learning Models. Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte,

Ramya Raghavendra, John Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei, and Yunfeng Zhang. *Journal of Machine Learning Research*, vol. 21, no. 130, p. 1-6, June 2020.

FactSheets: Increasing Trust in AI Services through Supplier's Declarations of Conformity. Matthew Arnold, Rachel K. E. Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, Aleksandra Mojsilović, Ravi Nair, Karthikeyan Natesan Ramamurthy, Alexandra Olteanu, David Piorkowski, Darrell Reimer, John Richards, Jason Tsay, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 63, no. 4/5, p. 6, July/September 2019.

AI Fairness 360: An Extensible Toolkit for Detecting and Mitigating Algorithmic Bias. Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. *IBM Journal of Research and Development*, vol. 63, no. 4/5, p. 4, July/September 2019.

Fairness GAN: Generating Datasets with Fairness Properties Using a Generative Adversarial Network. Prasanna Sattigeri, Samuel C. Hoffman, Vijil Chenthamarakshan, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 63, no. 4/5, p. 3, July/September 2019.

Teaching AI Agents Ethical Values Using Reinforcement Learning and Policy Orchestration. Ritesh Noothigattu, Djallel Bouneffouf, Nicholas Mattei, Rachita Chandra, Piyush Madan, Kush R. Varshney, Murray Campbell, Moninder Singh, and Francesca Rossi. *IBM Journal of Research and Development*, vol. 63, no. 4/5, p. 2, July/September 2019.

Think Your Artificial Intelligence Software is Fair? Think Again. Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Sameep Mehta, Aleksandra Mojsilović, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. *IEEE Software*, vol. 36, no. 4, p. 76–80, July–August 2019.

Confronting Data Sparsity to Identify Potential Sources of Zika Virus Spillover Infection Among Primates. Barbara A. Han, Subhabrata Majumdar, Flavio P. Calmon, Benjamin S. Glicksberg, Raya Horesh, Abhishek Kumar, Adam Perer, Elisa B. von Marschall, Dennis Wei, Aleksandra Mojsilović, and Kush R. Varshney. *Epidemics*, vol. 27, p. 59–65, June 2019.

Trustworthy Machine Learning and Artificial Intelligence. Kush R. Varshney. *ACM XRDS Magazine*, vol. 25, no. 3, p. 26–29, Spring 2019.

A Big Data Approach to Computational Creativity: The Curious Case of Chef Watson. Lav R. Varshney, Florian Pinel, Kush R. Varshney, Debarun Bhattacharjya, Angela Schörgendorfer and Yi-Min Chee. *IBM Journal of Research and Development*, vol. 63, no. 1, p. 7, January–February 2019.

Distribution-Preserving k-Anonymity. Dennis Wei, Karthikeyan Natesan Ramamurthy, and Kush R. Varshney. *Statistical Analysis and Data Mining*, vol. 11, no. 6, p. 253–270, December 2018.

Data Pre-Processing for Discrimination Prevention. Flavio P. Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R. Varshney. *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 5, p. 1106–1119, October 2018.

How to Foster Innovation: A Data-Driven Approach to Measuring Economic Competitiveness. Caitlin Kuhlman, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Aurélie C. Lozano, Lei Cao, Chandra Reddy, Aleksandra Mojsilović, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, p. 11, November–December 2017.

Dataflow Representation of Data Analyses: Towards a Platform for Collaborative Data Science. Evan Patterson, Robert McBurney, Hollie Schmidt, Ioana Baldini, Aleksandra Mojsilović, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, p. 9, November–December 2017.

Real-Time Understanding of Humanitarian Crises via Targeted Information Retrieval. Kien T. Pham, Prasanna Sattigeri, Amit Dhurandhar, Arpith C. Jacob, Maja Vukovic, Patrice Chataigner, Juliana Freire, Aleksandra Mojsilović, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, p. 7, November–December 2017.

Understanding the Ecospace of Philanthropic Projects. Hemank Lamba, Mary E. Helander, Moninder Singh, Nizar Lethif, Anuradha Bhamidipaty, Salman Baset, Aleksandra Mojsilović, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, p. 6, November–December 2017.

Effectiveness of Peer Detailing in a Diarrhea Program in Nigeria. Yumeng Tao, Debarun Bhattacharjya, Aliza R. Heching, Aditya Vempaty, Moninder Singh, Felix Lam, Jason Houdek, Mohammed Abubakar, Ahmad Abdulwahab, Tiwadayo Baraimoh, Nnenna Ihebuzor, Aleksandra Mojsilović, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, p. 1, November–December 2017.

On the Safety of Machine Learning: Cyber-Physical Systems, Decision Sciences, and Data Products. Kush R. Varshney and Homa Alemzadeh. *Big Data*, vol. 5, no. 3, p. 246–255, September 2017.

Signal Processing for Social Good. Kush R. Varshney. *IEEE Signal Processing Magazine*, vol. 34, no. 3, p. 112, 108, May 2017.

Decision Making with Quantized Priors Leads to Discrimination. Lav R. Varshney and Kush R. Varshney. *Proceedings of the IEEE*, vol. 105, no. 2, p. 241–255, February 2017.

Associative Algorithms for Computational Creativity. Lav R. Varshney, Jun Wang, and Kush R. Varshney. *Journal of Creative Behavior*, vol. 50, no. 3, p. 211–223, September 2016.

Olfactory Signal Processing. Kush R. Varshney and Lav R. Varshney. *Digital Signal Processing*, vol. 48, p. 84–92, January 2016.

Data Challenges in Disease Response: The 2014 Ebola Outbreak and Beyond. Kush R. Varshney, Dennis Wei, Karthikeyan Natesan Ramamurthy, and Aleksandra Mojsilović. *ACM Journal of Data and Information Quality*, vol. 6, no. 2–3, p. 5, June 2015.

Targeting Villages for Rural Development Using Satellite Image Analysis. Kush R. Varshney, George H. Chen, Brian Abelson, Kendall Nowocin, Vivek Sakhrani, Ling Xu, and Brian L. Spatocco. *Big Data*, vol. 3, no. 1, p. 41–53, March 2015.

Optimal Grouping for Group Minimax Hypothesis Testing. Kush R. Varshney and Lav R. Varshney. *IEEE Transactions on Information Theory*, vol. 60, no. 10, p. 6511–6521, October 2014.

Bounded Confidence Opinion Dynamics in a Social Network of Bayesian Decision Makers. Kush R. Varshney. *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, p. 576–585, August 2014.

Collaborative Kalman Filtering for Dynamic Matrix Factorization. John Z. Sun, Dhruv Parthasarathy, and Kush R. Varshney. *IEEE Transactions on Signal Processing*, vol. 62, no. 14, p. 3499-3509, July 15, 2014.

Sparsity-Driven Synthetic Aperture Radar Imaging: Reconstruction, Autofocusing, Moving Targets, and Compressed Sensing. Müjdat Çetin, Ivana Stojanović, N. Özben Önhon, Kush R. Varshney, Sadegh Samadi, W. Clem Karl, and Alan S. Willsky. *IEEE Signal Processing Magazine*, vol. 31, no. 4, p. 27–40, July 2014.

Practical Ensemble Classification Error Bounds for Different Operating Points. Kush R. Varshney, Ryan J. Prenger, Tracy L. Marlatt, Barry Y. Chen, and William G. Hanley. *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 11, p. 2590–2601, November 2013.

Sales-Force Performance Analytics and Optimization. Moritz Baier, Jorge E. Carballo, Alice J. Chang, Yingdong Lu, Aleksandra Mojsilović, M. Jonathan Richard, Moninder Singh, Mark S. Squillante, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 56, no. 6, November/December 2012.

Generalization Error of Linear Discriminant Analysis in Spatially-Correlated Sensor Networks. Kush R. Varshney. *IEEE Transactions on Signal Processing*, vol. 60, no. 6, p. 3295-3301, June 2012.

Bayes Risk Error is a Bregman Divergence. Kush R. Varshney. *IEEE Transactions on Signal Processing*, vol. 59, no. 9, p. 4470–4472, September 2011.

Business Analytics Based on Financial Time Series. Kush R. Varshney and Aleksandra Mojsilović. *IEEE Signal Processing Magazine*, vol. 28, no. 5, p. 83–93, September 2011.

Linear Dimensionality Reduction for Margin-Based Classification: High-Dimensional Data and Sensor Networks. Kush R. Varshney and Alan S. Willsky. *IEEE Transactions on Signal Processing*, vol. 59, no. 6, p. 2496–2512, June 2011.

Classification Using Geometric Level Sets. Kush R. Varshney and Alan S. Willsky. *Journal of Machine Learning Research*, vol. 11, p. 491–516, February 2010.

Postarthroplasty Examination Using X-Ray Images. Kush R. Varshney, Nikos Paragios, Jean-François Deux, Alain Kulski, Rémy Raymond, Phillipe Hernigou, and Alain Rahmouni. *IEEE Transactions on Medical Imaging*, vol. 28, no. 3, p. 469–474, March 2009.

Quantization of Prior Probabilities for Hypothesis Testing. Kush R. Varshney and Lav R. Varshney. *IEEE Transactions on Signal Processing*, vol. 56, no. 10, p. 4553–4562, October 2008.

Sparse Representation in Structured Dictionaries with Application to Synthetic Aperture Radar. Kush R. Varshney, Müjdat Çetin, John W. Fisher, III, and Alan S. Willsky. *IEEE Transactions on Signal Processing*, vol. 56, no. 8, p. 3548–3561, August 2008.

Conference and Workshop Papers

Calibrated Deferral Routing for Cost-Efficient Guardrails. Giandomenico Cornacchia, Inkit Padhi, Manish Nagireddy, Subhajit Chaudhury, Pierre Dognin, Tejaswini Pedapati, Ambrish Rawat, Mark Purcell, Kush R. Varshney, and Prasanna Sattigeri. *ICML Workshop on Agents in the Wild: Safety, Security, and Beyond*, Seoul, South Korea, July 2026.

AI Steerability 360: A Toolkit for Steering Large Language Models. Erik Miebling, Karthikeyan Natesan Ramamurthy, Praveen Venkateswaran, Ching-Yun Ko, Pierre Dognin, Moninder Singh, Tejaswini Pedapati, Avinash Balakrishnan, Matthew Riemer, Dennis Wei, Inge Vejsbjerg, Elizabeth M. Daly, and Kush R. Varshney. *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, San Diego, CA, July 2026.

The Shepherd Test: How Will Superintelligent Agents Balance Care and Control in Asymmetric Relationships? Djallel Bouneffouf, Matthew Riemer, and Kush R. Varshney. *AAAI Conference on Artificial Intelligence*, Singapore, January 2026.

Assessing AI Certification Validity via Knowledge Graphs for Use Case Similarity. Elizabeth M. Daly, Jasmina Gajcin, Inge Vejsbjerg, Delaram Golpayegani, Rahul Nair, and Kush R. Varshney. *EurIPS Workshop on Private AI Governance*, Copenhagen, Denmark, December 2025.

The Shepherd Test: How Will Superintelligent Agents Balance Care and Control in Asymmetric Relationships? Djallel Bouneffouf, Matthew Riemer, and Kush R. Varshney. *NeurIPS Workshop on Evaluating the Evolving LLM Lifecycle*, San Diego, CA, December 2025.

Exposing AI Bias by Crowdsourcing: Democratizing Critique of Large Language Models. Hangzhi Guo, Pranav Narayanan Venkit, Eunhae Jang, Mukund Srinath, Wenbo Zhang, Bonam Mingole, Vipul Gupta, Kush R. Varshney, Shyam Sundar, and Amulya Yadav. *AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*, Madrid, Spain, October 2025.

Story Arena: A Multi-Agent Environment for Envisioning the Future of Software Engineering. Justin D. Weisz, Michael Muller, and Kush R. Varshney. *Workshop on the End of Programming (as we know it) - Envisioning Radical Re-Conceptualizations of Co-Coding with AI, Decennial Aarhus Conference*, Aarhus, Denmark, August 2025.

Granite Guardian: Comprehensive LLM Safeguarding. Inkit Padhi, Manish Nagireddy, Giandomenico Cornacchia, Subhjit Chaudhury, Tejaswini Pedapati, Pierre Dognin, Keerthiram Murugesan, Erik Miebling, Martín Santillán Cooper, Kieran Fraser, Giulio Zizzo, Muhammad Zaid Hameed, Mark Purcell, Michael Desmond, Qian Pan, Inge Vejsbjerg, Elizabeth M. Daly, Michael Hind, Werner Geyer, Ambrish Rawat, Kush R. Varshney, and Prasanna Sattigeri. *Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, Albuquerque, NM, April–May 2025.

Evaluating the Prompt Steerability of Large Language Models. Erik Miebling, Michael Desmond, Karthikeyan Natesan Ramamurthy, Elizabeth M. Daly, Kush R. Varshney, Eitan Farchi, Pierre Dognin, Jesus Rios, Djallel Bouneffouf, Miao Liu, and Prasanna Sattigeri. *Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, Albuquerque, NM, April–May 2025.

Contextual Value Alignment. Pierre Dognin, Jesus Rios, Ronny Luss, Prasanna Sattigeri, Miao Liu, Inkit Padhi, Matthew D. Riemer, Manish Nagireddy, Kush R. Varshney, and Djallel Bouneffouf. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Hyderabad, India, April 2025.

Scopes of Alignment. Kush R. Varshney, Zahra Ashktorab, Djallel Bouneffouf, Matthew D. Riemer, and Justin D. Weisz. *AAAI 2025 Workshop on AI Governance: Alignment, Morality, and Law*, Philadelphia, PA, March 2025.

Adversarial Prompt Evaluation: Systematic Benchmarking of Guardrails Against Prompt Input Attacks on LLMs. Giulio Zizzo, Giandomenico Cornacchia, Kieran Fraser, Muhammad Zaid Hameed, Ambrish

Rawat, Beat Buesser, Mark Purcell, Pin-Yu Chen, Prasanna Sattigeri, and Kush R. Varshney. *NeurIPS Workshop on Safe Generative AI*, Vancouver, Canada, December 2024.

Attack Atlas: A Practitioner's Perspective on Challenges and Pitfalls in Red Teaming GenAI. Amrbrish Rawat, Stefan Schoepf, Giulio Zizzo, Giandomenico Cornacchia, Muhammad Zaid Hameed, Kieran Fraser, Erik Miebling, Beat Buesser, Elizabeth M. Daly, Mark Purcell, Prasanna Sattigeri, Pin-Yu Chen, and Kush R. Varshney. *NeurIPS Workshop on Red Teaming GenAI*, Vancouver, Canada, December 2024.

Value Alignment from Unstructured Text. Inkit Padhi, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Manish Nagireddy, Pierre Dognin, and Kush R. Varshney. *NeurIPS Workshop on Pluralistic Alignment*, Vancouver, Canada, December 2024.

Open-Weight Guardians and Open-Source Frameworks for Governing LLMs. Kush R. Varshney. *Conference on Frontier AI Safety Commitments*, Berkeley, CA, November 2024.

Value Alignment from Unstructured Text. Inkit Padhi, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Manish Nagireddy, Pierre Dognin, and Kush R. Varshney. *Conference on Empirical Methods in Natural Language Processing*, Miami, FL, November 2024.

Racial and Neighborhood Disparities in Legal Financial Obligations in Jefferson County, Alabama. Óscar Lara-Yejas, Aakanksha Joshi, Andrew Martinez, Leah Nelson, Skyler Speakman, Krysten Thompson, Yuki Nishimura, Jordan Bond, and Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*, San Jose, CA, October 2024.

Decolonial AI Alignment: Openness, Viśeṣa-Dharma, and Including Excluded Knowledges. Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*, San Jose, CA, October 2024.

Individual Fairness in Graphs Using Local and Global Structural Information. Yonas Sium, Qi Li, and Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*, San Jose, CA, October 2024.

ComVas: Contextual Moral Values Aligned System. Inkit Padhi, Pierre Dognin, Jesús M. Ríos Aliaga, Ronny Luss, Swapnaja Achintalwar, Matthew D. Riemer, Miao Liu, Prasanna Sattigeri, Manish Nagireddy, Kush R. Varshney, and Djallel Bouneffouf. *International Joint Conference on Artificial Intelligence*, Jeju, South Korea, August 2024.

Using Causal Inference to Investigate Contraceptive Discontinuation in Sub-Saharan Africa. Victor Akinwande, Megan MacGregor, Celia Cintas, Ehud Karavani, Dennis Wei, Kush R. Varshney, and Pablo Nepomnaschy. *International Joint Conference on Artificial Intelligence*, Jeju, South Korea, August 2024.

Towards Healthy AI: Large Language Models Need Therapists Too. Baihan Lin, Djallel Bouneffouf, Guillermo Cecchi, and Kush R. Varshney. *NAACL Workshop on Trustworthy Natural Language Processing*, Mexico City, Mexico, June 2024.

Evaluating the Impact of Skin Tone Representation on Out-of-Distribution Detection Performance in Dermatology. Assala Benmalek, Celia Cintas, Girmaw Abebe Tadesse, Roxana Daneshjou, Kush R. Varshney, and Cherifi Dalila. *IEEE International Symposium on Biomedical Imaging*, Athens, Greece, May 2024.

