U.S. LEADERSHIP IN AI:
A Plan for Federal Engagement in Developing Technical Standards and Related Tools

Prepared in response to Executive Order 13859
Submitted on August 9, 2019

NIST National Institute of Standards and Technology
U.S. Department of Commerce

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Emphasizing the importance of artificial intelligence (AI) to the future of the U.S. economy and national security, on February 11, 2019, the President issued an Executive Order (EO 13859) directing Federal agencies to ensure that the nation maintains its leadership position in AI. Among its objectives, the EO aims to *"Ensure that technical standards...reflect Federal priorities for innovation, public trust, and public confidence in systems that use AI technologies...and develop international standards to promote and protect those priorities."*

United States global leadership in AI depends upon the Federal government playing an active and purpose-driven role in AI standards development. That includes AI standards-related efforts needed by agencies to fulfill their missions by:

- supporting and conducting AI research and development,
- actively engaging in AI standards development,
- procuring and deploying standards-based products and services, and
- developing and implementing supportive policies, including regulatory policies where needed.

The EO directs the Secretary of Commerce, through the National Institute of Standards and Technology (NIST), to issue *"a plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies."* This plan was prepared with broad public and private sector input.

There are a number of cross-sector (horizontal) and sector-specific (vertical) AI standards available now and many others are being developed by numerous standards developing organizations (SDOs). Some areas, such as communications, have well-established and regularly maintained standards in widespread use, often originally developed for other technologies. Other aspects, such as trustworthiness, are only now being considered.

This plan identifies the following nine areas of focus for AI standards:

- Concepts and terminology
- Data and knowledge
- Human interactions
- Metrics
- Networking
- Performance testing and reporting methodology
- Safety
- Risk management
- Trustworthiness

Trustworthiness standards include guidance and requirements for accuracy, explainability, resiliency, safety, reliability, objectivity, and security.

It is important for those participating in AI standards development to be aware of, and to act consistently with, U.S. government policies and principles, including those that address societal and ethical issues, governance, and privacy. While there is broad agreement that these issues must factor into AI standards, it is not clear how that should be done and whether there is yet sufficient scientific and technical basis to develop those standards provisions.

Standards should be complemented by related tools to advance the development and adoption of effective, reliable, robust, and trustworthy AI technologies. These tools—which often have overlapping applications—include, but are not limited to:

- Data sets in standardized formats, including metadata for training, validation and testing of AI systems
- Tools for capturing and representing knowledge, and reasoning in AI systems
- Fully documented use cases that provide a range of data and information about specific applications of AI technologies and any standards or best practice guides used in making decisions about deployment of these applications
- Testing methodologies to validate and evaluate AI technologies' performance
- Metrics to quantifiably measure and characterize AI technologies
- Benchmarks, evaluations, and challenge problems to drive innovation
- AI testbeds
- Tools for accountability and auditing

U.S. government agencies should prioritize involvement in AI standards efforts that are: inclusive and accessible, open and transparent, consensus-based, globally relevant, and non-discriminatory. This plan provides guidance regarding important characteristics of standards to help agencies in their decision making about AI standards. It also groups potential agency involvement into four categories ranked from least- to most-engaged: monitoring, participation, influencing, and leading. The plan provides a series of practical steps for agencies to take as they decide about engaging in AI standards.

This plan recommends that the Federal government commit to deeper, consistent, long-term engagement in AI standards development activities to help the United States to speed the pace of reliable, robust, and trustworthy AI technology development. Specifically, the Federal government should:

1. **Bolster AI standards-related knowledge, leadership, and coordination among Federal agencies to maximize effectiveness and efficiency.**
   - The National Science and Technology Council (NSTC) Machine Learning/Artificial Intelligence (ML/AI) Subcommittee should designate a Standards Coordinator with responsibility to gather and share AI standards-related needs, strategies, roadmaps, terminology, use cases, and best practices in support of reliable, robust, and trustworthy AI in government operations, including:

     -planned and ongoing standards approaches and engagement activities, including ensuring a robust feedback loop with SDOs,

     -specific horizontal or vertical areas for prioritization,

> -requirements for input into proposed standards activities, and
>
> -analysis of whether ongoing standards activities meet Federal government needs and whether additional guidance is appropriate.

