

NCCoE Hybrid Workshop on Mitigating AI Bias in Context

August 31, 2022 from 11:00 AM to 3:50 PM EST

Speaker Biographies



Keynote Speaker: Dr. Itiel Dror

Itiel Dror (PhD Harvard) is a cognitive neuroscientist who is interested in the cognitive architecture that underpins expert decision making. Dror's research, published in over 150 research articles, demonstrated how contextual information can influence judgments and decision making of experts. He has worked in a variety of domains, from policing and aviation to medical experts

and bankers, showing that even hard working and competent experts can make unintentional errors in evaluating data. Dror has worked with many agencies in various countries to minimize error and enhance decision making.

Recently Science magazine wrote a feature describing Dr. Itiel Dro's research:

<https://www.science.org/content/article/forensic-experts-biased-scientists-claims-spark-outrage>

Further background reading:

<https://pubs.acs.org/doi/10.1021/acs.analchem.0c00704>



Dr. Alexander Amini – Themis AI

Alexander Amini is the co-founder and chief scientific officer of Themis AI, a CSAL MIT spinoff with a mission to empower the world to create, advance and deploy trustworthy AI. Amini is currently a Postdoctoral Research Associate at the [Massachusetts Institute of Technology \(MIT\)](#), in the [Computer Science and Artificial Intelligence Laboratory \(CSAIL\)](#), with [Prof. Daniela Rus](#). He completed his PhD (2022), Master of Science (2018), and Bachelor of Science (2017) in [Computer Science](#) at MIT, with a minor in Mathematics. The objective of his research is to develop the science and engineering of autonomy and its applications to safe decision

making for autonomous agents. His vision is a world with adaptive autonomous agents capable of learning to interact in complex, uncertain, and extreme scenarios, supporting people with cognitive and physical tasks.




Mislav Balunović - ETH Zurich

Mislav is a fourth year PhD student at ETH Zurich, advised by Prof. Martin Vechev. Before that he completed a master's degree in computer science at ETH Zurich in 2019. The main goal of his research is to develop techniques for training trustworthy machine learning models, more specifically models that make provably fair decisions, are robust to input changes, and use as little private user data as possible.



Laura Blattner – Stanford University

Laura Blattner is an Assistant Professor of Finance at the Stanford Graduate School of Business. Her research interests are banking, corporate finance, and macroeconomics. Laura recently earned her Ph.D. at Harvard University. She also holds a B.A. in Philosophy, Politics and Economics and an M.Phil. in Economics from Oxford University.



Jieshan Chen – CSIRO

Jieshan Chen is a research scientist of software engineering for AI research team at CSIRO's Data61, Australia. She finished her PhD degree in Computer Science at Australian National University in 2022. Prior to Ph.D., Jieshan received her Bachelor's Degree from Sun Yat-Sen University. Her work lies in the fields of software engineering, deep learning, and human-computer interaction. Her current research focus is software engineering for AI-based systems, responsible AI, UI semantic understanding for responsible UIs and accessibility enhancement in mobile applications.



Miroslav Dudík – Microsoft

Miroslav Dudík is a Senior Principal Researcher Manager in machine learning at Microsoft Research, NYC. His research focuses on combining theoretical and applied aspects of machine learning, statistics, convex optimization, and algorithms. Most recently he has worked on contextual bandits, reinforcement learning, and algorithmic fairness. He received his PhD from Princeton in 2007. He is a co-creator of the Fairlearn toolkit for assessing and improving the fairness of machine learning models and of the Maxent package for modeling species distributions, which is used by biologists around the world to design national parks, model the impacts of climate change, and discover new species.



Nakia Grayson – NIST

Nakia Grayson is an IT Security Specialist who leads Supply Chain Assurance & Autonomous Vehicle project efforts at the National Cybersecurity Center of Excellence (NCCoE), which is part of the National Institute of Standards and Technology (NIST). She is also a part of the Privacy Engineering Program at NIST, where she supports the development of privacy risk management best practices, guidance, and communications efforts. Nakia serves as the Contracting Officer Representative for several NIST cybersecurity contracts. She holds a bachelor's in criminal justice from University of Maryland-Eastern Shore and a master's in information technology, information assurance and business administration from the University of Maryland University College.



Cathy O'Neil – ORCAA

Cathy has been an independent data science consultant since 2012 and has worked for clients including the Illinois Attorney General's Office and Consumer Reports. She wrote the book *Doing Data Science* in 2013 and *Weapons of Math Destruction: How Big Data Increases Inequality And Threatens Democracy*, released in September 2016.



Jann Spiess- Stanford University

Jann works on integrating techniques and insights from machine learning into the econometric toolbox. His research brings together microeconomic methods, statistical decision theory, and mechanism design to clarify the use of flexible prediction algorithms in causal inference and data-driven decision-making. He is particularly interested in the role of human and machine decisions in replicable and robust inferences from big data.



Paula-Rose Stark- FinRegLab

Paula-Rose Stark is FinRegLab's Director of Machine Learning Research. In that role, she leads the organization's efforts to develop and execute policy-relevant research on the use of AI and machine learning in financial services. That includes ongoing research being conducted with researchers at the Stanford Graduate School of Business on the explainability and fairness of machine learning underwriting models. This research explores the capabilities and performance of available model diagnostic tools in helping lenders comply with a variety of regulations that apply irrespective of the kind of model a lender opts to use to make consumer credit decisions.