Keeping Up with the Language Models: Robustness-Bias Interplay in NLI Data and Models. Ioana Baldini, Chhavi Yadav, Payel Das, and Kush R. Varshney. *ACL Workshop on Trustworthy Natural Language Processing*, Toronto, Canada, July 2023.

Add-Remove-or-Relabel: Practitioner-Friendly Bias Mitigation via Influential Fairness. Brianna Richardson, Prasanna Sattigeri, Dennis Wei, Karthikeyan Natesan Ramamurthy, Kush R. Varshney, Amit Dhurandhar, and Juan E. Gilbert. *ACM Conference on Fairness, Accountability, and Transparency*, p. 736–752, Chicago, IL, June 2023.

Trustworthy AI and the Logics of Intersectional Resistance. Bran Knowles, Jasmine Fledderjohann, John T. Richards, and Kush R. Varshney. *ACM Conference on Fairness, Accountability, and Transparency*, p. 172–182, Chicago, IL, June 2023.

Foundation Model Platforms and Bottom-of-the-Pyramid Innovation. Kush R. Varshney. *ICLR Workshop on Practical Machine Learning for Developing Countries*, Kigali, Rwanda, May 2023.

What Is Missing in IRM Training and Evaluation? Challenges and Solutions. Yihua Zhang, Pranay Sharma, Parikshit Ram, Mingyi Hong, Kush R. Varshney, and Sijia Liu. *International Conference on Learning Representations*, Kigali, Rwanda, May 2023.

A Banal Account of a Safety-Creativity Tradeoff. Kush R. Varshney and Lav R. Varshney. *IUI Workshop on Designing for Safety in Human-AI Interactions*, Sydney, Australia, March 2023.

Equi-Tuning: Group Equivariant Fine-Tuning of Pretrained Models. Sourya Basu, Prasanna Sattigeri, Karthikeyan Natesan Ramamurthy, Vijil Chenthamarakshan, Kush R. Varshney, Lav R. Varshney, and Payel Das. *AAAI Conference on Artificial Intelligence*, p. 6788–6796, Washington, DC, February 2023.

Minimax AUC Fairness: Efficient Algorithm with Provable Convergence. Zhenhuan Yang, Yan Lok Ko, Kush R. Varshney, and Yiming Ying. *AAAI Conference on Artificial Intelligence*, p. 11909–11917, Washington, DC, February 2023.

Fair Infinitesimal Jackknife: Mitigating the Influence of Biased Training Data Points Without Refitting. Prasanna Sattigeri, Soumya Ghosh, Inkit Padhi, Pierre Dognin, and Kush R. Varshney. *Advances in Neural Information Processing Systems*, New Orleans, LA, November–December 2022.

On the Safety of Interpretable Machine Learning: A Maximum Deviation Approach. Dennis Wei, Rahul Nair, Amit Dhurandhar, Kush R. Varshney, Elizabeth M. Daly, and Moninder Singh. *Advances in Neural Information Processing Systems*, New Orleans, LA, November–December 2022.

The Empathy Gap: Why AI Can Forecast Behavior But Cannot Assess Trustworthiness. Jason R. D’Cruz, William Kidder, and Kush R. Varshney. *AAAI Fall Symposium Series Symposium on Thinking Fast and Slow and Other Cognitive Theories in AI*, Arlington, VA, November 2022.

Deciding Fast and Slow: The Role of Cognitive Biases in AI-Assisted Decision-Making. Charvi Rastogi, Yunfeng Zhang, Dennis Wei, Kush R. Varshney, Amit Dhurandhar, and Richard Tomsett. *ACM Conference on Computer-Supported Cooperative Work and Social Computing*, November 2022.

Humble Machines: Attending to the Underappreciated Costs of Misplaced Distrust. Bran Knowles, Jason D’Cruz, John T. Richards, and Kush R. Varshney. *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, Arlington, VA, October 2022.

Differentially Private SGDA for Minimax Problems. Zhenhuan Yang, Shu Hu, Yunwen Lei, Kush R. Varshney, Siwei Lyu, and Yiming Ying. *Conference on Uncertainty in Artificial Intelligence*, Eindhoven, Netherlands, August 2022.

Causal Feature Selection for Algorithmic Fairness. Sainyam Galhotra, Karthikeyan Shanmugam, Prasanna Sattigeri, and Kush R. Varshney. *ACM SIGMOD/PODS International Conference on Management of Data*, Philadelphia, PA, June 2022.

Out-of-Distribution Detection in Dermatology using Input Perturbation and Subset Scanning. Hannah Kim, Girmaw Abebe Tadesse, Celia Cintas, Skyler Speakman, and Kush R. Varshney. *IEEE International Symposium on Biomedical Imaging*, Kolkata, India, March 2022.

AI Explainability 360: Impact and Design. Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte, Ramya Raghavendra, John T. Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei and Yunfeng Zhang. *Conference on Innovative Applications of Artificial Intelligence*, February 2022.

Uncertainty Quantification 360. Soumya Ghosh, Q. Vera Liao, Karthikeyan Natesan Ramamurthy, Jiri Navratil, Prasanna Sattigeri, Kush R. Varshney, Yunfeng Zhang. *ACM India Joint International Conference on Data Science and Management of Data*, p. 333–335, January 2022.

CoFrNets: Interpretable Neural Architecture Inspired by Continued Fractions. Isha Puri, Amit Dhurandhar, Tejaswini Pedapati, Karthikeyan Shanmugam, Dennis Wei, and Kush R. Varshney. *Advances in Neural Information Processing Systems*, December 2021.

Blockchain and the Scientific Method. James A. Evans, Kweku Opoku-Agyemang, Krishna Ratakonda, Kush R. Varshney, and Lav R. Varshney. *ASCR Workshop on Cybersecurity and Privacy for Scientific Computing Ecosystems*, November 2021.

Out-of-Distribution Detection and Fairness Assessment in Dermatology. Hannah Kim, Girmaw Abebe Tadesse, Celia Cintas, Skyler Speakman, and Kush R. Varshney. *KDD Outlier Detection and Description Workshop*, August 2021.

An Empirical Study of Accuracy, Fairness, Explainability, Distributional Robustness, and Adversarial Robustness. Moninder Singh, Gevorg Ghalachyan, Kush R. Varshney, and Reginald E. Bryant. *KDD Workshop on Measures and Best Practices for Responsible AI*, August 2021.

A Research Framework for Understanding Education-Occupation Alignment with NLP Techniques. Renzhe Yu, Subhro Das, Sairam Gurajada, Kush R. Varshney, Hari Raghavan and Carlos X. Lastra-Anadon. *ACL-IJCNLP Workshop on NLP for Positive Impact*, August 2021.

Biomedical Interpretable Entity Representations. Diego Garcia-Olano, Yasumasa Onoe, Ioana Baldini, Joydeep Ghosh, Byron C. Wallace, and Kush R. Varshney. *Findings of ACL: ACL-IJCNLP*, August 2021.

Treatment Effect Estimation Using Invariant Risk Minimization. Abhin Shah, Kartik Ahuja, Karthikeyan Shanmugam, Dennis Wei, Kush R. Varshney, and Amit Dhurandhar. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, p. 5005–5009, June 2021.

Beyond Reasonable Doubt: Improving Fairness in Budget-Constrained Decision Making Using Confidence Thresholds. Michiel Bakker, Duy Patrick Tu, Krishna Gummadi, Alex ‘Sandy’ Pentland,

Kush R. Varshney, and Adrian Weller. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, May 2021.

Towards Interpreting Zoonotic Potential of Betacoronavirus Sequences with Attention. Kahini Wadhawan, Payel Das, Barbara A. Han, Ilya R. Fischhoff, Adrian Castellanos, Arvind Varsani, and Kush R. Varshney. *ICLR Workshop on Machine Learning for Preventing and Combating Pandemics*, May 2021.

Fairly Estimating Socioeconomic Status Under Costly Feature Acquisition. Ritika Brahmadesam and Kush R. Varshney. *ICLR Workshop on Practical Machine Learning for Developing Countries*, May 2021.

Empirical or Invariant Risk Minimization? A Sample Complexity Perspective. Kartik Ahuja, Jun Wang, Amit Dhurandhar, Karthikeyan Shanmugam, and Kush R. Varshney. *International Conference on Learning Representations*, May 2021.

Automated Meta-Analysis in Medical Research: A Causal Learning Perspective. Lu Cheng, Dmitriy Katz-Rogozhnikov, Kush R. Varshney, and Ioana Baldini. *ACM Conference on Health, Inference, and Learning Workshop*, April 2021.

Disparate Impact Diminishes Consumer Trust Even for Advantaged Users. Tim Draws, Zoltán Szilávik, Benjamin Timmermans, Nava Tintarev, Kush R. Varshney, and Michael Hind. *International Conference on Persuasive Technologies*, April 2021.

Automated Evaluation of Representation in Dermatology Educational Materials. Girmaw Abebe Tadesse, Hannah Kim, Roxana Daneshjou, Celia Cintas, Kush R. Varshney, Ademide Adelekun, Jules Lipoff, Ginikanwa Onyekab, Veronica Rotemberg, and James Zou. *AAAI Workshop on Trustworthy AI for Healthcare*, February 2021.

Exploring the Efficacy of Generic Drugs in Treating Cancer. Ioana Baldini, Mariana Bernagozzi, Sulbha Aggarwal, Mihaela Bornea, Saksham Chawla, Joppe Geluykens, Dmitriy A. Katz-Rogozhnikov, Pratik Mukherjee, Smruthi Ramesh, Sara Rosenthal, Jagrati Sharma, Kush R. Varshney, Catherine Del Vecchio Fitz, Pradeep Mangalath, and Laura B. Kleiman. *AAAI Conference on Artificial Intelligence*, February 2021.

AI Explainability 360 Toolkit. Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte, Ramya Raghavendra, John T. Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei, and Yunfeng Zhang. *ACM India Joint International Conference on Data Science and Management of Data*, p. 376–379, Bangalore, India, January 2021. (held virtually)

Identifying Factors Associated with Neonatal Mortality in Sub-Saharan Africa using Machine Learning. William Ogallo, Skyler Speakman, Victor Akinwande, Kush R. Varshney, Aisha Walcott-Bryant, Charity Wayua, Komminist Weldemariam, Claire-Helene Mershon, and Nosa Orobato. *American Medical Informatics Association Annual Symposium*, Chicago, IL, November 2020. (held virtually)

Fairness of Classifiers Across Skin Tones in Dermatology. Newton M. Kinyanjui, Timothy Odonga, Celia Cintas, Noel C. F. Codella, Rameswar Panda, Prasanna Sattigeri, and Kush R. Varshney. *International Conference on Medical Image Computing and Computer Assisted Intervention*, Lima, Peru, October 2020. (held virtually)

Trust and Transparency in Contact Tracing Applications. Stacy Hobson, Michael Hind, Aleksandra Mojsilović, and Kush R. Varshney. *KDD Workshop on Fragile Earth: Data Science for a Sustainable Planet*, San Diego, CA, August 2020. (held virtually)

Tutorial on Human-Centered Explainability for Healthcare. Prithwish Chakraborty, Bum Chul Kwon, Sanjoy Dey, Amit Dhurandhar, Daniel Gruen, Kenney Ng, Daby Sow, and Kush R. Varshney. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, p. 3547–3548, San Diego, CA, August 2020. (held virtually)

On the Equivalence of Bi-Level Optimization and Game-Theoretic Formulations of Invariant Risk Minimization. Kartik Ahuja, Karthikeyan Shanmugam, Kush R. Varshney, and Amit Dhurandhar. *ICML Workshop on Inductive Biases, Invariances and Generalization in RL*, Vienna, Austria, July 2020. (held virtually)

Invariant Risk Minimization Games. Kartik Ahuja, Karthikeyan Shanmugam, Kush R. Varshney, and Amit Dhurandhar. *International Conference on Machine Learning*, p. 145–155, Vienna, Austria, July 2020. (held virtually)

Is There a Trade-Off Between Fairness and Accuracy? A Perspective Using Mismatched Hypothesis Testing. Sanghamitra Dutta, Dennis Wei, Hazar Yueksel, Pin-Yu Chen, Sijia Liu, and Kush R. Varshney. *International Conference on Machine Learning*, p. 2803–2813, Vienna, Austria, July 2020. (held virtually)

Inspection of Blackbox Models for Evaluating Vulnerability in Maternal, Newborn, and Child Health. William Ogallo, Skyler Speakman, Victor Akinwande, Kush R. Varshney, Aisha Walcott-Bryant, Charity Wayua, and Komminist Weldemariam. *International Joint Conference on Artificial Intelligence–Pacific Rim International Conference on Artificial Intelligence*, p. 5282–5284, Yokohama, Japan, July 2020. (delayed and held virtually)

Characterization of Overlap in Observational Studies. Michael Oberst, Fredrik D. Johansson, Dennis Wei, Tian Gao, Gabriel Brat, David Sontag, and Kush R. Varshney. *International Conference on Artificial Intelligence and Statistics*, p. 788–798, Palermo, Italy, June 2020. (delayed and held virtually)

Preservation of Anomalous Subgroups on Variational Autoencoder Transformed Data. Samuel C. Maina, Reginald E. Bryant, William Ogallo, Kush R. Varshney, Skyler Speakman, Celia Cintas, Aisha Walcott-Bryant, Robert-Florian Samoilescu, and Komminist Weldemariam. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, p. 3627–3631, Barcelona, Spain, May 2020. (held virtually)

DADI: Dynamic Discovery of Fair Information with Adversarial Reinforcement Learning. Michiel Bakker, Duy Patrick Tu, Humberto Riverón Valdés, Krishna Gummadi, Kush R. Varshney, Adrian Weller and Alex ‘Sandy’ Pentland. *ICLR Workshop on Towards Trustworthy ML*, Addis Ababa, Ethiopia, April 2020. (held virtually)

Experiences with Improving the Transparency of AI Models and Services. Michael Hind, Stephanie Houde, Jacquelyn Martino, Aleksandra Mojsilović, David Piorkowski, John Richards, and Kush R. Varshney. *ACM CHI Conference on Human Factors in Computing Systems*, p. LBW229, Honolulu, HI, April 2020. (held virtually)

Interpretable Subgroup Discovery in Treatment Effect Estimation with Application to Opioid Prescribing Guidelines. Chirag Nagpal, Dennis Wei, Bhanukiran Vinzamuri, Monica Shekhar, Sara E. Berger, Subhro

Das, and Kush R. Varshney. *ACM Conference on Health, Inference, and Learning*, p. 19–29, Toronto, Canada, April 2020. (delayed and held virtually)

On Mismatched Detection and Safe, Trustworthy Machine Learning. Kush R. Varshney. *Conference on Information Sciences and Systems*, Princeton, NJ, March 2020. (not held, proceedings printed)

Event-Driven Continuous Time Bayesian Networks. Debarun Bhattacharjya, Karthikeyan Shanmugam, Tian Gao, Nicholas Mattei, Kush R. Varshney, and Dharmashankar Subramanian. *AAAI Conference on Artificial Intelligence*, p. 3259–3266, New York, NY, February 2020.

A Natural Language Processing System for Extracting Evidence of Drug Repurposing from Scientific Publications. Shivashankar Subramanian, Ioana Baldini, Sushma Ravichandran, Dmitriy A. Katz-Rogozhnikov, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Kush R. Varshney, Annmarie Wang, Pradeep Mangalath, and Laura B. Kleiman. *Conference on Innovative Applications of Artificial Intelligence*, p. 13376–13381, New York, NY, February 2020.

Data Augmentation for Discrimination Prevention and Bias Disambiguation. Shubham Sharma, Yunfeng Zhang, Jesús M. Ríos Aliaga, Djallel Bouneffouf, Vinod Muthusamy, and Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, p. 358–364, New York, NY, February 2020.

Joint Optimization of AI Fairness and Utility: A Human-Centered Approach. Yunfeng Zhang, Rachel K. E. Bellamy, and Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, p. 400–406, New York, NY, February 2020.

Fair Enough: Improving Fairness in Budget-Constrained Decision Making Using Confidence Thresholds. Michiel Bakker, Duy Patrick Tu, Humberto Riverón Valdés, Krishna Gummadi, Kush R. Varshney, Adrian Weller and Alex ‘Sandy’ Pentland. *AAAI Workshop on Safe Artificial Intelligence*, New York, NY, February 2020.

How Data Scientists Work Together with Domain Experts in Scientific Collaborations: To Find The Right Answer or to Ask the Right Question? Yaoli Mao, Dakuo Wang, Michael Muller, Kush R. Varshney, Ioana Baldini, Casey Dugan, and Aleksandra Mojsilović. *ACM International Conference on Supporting Group Work*, p. 237, Sanibel Island, FL, January 2020.

Estimating Skin Tone and Effects on Classification Performance in Dermatology Datasets. Newton M. Kinyanjui, Timothy Odonga, Celia Cintas, Noel C. F. Codella, Rameswar Panda, Prasanna Sattigeri, and Kush R. Varshney. *NeurIPS Fair Machine Learning for Health Workshop*, Vancouver, Canada, December 2019. (*Spotlight*)

Drug Repurposing for Cancer: An NLP Approach to Identify Low-Cost Therapies. Shivashankar Subramanian, Ioana Baldini, Sushma Ravichandran, Dmitriy A. Katz-Rogozhnikov, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Kush R. Varshney, Annmarie Wang, Pradeep Mangalath, and Laura B. Kleiman. *NeurIPS Workshop on Machine Learning for Health*, Vancouver, Canada, December 2019.