- Make maximum use of existing standards that are broadly adopted by industry sectors and can be used or evolved within the new context of AI solutions.

- Reinforce the importance of agencies' adherence to Federal policies for standards and related tools: for example, data access and quality.

- Be flexible in selecting AI standards for use in regulatory or procurement actions to meet agency objectives. Flexibility is required to adapt to the rapid pace of AI technology developments and standards and to reflect the increasing understanding about trustworthiness, accessibility, and human-centered implications of AI.

- Grow a cadre of Federal staff with the relevant skills and training. Develop and provide a clear career development and promotion path that values and encourages participation in and expertise in AI standards and standards development.

2. Promote focused research to advance and accelerate broader exploration and understanding of how aspects of trustworthiness can be practically incorporated within standards and standards-related tools.

- Plan, support, and conduct research and evaluation that underlies technically sound, fit-for-purpose standards and related tools in support of reliable, robust, and trustworthy AI.

- Develop metrics and data sets to assess reliability, robustness and other trustworthy attributes of AI systems, focusing on approaches that are readily understandable, available, and can be put on a path to standardization.

- Conduct research to inform standardization of risk management strategies, including monitoring and mitigating risks.

- Identify research needs, requirements and approaches that help advance scientific breakthroughs for reliable, robust, and trustworthy AI, give us confidence in AI technologies, and cultivate trust in the design, development, and use of AI.

3. Support and expand public-private partnerships to develop and use AI standards and related tools to advance reliable, robust, and trustworthy AI.

- Strategically increase participation in the development of technical AI standards in targeted venues and exercise a variety of engagement options ranging from monitoring to leading – especially at the early stage of standards development where major decisions can be made about scoping and leadership.

- Advance non-traditional collaborative models for standards development, such as open source efforts and Federal open data initiatives.

- Increase data discoverability and access to the Federal government data that is authorized for public use to enable more widespread training and use of AI technologies particularly as they relate to standards development.

- Spur benchmarking efforts to assess the reliability, robustness, and trustworthiness of AI systems. Ensure that these benchmarks are widely available, result in best practices, and improve AI evaluations and methods for verification and validation.

- Foster collaborative environments to promote creative problem solving through AI challenge problems and testbeds to advance standards development.

- Facilitate research and collaboration across scientific disciplines to increase understanding of how societal and ethical considerations relate to and can advance the development and use of standards for reliable, robust, and trustworthy AI technologies.

4. Strategically engage with international parties to advance AI standards for U.S. economic and national security needs.

- Champion U.S. AI standards priorities in AI standards development activities around the world.

- Accelerate the exchange of information between Federal officials and counterparts in like-minded countries through partnering on development of AI standards and related tools.

- Track and understand AI standards development strategies and initiatives of foreign governments and entities.

1. STANDARDS AND ARTIFICIAL INTELLIGENCE

A. WHY IS A PLAN FOR FEDERAL ENGAGEMENT IN AI TECHNICAL STANDARDS NEEDED?

Emphasizing the importance of artificial intelligence (AI) to the future of the U.S. economy and national security, on February 11, 2019, the President issued an Executive Order (EO 13859)[1] directing Federal agencies to take a variety of steps designed to ensure that the nation maintains its leadership position in AI.

Among its objectives, the EO aims to "Ensure that technical standards...reflect Federal priorities for innovation, public trust, and public confidence in systems that use AI technologies; and develop international standards to promote and protect those priorities." The EO also states that the United States must drive development of appropriate technical standards in order to enable the creation of new AI-related industries and the adoption of AI by today's industries. Technical standards will provide agreed-upon language and frameworks to underpin the development and deployment of technological innovations. With the goal of fulfilling their missions more effectively and efficiently, Federal agencies are major players in developing and using AI technologies.[2] Likewise, these agencies should be directly engaged in prioritizing and developing AI technical standards.

