

Testimony of

Elham Tabassi

Associate Director for Emerging Technologies

Information Technology Laboratory

National Institute of Standards and Technology United States

Department of Commerce

Before the

United States House of Representatives

Committee on Science, Space, and Technology

Subcommittee on Investigations and Oversight and

Subcommittee on Research & Technology

*Balancing Knowledge and Governance: Foundations for
Effective Risk Management of Artificial Intelligence*

October 18, 2023

Chairman Obernolte, Chairman Collins, Ranking Members Foushee and Stevens and members of the subcommittees, I am Elham Tabassi, Associate Director for Emerging Technologies at the Information Technology Laboratory in the Department of Commerce's National Institute of Standards and Technology (NIST). I also am the lead for NIST's trustworthy and responsible artificial intelligence (AI) program. We appreciate the committee's continued support of our work and thank you for the opportunity to testify today on NIST's efforts to improve the trustworthiness of AI.

NIST is home to five Nobel Prize winners, with programs focused on national priorities such as cybersecurity, advanced manufacturing, semiconductors, the digital economy, precision metrology, quantum information science, biosciences, and artificial intelligence. NIST's mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life.

In the NIST Information Technology Laboratory (ITL), we work to cultivate trust in information technology and metrology. Trust in the digital economy is built on key attributes of technology like cybersecurity, privacy, usability, interoperability, equity, and avoiding bias and increasing usefulness. NIST conducts fundamental and applied research and advances standards to understand and measure limits and capabilities of information technology. We develop tools to measure and evaluate those technologies. These technology standards and measurements – and the foundational and applied research that enables their development and use – are critical to advancing trust in digital products and services. They provide increased assurance and utility, enabling more secure, private, and rights-affirming technologies. More than ever before, our economy and society sorely need and depend on these kinds of standards and measurements. That is especially true when it comes to AI.

We are aware of the enormous recent increase in interest in AI. We know that the world is demanding collaborative, sophisticated frameworks for the development and deployment of trustworthy AI; infrastructure for evaluations of AI systems; and coordinated standards around which AI can be defined, built, and assessed. This is at the heart of NIST's work. We are maximizing the impact of our limited resources to meet the urgent needs of this moment.

NIST's Role in Artificial Intelligence

NIST contributes to the research, standards, measurements, and data required to realize the full promise of AI, while managing its risks, to enable American innovation, enhance economic security, and improve our quality of life.

As a non-regulatory agency, NIST prides itself on the strong partnerships it has cultivated with the public and private sectors. NIST seeks and relies on diverse stakeholder insights and feedback from government, industry, academia, and non-profit entities to develop and improve its resources.

The collaborative, transparent, and open processes NIST uses to prioritize, develop, and carry out its research and to produce its guidelines result in more effective and usable resources that are trusted and, therefore, widely used by federal agencies, as well as private sector organizations of all sizes, educational institutions, and state, local, tribal, and territorial governments. Our inclusion of international experts and organizations in most of these processes means that NIST's work is more likely to help shape the way that others around the globe approach issues related to AI standards and guidelines. This enhances the opportunity for alignment.

NIST’s team includes some of the top AI and standards experts in the world. Our staff has multidisciplinary backgrounds from industry, government, and academia with deep experience in various aspects of science and engineering related to AI.

Much of NIST’s AI effort¹ focuses on cultivating trust in the design, development, and use of AI technologies and systems. Working with the community, NIST is:

- conducting fundamental research to advance trustworthy AI technologies and understand and measure their capabilities and limitations
- applying AI research and innovation across NIST laboratory programs
- establishing benchmarks and developing data and metrics to evaluate AI technologies
- leading and participating in the development of technical AI standards
- contributing to discussions and development of AI policies, including supporting the National AI Advisory Committee².

NIST AI Risk Management Framework

Among its many AI-related activities, NIST developed the AI Risk Management Framework³ (AI RMF 1.0). Released in January 2023, the AI RMF is a voluntary framework that provides a flexible, structured, and measurable process to address AI risks purposefully and continually throughout the AI lifecycle. **AI risk management is about offering a path to minimize potential negative impacts of AI systems, as well as pointing to opportunities to maximize positive impacts and creating opportunities for innovation.** Identifying, mitigating, and minimizing risks and potential harms associated with AI technologies are essential steps towards the development of trustworthy AI systems and their appropriate and responsible use. Like NIST’s well-known Cybersecurity and Privacy Frameworks, the NIST AI RMF provides a set of outcomes that enable dialogue, understanding, and actions to manage AI risks.

