Building the NIST AI Risk Management Framework: Workshop #3

Day 1 Speakers

Elham Tabassi
National Institute of Standards and Technology
Elham Tabassi is the Chief of Staff in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST). She leads NIST Trustworthy and Responsible AI program that aims to cultivate trust in the design, development, and use of AI technologies by improving measurement science, standards, and related tools in ways that enhance economic security and improve quality of life. She has been working on various machine learning and computer vision research projects with applications in biometrics evaluation and standards since she joined NIST in 1999. She is a member of the National AI Resource Research Task Force, a senior member of IEEE, and a fellow of Washington Academy of Sciences.

Panel 1: How Does the AI RMF Fit In with Other AI Risk Management Approaches?

Mark Latonero
National Institute of Standards and Technology
Mark Latonero is a senior policy advisor for AI and international cooperation at the National Institute of Standards and Technology (NIST). He is a member of the OECD Network of Experts on AI and a research professor at Georgetown. Recently, he was a senior consultant at the UN Secretary General's Office, senior policy advisor at the Partnership on AI, and fellow at the Harvard Kennedy School. He created the data and human rights program at the Data & Society Research Institute and led the technology & human trafficking initiative at the USC Annenberg School, where he was a research director. Mark completed his PhD at the University of Southern California specializing in the social implications of emerging technologies and was a postdoctoral scholar at the London School of Economics.

Nozha Boujemaa
Global Vice President, Digital Ethics and Responsible AI - IKEA Retail (Ingka Group), Co-Chair of the OECD.AI Trustworthy
Nozha Boujemaa was Vice-Chair of the European Commission's High Level Expert Group on Artificial Intelligence, where she led the development of the "Ethical Guidelines for Trustworthy AI." She also coordinated the work on Trustworthy AI principles at the OECD, where today, as co-chair of the "OECD.AI Network of Experts working group on implementing trustworthy AI", she translates these principles into actions. She founded the interdisciplinary AI institute DATAIA and was the director of Inria Saclay Research Center, where she also supervised more than 25 PhDs in computer vision. Previously, Dr. Boujemaa served as Chief Science and Innovation Officer for AI in the health Industry and has held numerous positions leading teams in areas such as large-scale information retrieval in media archives, earth observation, biodiversity, as well as personalised medicine. Dr. Boujemaa holds a PhD and HDR in Computer Science, is a Knight of

the French National Order of Merit, and a proud mother of a data scientist and a data analyst, as well as the wife of a digital strategist.

Greg Cannon
Amazon Web Services
Greg has a Master's degree in electrical and computer engineering, and has worked as an engineer for 30 years, primarily specializing in signal and image processing, software architecture, and system design. Greg is the principal inventor for over 40 US and International patents, primarily developed in the fields of mobile communications and biometrics. Greg has been in the biometrics Industry for over 20 years. He has developed matching algorithms, multimodal fusion software, transaction routing servers, enrollment systems, and verification software for physical access control terminals and single sign-on applications. He has worked in standards development for the US and International Standards Bodies, serving as editor and liaison for various standards, and currently is convener for the ISO working group for the technical implementation of biometric systems. Greg is particularly interested in technology development, both from a corporate perspective and from an industry perspective. Greg serves as a technology innovation expert to help solve technical issues, identify use cases, and develop business cases that derive from the technology value adoption cycle.

Sri Krishnamurthy
Founder, QuantUniversity
Sri Krishnamurthy is the founder of www.QuantUniversity.com, an AI advisory and Quantitative Analysis Company, specializing in large-scale Machine Learning and AI solutions focused on adoption of responsible AI products in the industry. Prior to starting QuantUniversity, Sri has worked at Citigroup, Endeca, MathWorks and with more than 25 customers in the financial services and energy industries. He has trained more than 10,000 students in quantitative methods, analytics and big data in the industry and at Babson College, Northeastern University and Hult International Business School. Sri earned an MS in Computer Systems Engineering and another MS in Computer Science, both from Northeastern University specializing in robotics and agent-based systems and an MBA with a focus on Investments from Babson College.
Some of Sri's publications are available here: https://www.quantuniversity.com/publications.html