United States global leadership in AI will benefit from the Federal government playing an active and purpose-driven role in AI standards development. That includes AI standards-related efforts needed by agencies to fulfill their missions by:

- supporting and conducting AI research and development,
- actively engaging in AI standards development,
- procuring and deploying standards-based products and services, and
- developing and implementing supportive policies, including regulatory policies where needed.

The government's meaningful engagement in fulfilling that role is necessary—but not sufficient—for the nation to maintain its leadership in this competitive realm. Active involvement and leadership by the private sector, as well as academia, is required.

The order directs the Secretary of Commerce, through the National Institute of Standards and Technology (NIST), to issue "a plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies."

Prepared with extensive public and private sector input, this plan provides guidance for bolstering the engagement of Federal agencies in AI technical standards to promote continued U.S. leadership in AI. It focuses on the Federal government's role in advancing AI standards and priorities for research that support development of technically sound and fit for purpose standards.

Note: While definitions of AI vary, for purposes of this plan AI technologies and systems are considered

---

[1]Exec Order No. 13859 Maintaining American Leadership in Artificial Intelligence, 84 FR 3967 (February 14, 2019) https://www.govinfo.gov/content/pkg/FR-2019-02-14/pdf/2019-02544.pdf.

[2]https://www.whitehouse.gov/ai/ai-american-values/.

AI AND TRUSTWORTHINESS

Increasing trust in AI technologies is a key element in accelerating their adoption for economic growth and future innovations that can benefit society. Today, the ability to understand and analyze the decisions of AI systems and measure their trustworthiness is limited.

Among the characteristics that relate to trustworthy AI technologies are accuracy, reliability, resiliency, objectivity, security, explainability, safety, and accountability. Ideally, these aspects of AI should be considered early in the design process and tested during the development and use of AI technologies. AI standards and related tools, along with AI risk management strategies, can help to address this limitation and spur innovation.

to comprise software and/or hardware that can learn to solve complex problems, make predictions or undertake tasks that require human-like sensing (such as vision, speech, and touch), perception, cognition, planning, learning, communication, or physical action. Examples are wide-ranging and expanding rapidly. They include, but are not limited to, AI assistants, computer vision systems, biomedical research, unmanned vehicle systems, advanced game-playing software, and facial recognition systems as well as application of AI in both Information Technology (IT) and Operational Technology (OT).

B. WHAT ARE TECHNICAL STANDARDS AND WHY ARE THEY IMPORTANT?

For the purpose of this plan, "technical standards" refer to "documentary" standards. *ISO/IEC Guide 2:2004 Standardization and related activities—General vocabulary*[3] defines such a standard as "a document, established by consensus and approved by a recognized body, that provides for common and repeated use, rules, guidelines or characteristics for activities or their results, aimed at the achievement of the optimum degree of order in a given context." This plan refers to these simply as "standards." Standards can be developed in many types of organizations that cover a broad spectrum of formality, structure and approach.

Widespread use of standards facilitates technology advancement by defining and establishing common foundations for product differentiation, technological innovation, and other value-added services. Standards also promote an expanded, more interoperable, and efficient marketplace. AI standards that articulate requirements, specifications, guidelines, or characteristics can help to ensure that AI technologies and systems meet critical objectives for functionality, interoperability, and trustworthiness—and that they perform accurately, reliably, and safely.

In contrast, standards that are not fit-for-purpose, are not available when needed, or that are designed around less than ideal technological solutions may hamper innovation and constrain the effective or timely development and deployment of reliable, robust, and trustworthy AI technologies.

Global cooperation and coordination on AI standards will be critical for having a consistent set of "rules of the road" to enable market competition, preclude barriers to trade, and allow innovation to flourish.

---

[3]https://www.iso.org/standard/39976.html.

The U.S. government should ensure cooperation and coordination across Federal agencies and partner with private sector stakeholders to continue to shape international dialogues in regards to AI standards development.