As directed by the National Artificial Intelligence Initiative Act of 2020 (Division E of P.L. 116-283, the National Defense Authorization Act for Fiscal Year 2021), the AI RMF offers a resource to the organizations designing, developing, deploying, or using AI systems to help manage the many risks of AI and promote trustworthy and responsible development and use of AI systems. It has been exceptionally well-received nationally and internationally, where public and private entities in the U.S. and abroad are adopting or incorporating it in their responsible AI practices.⁴

The Framework is intended to be voluntary, rights-preserving, non-sector-specific, and use-case agnostic, providing flexibility to organizations of all sizes and in all sectors and throughout society to implement the approaches in the Framework. The AI RMF is designed to be practical, to adapt to the AI landscape as AI technologies continue to develop, and to be operationalized by organizations in varying degrees and capacities so society can benefit from AI while also being protected from its potential harms.

¹ <https://www.nist.gov/artificial-intelligence>

² <https://www.ai.gov/naiac/>

³ <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

⁴ <https://www.nist.gov/itl/ai-risk-management-framework/perspectives-about-nist-artificial-intelligence-risk-management>

The Framework was developed through a consensus-driven, open, transparent, and collaborative process. From the start of this initiative, NIST has offered a broad range of stakeholders the opportunity to take part in workshops⁵, respond to a Request for Information (RFI)⁶, and review draft reports⁷ and other documents including draft approaches⁸ and versions of the framework⁹.

NIST also has reached out directly to AI practitioners along with other stakeholders across a full spectrum of interests domestically and internationally. This outreach included companies, government agencies, academia, and not-for-profit organizations representing civil society, consumers, and industry. NIST actively encouraged others to provide direct input, and many organizations and individuals have contributed their insights to NIST. Those have included international organizations, with the goal of aligning the NIST Framework with standards and approaches being developed around the globe.

The AI RMF defines certain key characteristics of trustworthy AI systems and offers guidelines for mapping, measuring, and managing them. As defined in the AI RMF, trustworthy AI is valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, and privacy-enhanced, and fair with their harmful biases managed. AI systems are socio-technical in nature, meaning they are a product of the complex human, organizational, and technical factors involved in their design, development, and use. Many of the trustworthy AI characteristics – such as bias, fairness, interpretability, and privacy – are directly connected to societal dynamics and human behavior.

NIST also released a companion NIST AI RMF Playbook¹⁰ and a roadmap¹¹. The Playbook provides additional guidelines to organizations on the actions they can take to meet the outcomes included in the Framework. The roadmap identifies key activities for advancing the AI RMF that could be carried out by NIST in collaboration with private and public sector organizations – or by those organizations independently.

To support and operationalize the AI RMF and playbook, NIST established the online Trustworthy and Responsible AI Resource Center¹² in March 2023. The resource center is a platform to support people and organizations in government, industry, and academia driving technical and scientific innovation in AI. It serves as a one-stop-shop for foundational content, technical documents, and AI toolkits such as a glossary, a repository hub for measurement methods and metrics. In the coming months, NIST will add to the Resource Center a tracker of AI-related standards and a curated selection of public data sets, so that the public can engage with trustworthy and responsible AI technologies and standards.

In June of 2023, NIST launched a Generative AI Public Working Group¹³ of volunteers led by NIST staff to help NIST develop a profile of AI RMF for generative AI. To address the risk of generative AI technologies, this public working group is helping to develop four sets of

⁵ <https://www.nist.gov/itl/ai-risk-management-framework/ai-risk-management-framework-workshops-events>

⁶ <https://www.nist.gov/itl/ai-risk-management-framework/ai-rmf-development-request-information>

⁷ <https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf>

⁸ https://www.nist.gov/system/files/documents/2021/12/14/AI%20RMF%20Concept%20Paper_13Dec2021_posted.pdf

⁹ <https://www.nist.gov/itl/ai-risk-management-framework>

¹⁰ <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>

¹¹ <https://www.nist.gov/itl/ai-risk-management-framework/roadmap-nist-artificial-intelligence-risk-management-framework-ai>

¹² <https://airc.nist.gov/home>

¹³ <https://www.nist.gov/news-events/news/2023/06/biden-harris-administration-announces-new-nist-public-working-group-ai>

guidelines: pre-deployment verification and validation of generative AI models; digital content provenance; incident disclosure; and governance of generative AI systems. This Working Group has attracted about 1,000 participants who are providing NIST with substantive input that will lead to very practical guidelines on generative AI.

