NIST AI PROGRAM

Artificial Intelligence: The Vitals

The National Institute of Standards and Technology (NIST) aims to cultivate trust in the design, development, use, and governance of Artificial Intelligence (AI) technologies and systems in ways that enhance economic security and improve quality of life. NIST focuses on improving measurement science, technology, standards, and related tools – including evaluation and data.

With AI and Machine Learning (ML) already changing how society addresses challenges and opportunities, these technologies must be trustworthy and embedded in responsible practices. Trustworthy AI systems are those demonstrated to be valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed. Trustworthiness is critical for successful deployment of new technologies for genomics, image and video processing, materials, natural language processing, robotics, wireless spectrum monitoring, and more. NIST helps to deliver the needed measurements, standards, tools, and expertise. The agency's AI goals and activities are driven by its statutory mandates, Presidential Executive Orders, and the needs expressed by U.S. industry, other federal agencies, and the global AI research community.

NIST's AI goals include:

1. Conduct fundamental research to advance trustworthy AI technologies
2. Apply AI research and innovation across the NIST Laboratory Programs
3. Establish benchmarks, data, and metrics to evaluate AI technologies
4. Lead and participate in development of technical AI standards
5. Contribute technical expertise to discussions and development of AI policies

FUNDAMENTAL AI RESEARCH

NIST's AI portfolio includes fundamental research, development, and standards vital for trust in AI technologies. This work includes efforts on software, hardware, architectures, human-interaction and teaming, and all relevant intersections and interfaces.

Establishing and Promoting Technical Requirements for Trustworthy and Responsible AI

- NIST identifies and quantifies trustworthy and responsible AI in technical terms and develops tools and guidance so that designers, developers, and evaluators can act. This includes developing taxonomy, terminology, and testbeds for measuring AI risks, as well as standards for key technical characteristics of AI trustworthiness.

- Aided by workshops and other engagements, NIST develops documents, software, standards, and related tools that help AI actors better understand, identify, measure, and manage risks associated with AI systems.

- As mandated by Congress, NIST has developed a voluntary AI Risk Management Framework. The framework was worked on collaboratively with public and private stakeholders, many of whom are now implementing it and using its companion Playbook.

- NIST has established a Trustworthy and Responsible AI Resource Center, a "one stop shop" for foundational content, technical documents, and toolkits to enable responsible use of AI.

- Per an October 30, 2023 Presidential Executive Order, NIST is launching a variety of new initiatives to advance critical guidelines, benchmarks, and tools to evaluate AI trustworthiness.

- NIST has launched the U.S. AI Safety Institute and a related consortium to empower a new measurement science. These institutions will help identify proven, scalable, and interoperable techniques and metrics to promote responsible development and use of safe and trustworthy AI.

Hardware for AI: Creating New Measurements and Technical Approaches for New AI Chips

- As AI becomes more prevalent in our everyday lives, demand for faster, more energy-efficient information processing is growing exponentially. Conventional hardware cannot keep up with this demand. That is why researchers, taking inspiration from the brain, are considering alternatives—e.g., where massively connected networks of artificial neurons and synapses process

information with high speed, energy efficiency, scalability, and adaptive learning. NIST is helping develop [devices, circuits, systems, measurements, and theory](#) to support the evolution of new AI hardware from laboratory research to commercial applications. NIST is focusing on scalability, energy efficiency, hardware optimization, and architecture development. The new physical computation mechanisms and intelligent functionality call for new measurement techniques and protocols.

APPLIED AI RESEARCH

NIST's multidisciplinary laboratories offer an ideal environment to develop and apply AI. AI approaches are increasingly an essential component in new project conception and execution. [NIST scientists and engineers are drawing on a variety of ML and AI tools](#) to gain a deeper understanding of and insight into their research. At the same time, NIST laboratories' experiences with AI are assisting in better understanding AI's capabilities and limitations.

Exploring Research Frontiers by Incorporating AI and ML

- NIST researchers are developing partnerships across the Institute and with academia, other government laboratories, and industrial entities. They are integrating AI into the design, planning, and optimization of NIST's research efforts – including computer vision, engineering biology and biomanufacturing, image and video understanding, medical imaging, materials science, manufacturing, disaster resilience, energy efficiency, natural language processing, 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.

Producing Training Data, Algorithms, and Other Tools

- NIST's technical staff are working on key practices for documentating and characterizing datasets, and developing datasets that the broader community can use to test or train AI systems. The agency is expanding the availability of algorithms, test or training data, and other tools for domain-specific applications including advanced materials, computer vision, design, industrial robotics, natural language processing, spectrum management, and video processing.

MEASUREMENT & EVALUATION

With a long history of devising and revising metrics, measurement tools, standards, and testbeds, NIST increasingly is focusing on measurement and evaluation of technical characteristics of trustworthy AI.

NIST's actions under the [President's Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence](#) and its new [U.S. AI Safety Institute](#) are major initiatives designed to dramatically advance the ability to measure and evaluate the trustworthiness of rapidly advancing AI technologies. While NIST's activities have typically addressed measures of accuracy and robustness, other types of AI-related measurements and evaluations under investigation include safety, bias, interpretability, and transparency. Working with others, NIST aims to augment such efforts by:

- *Advancing measurement science for AI:* Defining, characterizing, and theoretically and empirically developing and analyzing quantitative and qualitative metrics and measurement methods for various characteristics of AI technologies.

- *Conducting rigorous evaluations of AI:* Designing and conducting evaluations of AI technologies – including developing tasks, challenge problems, testbeds, software tools, and helping to curate and characterize meaningful data sets – and identifying technical gaps and limitations in AI technologies and related measurements.

- *Developing best practices and technical guidance:* Sharing results and guidance to inform academic, industrial, and government programs.

- *Contributing to voluntary consensus-based standards for measuring and evaluating AI:* Leading or participating in standardization efforts to support the development, deployment, and evaluation of AI technologies.

TECHNICAL AI STANDARDS

NIST leads and participates in the development of technical [standards](#), including international standards, that promote innovation and public trust in systems that use AI.

Ensuring Awareness and Federal Coordination in AI Standards Efforts

- In its role as federal AI standards coordinator, NIST works across the government and with industry to identify critical standards development activities, strategies, and gaps. Based on priorities outlined in the NIST-developed [Plan for Federal Engagement in AI Standards and Related Tools](#), the agency is tracking AI standards development opportunities, periodically collecting and analyzing information about agencies' AI standards-related priority activities, and making recommendations through interagency processes to optimize engagement.

- NIST is coordinating federal agencies in the development and use of AI standards in part through the Interagency Committee on Standards Policy (ICSP), which it chairs. ICSP's AI Standards Coordination Working Group promotes effective and consistent federal policies to leverage and further development of AI standards, raise awareness, and foster agencies' use of AI. The group helps to coordinate government and private sector positions on AI international standards activities, a priority spelled out in the October 2023 Executive Order.

Participating in High-Priority AI Standards Activities

- NIST supports its staff's participation and leadership in standards development activities aligned with NIST technical expertise. Participation is based on national needs and an assessment of priority topics and gaps.

- NIST efforts include work through two prominent international standards developing organizations: ISO and IEEE.

AI POLICY CONTRIBUTIONS

NIST expertise in AI is helping national and global AI policy bodies shape the development and use of innovative, trustworthy AI. In the international arena, standards are a critical element of government's AI policies and industry pracices. NIST plays a key role as a neutral convenor of organizations and individuals with disparate views about AI matters. The NIST AI Risk Management Framework highlights how NIST's reputation for technical excellence and interacting impartially with the private and public sectors is influencing international approaches to AI.

Contributing to Federal Engagements that Explore or Determine AI-Related Policies

- NIST leads or 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, co-chairing the National Science and Technology Council's Machine Learning and Artificial Intelligence Subcommittee, co-chairing the Networking and Information Technology Research and Development's AI Working group, and founding and co-chairing the AI Standards Coordination Working Group under the Interagency Committee Standards Policy. NIST's AI lead also serves as Federal AI Standards Coordinator and was a member of the National AI Research Resource Task Force.

Participating in Major AI Forums

- NIST engages with private, public, and non-profit organizations – directly and via international forums and other dialogues – about AI-related policies that align with NIST's mission and technical contributions. NIST also convenes national and international stakeholders to ensure two-way communication on select AI-related issues.

- NIST supports U.S. and internatoinal AI policy efforts such as the U.S.-EU Trade and Technology Council, OECD, Global Partnership on AI, and a host of bilateral initiatives in Asia, Europe, and North America.

WORKING WITH NIST ON AI

NIST relies on and encourages robust interactions with companies, universities, nonprofits, and other government agencies in driving and carrying out its AI agenda. There are multiple ways to engage with NIST, including:

NIST Draft Reports: NIST counts on public and private sector review of draft reports on a variety of AI issues. Drafts typically are prepared based on inputs from private and public sector individuals and organizations and then offered for public review on NIST's AI website and via email alerts. Public comments help to improve them. Sign up for AI-related emails here.

Workshops: NIST convenes experts for single day, multi-day, and multi-week sessions to tackle key characteristics of AI trustworthiness and other AI-related topics. All workshops are virtual for now and readily accessible.

AI Visiting Fellows: Accomplished Visiting Fellows bring thought leadership to fundamental research for trustworthy and responsible AI, use-inspired AI research, and AI hardware research, as well as AI related standards and evaluations conducted in NIST laboratories.

AI Visiting Researchers: To advance their careers and assist NIST in the process, Visiting Researchers conduct their work based on NIST laboratory priorities with guidance and recommendations from NIST mentors and AI Visiting Fellows.

Student Programs: NIST offers a range of opportunities for students to engage with NIST on AI-related work. That includes the Professional Research Experience Program (PREP), which provides valuable laboratory experience and financial assistance to undergraduate, graduate, and post-graduate students.

Grants: NIST offers some financial assistance to support collaborative research, including AI projects.

Sign up for AI email alerts here. If you have questions or ideas about how to engage with us on AI topics, send us an email: ai-inquiries@nist.gov.