C. HOW ARE TECHNICAL STANDARDS DEVELOPED?

The standards development approaches followed in the United States rely largely on the private sector to develop voluntary consensus standards, with Federal agencies contributing to and using these standards. Typically, the Federal role includes contributing agency requirements to standards projects, providing technical expertise to standards development, incorporating voluntary standards into policies and regulations, and citing standards in agency procurements. This use of voluntary consensus standards that are open to contributions from multiple parties, especially the private sector, is consistent with the U.S. market-driven economy and has been endorsed in Federal statute and policy. (See "Maximizing Use of the Voluntary Consensus Standards Process" on Page 18.)

Some governments play a more centrally managed role in standards development-related activities—and they use standards to support domestic industrial and innovation policy, sometimes at the expense of a competitive, open marketplace. This merits special attention to ensure that U.S. standards-related priorities and interests, including those related to advancing reliable, robust, and trustworthy AI systems, are not impeded.

The timing of standards development can greatly influence the state of technologies. Premature efforts can result in standards that do not reflect the state of the affected technology or may not be supported by a critical mass of technological understanding. This can yield standards that are not fit-for-purpose and can impede innovation. Alternatively, development efforts timed too late may deliver standards that cannot gain market acceptance due to the built-up infrastructure and market-power exerted by incumbent technologies, which will also hinder innovation. Regular review and updating are key elements to ensure that standards reflect technological innovations and take into account changing economic and societal systems.

The development of standards for IT is integral to AI technologies and systems. IT encompasses all technologies for the capture, storage, retrieval, processing, display, representation, security, privacy and interchange of data and information. Worldwide, there are multiple Standards Development Organizations (SDOs) developing IT standards using different models to address varying standardization needs. The rapid innovation in IT has been accompanied by competition among SDOs in areas of market relevance (e.g., cloud computing, cybersecurity, and Internet of Things). This has encouraged SDOs to streamline their consensus-building processes to develop and approve timely, technically sound standards that meet current market needs.

Broadly, IT standards can address cross-sector or sector-specific needs. *Horizontal* IT standards can be used across many applications and industries. Standards developed for specific applications areas such as healthcare or transportation are *vertical* standards. Developers of horizontal standards often seek to establish collaborative working relationships (e.g., liaisons) with sector-specific (vertical) standards developers. These liaisons foster cooperation, establish or reinforce boundaries, and help to ensure that horizontal standards are relevant to other IT standardization efforts and vice versa.

D. WHAT AI TECHNICAL STANDARDS ARE NEEDED?

Systems using AI technologies are generally systems of systems, and AI standards should take this into account. AI standards encompass those specific to AI applications as well as standards for parts of an AI-driven system—and *both types of standards are needed.*

ONE PERSPECTIVE ON AI TECHNICAL STANDARDS

The Center for Data Innovation describes AI standards this way:

"Technical standards for AI can encompass a wide variety of issues, including safety, accuracy, usability, interoperability, security, reliability, data, and even ethics....Flexible, robust, common technical standards for AI will be critical to the successful development and deployment of the technology for two key reasons.

First, technical standards can provide developers clear guidelines for the design of AI systems to ensure that they can be easily integrated with other technologies, utilize best practices for cybersecurity and safety, and adhere to a variety of different technical specifications that maximize their utility.

Second, common standards can serve as a mechanism to evaluate and compare AI systems. For example, in some contexts, there may be a legal requirement for transparency for a decision-making process, such as judicial decision-making. However, without clear standards defining what algorithmic transparency actually is and how to measure it, it can be prohibitively difficult to objectively evaluate whether a particular AI system meets these requirements or expectations, or does so better than another similar system, which discourages the adoption of these technologies. For this reason, in many cases technical standards will be a key component of determining whether an AI system is appropriate for use in a particular context."[4]

A growing number of cross-sector (horizontal) and sector-specific (vertical) AI standards exist and many others are being developed by numerous SDOs.[5] These SDOs have established liaisons that facilitate information exchange and collaboration on standards development. Some areas, such as communications, have well established and regularly maintained standards in widespread use. Other aspects, especially as regards to trustworthiness, are only now being considered.

There are several existing technology standards applicable to AI that were originally developed for other technologies. Standards related to data formats, testing methodology, transfer protocols, cybersecurity, and privacy are examples. The utility and applicability of these standards should be considered before new standardization is initiated.