Lee Wan Sie
Infocomm Media Development Authority (IMDA), Singapore
Lee Wan Sie is Director for Development of Data-Driven Tech at Singapore's Infocomm Media Development Authority. In the area of AI, her responsibilities include driving Singapore's approach to AI governance, growing the trustworthy AI ecosystem in Singapore and collaborating with governments around the world to further the development of responsible AI. She is also responsible for encouraging greater use of emergent data technologies, such as privacy enhancing tech, to enable more trusted use and sharing of data in Singapore. Wan Sie has extensive experience in technology use in the public sector – prior to her current role, she developed Singapore's strategies for Digital Economy and Smart Nation, enabled tech experimentation and innovation as head of IDA Labs, and implemented government digital services that helped Singapore

become one of the easiest places in a world to do business. Outside of Singapore, she has worked with governments around the world in their national digitalisation efforts. Wan Sie actively supports Women in Tech in Singapore, to encourage young women to take up a career in technology.

Anneke Olvera
Standards Council of Canada
Anneke Olvera is the Director at the Standards Council of Canada (SCC) responsible for Programs and Operations in the Strategy and Stakeholder Engagement branch. Anneke has over 25 years of experience providing policy advice and guidance on standardization to governments, industry, academia and civil society. She currently leads various teams in implementing government funded programs focused on delivering standardization strategies and solutions in areas such as Infrastructure and Climate Change, Innovation and IP, Data Governance, and Artificial Intelligence, and Mental Health and Substance Use Health. Anneke is the Secretary of the Canadian Data Governance Standardization Collaborative, which published its first Roadmap mapping the standardization landscape for data governance in Canada. SCC will soon be launching its third Collaborative focused on Artificial Intelligence.

PANEL 2: EXACTLY WHAT IS AN AI RMF PROFILE?

Cherilyn Pascoe
National Institute of Standards and Technology
Cherilyn Pascoe is Senior Technology Policy Advisor at the National Institute of Standards and Technology (NIST), U.S. Department of Commerce. She advises NIST leadership on technology policy and strategy, including cybersecurity, privacy, and artificial intelligence. She also leads the NIST Cybersecurity Framework program and is active in the NIST AI Risk Management Framework development. Prior to joining NIST, she served more than a decade in staff leadership roles on the U.S. Senate Committee on Commerce, Science, and Transportation working for former Senator Hutchison (R-TX), Senator Thune (R-SD), and current Ranking Member Wicker (R-MS). Most recently, she served as Deputy Policy Director managing the Committee's Space and Science Subcommittee, which has legislative and oversight jurisdiction over science, technology, standards, and civil space policy. During her time on the Hill, she led efforts to develop and advance several notable pieces of legislation, including the U.S. Innovation and Competition Act, the AV Start Act, as well as three surface transportation reauthorization laws and ten cybersecurity laws. Pascoe received her M.A. in International Science and Technology Policy from the George Washington University and her B.S. Chem. with Highest Honors in Chemistry from the University of Michigan.

Tony Barrett
UC Berkeley
Dr. Barrett is a Visiting Scholar with the Center for Long-Term Cybersecurity (CLTC) AI Security Initiative (AISI) at UC Berkeley, and a Senior Policy Analyst at the Berkeley Existential Risk Initiative (BERI). He is also Co-Founder and Director of Research of the Global Catastrophic Risk Institute (GCRI). Dr. Barrett was recently Lead for the Advanced

Analytics – Machine Learning area in ABS Group's Product Development and Innovation (PD&I) Center, where his project clients included the US Department of Defense (DOD) and Department of Homeland Security (DHS). He has over 20 years experience on a range of topics, which includes machine learning / artificial intelligence systems and standards, software development, cybersecurity, nuclear security, risk assessment, and project management. Dr. Barrett's published work includes papers in the peer-reviewed journals Risk Analysis, Decision Analysis, Science & Global Security, and the Journal of Experimental & Theoretical Artificial Intelligence. He has a Bachelor's degree in engineering from UC San Diego, and a Ph.D. in engineering and public policy from Carnegie Mellon University; he also was an ASME/AAAS fellow to DHS and a Stanton nuclear security fellow at RAND. He is based in the Washington, D.C. area.