DADI: Dynamic Discovery of Fair Information with Adversarial Reinforcement Learning. Michiel Bakker, Duy Patrick Tu, Humberto Riverón Valdés, Krishna Gummadi, Kush R. Varshney, Adrian Weller and Alex ‘Sandy’ Pentland. *NeurIPS Workshop on Human-Centric Machine Learning*, Vancouver, Canada, December 2019.

Subgroup Preservation in Financial Data Anonymized by a Variational Autoencoder. Samuel C. Maina, Reginald E. Bryant, William Ogallo, Kush R. Varshney, Skyler Speakman, Celia Cintas, Aisha Walcott-

Bryant, and Robert-Florian Samoilescu. *NeurIPS Workshop on Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy*, Vancouver, Canada, December 2019.

Teaching AI Ethical Values Using Reinforcement Learning and Policy Orchestration. Ritesh Noothigattu, Djallel Bouneffouf, Nicholas Mattei, Rachita Chandra, Piyush Madan, Kush R. Varshney, Murray Campbell, Moninder Singh, and Francesca Rossi. *International Joint Conference on Artificial Intelligence*, p. 6377–6381, Macao, August 2019.

Event-Driven Continuous Time Bayesian Networks: An Application in Modeling Progression out of Poverty through Integrated Social Services. Debarun Bhattacharjya, Karthikeyan Shanmugam, Tian Gao, Nicholas Mattei, and Kush R. Varshney. *IJCAI Workshop on AI for Social Good*, Macau, August 2019.

On Fairness in Budget-Constrained Decision Making. Michiel Bakker, Alejandro Noriega Campero, Duy Patrick Tu, Prasanna Sattigeri, Kush R. Varshney, and Alex ‘Sandy’ Pentland. *KDD Workshop on Explainable Artificial Intelligence*, Anchorage, AK, August 2019.

Open Platforms for Artificial Intelligence for Social Good: Common Patterns as a Pathway to True Impact. Kush R. Varshney and Aleksandra Mojsilović. *ICML Workshop on AI for Social Good*, Long Beach, CA, June 2019.

Teaching AI to Explain its Decisions Using Embeddings and Multi-Task Learning. Noel C. F. Codella, Michael Hind, Karthikeyan Natesan Ramamurthy, Murray Campbell, Amit Dhurandhar, Kush R. Varshney, Dennis Wei, and Aleksandra Mojsilović. *ICML Workshop on Human in the Loop Learning*, Long Beach, CA, June 2019.

Topological Data Analysis of Decision Boundaries with Application to Model Selection. Karthikeyan Natesan Ramamurthy, Kush R. Varshney, and Krishnan Mody. *International Conference on Machine Learning*, Long Beach, CA, p. 5351–5360, June 2019.

Bias Mitigation Post-Processing for Individual and Group Fairness. Pranay K. Lohia, Karthikeyan Natesan Ramamurthy, Manish Bhide, Diptikalyan Saha, Kush R. Varshney, and Ruchir Puri. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brighton, UK, p. 2847–2851, May 2019.

Constructing and Compressing Frames in Blockchain-Based Verifiable Multi-Party Computation. Ravi Kiran Raman, Kush R. Varshney, Roman Vaculin, Nelson Kibichii Bore, Sekou L. Remy, Eleftheria K. Pissadaki, and Michael Hind. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brighton, UK, p. 7500–7504, May 2019.

Promoting Distributed Trust in Machine Learning and Computational Simulation. Nelson Kibichii Bore, Ravi Kiran Raman, Isaac M. Markus, Sekou L. Remy, Oliver Bent, Michael Hind, Eleftheria K. Pissadaki, Biplav Srivastava, Roman Vaculin, Kush R. Varshney, and Komminist Weldemariam. *IEEE International Conference on Blockchain and Cryptocurrency*, Seoul, Korea, May 2019.

A Scalable Blockchain Approach for Trusted Computation and Verifiable Simulation in Multi-Party Collaboration. Ravi Kiran Raman, Roman Vaculin, Michael Hind, Sekou L. Remy, Eleftheria K. Pissadaki, Nelson Kibichii Bore, Roozbeh Daneshvar, Biplav Srivastava, and Kush R. Varshney. *IEEE International Conference on Blockchain and Cryptocurrency*, Seoul, Korea, May 2019.

Fairness GAN: Generating Datasets with Fairness Properties using a Generative Adversarial Network. Prasanna Sattigeri, Samuel C. Hoffman, Vijil Chenthamarakshan, and Kush R. Varshney. *ICLR Workshop on Safe Machine Learning*, New Orleans, LA, May 2019.

Fair Transfer Learning with Missing Protected Attributes. Amanda Coston, Karthikeyan Natesan Ramamurthy, Dennis Wei, Kush R. Varshney, Skyler Speakman, Zairah Mustahsan, and Supriyo Chakraborty. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, Honolulu, HI, p. 91–98, January 2019.

TED: Teaching AI to Explain Its Decisions. Michael Hind, Dennis Wei, Murray Campbell, Noel C. F. Codella, Amit Dhurandhar, Aleksandra Mojsilović, Karthikeyan Natesan Ramamurthy, and Kush R. Varshney. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, Honolulu, HI, p. 123–129, January 2019.

Financial Forecasting and Analysis for Low-Wage Workers. Wenyu Zhang, Raya Horesh, Karthikeyan Natesan Ramamurthy, Lingfei Wu, Jinfeng Yi, Kryn Anderson, and Kush R. Varshney. *Data for Good Exchange*, New York, NY, September 2018.

Teaching Machines to Understand Data Science Code by Semantic Enrichment of Dataflow Graphs. Evan Patterson, Ioana Baldini, Aleksandra Mojsilović, and Kush R. Varshney. *KDD Workshop on Fragile Earth: Theory Guided Data Science to Enhance Scientific Discovery*, London, UK, p. 5, August 2018.

Semantic Representation of Data Science Programs. Evan Patterson, Ioana Baldini, Aleksandra Mojsilović, and Kush R. Varshney. *International Joint Conference on Artificial Intelligence and European Conference on Artificial Intelligence*, Stockholm, Sweden, p. 5847–5849, July 2018.

Why Interpretability in Machine Learning? An Answer Using Distributed Detection and Data Fusion Theory. Kush R. Varshney, Prashant Khanduri, Pranay Sharma, Shan Zhang, and Pramod K. Varshney. *ICML Workshop on Human Interpretability in Machine Learning*, p. 15–20, Stockholm, Sweden, July 2018.

The Effect of Extremist Violence on Hateful Speech Online. Alexandra Olteanu, Carlos Castillo, Jeremy Boy, and Kush R. Varshney. *International AAAI Conference on Weblogs and Social Media*, Palo Alto, CA, p. 221–230, June 2018.

False Positive Control with Concave Penalties using Stability Selection. Bhanukiran Vinzamuri and Kush R. Varshney. *IEEE Data Science Workshop*, Lausanne, Switzerland, p. 76–80, June 2018.

Assessing National Development Plans for Alignment with Sustainable Development Goals via Semantic Search. Jonathan Galsurkar, Moninder Singh, Lingfei Wu, Aditya Vempaty, Mikhail Sushkov, Devika Iyer, Serge Kapto, and Kush R. Varshney. *Conference on Innovative Applications of Artificial Intelligence*, New Orleans, LA, p. 7753–7758, February 2018.

Neurology-as-a-Service for the Developing World. Tejas Dharamsi, Payel Das, Tejaswini Pedapati, Gregory Bramble, Vinod Muthusamy, Horst Samulowitz, and Kush R. Varshney, Yuvaraj Rajamanickam, John Thomas, and Justin Dauwels. *NIPS Workshop on Machine Learning for the Developing World*, Long Beach, CA, December 2017.

Scalable Demand-Aware Recommendation. Jinfeng Yi, Cho-Jui Hsieh, Kush R. Varshney, Lijun Zhang, and Yao Li. *Advances in Neural Information Processing Systems*, Long Beach, CA, p. 2412–2421, December 2017.

Optimized Pre-Processing for Discrimination Prevention. Flavio P. Calmon, Dennis Wei, Karthikeyan Natesan Ramamurthy, Bhanukiran Vinzamuri, and Kush R. Varshney. *Advances in Neural Information Processing Systems*, Long Beach, CA, p. 3992–4001, December 2017.

Exploring the Causal Relationships between Initial Opioid Prescriptions and Outcomes. Jinghe Zhang, Vijay S. Iyengar, Dennis Wei, Bhanukiran Vinzamuri, Hamsa Bastani, Alexander R. Macalalad, Anne E. Fischer, Gigi Yuen-Reed, Aleksandra Mojsilović, and Kush R. Varshney. *AMIA Workshop on Data Mining for Medical Informatics*, Washington, DC, November 2017.

An End-To-End Machine Learning Pipeline That Ensures Fairness Policies. Samiulla Shaikh, Harit Vishwakarma, Sameep Mehta, Kush R. Varshney, Karthikeyan Natesan Ramamurthy, and Dennis Wei. *Data for Good Exchange*, New York, NY, September 2017.

The Limits of Abstract Evaluation Metrics: The Case of Hate Speech Detection. Alexandra Olteanu, Kartik Talamadupula, and Kush R. Varshney. *ACM Web Science Conference*, Troy, NY, p. 405–406, June 2017.

Statistical Analysis of Peer Detailing for Children’s Diarrhea Treatments. Yumeng Tao, Debarun Bhattacharjya, Aliza R. Heching, Aditya Vempaty, Moninder Singh, Felix Lam, Kush R. Varshney, and Aleksandra Mojsilović. *AAAI Spring Symposium on AI for Social Good*, Stanford, CA, p. 101–106, March 2017.

Machine Representation of Data Analyses: Towards a Platform for Collaborative Data Science. Evan J. Patterson, Ioana Baldini, Aleksandra Mojsilović, and Kush R. Varshney. *AAAI Spring Symposium on AI for Social Good*, Stanford, CA, p. 53–59, March 2017.

Information Retrieval, Fusion, Completion, and Clustering for Employee Expertise Estimation. Raya Horesh, Kush R. Varshney, and Jinfeng Yi. *IEEE International Conference on Big Data*, Washington, DC, p. 1385–1393, December 2016.

Stable Estimation of Granger-Causal Factors of Country-Level Innovation. Aurélie C. Lozano, Prasanna Sattigeri, Aleksandra Mojsilović, and Kush R. Varshney. *IEEE Global Conference on Signal and Information Processing*, Washington, DC, p. 1290–1294, December 2016.

Interpretable Machine Learning via Convex Cardinal Shape Composition. Kush R. Varshney. *Allerton Conference on Communication, Control, and Computing*, Monticello, IL, p. 327–330, September 2016.

Learning Sparse Two-Level Boolean Rules. Guolong Su, Dennis Wei, Kush R. Varshney, and Dmitry M. Malioutov. *IEEE Workshop on Machine Learning for Signal Processing*, Salerno, Italy, September 2016.

Dynamic Matrix Factorization with Social Influence. Aleksandr Y. Aravkin, Kush R. Varshney, and Liu Yang. *IEEE Workshop on Machine Learning for Signal Processing*, Salerno, Italy, September 2016.

Understanding Innovation to Drive Sustainable Development. Prasanna Sattigeri, Aurélie Lozano, Aleksandra Mojsilović, Kush R. Varshney, and Mahmoud Naghshineh. *ICML Workshop on #Data4Good: Machine Learning in Social Good Applications*. New York, NY, p. 21–25, June 2016.

Interpretable Two-Level Boolean Rule Learning for Classification. Guolong Su, Dennis Wei, Kush R. Varshney, and Dmitry M. Malioutov. *ICML Workshop on Human Interpretability in Machine Learning*, New York, NY, p. 66–70, June 2016.

Fidelity Loss in Distribution-Preserving Anonymization and Histogram Equalization. Lav R. Varshney and Kush R. Varshney. *Conference on Information Sciences and Systems*, Princeton, NJ, p. 30–35, March 2016.

Engineering Safety in Machine Learning. Kush R. Varshney. *Information Theory and Applications Workshop*, La Jolla, CA, February 2016.

Predictive Modeling of Customer Repayment for Sustainable Pay-As-You-Go Solar Power in Rural India. Hugo Gerard, Kamallesh Rao, Mark Simithraaratchy, Kush R. Varshney, Kunal Kabra, and G. Paul Needham. *Data for Good Exchange Conference*, New York, NY, September 2015.

Data Science of the People, for the People, by the People: A Viewpoint on an Emerging Dichotomy. Kush R. Varshney. *Data for Good Exchange Conference*, New York, NY, September 2015.

From Open Data Ecosystems to Systems of Innovation: A Journey to Realize the Promise of Open Data. Shubir Kapoor, Aleksandra Mojsilović, Jade Nguyen Strattnner, and Kush R. Varshney. *Data for Good Exchange Conference*, New York, NY, September 2015.

A Robust Nonlinear Kalman Smoothing Approach for Dynamic Matrix Factorization. Aleksandr Y. Aravkin, Kush R. Varshney, and Dmitry M. Malioutov. *Signal Processing with Adaptive Sparse Structured Representations Workshop*, Cambridge, United Kingdom, July 2015.

A Semiquantitative Group Testing Approach for Learning Interpretable Clinical Prediction Rules. Amin Emad, Kush R. Varshney, and Dmitry M. Malioutov. *Signal Processing with Adaptive Sparse Structured Representations Workshop*, Cambridge, United Kingdom, July 2015.

Optigrow: People Analytics for Job Transfers. Dennis Wei, Kush R. Varshney, and Marcy Wagman. *IEEE International Congress on Big Data*, New York, NY, p. 535–542, June–July 2015.

Health Insurance Market Risk Assessment: Covariate Shift and k-Anonymity. Dennis Wei, Karthikeyan Natesan Ramamurthy, and Kush R. Varshney. *SIAM International Conference on Data Mining*, Vancouver, Canada, p. 226–234, April–May 2015. (*Best Research Paper Honorable Mention*)

Learning Interpretable Classification Rules Using Sequential Row Sampling. Sanjeeb Dash, Dmitry M. Malioutov, and Kush R. Varshney. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brisbane, Australia, p. 3337–3341, April 2015.

Robust Binary Hypothesis Testing Under Contaminated Likelihoods. Dennis Wei and Kush R. Varshney. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brisbane, Australia, p. 3407–3411, April 2015.

Persistent Topology of Decision Boundaries. Kush R. Varshney and Karthikeyan Natesan Ramamurthy. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brisbane, Australia, p. 3931–3935, April 2015.

Targeting Direct Cash Transfers to the Extremely Poor. Brian Abelson, Kush R. Varshney, and Joy Sun. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, New York, NY, p. 1563–1572, August 2014. (*Best Social Good Paper Award*)

Predicting Employee Expertise for Talent Management in the Enterprise. Kush R. Varshney, Vijil Chenthamarakshan, Scott W. Fancher, Jun Wang, Dongping Fang, and Aleksandra Mojsilović. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, New York, NY, p. 1729–1738, August 2014.

An Analysis of Losing Unimportant Points in Tennis. Kush R. Varshney. *KDD Workshop on Large-Scale Sports Analytics*, New York, NY, August 2014.

Computing Persistent Homology Under Random Projection. Karthikeyan Natesan Ramamurthy, Kush R. Varshney, and Jayaraman J. Thiagarajan. *IEEE International Workshop on Statistical Signal Processing*, Gold Coast, Australia, p. 105–108, June–July 2014.

Food Steganography with Olfactory White. Kush R. Varshney and Lav R. Varshney. *IEEE International Workshop on Statistical Signal Processing*, Gold Coast, Australia, p. 21–24, June–July 2014.

Active Odor Cancellation. Kush R. Varshney and Lav R. Varshney. *IEEE International Workshop on Statistical Signal Processing*, Gold Coast, Australia, p. 25–28, June–July 2014.

Screening for Learning Classification Rules via Boolean Compressed Sensing. Sanjeeb Dash, Dmitry M. Malioutov, and Kush R. Varshney. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Florence, Italy, p. 3360–3364, May 2014.

Prescriptive Analytics for Allocating Sales Teams to Opportunities. Ban Kawas, Mark S. Squillante, Dharmashankar Subramanian, and Kush R. Varshney. *IEEE International Conference on Data Mining Workshops*, Dallas, TX, p. 211–218, December 2013.

Quantifying and Recommending Expertise When New Skills Emerge. Dongping Fang, Kush R. Varshney, Jun Wang, Karthikeyan Natesan Ramamurthy, Aleksandra Mojsilović, and John H. Bauer. *IEEE International Conference on Data Mining Workshops*, Dallas, TX, p. 672–679, December 2013.

Quantile Regression for Workforce Analytics. Karthikeyan Natesan Ramamurthy, Kush R. Varshney, and Moninder Singh. *IEEE Global Conference on Signal and Information Processing*, Austin, TX, p. 1134, December 2013.