AI research and development, as well as the standards landscape, are developing at an extraordinary pace. For that reason, the AI RMF and its related documents will evolve over time and reflect new knowledge, awareness, and practices. NIST intends to continue its robust engagement with stakeholders to keep the Framework up to date with AI trends and reflect experience based on the use of the AI RMF.

NIST’s Research on AI Trustworthiness Characteristics

To build on NIST’s work on the AI RMF and provide additional guidelines to organizations to advance trustworthy and responsible AI, NIST also conducts fundamental research on many of the AI trustworthiness characteristics, prioritizing its work based on its insights and the community’s stated needs.

» AI Trustworthiness Characteristics – Fair and Bias is Managed

While there are many approaches for ensuring technologies that we use every day are safe and secure, there is less research into how to advance systems that are fair with harmful bias managed.

Fairness in AI includes concerns for equality and equity by addressing issues such as bias and discrimination. While certain forms of discrimination may violate federal civil rights laws, standards of fairness in AI can be complex and difficult to define because perceptions of fairness differ among cultures and may shift depending on application and context of use.

NIST has significantly expanded its research efforts to identify, understand, measure, manage and mitigate bias, with a focus on a socio-technical approach. NIST published “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence” (NIST Special Publication 1270)¹⁴, which identifies the concepts and challenges associated with bias in AI and provides preliminary guidelines for addressing them.

NIST has identified three major categories of AI bias to be considered and managed: systemic, computational, and human, all of which can occur in the absence of prejudice, partiality, or discriminatory intent. Current approaches to address the harmful effects of AI bias remain focused largely on computational factors such as representativeness of datasets and fairness of machine learning algorithms. Human and systemic institutional and societal factors are significant sources of AI bias that are currently overlooked. Systemic bias can be present in AI datasets, the organizational norms, practices, and processes across the AI lifecycle, and the broader society that uses AI systems. Human biases relate to how an individual or group perceives and uses AI system information to make a decision or fill in missing information. NIST has made this aspect of trustworthiness a premier area of focus and has been recognized widely for doing so.

» AI Trustworthiness Characteristics – Explainable and Interpretable

Explainability and interpretability are important characteristics to ensure that users and operators of AI can understand the decisions or predictions made by AI, thus avoiding the “opaque system”

¹⁴ <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>

concept associated with AI. Explainability refers to a representation of the mechanisms underlying an algorithm’s operation, whereas interpretability refers to the meaning of an AI system’s output in the context of its designed functional purpose.

NIST has released two publications aimed at providing deeper understanding of the principles of explainability and interpretability: “Four Principles of Explainable Artificial Intelligence” (NISTIR 8312)¹⁵ and “Psychological Foundations of Explainability and Interpretability in Artificial Intelligence” (NISTIR 8367)¹⁶.

» *AI Trustworthiness Characteristics – Secure and Resilient*

AI systems that can withstand adversarial attacks and maintain confidentiality, integrity, and availability are resilient and secure systems.

NIST released the draft “A Taxonomy and Terminology of Adversarial Machine Learning”¹⁷ (NIST AI 100-2) to advance a taxonomy for securing applications of AI, specifically, adversarial machine learning. NIST’s Cybersecurity Framework¹⁸ is widely used to address the cybersecurity risks of organizations. NIST periodically updates that Framework to account for changes in the cybersecurity technology, standards, and risk landscape.

NIST has built an experimentation testbed called Dioptra¹⁹ to begin to evaluate adversarial attacks against ML algorithms. The testbed aims to facilitate security evaluations of ML algorithms under a diverse set of conditions. It is publicly available on GitHub, and we intend to continue maintaining and updating it for public use. To that end, the testbed has a modular design that enables researchers to easily swap in alternative datasets, models, attacks, and defenses. This should advance the measurement capabilities needed to ultimately help secure AI systems. We know that the need for these types of testbeds is growing.

» *AI Trustworthiness Characteristics – Privacy-enhanced*

Privacy safeguards the important human values of autonomy and dignity through methods that focus on providing individuals with anonymity, confidentiality, and control over various facets of their identities. These outcomes generally should guide choices for AI system design, development, and deployment. From a policy perspective, privacy-related risks may overlap with security, bias, and transparency.