Phillip Collett
American Express
Phil Collett has more than 15 years' experience in Financial Services technology and cybersecurity in a wide range of hands-on roles that have included network administration, production support, architecture, information security, and risk management. Phil is an active contributor within the Financial Services Information Security Community, presenting regularly in industry forums. He serves on the Board of Directors for the Cyber Risk Institute which developed the CRI Financial Sector Risk Profile. In addition to several technical and cybersecurity certifications, Phil holds a Bachelor's in Economics from the University of Utah, a Master of Information Security from Carnegie Mellon University. In his current role as VP Technology Risk, Phil is responsible for global IT/IS policies and standards, risk management, technology operational excellence, business information security officers, security control enforcement, and risk reporting across all American Express lines of business. His organization helps ensure safe and sound banking and business operations via embedded partnerships focused on tailored risk management for business units and legal entities. Phil is a voting member on the American Express Model Tech Risk Management Committee, Compliance Risk Committee, Privacy Risk Committee, and New Product Committee. Other responsibilities include management of operational risk events, regulatory readiness assessments, technology compliance, and cybersecurity assessments on all new products and major initiatives globally.

Carlos Ignacio Gutierrez
Future of Life Institute
Carlos Ignacio Gutierrez is an artificial intelligence (AI) policy researcher at the Future of Life Institute. His research agenda focuses on the examination of the mechanisms, issues, incentives, and effectiveness of hard and soft law approaches to the governance of AI methods and applications. His research on this area is published in peer-reviewed journals, Brookings, OECD, among others (see here). He has a Ph.D. in policy analysis from the Pardee RAND Graduate school.

Snigdha Sharma
National Fair Housing Alliance
Snigdha Sharma is a Tech Equity Analyst and FAIR Ops Team lead at the National Fair Housing Alliance (NFHA). Snigdha conducts research on

regulations, and laws applicable to algorithmic systems. She also reviews research articles, regulations and laws that are applicable to work within the Tech Equity department to assure solutions developed within the department are fair and equitable. She leads and manages projects that focus on Privacy, Security, Fairness, Explainability and Interpretability elements of Responsible AI ecosystem in addition to her work on environmental impact of technology. She oversees auditing services centered on the Purpose, Process and Monitoring (PPM) algorithmic auditing framework developed by NFHA.

Panel 3: Moderated Discussion: What We Heard and What We Hope to Hear

Courtney Lang
Information Technology Industry Council (ITI)
Courtney Lang is Senior Director of Policy for Trust, Data, and Technology at the Information Technology Industry Council. In this role, Courtney leads ITI's policy development and global advocacy on Artificial Intelligence and serves as an issue expert on ICT supply chain security policy, engaging with government in the U.S. and abroad to shape and guide implementation of supply chain security efforts. Courtney also focuses on other cybersecurity issues, including those related to 5G and next generation network security. Prior to joining ITI, Courtney was a trade policy analyst at the U.S. Department of Commerce's International Trade Administration (ITA), where she worked to understand the priorities of U.S. technology companies and helped them to overcome trade barriers in markets overseas. Her areas of focus included mobile & wireless communications (including 5G), cybersecurity, and the Internet of Things. During her time at ITA, Courtney worked to promote policies that would facilitate trade and support innovation. For example, she established and led the Trade in the Digital Economy Working Group under the U.S.-Brazil Commercial Dialogue, aimed at addressing digital trade barriers in Brazil, as well as a long-term project to promote risk-based approaches to cybersecurity in the Asia-Pacific. She also worked on developing and executing a strategy aimed at supporting U.S. competitiveness in the 5G telecom technology sector. Courtney holds an MA in Security Studies from Georgetown University and a BA in Political Science from Duke University.

Elham Tabassi
National Institute of Standards and Technology
Elham Tabassi is the Chief of Staff in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST). She leads the NIST Trustworthy and Responsible AI program that aims to cultivate trust in the design, development, and use of AI technologies by improving measurement science, standards, and related tools in ways that enhance economic security and improve quality of life. She has been working on various machine learning and computer vision research projects with applications in biometrics evaluation and standards since she joined NIST in 1999. She is a member of the National AI Resource Research Task Force, a senior member of IEEE, and a fellow of Washington Academy of Sciences.