Flavor Pairing in Medieval European Cuisine: A Study in Cooking with Dirty Data. Kush R. Varshney, Lav R. Varshney, Jun Wang, and Daniel Meyers. *IJCAI Workshop on Cooking with Computers*, Beijing, China, p. 3–12, August 2013.

Balancing Lifetime and Classification Accuracy of Wireless Sensor Networks. Kush R. Varshney and Peter M. van de Ven. *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Bengaluru, India, p. 31–38, July–August 2013.

Expertise Assessment with Multi-Cue Semantic Information. Jun Wang, Kush R. Varshney, Aleksandra Mojsilović, Dongping Fang, and John H. Bauer. *IEEE International Conference on Service Operations and Logistics, and Informatics*, Dongguan, China, p. 534–539, July 2013. (*Best Paper Award*)

Dose-Response Signal Estimation and Optimization for Salesforce Management. Kush R. Varshney and Moninder Singh. *IEEE International Conference on Service Operations and Logistics, and Informatics*, Dongguan, China, p. 328–333, July 2013.

Cognition as a Part of Computational Creativity. Lav R. Varshney, Florian Pinel, Kush R. Varshney, Angela Schörgendorfer, and Yi-Min Chee. *IEEE International Conference on Cognitive Informatics and Cognitive Computing*, New York, NY, p. 36–43, July 2013.

Predicting and Recommending Skills in the Social Enterprise. Kush R. Varshney, Jun Wang, Aleksandra Mojsilović, Dongping Fang, and John H. Bauer. *International AAAI Conference on Weblogs and Social Media Workshop on Social Computing for Workforce 2.0*, Cambridge, MA, p. 20–23, July 2013.

A Salesforce Control Theory Analysis of Enterprise Microblog Posts. Kush R. Varshney and N. Sadat Shami. *International AAAI Conference on Weblogs and Social Media Workshop on Social Computing for Workforce 2.0*, Cambridge, MA, p. 16–19, July 2013.

Exact Rule Learning via Boolean Compressed Sensing. Dmitry M. Malioutov and Kush R. Varshney. *International Conference on Machine Learning*, Atlanta, GA, p. 765–773, June 2013.

Opinion Dynamics with Bounded Confidence in the Bayes Risk Error Divergence Sense. Kush R. Varshney. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vancouver, Canada, p. 6600–6604, May 2013.

Interactive Visual Salesforce Analytics. Kush R. Varshney, Jamie C. Rasmussen, Aleksandra Mojsilović, Moninder Singh, and Joan M. DiMicco. *International Conference on Information Systems*, Orlando, FL, December 2012.

An Analytics Approach for Proactively Combating Voluntary Attrition of Employees. Moninder Singh, Kush R. Varshney, Jun Wang, Aleksandra Mojsilović, Alisia R. Gill, Patricia I. Faur, and Raphael Ezry. *IEEE International Conference on Data Mining Workshops*, Brussels, Belgium, p. 317–323, December 2012.

Deconvolving the Productivity of Salespeople via Constrained Quadratic Programming. Gautam K. Bhat and Kush R. Varshney. *National Systems Conference*, Annamalainagar, India, December 2012.

Decision Trees for Heterogeneous Dose-Response Signal Analysis. Kush R. Varshney, Moninder Singh, and Jun Wang. *IEEE International Workshop on Statistical Signal Processing*, Ann Arbor, MI, p. 916–919, August 2012.

Does Selection Bias Blind Performance Diagnostics of Business Decision Models? A Case Study in Salesforce Optimization. Jun Wang, Moninder Singh, and Kush R. Varshney. *IEEE International Conference on Service Operations and Logistics, and Informatics*. Suzhou, China, p. 416–421, July 2012.

Legislative Prediction via Random Walks over a Heterogeneous Graph. Jun Wang, Kush R. Varshney, and Aleksandra Mojsilović. *SIAM International Conference on Data Mining*, Anaheim, CA, p. 1095–1106, April 2012.

Dynamic Matrix Factorization: A State Space Approach. John Z. Sun, Kush R. Varshney, and Karthik Subbian. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Kyoto, Japan, p. 1897–1900, March 2012.

Estimating Post-Event Seller Productivity Profiles in Dynamic Sales Organizations. Kush R. Varshney, Moninder Singh, Mayank Sharma, and Aleksandra Mojsilović. *IEEE International Conference on Data Mining Workshops*, Vancouver, Canada, p. 1191–1198, December 2011.

A Risk Bound for Ensemble Classification with a Reject Option. Kush R. Varshney. *IEEE International Workshop on Statistical Signal Processing*, Nice, France, p. 773–776, June 2011.

Multilevel Minimax Hypothesis Testing. Kush R. Varshney and Lav R. Varshney. *IEEE International Workshop on Statistical Signal Processing*, Nice, France, p.109–112, June 2011.

MCMC Inference of the Shape and Variability of Time-Response Signals. Dmitriy A. Katz-Rogozhnikov, Kush R. Varshney, Aleksandra Mojsilović, and Moninder Singh. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, p. 3956–3959, Prague, Czech Republic, May 2011.

Spatially-Correlated Sensor Discriminant Analysis. Kush R. Varshney. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, p. 3680–3683, Prague, Czech Republic, May 2011.

Categorical Decision Making by People, Committees, and Crowds. Lav R. Varshney, Joong Bum Rhim, Kush R. Varshney, and Vivek K Goyal. *Information Theory and Applications Workshop*, La Jolla, CA, February 2011.

Class-Specific Error Bounds for Ensemble Classifiers. Ryan J. Prenger, Tracy D. Lemmond, Kush R. Varshney, Barry Y. Chen, and William G. Hanley. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, p. 843–852, Arlington, VA, July 2010.

Learning Dimensionality-Reduced Classifiers for Information Fusion. Kush R. Varshney and Alan S. Willsky. *International Conference on Information Fusion*, p. 1881–1888, Seattle, WA, July 2009. (*Best Student Paper Travel Award*)

Supervised Learning of Classifiers via Level Set Segmentation. Kush R. Varshney and Alan S. Willsky. *IEEE Workshop on Machine Learning for Signal Processing*, p. 115–120, Cancún, Mexico, October 2008.

Minimum Mean Bayes Risk Error Quantization of Prior Probabilities. Kush R. Varshney and Lav R. Varshney. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, p. 3445–3448, Las Vegas, NV, April 2008.

Geniş Açılı Radarda Görüntü Oluşturma ve Yönbağımlılık Tespiti için Seyrek Sinyal Temsiline Dayalı bir Yaklaşım (A Sparse Signal Representation-Based Approach to Image Formation and Anisotropy Determination in Wide-Angle Radar). Kush R. Varshney, Müjdat Çetin, John W. Fisher, III, and Alan S. Willsky. *IEEE Signal Processing and Communication Applications Conference*, Eskişehir, Turkey, June 2007.

Multi-View Stereo Reconstruction of Total Knee Replacement from X-Rays. Kush R. Varshney, Nikos Paragios, Alain Kulski, Remy Raymond, Phillippe Hernigou, and Alain Rahmouni. *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, p. 1148–1151, Arlington, VA, April 2007.

Wide-Angle SAR Image Formation with Migratory Scattering Centers and Regularization in Hough Space. Kush R. Varshney, Müjdat Çetin, John W. Fisher, III, and Alan S. Willsky. *Annual Workshop on Adaptive Sensor Array Processing*, Lexington, MA, June 2006.

Joint Image Formation and Anisotropy Characterization in Wide-Angle SAR. Kush R. Varshney, Müjdat Çetin, John W. Fisher, III, and Alan S. Willsky. *SPIE Defense and Security Symposium, Algorithms for Synthetic Aperture Radar Imagery XIII*, p. 62370D, Orlando (Kissimmee), FL, April 2006.

Conference Presentations and Posters

An Advaita Vedānta Approach to Agentic AI Identity and Governance. Kush R. Varshney. *International Conference on Vedānta*, Stony Brook, NY, May 2026.

A Retrospective Study of Contraceptive Discontinuation Across Multiple Sub-Saharan African Countries. Victor Akinwande, Celia Cintas, Ehud Karavani, Megan MacGregor, Dennis Wei, Kush R. Varshney, and Pablo Nepomnaschy. *Annual Meetings of the Human Biology Association*, Reno, NV, April 2023.

On the Duality of Transparency and Value Alignment and Their Mismatch. Kush R. Varshney. *ACM Conference on Fairness, Accountability, and Transparency CRAFT Session on Emerging Problems: New Challenges in FAccT from Research, to Practice, to Policy*, Seoul, Korea, June 2022.

Racial Representation Analysis in Dermatology Academic Materials. Girmaw Abebe Tadesse, Celia Cintas, Roxana Daneshjou, Kush R. Varshney, Peter Staar, Skyler Speakman, Kenya Andrews, Chinyere Agunwa, Justin Jia, Elizabeth Bailey, Ademide Adelekun, Jules B. Lipoff, Ginikanwa Onyekaba, Veronica Rottemberg, and James Zou. *AMIA Annual Symposium*, San Diego, CA, October–November 2021.

One Explanation Does Not Fit All: A Toolkit and Taxonomy of AI Explainability Techniques. Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte, Ramya Raghavendra, John Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei, and Yunfeng Zhang. *INFORMS Annual Meeting*, Anaheim, CA, October 2021.

Addressing the Design Needs of Implementing Fairness in AI via Influence Functions. Brianna Richardson and Kush R. Varshney. *INFORMS Annual Meeting*, Anaheim, CA, October 2021.

Quantifying Education-Occupation Alignment Through Natural Language Processing. Renzhe Yu, Subhro Das, Sairam Gurajada, Kush R. Varshney, Hari Raghavan and Carlos X. Lastra-Anadon. *International Conference on Computational Social Science*, July 2021.

Building Trustworthy AI Systems. Kush R. Varshney. *Department of Defense Artificial Intelligence Symposium*, June 2021.

Human Cognitive Biases in Interpreting Machine Learning. Charvi Rastogi, Yunfeng Zhang, Dennis Wei, Kush R. Varshney, Amit Dhurandhar, and Richard Tomsett. *INFORMS Annual Meeting*, National Harbor, MD, November 2020. (held virtually)

Empirical and Theoretical Relationships Among Pillars of Trustworthy Machine Learning. Kush R. Varshney and Moninder Singh. *INFORMS Annual Meeting*, National Harbor, MD, November 2020. (held virtually)

Drug Repurposing for Cancer: An NLP Approach to Identify Low-Cost Therapies. Shivashankar Subramanian, Ioana Baldini, Sushma Ravichandran, Dmitriy A. Katz-Rogozhnikov, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Kush R. Varshney, Annmarie Wang, Pradeep Mangalath, and Laura B. Kleiman. *New York Academy of Sciences Natural Language, Dialog and Speech Symposium*, New York, NY, November 2019.

Rigorous Analysis of Racial Bias in Gender Classification. Vidya Muthukumar, Tejaswini Pedapati, Nalini Ratha, Prasanna Sattigeri, Chai-Wah Wu, Brian Kingsbury, Abhishek Kumar, Samuel Thomas,

Aleksandra Mojsilović, and Kush R. Varshney. *Women in Machine Learning Workshop*, Montreal, Canada, December 2018.

SimplerVoice: A Key Message & Visual Description Generator System for Illiteracy. Minh N. B. Nguyen, Samuel Thomas, Kush R. Varshney, Sujatha Kashyap, and Anne E. Gattiker. *Data for Good Exchange*, New York, NY, September 2018.

Decision Support for Policymaking: Causal Inference Algorithm and Case Study. Bhanukiran Vinzamuri, Aleksandra Mojsilović, and Kush R. Varshney. *Workshop on Mechanism Design for Social Good*, Ithaca, NY, June 2018.

How to Foster Innovation: A Data-Driven Approach to Measuring Economic Competitiveness. Caitlin Kuhlman, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Aurélie C. Lozano, Lei Cao, Chandra Reddy, Aleksandra Mojsilović, and Kush R. Varshney. *Workshop on Mechanism Design for Social Good*, Ithaca, NY, June 2018.

Decision Making With Quantized Priors Leads to Discrimination. Lav R. Varshney and Kush R. Varshney. *Illinois Summit on Diversity in Psychological Science*, Champaign, IL, March, 2018.

A Social Good Program at an Industrial Research Laboratory. Kush R. Varshney. *American Association for the Advancement of Science Annual Meeting*, Austin, TX, February, 2018.

SimplerVoice: A Key Message & Visual Description Generator System for Illiteracy. Minh N. B. Nguyen, Samuel Thomas, Anne E. Gattiker, Sujatha Kashyap, and Kush R. Varshney. *Women in Machine Learning Workshop*, Long Beach, CA, December 2017.

Machine Learning from Health Insurance Administrative Data: Opioids, Obamacare, and Other Applications. Kush R. Varshney. *Administrative Data Research Facilities Network Conference*, Washington, DC, November 2017.

Interpretable Two-Level Boolean Rule Learning. Guolong Su, Dennis Wei, Kush R. Varshney, and Dmitry M. Malioutov. *INFORMS Annual Meeting*, Houston, TX, October 2017.

Semantic Searching for Efficient Assessment of Sustainable Development in National Plans. Jonathan Galsurkar, Aditya Vempaty, Kush R. Varshney, Lingfei Wu, Mikhail Sushkov, Moninder Singh, Devika Iyer, and Serge Kaptov. *Data Science for Social Good Conference*, Chicago, IL, September 2017.

Cognitive Disease Hunter: Developing Automated Pathogen Feature Extraction from Scientific Literature. Timothy NeCamp, Prasanna Sattigeri, Dennis Wei, Emily Ray, Youssef Drissi, Ananya Poddar, Diwakar Mahajan, Sarah Bowden, Barbara A. Han, Aleksandra Mojsilović, and Kush R. Varshney. *Data Science for Social Good Conference*, Chicago, IL, September 2017.

Demystifying Social Entrepreneurship: An NLP Based Approach to Finding a Social Good Fellow. Aditya Garg, Alexandra Olteanu, Richard B. Segal, Dmitriy A. Katz-Rogozhnikov, Keerthana Kumar, Joana Maria, Liza Mueller, Ben Beers, and Kush R. Varshney. *Data Science for Social Good Conference*, Chicago, IL, September 2017.

Decision Support for Policymaking: Causal Inference Algorithm and Case Study. Bhanukiran Vinzamuri, Aleksandra Mojsilović, and Kush R. Varshney. *Data Science for Social Good Conference*, Chicago, IL, September 2017.

A Bayesian Approach for Predicting Neotropical Primate Reservoirs of Zika Virus. Subhabrata Majumdar, Dennis Wei, Adam Perer, Benjamin Glicksberg, Aleksandra Mojsilović, Kush R. Varshney, and Barabara A. Han. *International Symposium on Business and Industrial Statistics*, Yorktown Heights, NY, June 2017.

Employee Expertise Estimation by Information Retrieval, Fusion, Completion, and Clustering. Raya Horesh, Kush R. Varshney, and Jinfeng Yi. *SDM Workshop on Women in Data Science*, Houston, TX, April 2017.

Persistent Homology of Classifier Decision Boundaries. Kush R. Varshney and Karthikeyan Natesan Ramamurthy. *STOC/SoCG Workshop on Geometry and Machine Learning*, Cambridge, MA, June 2016.

Learning Interpretable Clinical Prediction Rules using Threshold Group Testing. Amin Emad, Kush R. Varshney, and Dmitry M. Malioutov. *NIPS Workshop on Machine Learning in Healthcare*, Montréal, Canada, December 2015.

Assessing Expertise in the Enterprise: The Recommender Point of View. Aleksandra Mojsilović and Kush R. Varshney. *ACM Recommender Systems Conference*, p. 231, Vienna, Austria, September 2015.

Learning Interpretable Classification Rules via Boolean Compressed Sensing. Dmitry M. Malioutov, Kush R. Varshney, and Sanjeeb Dash. *Joint Statistical Meetings*, Seattle, WA, August 2015.

Personalization of Product Novelty Assessment via Bayesian Surprise. Nan Shao, Kush R. Varshney, Lav R. Varshney, Florian Pinel, Anshul Sheopuri, and Pavankumar Murali. *Joint Statistical Meetings*, Boston, MA, August 2014.

Talent Analytics to Predict Employee Job Roles and Skill Sets Using Diverse Data Sources. Kush R. Varshney, Vijil Chenthamarakshan, Jun Wang, Dongping Fang, and Aleksandra Mojsilović. *International Symposium on Business and Industrial Statistics / Conference of the ASA Section on Statistical Learning and Data Mining*, Durham, NC, June 2014.

Sales Team Effort Allocation Analytics. Ban Kawas, Mark S. Squillante, Dharmashankar Subramanian, and Kush R. Varshney. *INFORMS Annual Meeting*, Minneapolis, MN, October 2013.

Dynamic Factor Modeling via Robust Subspace Tracking. Aleksandr Y. Aravkin, Kush R. Varshney, and Dmitry M. Malioutov. *Industrial-Academic Workshop on Optimization in Finance and Risk Management*, Toronto, Canada, September 2013.