NIST’s Privacy Risk Assessment Methodology²⁰, developed in 2016 and NIST’s Privacy Framework²¹, issued in 2020, are voluntary tools that organizations from industry sectors across the world are using to identify and manage privacy risks in the systems, products, and services they develop and deploy, improve their privacy programs, and better comply with privacy regulation.

NIST is also conducting research on privacy-enhancing technologies (PETs) to advance data-driven, innovative solutions to preserve the right to privacy, including hosting the Privacy

¹⁵ <https://www.nist.gov/publications/four-principles-explainable-artificial-intelligence>

¹⁶ <https://www.nist.gov/publications/psychological-foundations-explainability-and-interpretability-artificial-intelligence>

¹⁷ <https://www.nccoe.nist.gov/ai/adversarial-machine-learning>

¹⁸ <https://www.nist.gov/cyberframework>

¹⁹ <https://pages.nist.gov/dioptra/>

²⁰ <https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources>

²¹ <https://www.nist.gov/privacy-framework/privacy-framework>

Engineering Collaboration Space²², a virtual public platform that serves as a clearinghouse for open-source tools and PETs use cases. In coordination with the National Science Foundation (NSF) and the White House Office of Science and Technology Policy (OSTP), NIST co-sponsored the U.S.-U.K. prize competition on PETs²³. First announced at the Summit for Democracy in December 2021, the winning solutions competed for a combined U.S.-U.K. prize pool of \$1.6 million and the winners were showcased at the second Summit for Democracy in March 2023.

Research on Applications of AI

NIST's multidisciplinary laboratories and varied fields are an ideal environment to develop and apply AI²⁴. Various AI techniques are being used to support NIST scientists and engineers, drawing on machine learning (ML) and AI tools to gain a deeper understanding of, and insight into, our research.

NIST is integrating AI into the design, planning, and optimization of our research efforts – including hardware for AI²⁵, computer vision, engineering biology and biomanufacturing, image and video understanding, medical imaging, materials science, manufacturing, disaster resilience, energy efficiency, natural language processing, biometrics, quantum science, robotics, and advanced communications technologies. Key focus areas include innovative measurements using AI/ML techniques, predictive systems using AI/ML models, and enabling and reducing the barriers to autonomous measurement platforms.

Just in the last year, NIST's scientists have designed new circuits for supercomputing,²⁶ developed a deep learning algorithm to interpret breathing patterns,²⁷ engineered a new AI-enabled method to digitally simulate hurricanes,²⁸ and used AI to detect disease in human breath molecules.²⁹

AI Measurement and Evaluation

NIST has a long history of devising appropriate metrics, measurement tools, and challenge problems to support technology development. NIST first started the measurement and evaluation of automated fingerprint identification systems in the 1960s. Evaluations strengthen research communities, establish research methodology, support the development of standards, and facilitate technology transfer. NIST is looking to bring these benefits of community evaluations to bear on the problem of constructing trustworthy AI systems. These evaluations will begin with community input to identify potential harms of selected AI technologies in context, and the data requirements for AI evaluations. NIST also hosts a biweekly AI metrology colloquia series³⁰, where leading researchers share current work on AI measurement and evaluation.

NIST has been engaged in focused efforts to establish common terminologies, definitions, and

²² <https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/collaboration-space>

²³ <https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/collaboration-space/prize-challenges>

²⁴ <https://www.nist.gov/applied-ai>

²⁵ <https://www.nist.gov/artificial-intelligence/hardware-ai>

²⁶ <https://www.nist.gov/news-events/news/2022/10/nists-superconducting-hardware-could-scale-brain-inspired-computing>

²⁷ <https://www.nist.gov/news-events/news/2022/12/wi-fi-could-help-identify-when-youre-struggling-breathe>

²⁸ <https://www.nist.gov/news-events/news/2023/03/ai-could-set-new-bar-designing-hurricane-resistant-buildings>

²⁹ <https://www.nist.gov/news-events/news/2023/04/jilas-frequency-comb-breathalyzer-detects-covid-19-excellent-accuracy>

³⁰ <https://www.nist.gov/programs-projects/ai-measurement-and-evaluation/ai-metrology-colloquia-series>

taxonomies of concepts pertaining to characteristics of AI technologies in order to form the necessary underpinnings for trustworthy AI systems. Each of these characteristics also requires its own portfolio of measurements and evaluations. For each characteristic, NIST aims to document and improve the definitions, applications, and strengths and limitations of metrics and measurement methods in use or being proposed. NIST's current efforts represent a vital – but only a small – portion of the research that will be required to test and evaluate trustworthy AI systems.