Mark Latonero
National Institute of Standards and Technology
Mark Latonero is a senior policy advisor for AI and international cooperation at the National Institute of Standards and Technology (NIST). He is a member of the OECD Network of Experts on AI and a research professor at Georgetown. Recently, he was a senior consultant at the UN Secretary General's Office, senior policy advisor at the Partnership on AI, and fellow at the Harvard Kennedy School. He created the data and human rights program at the Data & Society Research Institute and led the technology & human trafficking initiative at the USC Annenberg School, where he was a research director. Mark completed his PhD at the University of Southern California specializing in the social implications of emerging technologies and was a postdoctoral scholar at the London School of Economics.

Cherilyn Pascoe
National Institute of Standards and Technology
Cherilyn Pascoe is Senior Technology Policy Advisor at the National Institute of Standards and Technology (NIST), U.S. Department of Commerce. She advises NIST leadership on technology policy and strategy, including cybersecurity, privacy, and artificial intelligence. She also leads the NIST Cybersecurity Framework program and is active in the NIST AI Risk Management Framework development. Prior to joining NIST, she served more than a decade in staff leadership roles on the U.S. Senate Committee on Commerce, Science, and Transportation working for former Senator Hutchison (R-TX), Senator Thune (R-SD), and current Ranking Member Wicker (R-MS). Most recently, she served as Deputy Policy Director managing the Committee's Space and Science Subcommittee, which has legislative and oversight jurisdiction over science, technology, standards, and civil space policy. During her time on the Hill, she led efforts to develop and advance several notable pieces of legislation, including the U.S. Innovation and Competition Act, the AV Start Act, as well as three surface transportation reauthorization laws and ten cybersecurity laws. Pascoe received her M.A. in International Science and Technology Policy from the George Washington University and her B.S. Chem. with Highest Honors in Chemistry from the University of Michigan.

Lori Perine
National Institute of Standards and Technology
Lori A. Perine is an associate researcher in the Information Technology Laboratory at NIST and a doctoral candidate at the University of Maryland's iSchool. Her research explores sociotechnical approaches to AI design, development, and deployment in various domains, as well as associated public policy. Ms. Perine has had a substantive career as a STEM policy-maker, alliance executive, and consultant, focused on translating innovation to advance societal goals. She has served on a Presidential Transition Team, held executive positions at the White House Office of Science and Technology Policy and in international technology alliances, and sat on intra- and intergovernmental councils and advisory groups for STEM and energy. A dedicated advocate for women and underrepresented minorities in STEM fields, Lori also consults as a strategic advisor to AnitaB.org.

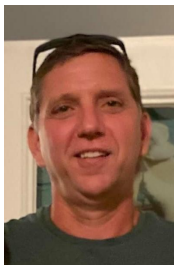

Mark Przybocki
National Institute of Standards and Technology
Mark Przybocki is the Chief of the Information Access Division (IAD) in the Information Technology Laboratory at NIST. He joined NIST in 1993 and has led evaluation-driven research projects for speech transcription, speaker and language recognition, machine translation and information extraction. IAD has a strong focus on AI measurement and evaluation, and is home to high visibility programs such as TREC (Text REtrieval Conference) and FRVT (Face Recognition Vendor Test). Mark has over 30 years of federal experience in test and evaluation for human language technologies, and is leading IAD as we launch new evaluation efforts to promote trustworthy AI.

Reva Schwartz
National Institute of Standards and Technology
Reva Schwartz is a research scientist in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST). 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 expertise and expert judgment in socio-technical systems. She has advised federal agencies about how experts interact with automation to make sense of information in high-stakes settings. Her background includes a forensic science posting for almost 15 years at the United States Secret Service, advising forensic science practice while at NIST, a temporary duty assignment at the National Security Agency, and adjunct researcher at the Johns Hopkins University Human Language Technology Center of Excellence.

Building the NIST AI Risk Management Framework: Workshop #3

Day 2 Speakers

Elham Tabassi
National Institute of Standards and Technology
Elham Tabassi is the Chief of Staff in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST). She leads NIST Trustworthy and Responsible AI program that aims to cultivate trust in the design, development, and use of AI technologies by improving measurement science, standards, and related tools in ways that enhance economic security and improve quality of life. She has been working on various machine learning and computer vision research projects with applications in biometrics evaluation and standards since she joined NIST in 1999. She is a member of the National AI Resource Research Task Force, a senior member of IEEE, and a fellow of Washington Academy of Sciences.