More Contentious Issues Lead to More Factions: Bounded Confidence Opinion Dynamics of Bayesian Decision Makers. Kush R. Varshney. *Interdisciplinary Workshop on Information and Decision in Social Networks*, Cambridge, MA, November 2012.

Proactive Retention. Yingdong Lu, Aleksandra Mojsilović, Moninder Singh, Mark S. Squillante, Kush R. Varshney, and Jun Wang. *INFORMS Annual Meeting*, Phoenix, AZ, October 2012.

Heterogeneous Graphs in Social Business. Kush R. Varshney, Jun Wang, and Aleksandra Mojsilović. *Graph Exploitation Symposium*, Dedham, MA, April, 2012.

Map of the Marketplace: Visualizing the Relationships in the IT Services Marketplace. Aleksandra Mojsilović, Kush R. Varshney, and Jun Wang. *INFORMS Annual Meeting*, Charlotte, NC, November 2011.

Identifying Optimal Sales Team Composition for Business Opportunities. Aleksandra Mojsilović, Moninder Singh, Kush R. Varshney, and Jun Wang. *INFORMS Annual Meeting*, Charlotte, NC, November 2011.

Classification of IT Service Tickets for Defect Prevention. Kush R. Varshney, Dongping Fang, Aliza R. Heching, Aleksandra Mojsilović, and Moninder Singh. *INFORMS Annual Meeting*, Charlotte, NC, November 2011.

A Framework for Sales Force Productivity Profile Estimation. Mayank Sharma, Aleksandra Mojsilović, Moninder Singh, and Kush R. Varshney. *INFORMS Annual Meeting*, Austin, TX, November 2010.

Random Forest Classifier Performance in Low False Alarm and Low Missed Detection Regimes. Kush R. Varshney. *Lawrence Livermore National Laboratory Institutional Summer Student Poster Symposium*, Livermore, CA, August 2009.

Block-Segmentation and Classification of Grayscale Postal Images. Kush R. Varshney. *Bits on our Mind Symposium*, Ithaca, NY, March 2004.

Theses and Technical Reports

Go Further, Faster with AI: How Governance Increases Velocity. Heather Gentile, Brian Goehring, Betsy Greytok, Wouter Oosterbosch, and Kush Varshney. *Research Insight*. IBM Institute for Business Value, December 2025.

A Line in the Sand: Artificial Intelligence and Human Liberty. Julian Adler, Jethro Antoine, and Kush R. Varshney. *Policy Brief*. Center for Justice Innovation, March 2025.

Trustworthiness and Responsibility in AI – Causality, Learning, and Verification. Vaishak Belle, Hana Chockler, Shannon Vallor, Kush R. Varshney, Joost Vennekens, and Sander Beckers. *Report from Dagstuhl Seminar 24121*, March 2024.

Trustworthy AI Practitioner Requirements. Phaedra Boinodiris, Mayan Murray, Martin Wiertz, Bernard Freund, Graham White, Kim Holmes, Heather Hagerty, Michael Vössing, Heather Domin, Ulrike Zeilberger, Sheri Feinzig, James Winters, Michael Flores, Boz Handy Bosma, Georg Olowson, Herman Colquhoun, Jr., Jill Villany, Zuzana Leova, Kush Varshney, Bill Dusch, Dai Lins, and Jurriaan Parie. *Technical Report ARB 211320a*. IBM Academy of Technology, September 2022.

Frugal Hypothesis Testing and Classification. Kush R. Varshney. *Ph.D. Thesis*. Massachusetts Institute of Technology, Cambridge, MA, February 2010.

The Feature Analysis of Flow-Based Statistics for Network Traffic Classification. Kelley B. Herndon Ford, Kush R. Varshney, Tracy D. Lemmond, William G. Hanley, Barry Y. Chen, and William C. Kallander. *Technical Report*. Lawrence Livermore National Laboratory, Livermore, CA, September 2009.

Surface Evolution for 3-D Shape Reconstruction of Knee Joint Prosthetics and Bones. Kush R. Varshney and Nikos Paragios. *Technical Report 0703*. Laboratoire de Mathématiques Appliquées aux Systèmes, École Centrale Paris, Châtenay-Malabry, France, February 2007.

Joint Anisotropy Characterization and Image Formation in Wide-Angle Synthetic Aperture Radar. Kush R. Varshney. *Master's Thesis*. Massachusetts Institute of Technology, Cambridge, MA, May 2006.

Prefaces

Sanmay Das, Brian Patrick Green, Kush R. Varshney, Marianna Ganapini, and Andrea Renda. *Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*, San Jose, CA, October 2024.

AI Ethics. Sameep Mehta, Francesca Rossi, and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 63, no. 4/5, July/September 2019.

Advances in Computational Creativity Technology. Richard Goodwin, Kush R. Varshney, and Jinjun Xiong. *IBM Journal of Research and Development*, vol. 63, no. 1, January–February 2019.

Data Science for Social Good. Aleksandra Mojsilović and Kush R. Varshney. *IBM Journal of Research and Development*, vol. 61, no. 6, November–December 2017.

Online Contributions

Expert, Mixture of Experts Podcast.

- Episode 107. (<https://www.youtube.com/watch?v=YCWwh70FZtQ>), 15 May 2026.
- Episode 104. (<https://www.youtube.com/watch?v=vCP7C8CVw28>), 24 April 2026.
- Episode 95. (<https://www.youtube.com/watch?v=3Wb6cA4CBIQ>), 20 February 2026.
- Episode 94. (<https://www.youtube.com/watch?v=Lw5kD9xb9Ic>), 13 February 2026.
- Episode 86. (<https://www.youtube.com/watch?v=dJTwev1O-mE>), 19 December 2025.
- Episode 75. (<https://www.youtube.com/watch?v=LAXAmXHNGeM>), 3 October 2025.
- Episode 66. (<https://www.youtube.com/watch?v=EB3rWZPz0gk>), 1 August 2025.
- Episode 62. (<https://www.youtube.com/watch?v=myIre7iONII>), 4 July 2025.
- Episode 53. (<https://www.youtube.com/watch?v=yeutlpU13YM>), 2 May 2025.
- Episode 48. (<https://www.youtube.com/watch?v=CgqHN38l6Ko>), 28 March 2025.
- Episode 36. (<https://www.youtube.com/watch?v=QzERUfTbKQw>), 3 January 2025.
- Episode 35. (<https://www.youtube.com/watch?v=l8plyR8aqVQ>), 27 December 2024.
- Episode 26. (<https://www.youtube.com/watch?v=5-xMSQZ9xx0>), 25 October 2024.
- Episode 18. (<https://www.youtube.com/watch?v=IKaj8gATzoY>), 30 August 2024.
- Episode 15. (<https://www.youtube.com/watch?v=1tTXTDzaPyY>), 9 August 2024.
- Episode 8. (<https://www.youtube.com/watch?v=G8lnNR1-rsw>), 21 June 2024.
- Episode 2. (<https://www.youtube.com/watch?v=K83tTEeGCBC>), 10 May 2024.
- Episode 1. (<https://www.youtube.com/watch?v=hwNkFnR1U0I>), 3 May 2024.

IBM Research Video and Written Q&A (<https://research.ibm.com/blog/map-measure-manage-gen-ai>), 26 September 2025.

IEEE Signal Processing Society Our Digital Life, Episode 6. (<https://www.youtube.com/watch?v=PXVbCo8mZGU>), 9 September 2025.

Granite Guardian with Kush Varshney. Episode 2, What's New at IBM (<https://www.ibm.com/new/videos/whats-new-granite-episode-2>), 24 October 2024.

On Apprenticeship with Kush Varshney and Skyler Speakman. On Not Knowing (https://w3.ibm.com/w3publisher/ibmresearch_growthmindset/learn-from-our-peers), 20 November 2024.

Where to Begin: 3 IBM Leaders Offer Guidance to Newly Appointed Chief AI Officers. *IBM Blog* (<https://www.ibm.com/blog/caio-first-days-guidance>), 14 May 2024.

Pioneering AI Fairness and Decolonial Perspectives with Kush Varshney. *Humans of AI Podcast* (<https://humansofai.transistor.fm/episodes/24-pioneering-ai-fairness-and-decolonial-perspectives-with-kush-varshney>), 1 December 2023.

The Latest AI Safety Method is a Throwback to Our Maritime Past. *IBM Research Blog* (<https://research.ibm.com/blog/AI-governance-explained>), 16 November 2023.

Trust, Transparency, and Governance in the Age of Generative AI. *AI Academy* (<https://www.youtube.com/watch?v=odrD0OLPeiY>), 16 November 2023.

Guest, ACM ByteCast (<https://learning.acm.org/bytecast/ep43-kush-varshney>), 14 September 2023.

Behind the Paper: Skin Tone Analysis for Representation in Educational Materials (STAR-ED). *Springer Nature Research Communities* (<https://communities.springernature.com/posts/skin-tone-analysis-for-representation-in-educational-materials-star-ed>), 18 August 2023.

Guest, MLOps Community Podcast (https://mlops.community/watch/trustworthy-machine-learning_mzL4fCH1ok8UD7), 21 September 2022.

Guest, Computer Science Chats (https://www.youtube.com/watch?v=tg_2VYbOO44), 27 April 2022.

Guest, Pipeline Conversations Podcast, ZenML (<https://podcast.zenml.io/trustworthy-ml-kush-varshney>), 14 April 2022.

The Toolkit Approach to Trustworthy AI. *Open Data Science Blog* (<https://opendatascience.com/the-toolkit-approach-to-trustworthy-ai/>), 13 April 2022.

The Watsons Meet Watson: A Call for Carative AI. *Montreal AI Ethics Institute Blog* (<https://montrealaiethics.ai/the-watsons-meet-watson-a-call-for-carative-ai>), 23 March 2022.

Guest, Episode 82, Artificial Intelligence and You (<https://humancusp.com/podcast.html>), 10 January 2022.

The Lifecycle View of Trustworthy AI. *IBM Data Science in Practice Blog* (<https://medium.com/ibm-data-ai/the-lifecycle-view-of-trustworthy-ai-a0006339a4cd>), 22 October 2021.

AIX360 Meets CodeFlare: How to Scale Explainability Using CodeFlare Pipelines. *CodeFlare Blog* (<https://medium.com/codeflare/aix360-meets-codeflare-how-to-scale-explainability-using-codeflare-pipelines-c6bd8a7b0090>), 29 September 2021.

AI Governance: Gain Control Over the AI Lifecycle. *Linux Foundation AI & Data Blog* (<https://lfaidata.foundation/blog/2021/09/22/ai-governance-gain-control-over-the-ai-lifecycle/>), 22 September 2021.

Foundations of Trustworthy AI: How to Conduct Trustworthy AI Assessment and Mitigation. *Watson Blog* (<https://www.ibm.com/blogs/watson/2021/06/trustworthy-ai-assessment-mitigation/>), 3 June 2021.

What You Cannot Miss in Any AI Implementation: Fairness. *Linux Foundation AI & Data Blog* (<https://lfaidata.foundation/blog/2021/05/19/what-you-cannot-miss-in-any-ai-implementation-fairness/>), 19 May 2021.

IBM Research LIVE: Responsible AI (<https://www.youtube.com/watch?v=q67JXxcfAC4>), 13 May 2021.

Guest, Can We Trust AI, *In AI We Trust?* (<https://www.equalai.org/newsroom/>), 30 April 2021.

Guest, Machine Learning and Trustworthy AI, *Making Data Simple* (<https://www.ibm.com/blogs/journey-to-ai/2020/12/making-data-simple-ai-martin-and-kush-varshney-discuss-machine-learning-and-trustworthy-ai/>), 9 December 2020.

Guest, AI for Good and the Real World, Season 4, Episode 19, *Talking Machines* (<https://www.thetalkingmachines.com/episodes/ai-good-and-real-world>), 18 October 2018.

Introducing AI Fairness 360. Kush R. Varshney. *IBM Research Blog* (<https://www.ibm.com/blogs/research/2018/09/ai-fairness-360>), 19 September 2018.

AI for Code Encourages Collaborative, Open Scientific Discovery. Kush R. Varshney. *IBM Research Blog* (<https://www.ibm.com/blogs/research/2018/08/ai-code>), 15 August 2018.

New AI Algorithm Recommends Right Products at the Right Time. Kush R. Varshney. *IBM Research Blog* (<https://www.ibm.com/blogs/research/2017/12/ai-algorithm-recommends-products>), 8 December 2017.

Reducing Discrimination in AI with New Methodology. Kush R. Varshney. *IBM Research Blog* (<https://www.ibm.com/blogs/research/2017/12/ai-reducing-discrimination>), 6 December 2017.

Social Good, Meet Data Science. Kush R. Varshney. *IBM Big Data & Analytics Hub Blog* (<http://www.ibmbigdatahub.com/blog/social-good-meet-data-science>), 16 February 2016.

For Analytics to Have an Impact, Keep it Simple. Tyler H. McCormick, Cynthia Rudin, Dmitry Malioutov, and Kush Varshney. *Data Informed* (<http://data-informed.com/for-analytics-to-have-an-impact-keep-it-simple>), 3 August 2015.

Statistical Analysis of Serving and Breaking First in Tennis. Kush Varshney. *Jon Wertheim's Tennis Mailbag* (<http://sportsillustrated.cnn.com/tennis/news/20131030/serena-williams-mailbag>), 30 October 2013.

Patents:

Issued

Interpretable Neural Architecture Using Continued Fractions. Amit Dhurandhar, Tejaswini Pedapati, Isha Puri, Karthikeyan Shanmugam, Kush R. Varshney, and Dennis Wei. US 12,619,871, granted May 5, 2026.

Multi-Agent Platform for Trusted Computation of Models and Data. Nelson Kibichii Bore, Michael Hind, Eleftheria K. Pissadaki, Ravi Kiran Raman, Sekou Lionel Remy, Roman Vaculin, and Kush R. Varshney. US 12,567,948, granted March 3, 2026.

Mitigating the Influence of Biased Training Instances Without Refitting. Pierre L. Dognin, Soumya Ghosh, Inkit Padhi, Prasanna Sattigeri, and Kush R. Varshney. US 12,561,612, granted February 24, 2026.

Multi-Agent Platform for Trusted Computation of Models and Data. Nelson Kibichii Bore, Michael Hind, Eleftheria K. Pissadaki, Ravi Kiran Raman, Sekou Lionel Remy, Roman Vaculin, and Kush R. Varshney. US 12,530,339, granted January 20, 2026.

Distributed Platform for Computation and Trusted Validation. Nelson Kibichii Bore, Michael Hind, Eleftheria K. Pissadaki, Ravi Kiran Raman, Sekou Lionel Remy, Roman Vaculin, and Kush R. Varshney. US 11,940,978, granted March 26, 2024.

Initializing Optimization Solvers. Kartik Ahuja, Amit Dhurandhar, Karthikeyan Shanmugam, and Kush R. Varshney. US 11,915,131, granted February 27, 2024.

Input Encoding for Classifier Generalization. Brian E. D. Kingsbury, Kush R. Varshney, and Hazar Yueksel. US 11,914,678, granted February 27, 2024.

System and Method for Post Hoc Improvement of Instance-Level and Group-Level Prediction Metrics. Manish Bhide, Pranay Lohia, Karthikeyan Natesan Ramamurthy, Ruchir Puri, Diptikalyan Saha, and Kush R. Varshney. JP 7,289,086, granted May 31, 2023. US 11,734,585, granted August 22, 2023.

Mitigating Statistical Bias in Artificial Intelligence Models. Rachel K. E. Bellamy, Kush R. Varshney, and Yunfeng Zhang. US 11,586,849, February 21, 2023.

Semantic Queries Based on Semantic Representation of Programs and Data Source Ontologies. Ioana M. Baldini Soares, Aleksandra Mojsilović, Evan J. Patterson, and Kush R. Varshney. US 11,520,830, December 6, 2022.

Artificial Intelligence Certification of Factsheets Using Blockchain. Samuel C. Hoffman, Kalapriya Kannan, Pranay K. Lohia, Sameep Mehta, and Kush R. Varshney. US 11,483,154, October 25, 2022.

Enhancing Fairness in Transfer Learning for Machine Learning Models with Missing Protected Attributes in Source or Target Domains. Supriyo Chakraborty, Amanda Coston, Zairah Mustahsan, Karthikeyan Natesan Ramamurthy, Skyler Speakman, Kush R. Varshney, and Dennis Wei. US 11,443,236, September 13, 2022.

Generation and Management of an Artificial Intelligence (AI) Model Documentation Throughout Its Life Cycle. Matthew R. Arnold, Rachel K. E. Bellamy, Kaoutar El Maghraoui, Michael Hind, Stephanie Houde, Kalapriya Kannan, Manish Kesarwani, Sameep Mehta, Aleksandra Mojsilović, Ramya Raghavendra, Darrell C. Reimer, John T. Richards, David J. Piorkowski, Jason Tsay, and Kush R. Varshney. US 11,263,188, March 1, 2022.

Tracking the Evolution of Topic Rankings from Contextual Data. Mary E. Helander, Hemank Lamba, Nizar Lethif, Joana S. B. T. Maria, Emily A. Ray, and Kush R. Varshney, US 11,244,013, February 8, 2022.

Distributed Platform for Computation and Trusted Validation. Nelson K. Bore, Michael Hind, Eleftheria K. Pissadaki, Ravi Kiran Raman, Sekou L. Remy, Roman Vaculin, and Kush R. Varshney. US 11,212,076, December 28, 2021.

Distributed Platform for Computation and Trusted Validation. Nelson K. Bore, Michael Hind, Eleftheria K. Pissadaki, Ravi Kiran Raman, Sekou L. Remy, Roman Vaculin, and Kush R. Varshney. US 11,032,063, June 8, 2021.

Representation of a Data Analysis Using a Flow Graph. Ioana M. Baldini Soares, Aleksandra Mojsilović, Evan J. Paterson, and Kush R. Varshney. US 10,891,326, January 12, 2021.

Node Relevance Determination in an Evolving Network. Mary E. Helander, Hemank Lamba, Nizar Lethif, Joana S. B. T. Maria, Emily A. Ray, and Kush R. Varshney, US 10,756,977, August 25, 2020.

Humanitarian Crisis Analysis Using Secondary Information Gathered by a Focused Web Crawler. Ioana M. Baldini Soares, Amit Dhurandhar, Abhishek Kumar, Aleksandra Mojsilović, Kien T. Pham, Kush R. Varshney, and Maja Vukovic. US 10,740,860, August 11, 2020.

Recommendation Prediction Based on Preference Elicitation. Aleksandra Mojsilović, Kush R. Varshney, Jun Wang, and Jinfeng Yi. US 10,678,800, June 9, 2020.

Method, System and Computer Program Product for Automating Expertise Management Using Social and Enterprise Data. John H. Bauer, Dongping Fang, Aleksandra Mojsilović, Karthikeyan Natesan Ramamurthy, Kush R. Varshney, and Jun Wang. US 10,643,140, May 5, 2020.

Generating Semantic Flow Graphs Representing Computer Programs. Ioana M. Baldini Soares, Aleksandra Mojsilović, Evan J. Paterson, and Kush R. Varshney. US 10,628,282, April 21, 2020.

Generating Work Products Using Work Product Metrics and Predicted Constituent Availability. Debarun Bhattacharjya, Kush R. Varshney, and Lav R. Varshney. US 10,467,638, November 5, 2019.

Accelerating Data-Driven Scientific Discovery. Flavio P. Calmon and Kush R. Varshney. US 10,388,039, August 20, 2019.

Association-Based Product Design. Kush R. Varshney, Lav R. Varshney, and Jun Wang. US 10,019,689, July 10, 2018.

Computing Personalized Probabilistic Familiarity Based on Known Artifact Data. Florian Pinel, Nan Shao, Kush R. Varshney, and Lav R. Varshney. US 9,852,380, December 26, 2017.

Active Odor Cancellation. Kush R. Varshney and Lav R. Varshney. US 9,600,793, March 21, 2017.

Food Steganography. Kush R. Varshney and Lav R. Varshney. US 9,417,221, August 16, 2016.

Predictive and Descriptive Analysis on Relations Graphs with Heterogeneous Entities. Aleksandra Mojsilović, Kush R. Varshney, and Jun Wang. US 9,195,941, November 24, 2015.

Applied

Shepherd Test Framework. Djallel Bouneffouf, Matthew D. Riemer, and Kush R. Varshney. 19/467,148, filed February 2, 2026.

Aligning Models Based on Rules Associated with Harm Perceptions and Scopes of Alignment. Zahra Ashktorab, Djallel Bouneffouf, Matthew D. Riemer, Kush R. Varshney, and Justin D. Weisz. 19/449,865, filed January 15, 2026.

Contextual Moral Value Alignment through Context-Based Aggregation. Djallel Bouneffouf, Pierre L. Dognin, Miao Liu, Ronny Luss, Manish Nagireddy, Inkit Padhi, Matthew D. Riemer, Jesus M. Rios Aliaga, Prasanna Sattigeri, and Kush R. Varshney. 18/925,251, filed October 24, 2024.

System Development Incorporating Ethical Context. Venkata N. Pavuluri, John Rofrano, Kush R. Varshney, and Maja Vukovic. 18/392,405, filed December 21, 2023.

Sufficiency Assessment of Machine Learning Models Through Maximum Deviation. Elizabeth Daly, Amit Dhurandhar, Michael Hind, Rahul Nair, Moninder Singh, Kush R. Varshney, and Dennis Wei. 17/931,803, filed September 13, 2022.

Input-Encoding with Federated Learning. Pradip Bose, Supriyo Chakraborty, Brian E. D. Kingsbury, Kush R. Varshney, Augusto Vega, Dinesh C. Verma, Ashish Verma, Shiqiang Wang, and Hazar Yueksel. 17/239,812, filed April 26, 2021.

Learning Robust Predictors Using Game Theory. Kartik Ahuja, Amit Dhurandhar, Karthikeyan Shanmugam, and Kush R. Varshney. 17/115,489, filed December 8, 2020.

Modeling External Event Effects Upon System Variables. Debarun Bhattacharjya, Tian Gao, Nicholas S. Mattei, Karthikeyan Shanmugam, Dharmashankar Subramanian, and Kush R. Varshney. 16/942,842, filed July 30, 2020.

Method for Market Risk Assessment for Healthcare Applications. Shilpa Mahatma, Aleksandra Mojsilović, Karthikeyan Natesan Ramamurthy, Kush R. Varshney, Dennis Wei, and Gigi Yuen-Reed. 14/699,482, filed April 29, 2015.

Nonparametric Tracking and Forecasting of Multivariate Data. Aleksandr Y. Aravkin, Dmitry M. Malioutov, and Kush R. Varshney. 14/480,704, filed September 9, 2014, abandoned July 24, 2018.

Tutorials, Talks, Panels, and Demonstrations:

Tutorials

Safe and Trustworthy Large Language Models.

- IEEE International Conference on Acoustics, Speech, and Signal Processing, Seoul, Korea, April 14, 2024.

Trustworthy Machine Learning.

- Trinity College Dublin + IBM Research AI Masterclass, Dublin, Ireland, October 17, 2023.

AI Fairness 360.

- Immersive AI Session, Open Data Science Conference, New York, NY, June 27, 2019.
- O'Reilly Artificial Intelligence Conference, New York, NY, April 16, 2019.
- ACM Conference on Fairness, Accountability, and Transparency, Atlanta, GA, January 29, 2019.

Artificial Intelligence and Machine Learning.

- United Nations Data Innovation Lab, Nairobi, Kenya, March 13, 2017.

Introduction to Business Analytics.

- IEEE International Conference on Acoustics, Speech, and Signal Processing, Kyoto, Japan, March 25, 2012.

Invited Talks

Individual and Collective Human Agency in the Face of 'AI.'

- Society Seminar, Lucy Family Institute for Data & Society, University of Notre Dame, South Bend, IN, April 2, 2026.
- Center for Statistics and the Social Sciences Seminar, University of Washington, Seattle, March 4, 2026.

Azadi in the age of AI.

- India Science Festival, Pune, India, January 11, 2026.

A Synergistic Future for AI and Ecology.

- *Keynote*, IEEE International Conference on BioInspired Processing, San Isidro de El General, Costa Rica, December 3, 2025.

Human Agency in the Face of AI Agency.

- *Keynote*, NeurIPS Workshop on Socially Responsible and Trustworthy Foundation Models, Mexico City, Mexico (virtual), November 30, 2025
- *Keynote*, Deep Learning Indaba Workshop on Human-AI Interaction in the Global Majority, Kigali, Rwanda, August 22, 2025.

Paisa Fund AI Works: Enabling Cooperatives to Decolonialize AI.

- *Keynote*, Deep Learning Indaba Workshop on Trust AI, Kigali, Rwanda, August 21, 2025.

Introducing Granite Guardian: Safe and Responsible AI Use from GenAI Risks.

- Packt Machine Learning Summit, July 16, 2025.

Toward a Systems Theory for Human-Centered Trustworthy Agentic AI.

- Machine Learning Research Network, University of Sydney, Australia, April 14, 2025.
- Electrical and Computer Engineering Distinguished Colloquium Series, University of Illinois at Urbana-Champaign, Urbana, IL, February 20, 2025.

Harm Detectors and Guardian Models for LLMs: Implementations, Uses, and Limitations.

- *Keynote*, IEEE Big Data Conference Workshop on Trustworthy Machine Learning for Healthcare, December 18, 2024.
- *Keynote*, NeurIPS Workshop on Algorithmic Fairness through the Lens of Metrics and Evaluation, December 14, 2024.

Artificial Intelligence and Glimpses of Moral Psychology.

- Brotherhood of St. Andrew Webinar, December 4, 2024.

The Alignment Problem, Solutions, and Limits.

- Neuro-Symbolic AI Seminar, Alan Turing Institute, November 27, 2024.
- *Keynote*, IEEE New Trends in Signal Processing Conference, Rajkot, India, July 25, 2024.
- Indian Institute of Technology Delhi, New Delhi, India, July 18, 2024.
- *Keynote*, IEEE ISIT Workshop on Information-Theoretic Methods for Trustworthy Machine Learning, Athens, Greece, July 7, 2024.

Trustworthy Machine Learning.

- NIH Bridge2AI Discussion Forum on Emerging Ethical, Legal, and Social Issues, September 17, 2024.
- Data Science BK21xERC Seminar, Seoul National University, Seoul, Korea, April 16, 2024.

- Laboratory for Applied Mathematics, Numerical Software, and Statistics, Argonne National Laboratory, April 3, 2024.
- IEEE Hawaii Section, July 12, 2023.
- Computer Science Department Colloquium, Iowa State University, March 27, 2023.
- Distinguished Lecture Series, Computer Science Research Institute, Sandia National Laboratory, Albuquerque, NM, February 20, 2023.
- Dalian Maritime University, November 23, 2002.
- Quant Seminar, U. S. Securities and Exchange Commission, September 14, 2022.
- *Keynote*, IEEE International Workshop on Machine Learning for Signal Processing, August 24, 2022.
- Innovation Forum, Office of Computational Science, US Food and Drug Administration, August 23, 2022.
- Jour-Fixe Talk, Prosus AI, June 30, 2022.
- QuantUniversity, March 29, 2022.

A Carative Approach to AI Governance.

- IIT Kharagpur Student Chapter, IEEE Signal Processing Society, May 3, 2024.
- North Jersey Chapter, IEEE Signal Processing Society, March 27, 2024.
- Distinguished Lecture, Department of Electrical and Computer Engineering, University of California, Santa Barbara, February 23, 2024.
- Expert Talk Series, Merlyn Mind, December 6, 2023.
- Center for Artificial Intelligence in Society Seminar, University of Southern California, November 9, 2023.
- Water Resources Mission Area Seminar, United States Geologic Survey, June 7, 2023.
- Heinz College, Carnegie Mellon University, Pittsburgh, PA, November 7, 2022.
- AI COI and AIME Joint Seminar on Artificial Intelligence and Machine Learning, National Institute of Standards and Technology, May 3, 2022.
- *Keynote*, ICLR Workshop on Practical ML for Developing Countries, April 29, 2022.
- Institute for Assured Autonomy, Johns Hopkins University, April 19, 2022.
- Annual Winter Workshop, Department of Statistics, University of Florida, January 15, 2022.
- Responsible AI Systems and Experiences Seminar, University of Washington, December 3, 2021.

A Pattern for Making Impact in Trustworthy Machine Learning through Basic Research, Open Source, Product, and Deployment.

- *Spotlight*, ICASSP Industry Program, April 15, 2024.

Sociocultural Diversity and the Machine Learning Development Lifecycle.

- ICASSP Promoting Diversity in Signal Processing (PROGRESS) Workshop, April 14, 2024.

Safe and Trustworthy Foundation Models.

- Ridgefield Library, Ridgefield, Connecticut, October 25, 2023.
- CTO Data Science Speaker Series, Bloomberg, New York, NY, July 18, 2023.
- *Keynote*, RBCDSAI Conference on Deployable AI, June 6, 2023.

Decolonial AI Alignment.

- Trinity College Dublin, Ireland, October 18, 2023.

Do Asian-American Machine Learning Researchers Work on Social Justice?

- Math Department and Asian Coalition of Professionals Colloquium, University at Albany, State University of New York, Albany, NY, October 14, 2022.

The Watsons Meet Watson: A Call for Svadharma and Carative AI.

- University at Albany, State University of New York, Albany, NY, September 22, 2022.

On Algorithmic Disgorgement.

- Future of Privacy Forum AI Working Group, September 12, 2022.

Precautionary Human and Machine Decision Making, and Humble Trust.

- Data Science Institute, Lancaster University, June 22, 2022.

Interpretable Machine Learning for Safety and Teaming

- Interpretability in Artificial Intelligence (22w5055), Banff International Research Station for Mathematical Innovation and Discovery, May 3, 2022.

A Unified View of Trustworthy AI with the 360 Toolkits

- *Keynote*, Open Data Science Conference East, Boston, MA, April 21, 2022.

AI Fairness 360.

- Fraunhofer Institut für Arbeitswirtschaft und Organisation, April 12, 2022.
- Reinforce AI Conference, Budapest, Hungary, March 22, 2019.
- Budapest Data Science Meetup, Budapest, Hungary, March 20, 2019.
- Cognitive Systems Institute Group Speaker Series Webinar, September 20, 2018.

How Can Universities Collaborate with External Organizations and Local Communities to Address Questions of Bias and Discrimination in AI Enabled Technologies?

- Symposium on Artificial Intelligence and Social Responsibility, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, March 23, 2022.

Results in Trustworthy Machine Learning That Go Against Conventional Wisdom.

- *Keynote*, WACV Workshop on Demographic Variations in Performance of Biometrics and Related Technology, January 8, 2022.
- *Keynote*, IJCAI International Workshop on Mining Actionable Insights from Social Networks, August 21, 2021.
- Computation and Society Initiative Colloquium, Yale University, New Haven, CT, February 21, 2020.

Characterizing Threats to Validity in Trustworthy Machine Learning Using Information Theory.

- *Keynote*, ICML Workshop on Information-Theoretic Methods for Rigorous, Responsible, and Reliable Machine Learning, July 24, 2021.

Artificial Intelligence, Data Science, and Global Pandemics.

- Summer Undergraduate Research Program, New York University, July 12, 2021.
- Artificial Intelligence Summit, School of Informatics and Telecommunications, Duoc UC, Santiago, Chile, June 1, 2021. (held virtually)
- COVID-19 Policy Discussion Series, Cornell Institute for Public Affairs, May 28, 2020.

AI Safety in Cyber-Physical Systems, Decision Sciences, and Data Products.

- International Workshop on Signal and Information Intelligent Learning and Processing, Xi'an, China, July 11, 2021. (held virtually)
- QuantPort, New York, NY, September 13, 2018.
- Inside the Black Box Conference, New York, NY, June 16, 2018.
- Uehiro-Carnegie-Oxford Ethics Conference, New York, NY, May 17, 2018.

Lessons from Bottom of the Pyramid Innovation for AI for Social Good

- Catalyst Tech Start-Up Accelerator Talk, Oxford Brookes University, July 5, 2021.
- Center for Research for Computation and Society, Harvard University, March 29, 2021.

Developing Trust in Artificial Intelligence and Machine Learning for High-Stakes Applications.

- Goergen Institute for Data Science, University of Rochester, June 23, 2021.
- *Keynote Address*, Enterprise Computing Community Conference, June 14, 2021.
- Kaiser Permanente DS&E Invited Talk Series, May 21, 2021.
- IBM presents: What's next in AI? Webinar Series, May 19, 2021.
- *Keynote Address*, IEEE Green Mountain Section Annual Meeting, Burlington, VT, November 20, 2020. (held virtually)
- International Workshop on Signal and Information Intelligent Learning and Processing, Xi'an, China, August 15, 2020. (held virtually)
- Mind Food Talk, Ajua (formerly mSurvey), Nairobi, Kenya, September 27, 2019.
- Carnegie Mellon University Africa, Kigali, Rwanda, September 19, 2019.
- *Keynote Address*, IEEE Syracuse Section Fellows Night, Syracuse, NY, April 11, 2019.
- Electrical Engineering and Computer Science Seminar, Syracuse University, Syracuse, NY, April 10, 2019.
- Princeton University, Princeton, NJ, March 8, 2019.
- Electrical and Computer Engineering Seminar, University of Virginia, Charlottesville, VA, October 19, 2018.
- New York Artificial Intelligence Meetup, New York, NY, September 27, 2018.

Computational Intelligence, Creativity, and Wisdom for Man and Mankind.

- IDInsight, Nairobi, Kenya, October 30, 2019.
- New England Complex Systems Institute Winter School, Cambridge, MA, January 11, 2018.

Unequal Performance Across Groups in Face Image Classification.

- SAMSI Discussion Series on Controversial Topics in Machine Learning and Precision Medicine, Research Triangle Park, NC, February 27, 2019.

On Making Machine Learning Safe.

- PDT Partners, New York, NY, June 8, 2017.
- IEEE International Conference on Signal Processing and Integrated Networks, Noida, India, February 3, 2017.
- Computer Science Colloquium, University at Albany, Albany, NY, November 7, 2016.
- *Keynote Lecture*, INFORMS Southeast Michigan Symposium, Rochester, MI, May 14, 2016.

Data for Good.

- School of Social Policy and Practice, University of Pennsylvania, Philadelphia, PA, March 30, 2017.

Learning Classification Rules via Boolean Compressed Sensing with Application to Workforce Analytics.

- Optimization Seminar, Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, January 22, 2015.
- Communications Seminar, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, November 3, 2014.
- Interdisciplinary Distinguished Seminar Series, Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC, October 4, 2013.

Detecting Poverty with Satellite Imagery.

- Nielsen Machine Learning Webinar, December 16, 2016.
- *Keynote Address*, Big Data Summit, University of Illinois Research Park, Champaign, IL, November 5, 2014.
- Data Analysts for Social Good Webinar, October 9, 2014.

Proactive Retention among IBM Growth Market Employees.

- Using Analytics to Optimize Your Workforce Seminar, Analytics Solutions Center, Washington, DC, March 12, 2013.

Margin-Based Classification and Dimensionality Reduction Using Geometric Level Sets.

- École Centrale Paris, Châtenay-Malabry, France, July 1, 2011.
- General Electric Global Research Center, Niskayuna, NY, January 21, 2010.
- MITRE Corporation, Bedford, MA, January 14, 2010.
- NASA Jet Propulsion Laboratory, Pasadena, CA, December 10, 2009.
- MIT Lincoln Laboratory, Lexington, MA, November 24, 2009.
- IBM Thomas J. Watson Research Center, Yorktown Heights, NY, November 16, 2009.
- Naval Research Laboratory, Washington, DC, November 10, 2009.
- National Institute of Standards and Technology, Gaithersburg, MD, November 9, 2009.
- Information, Systems and Networks Seminar, Cornell University, Ithaca, NY, November 6, 2009.
- Schlumberger-Doll Research Center, Cambridge, MA, November 2, 2009.

Supervised Classification in Sensor Networks.

- Communications, Networking, Signal and Image Processing Seminar, Purdue University, West Lafayette, IN, October 6, 2010.
- Business Analytics and Mathematical Sciences Department Seminar, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, October 4, 2010.

Frugal Hypothesis Testing and Classification.

- Massachusetts Institute of Technology, Cambridge, MA, February 8, 2010.

Level Set Margin-Based Classification.

- National Security Engineering Division Seminar, Lawrence Livermore National Laboratory, Livermore, CA, August 11, 2009.
- Decision Modeling Research Initiative Seminar, MIT Lincoln Laboratory, Lexington, MA, January 16, 2009.
- Johns Hopkins University Applied Physics Laboratory, Laurel, MD, January 13, 2009.

Wide-Angle SAR Image Formation with Sparsifying Regularization.

- Decision Modeling Research Initiative Seminar, MIT Lincoln Laboratory, Lexington, MA, January 30, 2007.

- Dipartimento di Ingegneria dell'Informazione, Facoltà di Ingegneria, Università di Pisa, Italy, July 17, 2006.

Panels

Panelist, Beyond Accuracy: Building Trustworthy and Responsible AI, Amazon Trusted AI Symposium, New York, NY, January 21, 2026.

Panelist, AAAI Panel Discussion on Factuality and Trustworthiness, December 11, 2025.

Panelist, Modern Artificial Intelligence in the Research Process, IEEE International Conference on BioInspired Processing, San Isidro de El General, Costa Rica, December 3, 2025.

Moderator, Developing AI with Integrity: Strategic Data and Model Decisions, Raise the Bar, Notre-Dame-IBM Technology Ethics Lab, New York, NY, October 29, 2025.

Keynote panelist, Navigating the Confusion of AI, IAPP Navigate, Portsmouth, NH, June 26, 2025.

Panelist, AI in the Boardroom: Navigating the Challenge of Ethical, Regulatory, and Operational Risks & Opportunities, Private Directors Association, New York, NY, April 21, 2025.

Panelist, Mapping Opportunities and Risks of Generative AI, New York Fed Innovation Conference, Federal Reserve Bank of New York, New York, NY, April 4, 2025.

Panelist, The Future of AI: Ensuring U.S. National Competitiveness Through Strategic (Cross-Sectoral) Partnerships, AI and National Competitiveness Innovation Summit, Albany, NY, February 28, 2025.

Panelist, Responsible AI, *Keynote Panel*, PyTorch Conference, San Francisco, CA, September 18, 2024.

Panelist, Operationalizing Responsible AI, Global Digital Development Forum, June 12, 2024.

Panelist, XAI Beyond Academia: AI Governance and Policymaking, ACM CHI Workshop on Human-Centered Explainable AI, May 12, 2024.

Panelist, ACM Panel on Ethical Quandaries in AI-ML: Facing the Tough Questions, December 14, 2023.

Panelist, Entrepreneurship and Innovation, The Benefits and Challenges of Open Source: Fostering Innovation, Igniting Growth, and the Promise of Equitable Design, New York, NY, September 29, 2023.

Panelist, Generative AI in the Enterprise: Risks and Opportunities, EAIGG: Ethical AI Governance Group, September 6, 2023.

Panelist, AI Africa Mini-Conference, Nairobi, Kenya, May 9, 2023.

Panelist, ChatGPT, LLMs and AI: Opportunities and Ethical Challenges, Grove School of Engineering, City College of New York, April 27, 2023.

Panelist, Robust and Fair AI, IEEE International Conference on Cognitive Machine Intelligence, December 16, 2022.

Panelist, Responsible AI, International Conference on AI-ML Systems, Bangalore, India, October 15, 2022.

Panelist, Context is Everything: Requirements for Test, Evaluation, Validation and Verification (TEVV), Building the NIST AI Risk Management Framework: Workshop #2, March 31, 2022.

Panelist, Building a Career in AI/ML, New York University AI School, January 11, 2022.

Panelist, Bias, Fairness, and Ethics in Biometrics, WACV Workshop on Demographic Variations in Performance of Biometrics and Related Technology, January 8, 2022.

Panelist, AI + Dermatology, AI + Health Conference, Stanford University, December 8, 2021.

Panelist, Interdisciplinary Research in Trustworthy ML -- Challenges and Way Forward, The Invisible Elephant in the Room: The Trustworthy ML Un-Symposium, October 28, 2021.

Moderator, Trustworthy AI and Lending, Ethics and Explainability for Responsible Data Science Workshop, University of Johannesburg, October 27, 2021.

Panelist, Technical Characteristics of AI System Trustworthiness, NIST AI Risk Management Framework Workshop, October 20, 2021.

Panelist, KDD Workshop on Measures and Best Practices for Responsible AI, August 15, 2021.

Panelist, ICML Workshop on Information-Theoretic Methods for Rigorous, Responsible, and Reliable Machine Learning, July 24, 2021.

Panelist, Deployment & Impact of AI for Social Good, CMU Symposium on Artificial Intelligence for Social Good, April 30, 2021.

Panelist, Algorithmic Fairness and Social Justice, Ingram Olkin Forum Series, National Institute of Statistical Sciences, September 25, 2020.

Panelist, Algorithmic Bias is in the Question, Not the Answer: Measuring and Managing Bias Beyond Data, NIST Bias in AI Workshop, August 18, 2020.

Panelist, Model Explainability Forum, TWIML AI Podcast, August 11, 2020.

Panelist, Data Privacy and Ethics, Yale Economic Development Symposium, New Haven, CT, February 21, 2020.

Moderator, AI for Social Good, Yale Economic Development Symposium, New Haven, CT, February 21, 2020.

Panelist, AAI Artificial Intelligence Diversity, Belonging, Equity, and Inclusion Workshop, New York, NY, February 7, 2020.

Panelist, NetHope Workshop, AI for Good Global Summit, Geneva, Switzerland, May 31, 2019.

Panelist, Data and Partnerships, AI XPRIZE Team and Red Judge Summit, Geneva, Switzerland, May 31, 2019.

Panelist, ‘Living with AI’ Future Trends Forum, Fundación Innovación Bankinter, Madrid, Spain, May 28–30, 2019.

Panelist, Internet Policy Revolution, AI and Transparency, IBM IT Services Legal Summit, New York, NY, October 15, 2018.

Panelist, SDGs and Beyond: Leveraging Technology for Development, The Youth Assembly, New York, NY, August 12, 2018.

Moderator, Machine Learning that Matters Now, Stockholm, Sweden, July 14, 2018.

Panelist, Ethics and Fairness in Artificial Intelligence Think Tank, IBM Think, Las Vegas, NV, March 22, 2018.

Panelist, IBM Watson Health Opioid Summit, Cambridge, MA, January 17, 2018.

Panelist, Human Interpretability in Machine Learning, Sydney, Australia, August 10, 2017.

Moderator, Rapid-Fire Introduction to Data Science for Social Good Organizations and Opportunities, New York, NY, June 24, 2016.

Panelist, Stories from the Frontlines of Data Science for Good, New York, NY, October 27, 2015.

Panelist, IEEE Spectrum Forecasters Panel, 2013–2015.

Demonstrations

Automatic Generation of Factsheets for Trusted AI in a Runtime Environment.

- NeurIPS Expo, Montreal, Canada, December 2, 2018.

SellerScope: Interactive Salesforce Analytics.

- IBM Information on Demand EXPO, Las Vegas, NV, October 21–24, 2012.
- IBM Investor Briefing, Yorktown Heights, NY, May 9, 2012.
- IBM Innovation Lab, Lotusphere, Orlando, FL, January 15–19, 2012.

Teaching and Mentoring Experience:

Industry faculty, IEX Labs, *IE University*, since May 2026.

Adjunct project advisor, DS-GA 1003 Machine Learning and Computational Statistics, *New York University*, Spring 2015, Spring 2016.

Doctoral committee member

- Saleh Afroogh, *University at Albany, State University of New York*
- Lu Cheng, *Arizona State University*
- Sanghamitra Dutta, *Carnegie Mellon University*
- Shubham Sharma, *University of Texas*
- Yihua Zhang, *Michigan State University*

Mentor or Manager

- Zijian (Jason) Ding, Ph.D. candidate, *University of Maryland*, Intent-Based Outcome Specification, Summer 2024.
- Veronika Spielmann, B.S. candidate, *Duale Hochschule Baden-Wuerttemberg Stuttgart*, Exploratory Network Analysis of Risk Diffusion in Open Foundation Model Ecosystems, Spring 2024.
- Caroline Johnston, Ph.D. candidate, *University of Southern California*, Social Choice in Aligning Large Language Models, Summer 2023.
- Lucas Monteiro Paes, Ph.D. candidate, *Harvard University*, In-Context Source Attribution in Large Language Models, Summer 2023.
- Ikechukwu Daniel Adebisi, Ph.D. candidate, *University of Texas*, Reinforcement Learning for Generative AI to Avoid Taboo Topics, Summer 2023.
- Matthew Pisano, M.S. candidate, *Rensselaer Polytechnic Institute*, Aligning Large Language Models Using Text-Based Games, Summer 2023.
- Shrey Jain, M.S. candidate, *Rensselaer Polytechnic Institute*, Red-Teaming Platform for Auditing Large Language Models, Summer 2023.
- Julia Clare Arrington, Ph.D. candidate, *Rensselaer Polytechnic Institute*, Auditing Large Language Models for Social Biases, Summer 2023.
- Chhavi Yadav, Ph.D. candidate, *University of California San Diego*, Bias and Robustness in Natural Language Inference, Summer 2022.
- Katy Ilonka Gero, Ph.D. candidate, *Columbia University*, Data Curation Practices for Large Language Models, Summer 2022. (Now with Massachusetts Institute of Technology.)
- Sourya Basu, Ph.D. candidate, *University of Illinois*, Equivariance and Fine-Tuning of Large Language Models, Summer 2022.
- Siddharth Vishwanath, Ph.D. candidate, *Pennsylvania State University*, Topology-Regularized Machine Learning, Summer 2022.
- Yonas Sium, Ph.D. candidate, *Iowa State University*, Graph Representation Learning, Summer 2022.
- Manish Nagireddy, B.S. candidate, *Carnegie Mellon University*, Function Composition in Trustworthy Machine Learning, Summer 2022. (Now with IBM.)
- Brianna Richardson, Ph.D. candidate, *University of Florida*, Explaining Fairness Metrics with Influence Functions, Summer 2021, Summer 2022. (Now with Humana.)
- Kofi Arhin, Ph.D. candidate, *Rensselaer Polytechnic Institute*, Challenges with Annotating Toxic Text Datasets, Summer 2021, Summer 2022.
- Isha Puri, B.S. candidate, *Harvard University*, Interpretable Machine Learning Architectures, Spring 2021, Summer 2022.
- Ezinne Nwankwo, Ph.D. candidate, *University of California Berkeley*, Bandit Algorithms for Food, Financial, and Housing Security Among Vulnerable Populations, Summer 2021.
- Aditya Mate, Ph.D. candidate, *Harvard University*, Bandit Algorithms for Food, Financial, and Housing Security Among Vulnerable Populations, Summer 2021. (Now with Microsoft.)
- Conor Artman, Ph.D. candidate, *North Carolina State University*, Bandit Algorithms for Food, Financial, and Housing Security Among Vulnerable Populations, Summer 2021.
- Devleena Das, Ph.D. candidate, *Georgia Institute of Technology*, Explainable Artificial Intelligence, Summer 2021.
- Blossom Metevier, Ph.D. candidate, *University of Massachusetts*, Algorithmic Fairness in Dynamic Environments, Summer 2021.
- Charvi Rastogi, Ph.D. candidate, *Carnegie Mellon University*, The Role of Cognitive Biases and Uncertainty in AI-Assisted Decision-Making, Summer 2020, Summer 2021.
- Hannah Kim, Ph.D. candidate, *Duke University*, Out-of-Distribution Detection in Dermatology, Summer 2020.

- Ifrah Idrees, Ph.D. candidate, *Brown University*, Anomalous Patterns in Maternal, Newborn and Child Health, Summer 2020.
- Ching-Yao Chuang, Ph.D. candidate, *Massachusetts Institute of Technology*, Fairness via Interpolation, Summer 2020.
- Abhin Shah, Ph.D. candidate, *Massachusetts Institute of Technology*, Treatment Effect Estimation Using Invariant Risk Minimization, Summer 2020.
- Jun Wang, Ph.D. candidate, *Rensselaer Polytechnic Institute*, Theoretical Analysis of Invariant Risk Minimization, Summer 2020. (Now with United Imaging Intelligence.)
- Renzhe Yu, Ph.D. candidate, *University of California Irvine*, Understanding Education-Occupation Alignment with NLP Techniques, Summer 2020. (Now with Columbia University.)
- Lu Cheng, Ph.D. candidate, *Arizona State University*, Automated Meta-Analysis, Summer 2020. (Now with University of Illinois Chicago.)
- Diego Garcia-Olano, Ph.D. candidate, *University of Texas*, Biomedical Interpretable Entity Representations, Summer 2020. (Now with Meta.)
- Saksham Chawla, M.S. candidate, *Rensselaer Polytechnic Institute*, AI Platforms for the Bottom of the Pyramid, Summer 2020. (Now with Macy's.)
- Jagrati Sharma, M.S. candidate, *Rensselaer Polytechnic Institute*, AI Platforms for the Bottom of the Pyramid, Summer 2020. (Now with Morningstar.)
- Sulbha Aggarwal, B.S. candidate, *Queens College*, AI Platforms for the Bottom of the Pyramid, Summer 2020. (Now with Bank of America.)
- Smruthi Ramesh, M.S. candidate, *Northeastern University*, Extracting Evidence of Drug Repurposing from Scientific Papers, Summer 2020. (Now with Verisk.)
- Sejal Dua, B.S. candidate, *Tufts University*, Extracting Evidence of Drug Repurposing from Scientific Papers, Summer 2020. (Now with National Basketball Association.)
- Anastasia Spangler, A.A. candidate, *Bellevue College*, Extracting Evidence of Drug Repurposing from Scientific Papers, Summer 2020.
- Gevorg Ghalachyan, M.S. candidate, *Yerevan State University*, Trustworthy AI Pentathlon, Summer–Fall 2019. (Now with DataArt.)
- Peter Yefi, M.S. candidate, *Carnegie Mellon University – Africa*, Trustworthy AI Pentathlon, Summer 2019. (Now with Complete Farmer.)
- Timothy Odonga, M.S. candidate, *Carnegie Mellon University – Africa*, Fairness in Dermatology, Summer 2019. (Now with Dalberg.)
- Newton Kinyanjui, M.S. candidate, *Carnegie Mellon University – Africa*, Fairness in Dermatology, Summer 2019.
- Shubham Sharma, Ph.D. candidate, *University of Texas*, Eliciting Fairness Definitions from Policymakers, Summer 2019. (Now with JPMorgan Chase.)
- Sanghamitra Dutta, Ph.D. candidate, *Carnegie Mellon University*, Detection Theory-Based Analysis of Accuracy, Fairness, and Explainability, Summer 2019. (Now with University of Maryland.)
- Shivashankar Subramanian, Ph.D. candidate, *University of Melbourne*, Extracting Evidence of Drug Repurposing from Scientific Papers, Summer 2019. (Now with Amazon.)
- Akanksha Atrey, Ph.D. candidate, *University of Massachusetts*, Causal Inference to Address the Opioid Crisis, Summer 2019. (Now with Nokia Bell Labs.)
- Vivek Gupta, Ph.D. candidate, *University of Utah*, Contrastive Explanations for Natural Language Processing, Summer 2019. (Now with University of Pennsylvania.)
- Xiufan Yu, Ph.D. candidate, *Pennsylvania State University*, Optimal Pathways out of Poverty, Summer 2019. (Now with University of Notre Dame.)
- Oscar Chang, Ph.D. candidate, *Columbia University*, Creating New Antibiotics of Last Resort, Summer–Fall 2018. (Now with Google.)