A significant challenge in the evaluation of trustworthy AI systems is that context (the specific use case) matters; accuracy measures alone will not provide enough information to determine if deploying a system is warranted. The accuracy measures must be balanced by the associated risks or societal harms that could occur. The tolerance for error drops as the potential impacts of risk rise.

New NIST efforts in AI evaluation will focus on other socio-technical aspects of system performance in addition to accuracy. In particular, the evaluations have the goal of identifying risks and harms of systems before they are deployed, and to define data sets and evaluation infrastructure that will allow system builders to detect the extent to which their system exhibits those harms.

Examples of NIST AI measurement and evaluation projects³¹ include:

- *Biometrics*: Over that past sixty years, NIST has been testing and evaluating biometric recognition technologies, including face recognition, fingerprint, biometric quality, iris recognition, and speaker recognition.
- *Computer vision*: NIST's computer vision program includes several activities contributing to the development of technologies that extract information from image and video streams through systematic, targeted annual evaluations and metrology advances, including the Open Media Forensics Challenge, Activities in Extended Video (ActEV), and handwriting recognition and translation evaluation.
- *Information retrieval*: The information retrieval research uses large, human-generated text, speech, and video files to create test collections through the Text Retrieval (TREC), TREC Video Retrieval Evaluation (TRECVID), and Text Analysis (TAC) Conferences. The Text Retrieval Conference is responsible for significant advances in search technology. A 2010 NIST study³² estimated that, without TREC, U.S. internet users would have spent an estimated 3.5 billion additional hours using search engines between 1999 and 2009.

AI Standards

NIST plays a critical role in the standards process as the nation's measurement laboratory and has a unique role relating to standards in the Federal enterprise. Our coordination function, currently defined under the National Technology Transfer and Advancement Act and the NIST Organic Act, has yielded benefits to the nation ever since the Institute was established by Congress as the National Bureau of Standards in 1901. NIST's strong ties to industry and the standards development community have enabled NIST to take on critical standards-related challenges and deliver timely and effective solutions. That role was reinforced recently when NIST was charged by the Administration to lead the National Standards Strategy for Critical

³¹ <https://www.nist.gov/programs-projects/ai-measurement-and-evaluation/nist-ai-measurement-and-evaluation-projects>

³² <https://trec.nist.gov/pubs/2010.economic.impact.pdf>

and Emerging Technology³³. We have recently begun seeking public input on how to best implement that strategy; we will begin reviewing comments at the close of public comment in early November.³⁴

NIST works to support the development of AI standards that promote innovation and public trust in systems that use AI. Pursuant to U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools³⁵, NIST seeks to bolster AI standards-related knowledge, leadership, and coordination; conduct research to support development of technically sound standards for trustworthy AI; promote partnerships to develop and use standards; and engage internationally to advance AI standards.

I serve as the Federal AI Standards Coordinator, working across the government and with industry stakeholders to gather and share information on AI standards-related needs, strategies, and best practices.

NIST facilitates federal agency coordination in the development and use of AI standards in part through the Interagency Committee on Standards Policy (ICSP) AI Standards Coordination Working Group³⁶. This working group seeks to foster agency interest and participation in AI standards and conformity assessment activities, facilitate coordination of U.S. government positions on draft standards, identify effective means of coordinating with, and contributing towards, voluntary consensus bodies, align U.S. government activities with those of the private sector on AI standards development activities, promote effective and consistent federal policies leveraging AI standards, and raise awareness of federal agencies' use of AI that contributes to standards activities.

NIST also engages internationally through bilateral and multilateral work on AI. The United States championed development of the first international principles for the responsible use of AI at the Organisation for Economic Co-operation and Development, or OECD. The U.S. also serves as a founding member of the Global Partnership on AI, which includes all members of the G7 and other nations, including Brazil and India, to coordinate R&D AI initiatives. NIST supports the US- EU Trade and Technology Council (TTC) in building common approaches for trustworthy AI. Under the TTC, the U.S. and EU launched a new AI sub-working group which developed a joint AI roadmap on Evaluation and Measurement Tools for Trustworthy AI³⁷ in December 2022, through which NIST is working towards common definitions for AI risk management and developing metrics and methodologies for measuring AI trustworthiness. Progress in implementation of the joint map was reported to the TTC in May 2023, including the launch of three expert groups to focus on AI terminology and taxonomy, standards and tools for trustworthy AI and risk management, and monitoring and measuring AI risks. The groups have:

- (i) issued a list of 65 key AI terms essential to understanding risk-based approaches to AI³⁸, and
- (ii) mapped the involvement of the United States and the European Union in standardization activities with the goal of identifying relevant AI-related standards of mutual interest.