Panel 4: Mapping the Path from Current to Future AI Risk Management Practice

Reva Schwartz
National Institute of Standards and Technology
Reva Schwartz is a research scientist in the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST). 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 expertise and expert judgment in socio-technical systems. She has advised federal agencies about how experts interact with automation to make sense of information in high-stakes settings. Her background includes a forensic science posting for almost 15 years at the United States Secret Service, advising forensic science practice while at NIST, a temporary duty assignment at the National Security Agency, and adjunct researcher at the Johns Hopkins University Human Language Technology Center of Excellence.

Stevie Bergman
DeepMind
Stevie is a Senior Research Scientist on the Ethics Research team at DeepMind, where she explores questions at the intersection of technology and society. She has a global, human rights focus in her work and typically investigates questions as to the impacts of technology on marginalized and in-conflict communities outside the US and Europe. Her current work is on the topics of data governance and representation, meaningful participation, and the effective evaluation of AI systems. Stevie's research often has direct implications for both ML practitioners and policy. Stevie has further expertises in AI fairness and annotation from her prior work, including that on Meta's Responsible AI Fairness team. At Meta, in addition to analyzing "Integrity" and Arabic systems for capability fairness, she developed and taught company-wide seminars on AI fairness. Stevie received her technical research training during her physics doctorate, and she gained years of vital, on-the-ground experience working directly with communities as a Fulbright scholar (Indonesia) and Peace Corps volunteer (Uganda). Stevie has over a decade of experience in science and tech communication including a podcast series on AI & Human Rights and 100+ interviews on WPRB 103.3FM.

Sina Fazelpour
Northeastern University and NIST AI Visiting Fellow
Sina Fazelpour is an Assistant Professor of Philosophy and Computer Science at Northeastern University. His research focuses on issues of justice, diversity, and reliability in data-driven and artificial intelligence technologies. He also works on understanding the concepts and consequences of diversity in social groups and networks. He is a core member of the Institute for Experiential AI and NULab for Texts, Maps, and Networks at Northeastern University, as well as the 2020-21 Council Fellow on the World Economic Forum's Global Future Council on Data Policy. Before joining Northeastern, he was a SSHRC Postdoctoral Fellow in the Department of Philosophy at Carnegie Mellon University, with a secondary affiliation with the Machine Learning Department. In addition to a Ph.D. in Philosophy, he holds an M.Sc in medical biophysics and a B.Eng in electrical and biomedical engineering.

Bogdana Rakova
Mozilla Foundation
Bogdana Rakova is a Senior Trustworthy AI Fellow at Mozilla Foundation. She is an engineer, researcher, inventor, and organizer. Her work investigates the complex issues at the intersection of people, technology, trust, transparency, accountability, and social and environmental justice. Previously, she was a research manager at Accenture's Responsible AI team, a mentor at the Assembly Ethics and Governance of AI program led by Harvard's Berkman Klein Center for Internet and Society, and has also held positions with Partnership on AI, Samsung Research America, and others. She was part of the core team at the IEEE P7010 working group developing a recommended practice for assessing the impact of AI on human well-being and consequently joined the board at the Happiness Alliance nonprofit. Her work weaves together inspirations from feminist studies and speculative fiction in building tangible hyperlocal prototypes.

Harini Suresh
Massachusetts Institute of Technology (MIT)
Harini is a 6th year doctoral student in Computer Science at MIT, where she is part of Clinical and Applied Machine Learning Group, the Visualization Group, and the Data + Feminism Lab. Her research sits at the intersection of machine learning (ML) and human-computer interaction (HCI), and broadly aims to make the ML lifecycle more participatory and grounded in context. Her work includes both higher-level frameworks---for example, that identify important choices and potential sources of harm throughout the ML lifecycle---as well as concrete tools that give diverse stakeholders the agency to shape, probe, and contest ML-based systems. Her recent work asks how to shape the ML process from the start, and involves applying participatory methods with activist organizations throughout development and deployment, to co-design datasets and ML models that prioritize on-the-ground expertise and context. Previously, Harini completed her B.Sc. and M.Eng. at MIT in Computer Science, studying machine learning systems in healthcare contexts. She is also passionate about education, data/ML literacy, and community involvement. She is the co-creator of the MIT AI Ethics Discussion group, a forum to facilitate conversation and collaboration across disciplines, and ML Tidbits, an educational non-profit intended to empower non-experts to understand and discuss ML and its societal consequences. As part of the effort to integrate societal and ethical considerations into CS curricula, she has developed course modules for ML classes to teach topics such as subjectivity in categorization schemes and annotation.

Panel 5: What Does Good Governance for AI Risk Management Look Like?

Kathy Baxter
Salesforce and NIST AI Visiting Fellow
As a Principal Architect of Ethical AI Practice at Salesforce, Kathy develops research-informed best practices to educate Salesforce employees, customers, and the industry on the development of responsible AI. She collaborates and partners with external AI and ethics experts to continuously evolve Salesforce policies, practices, and products. She is a member of Singapore's Advisory Council on the Ethical Use of AI and Data, Visiting AI Fellow at NIST, and on the Board of EqualAI. Prior to

Salesforce, she worked at Google, eBay, and Oracle in User Experience Research. She is the co-author of "Understanding Your Users: A Practical Guide to User Research Methodologies." She received her MS in Engineering Psychology and BS in Applied Psychology from the Georgia Institute of Technology. You can follow her on Twitter at @baxterkb and read about the Ethical AI Practice Team's current research.

Rumman Chowdhury
Twitter
Dr. Rumman Chowdhury's passion lies at the intersection of artificial intelligence and humanity. She is a pioneer in the field of applied algorithmic ethics, creating cutting-edge socio-technical solutions for ethical, explainable and transparent AI. She is currently the Director of META (ML Ethics, Transparency, and Accountability) team at Twitter, leading a team of applied researchers and engineers to identify and mitigate algorithmic harms on the platform. Previously, she was CEO and founder of Parity, an enterprise algorithmic audit platform company. She formerly served as Global Lead for Responsible AI at Accenture Applied Intelligence. In her work as Accenture's Responsible AI lead, she led the design of the Fairness Tool, the first-in-industry enterprise algorithmic tool to identify and mitigate bias in AI systems. Dr. Chowdhury is dedicated to cultivating and growing the next wave of practitioners enabling the responsible use of emerging technologies. She is a General Partner (and founder) of the Parity Responsible Innovation venture capital fund. Her new venture, Bias Buccaneers, is a nonprofit organization that develops and hosts algorithmic bias bounties. Their first challenge, 8-bit bias bounty, launches tomorrow at www.biasbounty.ai.

Jen Gennai
Google
Jen founded, and now leads, Google's Responsible Innovation team, which operationalizes Google's AI Principles to ensure that Google's products have fair and ethical outcomes on individuals and society broadly. Her team works with product and engineering teams, leveraging a multidisciplinary group of experts in ethics, human rights, user research, and racial justice to validate that outputs align with our commitments to fairness, privacy, safety, societal benefit, and accountability to people. Before she co-authored the AI Principles with a multi-disciplinary group of experts in 2017 and established the Responsible Innovation team in 2018, Jen led user research in Trust and Safety and founded the product fairness testing team. She has been in Google for 15 years.

Sakshi Jain
LinkedIn
Sakshi Jain is a Senior Engineering Manager at LinkedIn and leads their Responsible AI efforts, focused on making LinkedIn fair, explainable and trustworthy. Prior to this, Sakshi spent 7 years in the Anti Abuse space, building new age machine learning algorithms & systems to protect the platform from malicious activities. She completed her undergraduate from IIT Bombay in Engineering Physics and her Masters from UC Berkeley in Network Security using AI. She is passionate about subjects at the intersection of social impact and AI within and outside work.

David Marcos
Microsoft
David Marcos leads the governance and enablement team within Microsoft's Office of Responsible AI, driving cross-company efforts to institutionalize AI governance, awareness, and training. Prior to his current position, Mr. Marcos led the development of Microsoft's Responsible AI compliance capabilities as part of the Ethics & Society team in Microsoft's Cloud & Artificial Intelligence division. Mr. Marcos was also previously Chief Privacy Officer of Microsoft's Cloud & Artificial Intelligence division, driving governance and privacy engineering solutions for GDPR. Before employment with Microsoft, Mr. Marcos worked for the National Security Agency, holding a variety of positions, including technical director of the NSA Office of Civil Liberties and Privacy, deputy technical director of the NSA Office of the Director of Compliance, and privacy research lead in the NSA Research Directorate. David specializes in governance, privacy, and compliance, focusing on legal automation and ethical computation in cloud technologies. Mr. Marcos holds a B.S. in Computer Engineering from Penn State and an M.S. in Strategic Intelligence from the National Intelligence University. Mr. Marcos is both a Certified Information Privacy Manager and Technologist (CIPM/CIPT). /in/davidjamesmarcos