Prior to joining FinRegLab, P-R served as an advisor to bank and nonbank financial institutions, most recently in the Finance and Risk Practice of Oliver Wyman. Her practice focused on helping firms operationalize compliance with the full range of federal consumer protection laws and regulations, including prohibitions on unfair, deceptive, or abusive acts or practices. P-R was also among the first employees of the Consumer Financial Protection Bureau, where she, among other things, she shaped early matters in the Enforcement Division and developed strategic priorities and shared frameworks for evaluating market risks for the division of Supervision, Enforcement, and Fair Lending.

Paula-Rose is a graduate of Princeton, Oxford, and Harvard Universities.



Nicholas Schmidt - SolasAI

Nicholas Schmidt is the CEO of SolasAI, a compliance-focused AI software platform that identifies and mitigates bias and discrimination in algorithmic decisioning. He is also the Artificial Intelligence Practice Leader at BLDS, LLC, where he provides expert guidance in the application of economics and statistics to questions of law and regulation. As head of the AI practice, Nick focuses on algorithmic fairness, explainable AI, and ensuring robust model governance practices. In addition to working with many of the largest U.S. lenders, FinTechs, and insurance companies, Nick regularly advises regulatory agencies in addressing questions relating to discrimination and innovation in AI.



Reva Schwartz – NIST

Reva Schwartz is a research scientist in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) where she serves as Principal Investigator on Bias in Artificial Intelligence for NIST's Trustworthy and Responsible AI program. Her research focuses on the role of context in human language and behavior and the nature of expertise and expert judgment in socio-technical systems. Reva has advised federal agencies about how experts interact with automation to make sense of information under high-risk and high-uncertainty operational conditions.



Kevin Stine – NIST

Mr. Kevin Stine is the Chief of the Applied Cybersecurity Division in the National Institute of Standards and Technology's Information Technology Laboratory (ITL). He is also NIST's Chief Cybersecurity Advisor and Associate Director for Cybersecurity in NIST's ITL. In these roles, he leads NIST collaborations with industry, academia, and government to improve cybersecurity and privacy risk management through the effective application of standards, best practices, and technologies. The Applied Cybersecurity Division develops cybersecurity and privacy guidelines, tools, and reference architectures in diverse areas such as public safety communications; health information technology; smart grid, cyber physical, and industrial control systems; and programs focused on outreach to small businesses and federal agencies. The Division is home to several priority programs including the National Cybersecurity Center of Excellence, Cybersecurity Framework, Cybersecurity for IoT, Identity and Access Management, Privacy Engineering and Risk Management, and the National Initiative for Cybersecurity Education.



Dr. Kush R. Varshney – IBM

Kush R. Varshney was born in Syracuse, NY in 1982. He received the B.S. degree (magna cum laude) in electrical and computer engineering with honors from Cornell University, Ithaca, NY, in 2004. He received the S.M. degree in 2006 and the Ph.D. degree in 2010, both in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge. While at MIT, he was a National Science Foundation Graduate Research Fellow.

Dr. Varshney is a distinguished research scientist and manager with IBM Research at the Thomas J. Watson Research Center, Yorktown Heights, NY, where he leads the machine learning group in the Foundations of Trustworthy AI department. He was a visiting scientist at IBM Research - Africa, Nairobi, Kenya in 2019. He is the founding co-director of the IBM Science for Social Good initiative. He applies data science and predictive analytics to human capital management, healthcare, olfaction, computational creativity, public affairs, international development, and algorithmic fairness, which has led to recognitions such as the Gerstner Award for Client Excellence for contributions to the WellPoint team, the Extraordinary IBM Research Technical Accomplishment for contributions to workforce innovation and enterprise transformation, and Falling Walls Science and Innovation winner for AI Fairness 360. He conducts academic research on the theory and methods of trustworthy machine learning. His work has been recognized through best paper awards at the Fusion 2009, SOLI 2013, KDD 2014, and SDM 2015 conferences and the 2019 Computing Community Consortium / Schmidt Futures Computer Science for Social Good White Paper Competition. He published the book entitled 'Trustworthy Machine Learning' (<http://www.trustworthymachinelearning.com/>) in 2022. He is a senior member of the IEEE.



Dr. Apostol Vassilev – NIST

Dr. Apostol Vassilev supervises a Research Team at NIST. His team focuses on a wide range of AI problems: AI bias identification and mitigation, meta learning with large language models for various NLP tasks, robustness and resilience of AI systems, AI applications to cybersecurity. Apostol is a keynote speaker and has authored over fifty peer-reviewed publications and multiple patents. He leads the development of standards on AI and cybersecurity. He is a coauthor of the recent NIST SP 1270: Towards a Standard for Identifying and Managing Bias in Artificial Intelligence, <https://doi.org/10.6028/NIST.SP.1270> . Apostol is a technologist by scientific background, but he is also interested in social aspects of using AI technology and advocates for a comprehensive socio-technical approach to evaluating AI's impact on individuals and society.



Sam Yoon – Harvard Kennedy School of Government

Sam Yoon is a Fulbright Scholar doing his master's in public policy at the Harvard Kennedy School of Government. His work focuses on AI system standards and audits. His relevant experiences include working as a Public Policy Manager at Meta and being a research associate at BCG with their responsible AI team.