- Amanda Coston, Ph.D. candidate, *Carnegie Mellon University*, Fair Financial Inclusion and Transfer Learning, Summer 2018.
- Niccolò Dalmasso, Ph.D. candidate, *Carnegie Mellon University*, Pruning Neural Networks for Multiple Objectives, Summer 2018. (Now with JPMorgan Chase.)
- Vidya Muthukumar, Ph.D. candidate, *University of California, Berkeley*, Color Blind Neural Networks, Summer 2018. (Now with Georgia Institute of Technology.)
- Chirag Nagpal, Ph.D. candidate, *Carnegie Mellon University*, Causal Inference to Address the Opioid Crisis, Summer 2018. (Now with Google.)
- Ritesh Noothigattu, Ph.D. candidate, *Carnegie Mellon University*, Winning Video Games Morally, Summer 2018. (Now with Hudson River Trading.)
- Ravi Kiran Raman, Ph.D. candidate, *University of Illinois*, Trusted Models and Results, Summer 2018. (Now with Analog Devices.)
- Evan Patterson, Ph.D. candidate, *Stanford University*, Open Discovery Platform for a Multiple Sclerosis Cure and Semantic Understanding of Data Science Code, Summer 2016; Summer 2017–Spring 2018. (Now with Topos Institute.)
- Tejas Dharamsi, M.S. candidate, *Columbia University*, Neurology-as-a-Service for the Developing World, Fall 2017. (Now with Twitter.)
- Jonathan Galsurkar, M.S. candidate, *Columbia University*, Smarter Sustainable Development, Summer–Fall 2017. (Now with eBay.)
- Anand Doshi, M.S. candidate, *University of Michigan*, Emergency Food Best Practice, Summer 2017. (Now with Square.)
- Aditya Garg, M.S. candidate, *Columbia University*, Demystifying Social Entrepreneurship, Summer 2017. (Now with Tesla.)
- Bernat Guillén Pegueroles, Ph.D. candidate, *Princeton University*, Cognitive Policy Advisor, Summer 2017. (Now with Google.)
- Yaoli Mao, Ph.D. candidate, *Columbia University*, User Experience of Open Discovery Platforms, Summer 2017. (Now with Autodesk.)
- Bryce Melvin, B.S. candidate, *University of Colorado*, Changing Behaviors to Conserve Energy, Summer 2017. (Now with DreamWorks Animation.)
- Timothy NeCamp, Ph.D. candidate, *University of Michigan*, Cognitive Disease Hunter, Summer 2017. (Now with Data Bloom.)
- Minh Nguyen, Ph.D. candidate, *University of Southern California*, SimplerVoice: Overcoming Illiteracy, Summer 2017. (Now with Formosa Interactive.)
- Hrishikesh Rao, Ph.D. candidate, *University of Michigan*, Emergency Food Best Practice, Summer 2017.
- Jinghe Zhang, Ph.D. candidate, *University of Virginia*, Combating the Opioid Crisis, Summer 2017. (Now with Apple.)
- Wenyu (Wendy) Zhang, Ph.D. candidate, *Cornell University*, Cognitive Financial Advisor for Low-Wage Workers, Summer 2017. (Now with Institute for Infocomm Research.)
- Yuanshuo (David) Zhao, Ph.D. candidate, *Georgia Institute of Technology*, Ask Nature for Design Inspiration, Summer 2017. (Now with Apple.)
- Caitlin Kuhlman, Ph.D. candidate, *Worcester Polytechnic Institute*, How to Foster Innovation, Summer 2016. (Now with Apple.)
- Hemank Lamba, Ph.D. candidate, *Carnegie Mellon University*, Understanding the Ecospace of Philanthropic Projects, Summer 2016. (Now with Dataminr.)
- Subhabrata Majumdar, Ph.D. candidate, *University of Minnesota*, Hunting Zika with Machine Learning, Summer 2016. (Now with Splunk.)
- Kien Pham, Ph.D. candidate, *New York University*, Real-Time Understanding of Humanitarian Crises, Summer 2016. (Now with Meta.)

- Yumeng Tao, Ph.D. candidate, *University of California Irvine*, Disseminating the Best Treatment for Diarrhea, Summer 2016. (Now with PyTorch.)
- Guolong Su, Ph.D. candidate, *Massachusetts Institute of Technology*, Interpretable Two-Level Rule Learning, Summer 2015. (Now with Google.)
- Michael Wattendorf, B.S. candidate, *Princeton University*, Predicting Outcomes of Individual Tennis Points, Spring 2015–Spring 2016. (Now with Waymo.)
- Ben Quazzo, B.S. candidate, *Princeton University*, Statistical Analysis of Professional Tennis, Spring 2015. (Now with Accel Partners.)
- Shaobo Han, Ph.D. candidate, *Duke University*, Nonparametric Bayesian Modeling of Legislation for Exploratory Analysis, Summer 2014. (Now with NEC Laboratories America.)
- Jinfeng Yi, Ph.D. candidate, *Michigan State University*, Expertise Assessment Recommendation, Summer 2013. (Now with JD AI Research.)
- Alex Gittens, Ph.D. candidate, *California Institute of Technology*, Voluntary Attrition Modeling, Summer 2012. (Now with Rensselaer Polytechnic Institute.)
- Gautam K. Bhat, M.Sc. candidate, *Karnataka State Open University*, Deconvolving the Productivity of Salespeople, Fall 2011–Summer 2012. (Now with Deutsche Bank.)
- John Z. Sun, Ph.D. candidate, *Massachusetts Institute of Technology*, Dynamic Matrix Factorization and Collaborative Filtering based on the Kalman Filter, Summer 2011. (Now with PDT Partners.)
- Ankan Saha, Ph.D. candidate, *University of Chicago*, Semi-Supervised Multi-Task Learning with Task-Dependent Regularization, Summer 2010. (Now with LinkedIn.)

Guest lecturer

- PADM-GP 4322 Data and AI Strategies for Social Impact Organizations, *New York University*, Spring 2022, Spring 2023, Spring 2024, Spring 2025, Spring 2026.
- CSCE 581 Trusted AI, *University of South Carolina*, Spring 2025, Spring 2026.
- COMS 5710X Responsible AI: Risk Management in Data Driven Discovery, *Iowa State University*, Spring 2026.
- Detecting and Addressing Bias in Data, Humans, and Institutions Research Practicum, *Stony Brook University*, Spring 2025.
- 6.S093 How to Ship Almost Anything with AI, *Massachusetts Institute of Technology*, Independent Activities Period (Winter) 2025.
- PADM 5472 Leveraging Information Technology for Public and Nonprofit Management, *Cornell University*, Spring 2019, Spring 2020, Spring 2021, Spring 2022, Spring 2024.
- LAW 644 National Security Lawyering, *Syracuse University*, Spring 2022.
- CIS 731 Artificial Neural Networks, *Syracuse University*, Fall 2020.
- HON 137 Honors Seminar: Understanding Algorithmic Bias, *State University of New York at Plattsburgh*, Fall 2020.
- MSFS 709 Tech4Development: Applying Technology to Sustainability Challenges, *Georgetown University*, Fall 2020.
- INFO-UB 0001 Information Technology in Business and Society, *New York University*, Spring 2019.
- CS-GY 9223 Foundations of Data Science, *New York University*, Spring 2016.
- 01:090:101 Data: What is it Good For? (Absolutely Something), *Rutgers University*, Fall 2015.
- ORF 360 Decision Modeling in Business Analytics, *Princeton University*, Spring 2015.

Teaching assistant, 6.972 Algorithms for Estimation and Inference, *Massachusetts Institute of Technology*, Fall 2006.

Group tutor, ECE 310 Introduction to Probability and Random Phenomena, *Cornell University*, Spring 2004.

Guest instructor, How Television Works, *Tompkins Seneca Tioga BOCES Community (alternative) School*, Spring 2004.

Service and Professional Activities:

Founding co-director: IBM Science for Social Good initiative, 2015–2023.

Board membership: External Advisory Committee, Goergen Institute for Data Science and Artificial Intelligence, University of Rochester, since 2025; Board of Directors, AI Risk and Vulnerability Alliance, since 2023; Board of Directors, CriticaLink, 2016.

Data ambassador: DataKind, Brooklyn, NY, USA, since August 2013.

Expert: AI & Faith, since May 2026.

Advisor: ALCF Concept to Clinic Lung Cancer Early Detection Challenge Hosted by DrivenData, 2017.

Professional societies: IEEE, since 2000 (fellow since 2024; senior member since 2015); ACM, 2013–2016; INFORMS, 2012; ISIF, 2009–2011; SIAM, 2006–2010; SPIE, 2006–2007; IEE/IET, 2005–2006.

Elected professional committees and working groups: ACM Conference on Fairness, Accountability and Transparency (FAccT) Steering Committee, since 2023; Partnership on AI Working Group on Safety-Critical AI, 2018–2019; IEEE Signal Processing Theory and Methods Technical Committee, 2015–2017; IEEE Machine Learning for Signal Processing Technical Committee, 2014–2016.

Journal editorial board: Digital Signal Processing: A Review Journal, 2013–2016.

Proposal review panel: Directorate for Computer and Information Science and Engineering, National Science Foundation, 2019.

Workshop and symposium organizer: 2020–2025 IEEE CVPR Workshop on Fair, Data Efficient and Trusted Computer Vision; 2023 ACM FAccT Doctoral Colloquium; 2021 INFORMS Annual Meeting Session on Participatory Specification of Trustworthy Machine Learning; 2016–2018, 2020 ICML Workshop on Human Interpretability (WHI) in Machine Learning; 2019 ICML Workshop on Human in the Loop Learning (HILL); 2019 AAAS Annual Meeting Symposium on Blockchain and the Scientific Method; 2017 IEEE GlobalSIP Symposium on Signal and Information Processing for Finance and Business; 2017 ACM SIGKDD Workshop on Machine Learning for Creativity; 2016 ICML Workshop on #data4good: Machine Learning in Social Good Applications.

General chair: AAAI/ACM Conference on Artificial Intelligence, Ethics and Society (AIES), 2025.

Program chair: AAAI/ACM Conference on Artificial Intelligence, Ethics and Society (AIES), 2024.

Track chair: ACM Conference on Fairness, Accountability and Transparency (FAT*), 2020.

Senior area chair: ACM Conference on Fairness, Accountability and Transparency (FAccT), 2023–2026; AAAI Conference on Artificial Intelligence (Social Impact Track), 2023.

Journal reviewer: ACM Computing Surveys; ACM Transactions on Interactive Intelligent Systems; Communications of the ACM; Computer Methods in Biomechanics and Biomedical Engineering; Computers in Human Behavior; Data Mining and Knowledge Discovery; Digital Signal Processing (2 articles); Ecology Letters; Electronic Journal of Statistics; Entropy; Future Business Journal; IBM Journal of Research and Development (2 articles); IEE Proceedings Vision, Image and Signal Processing; IEEE Geoscience and Remote Sensing Letters; IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing; IEEE Journal of Selected Topics in Signal Processing (5 articles); IEEE Potentials Magazine (several articles); IEEE Signal Processing Letters (5 articles); IEEE Signal Processing Magazine (3 articles); IEEE Transactions on Aerospace and Electronic Systems; IEEE Transactions on Communications; IEEE Transactions on Cybernetics; IEEE Transactions on Image Processing (2 articles); IEEE Transactions on Multimedia; IEEE Transactions on Information Theory (2 articles); IEEE Transactions on Signal Processing (9 articles); Information Sciences; Inverse Problems and Imaging; Journal of Artificial Intelligence Research; Journal of Machine Learning Research (11 articles); Mayo Clinic Proceedings: Digital Health; Nature Machine Intelligence (3 articles); Pattern Recognition; Remote Sensing Letters; Service Science (2 articles); Signal Processing (2 articles); Socio-Economic Planning Sciences; Stochastic Models; Technology in Society.

Conference program committee and reviewer: IEEE International Conference on Acoustics, Speech, and Signal Processing, 2014–2017, 2019–2020, 2026; IEEE International Workshop on Statistical Signal Processing, 2011, 2025; AAAI AI Alignment Track, 2025; International Conference on Machine Learning, 2012–2014, 2018, 2024 (ethics reviewer); ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2015–2018, 2024; ACM Web Conference, 2024; Conference on Neural Information Processing Systems, 2016–2018, 2020, 2021–2022 (ethics reviewer), 2023; AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society, 2023; International Joint Conference on Artificial Intelligence, 2023; IEEE International Workshop on Statistical Signal Processing, 2012, 2016, 2023; ACM Conference on Fairness, Accountability and Transparency, 2019, 2021–2022; ICLR Workshop on Practical ML for Developing Countries, 2020–2022; IEEE International Symposium on Information Theory, 2014, 2022; AAAI Workshop on Interactive Machine Learning, 2022; NeurIPS Workshop on Algorithmic Fairness through the Lens of Causality and Interpretability, 2020; NeurIPS Workshop on Machine Learning for Development, 2019; KDD Workshop on Adversarial Learning Methods for Machine Learning and Data Mining, 2019; IEEE/IFIP DSN Workshop on Dependable and Secure Machine Learning, 2018–2019; ICLR Workshop on Debugging Machine Learning Models, 2019; ICLR Workshop on Safe Machine Learning: Specification, Robustness, and Assurance, 2019; International Conference on Artificial Intelligence and Statistics, 2019; ACM SIGCAS Conference on Computing and Sustainable Societies, 2018; International Conference on Learning Representations, 2018; IEEE ICDM Workshop on Data Science for Human Capital Management, 2017; IEEE International Workshop on Machine Learning for Signal Processing, 2014–2017; IEEE International Conference on Intelligent Transportation Systems, 2017; International Conference on Signal Processing and Integrated Networks, 2015–2017; NIPS Workshop on Interpretable Machine Learning in Complex Systems, 2016; Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing, 2015; IEEE Sensor Array and Multichannel Signal Processing Workshop, 2014; IEEE International Conference on Service Operations and Logistics, and Informatics, 2013; IEEE GlobalSIP Symposium on Signal and Information Processing in Finance and Economics, 2013; European Signal Processing Conference, 2009, 2012; International Conference on Industrial, Engineering, and Other Applications of Applied Intelligent Systems, 2011.

Book proposal reviewer: Springer Nature, 2022; MIT Press, 2021; Manning Publications, 2020.

Conference session chair: ACM Conference on Fairness, Accountability and Transparency, 2023; AAAI Conference on Artificial Intelligence, 2023; IEEE International Workshop on Machine Learning for

Signal Processing, 2016; IEEE International Workshop on Statistical Signal Processing, 2012; INFORMS Annual Meeting, 2011.

Writer: Information Ashvins Blog (<http://informationashvins.wordpress.com>), since 2010; LIDS-Blog (<http://lidsblog.typepad.com>), 2007–2010.

U.S. Co-chair, Trustworthy AI Standing Committee, Indo-U.S. Science and Technology Forum, 2021.

Technical advisor, Fighter Steel Productions LLC, 2018.

Network member, *Council for Big Data, Ethics, and Society*, January–June 2016.

Organizing committee member, *Ebola Open Data Jam*, October 2014 and February 2015.

Summer intern coordinator, *Business Analytics and Mathematical Sciences Department*, IBM Thomas J. Watson Research Center, May–October 2013.

Organizing committee chair, *Stochastic Systems Symposium in Honor of Alan Willsky on the Occasion of his Sixtieth Birthday*, May 2008.

Organizer, *Stochastic Systems Group Seminars*, Fall 2007–Spring 2008.

Executive committee member, *MIT EECS Graduate Students Association*, 2006–2007.

Conference committee member, *LIDS Student Conference*, January 2005 and 2006.

Selected Coursework:

IBM Research: Technical Leadership; MicroMBA Program I & II.

Massive Open Online Courses: Quantitative Methods in Clinical and Public Health Research (HarvardX); The Challenges of Global Poverty (MITx).

Massachusetts Institute of Technology: Detection, Estimation, and Stochastic Processes; Discrete-Time Signal Processing; Optimization Methods; Machine Learning (listener); Representation and Modeling for Image Analysis; Computational Biology: Genomes, Networks, Evolution; Language and Its Structure II: Syntax; Language and Its Structure III: Semantics and Pragmatics.

Cornell University: Introduction to Computer Graphics; Signal Representation and Modeling; Digital Signal Processing; Telecommunications Systems I & II; Introduction to Scientific Computation; Geographic Information Systems; Introduction to Scanning Electron Microscopy; Computer Organization.

References:

Available upon request.