Panel 6: How to Measure AI Risk across the AI lifecycle

Jeanna Matthews
Clarkson University and NIST Faculty Fellow
Jeanna Matthews is a professor of computer science at Clarkson University. She is a founding Chair of the ACM Technology Policy Subcommittee on Artificial Intelligence and Algorithmic Accountability, a Chair of Institute of Electrical and Electronics Engineers (IEEE) - USA AI Policy Committee, and a member of the ACM Technology Policy Committee. She is an affiliate of the Data and Society Research Institute. She has been a member of the ACM Council (2015-2022), chair of the ACM Special Interest Group Governing Board (2016-2018), the chair of the ACM Special Interest Group on Operating Systems (SIGOPS) (2011-2015), an ACM Distinguished Speaker and an Fulbright Scholar. She has published work in a broad range of systems topics from virtualization and cloud computing to social media security and distributed file systems. She has been a four-time presenter at DEF CON on topics including security vulnerabilities in virtual environments (2015 and 2016), adversarial testing of criminal justice software (2018) and trolling (2018). Her current work focuses on securing societal decision-making processes and supporting the rights of individuals in a world of automation.

Jack Clark
Anthropic
Jack is a co-founder of Anthropic, an AI safety and research company. He is also co-chair of the AI Index at Stanford University, and writes a detailed newsletter about AI research called ImportAI (ImportAI.net).

David Danks
UC San Diego
David Danks is Professor of Data Science & Philosophy and affiliate faculty in Computer Science & Engineering at University of California, San Diego. His research interests range widely across philosophy, cognitive science, and machine learning, including their intersection. Danks has examined the ethical, psychological, and policy issues around AI and robotics in transportation, healthcare, privacy, and security. He has also done significant research in computational cognitive science and developed multiple novel causal discovery algorithms for complex types of observational and experimental data. Danks is the recipient of a James S. McDonnell Foundation Scholar Award, as well as an Andrew Carnegie Fellowship. He currently serves on multiple advisory boards, including the National AI Advisory Committee.

Ian Eisenberg
Credo AI
Ian Eisenberg is Head of Data Science at Credo AI, where he leads the development of Lens, an open-source Responsible AI Assessment Framework. Ian believes safe AI requires systems level approaches that draw on technical, social and regulatory advancements. He is currently focused on building the tools and processes that operationalize Responsible AI in an accessible manner. His interest in AI started as a neuroscientist at Stanford, which developed into a focus on reducing AI risk through his involvement with the Effective Altruism movement and as a machine learning engineer in the hiring space. Ian has been a researcher at Stanford, the NIH, Columbia and Brown University. He received his PhD from Stanford University, and BS from Brown University.

Nazneen Rajani
HuggingFace
Nazneen is a Research Lead at HuggingFace, a startup with a mission to democratize ML, leading data-centric ML research which involves systematically analyzing, curating, and automatically annotating data. Before HF, she worked at Salesforce Research with Richard Socher and led a team of researchers focused on building robust natural language generation systems based on LLMs. She completed her Ph.D. in CS at UT-Austin with Prof. Ray Mooney. Nazneen has over 30 papers accepted at ACL, EMNLP, NAACL, NeurIPs, ICLR and has her research covered by Quanta magazine, VentureBeat, SiliconAngle, ZDNet, and Datanami. She is also teaching a course on Interpreting ML models with CoRise. More details about her work can be found here.

Panel 7: Managing AI Risk

Daniel Castro
Information Technology and Innovative Foundation (ITIF)
Daniel Castro is vice president at the Information Technology and Innovation Foundation (ITIF) and director of ITIF's Center for Data Innovation. Castro writes and speaks on a variety of issues related to information technology and internet policy, including privacy, security, intellectual property, Internet governance, e-government, and accessibility for people with disabilities. His work has been quoted and cited in numerous media outlets, including The Washington Post, The Wall Street Journal, NPR, USA Today, Bloomberg News, and Bloomberg Businessweek. In 2013, Castro was named to FedScoop's list of the "top 25 most influential people under 40 in government and tech." In 2015, U.S. Secretary of Commerce Penny Pritzker appointed Castro to the Commerce Data Advisory Council. Castro previously worked as an IT analyst at the Government Accountability Office (GAO) where he audited IT security and management controls at various government agencies. He contributed to GAO reports on the state of information security at a variety of federal agencies, including the Securities and Exchange Commission and the Federal Deposit Insurance Corporation. In addition, Castro was a visiting scientist at the Software Engineering Institute in Pittsburgh, PA, where he developed virtual training simulations to provide clients with hands-on training of the latest information security tools. He has a B.S. in foreign service from Georgetown University and an M.S. in information security technology and management from Carnegie Mellon University.

Veena Calambur
Workday
Veena Calambur is a Senior Program Manager at Workday's Machine Learning Trust Program. In the past Veena has worked as data scientist at Pfizer where she pioneered their Responsible AI strategy and toolkit and as a decision analytics associate at ZS Associates. Veena has her bachelor's degree in Information Science and Statistics from Cornell University and is pursuing a part-time PhD in Information Science studying the intersection of human-computer interaction and artificial intelligence at Drexel University.

Trey Causey
Indeed
Trey Causey is Head of AI Ethics and Director of Data Science at Indeed, where he and his team work to increase fairness in Indeed's machine learning and artificial intelligence systems. As the owner of Indeed's responsible AI strategy, he leads an interdisciplinary group of data scientists, engineers, researchers, and PMs in tackling the sociotechnical issues of algorithmic, human, and systemic bias. He also is a member of Indeed's Environmental, Social, and Governance leadership team, where he and his colleagues are changing the way the world hires, for good. Prior to Indeed, Trey has led data science teams, machine learning engineering teams, and product teams across the technology industry. Trained as a sociologist at the University of Washington, he confronts challenges in AI ethics with both a sociological and a technological approach.

Ali Shah
Accenture

Ali leads Accenture's efforts to ensure AI is a force for good, working in the interests of citizens and society. Ali's expertise in AI, regulation, policy, data ethics, privacy, and human rights helps ensure Accenture and its clients can understand and practically address the challenges AI and other emerging technologies present. Prior to joining Accenture, Ali was Head of Technology Policy at the UK Information Commissioner's Office (ICO), where he was accountable for regulatory policy related to emerging technologies. His leadership and expertise focused on AI, data ethics, national and international cross regulatory collaboration and investigations, privacy by design engineering, and award-winning work helping designers and engineers re-imagine the internet with children's safety at the core. Ali has also held roles at the BBC, related to emerging technology strategy, engineering, and innovation, and led the development of the BBC's position on AI ethics. Ali has written and spoken extensively about the intersections of AI technology, ethics, regulation, and human values at fora including UK Parliament, the Global Partnership on AI, the Global Privacy Assembly, and Start-up Grind. He is also an invited member of the Council on Extended Intelligence established by the IEEE and MIT. Ali believes we can only realize the positive benefits of AI when Responsibility to society is embedded at the heart of AI design. He is passionate about creating practical and simple steps to ensure AI is a force for good.

Mona Sloane
New York University (NYU)

Mona Sloane, Ph.D. is a sociologist working on design and inequality, specifically in the context of AI design and policy. She is a Research Assistant Professor at NYU's Tandon School of Engineering, a Senior Research Scientist at the NYU Center for Responsible AI, a Fellow with NYU's Institute for Public Knowledge (IPK) and The GovLab, and the Director of the *This Is Not A Drill* program on technology, inequality and the climate emergency at NYU's Tisch School of the Arts. She is principal investigator on multiple research projects on AI and society, and holds an affiliation as postdoctoral scholar with the Tübingen AI Center at the University of Tübingen in Germany where she leads a 3-year federally funded research project on the operationalization of ethics in German AI startups. Mona founded and runs the IPK Co-Opting AI series at NYU and currently serves as editor of the technology section at Public Books. She holds a PhD in Sociology from the London School of Economics and Political Science. Follow her on Twitter @mona_sloane.