³³ <https://www.nist.gov/standardsgov/usg-nss>

³⁴ <https://www.nist.gov/news-events/news/2023/09/nist-seeks-input-implementation-national-standards-strategy-critical-and>

³⁵ https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

³⁶ <https://www.nist.gov/standardsgov/icsp-ai-standards-coordination-working-group-aiscwg-charter>

³⁷ https://www.nist.gov/system/files/documents/2022/12/04/Joint_TTC_Roadmap_Dec2022_Final.pdf

³⁸ <https://www.nist.gov/document/eu-us-terminology-and-taxonomy-artificial-intelligence>

Interagency Coordination

NIST leads and participates in several federal AI policymaking efforts and engages with many other federal offices and interagency groups. This includes administering the National Artificial Intelligence Advisory Committee (NAIAC)³⁹, on behalf of the Department of Commerce. Established by the National Artificial Intelligence Initiative Act of 2020, NAIAC is tasked with advising the President and the National AI Initiative Office. NIST provides executive support to this advisory committee. The Secretary of Commerce appointed the 26 members in April 2022. NAIAC held its first meeting in May 2022. In its year one report⁴⁰, the NAIAC presented findings and recommendations related to promoting leadership in trustworthy artificial intelligence and leadership in research and development; supporting the U.S. workforce and providing opportunity; and international collaboration. Subsequent publications have included instructive explainers such as the FAQs on Foundation Models and Generative AI⁴¹ document.

NIST also co-chairs the National Science and Technology Council's Machine Learning and Artificial Intelligence Subcommittee⁴², the Networking and Information Technology Research and Development's (NITRD) AI Working group⁴³, and the NITRD Fast Track Action Committee,⁴⁴ which drafted a national strategy to advance privacy-preserving data sharing and analytics. NIST founded and is co-chairing the AI Standards Coordination Working Group (AISCWG) under the Interagency Committee on Standards Policy (ICSP). As NIST's AI lead, I also serve as Federal AI Standards Coordinator and a member of the National AI Research Resource Task Force,⁴⁵ which released its final report in January 2023⁴⁶.

Conclusion

Advancing AI research and standards that contribute to a secure, private, interoperable, innovative, and world-leading digital economy is a top priority for NIST. Our economy is increasingly global, complex, and interconnected. It is characterized by rapid advances in technology. The timely availability of AI standards and guidelines is a dynamic and critical challenge. Through robust collaboration with stakeholders across government, industry, and academia in the U.S. and elsewhere, NIST is cultivating trust and fostering an environment that enables AI innovation on a global scale – ensuring that we can all gain the benefits of AI while managing its risks.

NIST's team includes some of the top AI and standards experts in the world. Our staff has multidisciplinary backgrounds from industry, government, and academia with expertise and experience in a wide range of science and engineering related to AI. Working with our partners in other federal agencies, the private sector, academia, and allied countries, and with the support of Congress, we have made important progress and will continue to work tirelessly to address current and future challenges.

Thank you for the opportunity to present on NIST activities to improve AI trustworthiness. I look forward to your questions.

³⁹ <https://www.nist.gov/artificial-intelligence/national-artificial-intelligence-advisory-committee-naiac>

⁴⁰ <https://www.ai.gov/wp-content/uploads/2023/05/NAIAC-Report-Year1.pdf>

⁴¹ <https://www.ai.gov/wp-content/uploads/2023/09/FAQs-on-Foundation-Models-and-Generative-AI.pdf>

⁴² https://www.ai.gov/about/#MLAI-SC_Machine_Learning_and_AI_Subcommittee

⁴³ <https://www.ai.gov/a-new-nitrd-iwg-for-artificial-intelligence-ai-rd/>

⁴⁴ <https://www.nitrd.gov/coordination-areas/privacy-rd/appdsa/>

⁴⁵ <https://www.ai.gov/nairr/>

⁴⁶ <https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf>