

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 the way in which society addresses economic and national security challenges and opportunities, these technologies must be developed and used in a trustworthy and responsible manner. Characteristics which support trustworthiness include accuracy, explainability and interpretability, privacy, reliability, robustness, safety, and security (resilience) and mitigation of harmful bias. Principles such as transparency, fairness, and accountability should be considered, especially during deployment and use. Trustworthy data, standards, and evaluation, validation, and verification are critical for the successful deployment of new technologies for genomics, image and video processing, materials, natural language processing, robotics, wireless spectrum monitoring, and more. Delivering the needed measurements, standards, and other tools is a primary focus for NIST's portfolio of AI efforts. It is an area in which the agency has special responsibilities and expertise – and where others often turn to NIST. The agency's AI goals and activities are prioritized and informed by its statutory mandates, White House directions, 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 and develop data and metrics to evaluate AI technologies
4. Lead and participate in development of technical AI standards
5. Contribute to discussions and development of AI policies

FUNDAMENTAL AI RESEARCH

NIST's AI portfolio includes [fundamental research, development, and standards for AI technologies](#) – including software, hardware, architectures, human-interaction and teaming, and all relevant intersections and interfaces – vital for AI computational trust.

Establishing and Promoting Technical Requirements for Trustworthy and Responsible AI

- Through targeted research investments, NIST identifies and quantifies trustworthy and responsible AI in technical terms and develops tools and guidance so that designers, developers, and evaluators can take appropriate actions. These include developing taxonomy, terminology, and testbeds for measurements of AI risks – as well as standards needed for key technical characteristics of AI trustworthiness. They address accuracy, explainability and interpretability, privacy, reliability,

robustness, safety, security (resilience), and harmful bias mitigation.

- By convening stakeholder workshops and via other engagements, NIST is developing a resource center of documents, software, standards, and related tools that contribute to better understanding, identifying, measuring, and managing various risks associated with AI systems.
- As mandated by Congress, NIST is developing a voluntary [AI risk management framework](#) through collaboration with stakeholders across public and private sectors.

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

- Demand for faster, more energy-efficient information processing is growing exponentially as AI becomes more prevalent in our everyday lives. Conventional digital processing hardware cannot keep up with this demand. That is why researchers, taking inspiration

from the brain, are considering alternatives where massively connected networks of artificial neurons and synapses process information with high speed, energy efficiency, scalability, and adaptive learning capabilities. To that end, NIST is helping to develop [devices, circuits, systems, measurements, and theory](#) to support the evolution of AI hardware technology from laboratory research to commercial applications. It is focusing on scalability, energy efficiency, hardware optimization, and architecture development. The new hardware leverages the physics of devices to perform computations, while the architectures and algorithms have entirely new intelligent functionality. This calls for a new system of measurement techniques and protocols.

APPLIED AI RESEARCH

NIST's multidisciplinary laboratories and varied fields are an ideal environment to develop and apply AI. AI approaches are increasingly an essential component in new project conception and execution. Various AI techniques are being used to support [NIST scientists and engineers, drawing on ML and AI tools](#) to gain a deeper understanding of and insight into their research. At the same time, NIST laboratory 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 is working on data characterization, key practices for documentation of datasets, and datasets that the broader community can use to test or train AI systems. The agency is expanding the availability of targeted test or training data, algorithms, and other tools for use in domain-specific applications including advanced materials, computer vision, design, industrial robotics, natural

language processing, spectrum management, and video processing.

BENCHMARKS & METRICS

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

Advancing Measurements and Evaluating AI Technologies

- Benchmarks – made up of data, tests, and evaluations – provide quantitative measures for developing standards and assessing conformance with standards as well as examining limits and capabilities of AI technologies. Benchmarks drive innovation by measurable advancements aimed at addressing strategically selected scenarios; they also provide objective data to track the evolution of AI science and technologies.
- NIST benchmarks are among the key elements of the testing methodologies and metrics that the agency develops to effectively evaluate AI technologies. This includes testing methods that prescribe protocols and procedures for assessing, comparing, and managing the performance and functionality of AI technologies. NIST is defining quantifiable measures to characterize AI technologies – including measuring accuracy, complexity, explainability and interpretability, privacy, reliability, robustness, safety, security, and bias – and to enable comparisons to human performance.

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. A broad spectrum of [standards for AI data, performance, and governance](#) are – and increasingly will be – a priority for trustworthy and responsible 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 stakeholders 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](#), NIST is tracking AI standards development opportunities, periodically collecting and analyzing information about agencies' AI standards-related priority activities and making recommendations

through the interagency process to optimize engagement.

- NIST is facilitating federal agency coordination in the development and use of AI standards in part through the Interagency Committee on Standards Policy (ICSP), which it chairs. A new ICSP AI Standards Coordination Working Group (AISCWG) will promote effective and consistent federal policies leveraging AI standards, raise awareness, and foster agencies' use of AI to inform the development of standards. The group will help to coordinate government and private sector positions regarding AI international standards activities.

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.

Developing Best Practices for Datasets Used to Train or Test AI Systems

- Recognizing the importance of data used in training and benchmarking the performance of AI systems, Congress directed NIST to develop best practices for metadata standards that describe the properties of datasets, along with standards for privacy and security of datasets with human characteristics.

AI POLICY CONTRIBUTIONS

NIST expertise in AI applications as well as in foundational AI enables the agency to make important technical contributions to the development of policies. NIST plays a key role as a neutral convener of organizations and individuals with disparate views about AI issues. The agency's active participation in national and global discussions can help to shape the development of trustworthy and responsible AI.

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

- Based on its specific technical expertise and reputation for being an unbiased source of information, NIST leads and participates in key federal AI policymaking efforts and engages with many other federal offices and interagency groups. This includes establishing and 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 (NITRD) AI Working group, and founding and co-chairing the AI Standards Coordination Working Group (AISCWG) under the

Interagency Committee on Standards Policy (ICSP). NIST's AI lead also serves on 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 communications on select AI-related issues.

Developing Guidance for Sharing AI Data

- Congress assigned NIST with developing guidance to facilitate voluntary data sharing arrangements among industry, federally funded research centers, and federal agencies to advance AI research and technologies. That includes offering options for partnership models between government entities, industry, universities, and nonprofits that incentivize each to share the data they have collected.

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 stakeholders to review drafts of 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 foundational research for trustworthy and responsible AI, use-inspired AI research, 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