Tables 1 and 2 of this Plan reflect the present state of AI-relevant standardization[6] based on stakeholder

---

[4]https://www.nist.gov/sites/default/files/documents/2019/05/10/nist-ai-rfi-ctr-for-data-innovation-001.pdf.

[5]See Appendix II for a list of SDOs that are developing AI standards.

[6]Note that some aspects of AI development do not lend themselves to standardization: for example, the mathematical foundations of AI, which focuses mainly on proving theorems in mathematical problems motivated by AI. This includes, but is not limited to areas such as statistics, learning theory and optimization.

input to the NIST Request for Information,[7] the NIST AI Standards Workshop, contacts and discussions with members of Federal agencies involved with artificial intelligence and related topics, comments on a draft of this plan, and other stakeholder interactions.[8] These tables reflect high-level, point-in-time snapshots of the AI related standards development efforts. Areas of standards listed are not mutually exclusive. Often, guidance and requirements in one standard are referenced in others. *Lastly, **even where standards are noted as available or being developed, each area could likely benefit from additional standards** to advance or keep pace with AI technologies, and their widespread use, in a reliable, robust, and trustworthy manner.*

TABLE 1. STATUS OF AI STANDARDIZATION

Note that standards efforts for AI systems recently have been initiated in all of these areas. In addition to developing new AI standards, each of these current or developing standards likely will need to be revised and expanded to address rapidly changing technological innovations, needs, and understanding.

AI STANDARDS	AVAILABLE[9]	BEING DEVELOPED[10]
Concepts and terminology	●	●
Data[11] and knowledge[12]	●	●
Human interaction[13]	●	●
Metrics	●	●
Networking[14]	●	●
Performance testing and reporting methodology[15]	●	●
Safety	●	●
Risk management[16]	●	●
Trustworthiness		●

---

[7]Artificial Intelligence Standards, 84 FR 18490 (May 1, 2019) https://www.govinfo.gov/content/pkg/FR-2019-05-01/pdf/2019-08818.pdf.

[8]See Appendix IV for the process NIST used to develop this plan.

[9]"Available" indicates that some SDO approved standards are available. These are subject to revision based upon feedback from testing and market deployments, as well as improvements in technology.

[10]"Being Developed" indicates that needed standards are still under development and/or that new standards development projects need to be initiated.

[11]Data standards include guidance and requirements for: big data analytics; data exchange; data quality; data accessibility, and data privacy.

[12]Knowledge standards include standards for knowledge representation and querying, such as the W3C Web Ontology Language (OWL) and the ISO Common Logic language (ISO/IEC 24707:2007), as well as standard ontologies formulated in such languages.

[13]Human interactions includes usability and accessibility.

[14]For example, see IEEE 1609.0: 2013—IEEE Draft Guide for Wireless Access in Vehicular Environments (WAVE) Architecture.

[15]Performance Testing and Reporting Methodology standards include testing guidance and requirements at the technology, prototype, and AI operational system levels.

[16]For example, see ISO/IEC 27005:2018 Information technology—Security techniques—Information security risk management.

Some categories in Table 1 are more primed for standards development than others. These include standards for concepts and terminology, data, human interaction, metrics, networking, performance testing and reporting methodology, as well as standards targeted to specific vertical domains. Standardization of AI safety, risk management, and some aspects of trustworthiness such as explainability or security, are in their formative stages and especially would benefit from research to provide a strong technical basis for development.

By defining common vocabularies, establishing the essential characteristics of reliable, robust, and trustworthy AI technologies, and identifying best practices within the life cycle of an AI system, these standards can accelerate the pace of innovation. Similarly, human interaction and performance testing standards spur innovation by establishing the "rules of the game" and forming a baseline from which new technologies emerge. Trustworthiness standards include guidance and requirements for: accuracy, explainability, resiliency, safety, reliability, objectivity, and security.[17] Aspects of trustworthiness also intersect with, and are addressed in, additional areas displayed in this table.

TABLE 2. NON-TECHNICAL STANDARDS TO INFORM POLICY DECISIONS

AI STANDARDS	AVAILABLE[18]	BEING DEVELOPED[19]
Societal and ethical considerations[20]		●
Governance[21]	●	●
Privacy[22]	●	●

Input to the development of this Federal engagement plan suggests that it is important for those participating in AI standards development to be aware of, and to act consistently with, policies and principles. That includes those set by public and private entities that address societal and ethical issues, governance, and privacy as mentioned in Section 1(F). Table 2 lists AI-related standards activities that may inform risk management and policy decisions.

Societal and ethical considerations in IT consist of the analysis of the nature and social impact of IT and the corresponding formulation and justification of policies for the appropriate use of such technology. In terms of developing standards for societal and ethical considerations, it is important to distin-

---

[17]For greater detail about relevant areas of cybersecurity, see NISTIR 8200, Interagency Report on the Status of International Cybersecurity Standardization for the Internet of Things (IoT), https://csrc.nist.gov/publications/detail/nistir/8200/final.

[18]"Available" indicates that some SDO approved standards are available. These standards are subject to revision based upon feedback from testing and market deployments, as well as improvements in technology.

[19]"Being Developed" indicates that needed standards are still under development and/or that new standards development projects need to be initiated.

[20]For example, see draft IEEE P7000 - Model Process for Addressing Ethical Concerns During System Design.

[21]For example, see ISO/IEC 38500:2015 Information technology—Governance of IT for the organization.

[22]For example, see ISO/IEC 29101:2018 Information technology—Security techniques—Privacy architecture framework. And ISO/IEC 27018 and ISO/IEC 27701. Any additional AI related standards or tools addressing privacy should be consistent with NIST Privacy Framework.

HELP WANTED: DATA STANDARDS AND DATA SETS

Data standards make the training data needed for machine learning applications more visible and more usable to all authorized users. Descriptions of data that define authorized use are important elements of data standards. These attributes include but are not limited to: Federal government security classification, the presence of law enforcement sensitive data, proprietary data, acquisition-sensitive data, personally identifiable information (to include biographic, biometric and contextual data for individuals), Freedom of Information Act (FOIA) exemptions, and fees that might be required for data access. This information can help potential users to rapidly evaluate the value and utility of the data before investing time seeking access.

guish between technical and non-technical standards. Not all societal and ethical issues of AI can be addressed by developing technical standards.[23] Non-technical standards can inform policy and human decision-making.[24]

There are information technology governance standards that are equally applicable to AI as they are to other technologies. It may be necessary to augment governance standards with AI considerations or to create new standards in this regard, but any such work should not ignore the existing applicable standards.

Likewise, privacy standards may or may not be specific to AI. They can encompass IT-related issues and operations and may be much broader with a focus on an organization's overall approach to consider potential problems individuals could experience arising from system, product, or service operations with data.

E. WHAT AI STANDARDS-RELATED TOOLS ARE NEEDED?

Standards should be complemented by an array of related tools to advance the development and adoption of effective, reliable, robust, and trustworthy AI technologies. These tools—which often have overlapping applications—include, but are not limited to:[25]

- *Data sets in standardized formats, including metadata* for training, validation and testing of AI systems. Data standards are vital in measuring and sharing information relating to the quality, utility and access of data sets.[26] They can preserve privacy, ensure accessibility, assist potential users in making informed decisions about the data's applicability to their purpose, and help prevent misuse.

---

[23]In considering the applications of AI in areas such as criminal justice and health care, organizations should design, build and deploy AI systems that leverage human judgment and responsibility where they are needed.

[24]Ethics, a mostly qualitative science, is critical to informing an organization's strategy for its technology deployments. Organizations should consider and articulate the values that should be encompassed in their product or service, and how these values might vary across different demographics. Organizations can then proceed by developing AI systems that incorporates those values.

[25]Text for several of the needed tools described in this section stems from discussions among member agencies of the National Science and Technology Council (NSTC) Machine Learning/Artificial Intelligence (ML/AI) Subcommittee.

[26]AI data-related needs also are addressed in other sections of the EO.

- ***Tools for capturing and representing knowledge and reasoning in AI systems*** to promote consistent formulation of, reasoning with, and sharing of knowledge, thereby promoting interoperability of AI systems and minimizing their misunderstandings and inferential errors.

- ***Fully documented use cases*** that provide a range of data and information about specific applications of AI technologies and any standards or best practice guides used in making decisions about deployment of these applications. For the use cases to be of real value, they must be accompanied not only by explicit information about the parameters of use, but also by the practical implications of such uses for persons who may be affected by AI deployments.

EXAMPLES OF AI STANDARDS-RELATED TOOLS

One successful example of a high-impact, community-based, AI-relevant benchmark program is the Text Retrieval Conference (TREC).[27] Started by NIST in 1992, TREC provides the infrastructure necessary for large-scale evaluation of information retrieval methodologies. More than 250 groups have participated in TREC, including academic and commercial organizations, both large and small. The standardized, widely available, and carefully constructed set of data put forth by TREC has been credited with revitalizing research on information retrieval.[28]

NIST also developed a comprehensive set of standard test methods and associated performance metrics to assess key capabilities of emergency response robots, including ground and aerial vehicles. The objective is to facilitate quantitative comparisons of different robot models by capturing data on robot capabilities using standard test methods. These comparisons guide purchasing decisions and help developers to understand the robots' capabilities. Resulting test methods are being standardized through the ASTM International Standards Committee on Homeland Security Applications for robotic operational equipment. Versions of the test methods are used to challenge the research community through the RoboCup Rescue Robot League competitions,[29] which emphasize autonomous capabilities such as mapping and navigation in unknown environments with difficult terrains.

Another example is the Agile Robotics for Industrial Automation Competition (ARIAC).[30] This competition promotes robot agility using the latest advances in AI. A core focus is to test the agility of industrial robot systems, with the goal of enabling them to be more productive and autonomous.

While these efforts provide a strong foundation for driving AI evaluations forward, they are limited by being domain-specific. Additional metrics, testing requirements, testbeds, and benchmarks are needed across a broader range of domains to ensure that AI solutions are broadly applicable and widely adopted.

---

[27]http://trec.nist.gov.

[28]E. M. Voorhees and D. K. Harman, TREC Experiment and Evaluation in Information Retrieval (Cambridge: MIT Press, 2005), Economic Impact Assessment of NIST's Text REtrieval Conference (TREC) Program, July 2010, Brent R. Rowe, Dallas W. Wood, Albert N. Link, Diglio A. Simoni, RTI International, https://trec.nist.gov/pubs/2010.economic.impact.pdf.

[29]https://www.robocup.org.

[30]http://www.nist.gov/ariac.

- **Benchmarks**, evaluations and challenge problems to drive innovation by promoting advancements aimed at addressing strategically selected scenarios; they also provide objective data to validate and track the evolution of AI technologies.

- **Testing methodologies** to validate and evaluate AI technologies' performance, especially to prescribe protocols and procedures. These tools are needed for specifying, assessing, comparing, and managing the performance and trustworthiness of AI technologies. Among other things, applications include testing for conformance, interoperability, security, and comparing AI systems to human performance. Testing methodologies have utility in multiple stages of AI deployment, from training, to production roll-out, to operational monitoring. They also come into play as part of on-going technology and system assessments and whenever existing AI systems and their outputs are repurposed for a new task that is outside of their original intent.

- **Metrics** to quantifiably measure and characterize AI technologies, including but not limited to aspects of hardware and its performance (at device, circuit, and system levels), trustworthiness (e.g., accuracy, explainability, resiliency, safety, reliability, objectivity, and security), complexity, domain-specific and context-dependent risk, and uncertainty.

- **AI testbeds** "so that researchers can use actual operational data to model and run experiments on real-world system[s] … and scenarios in good test environments."[31]

- **Tools for accountability and auditing** to enable examination of an AI system's output (e.g., decision-making or prediction). These tools can improve traceability by providing a record of events and information regarding technologies' implementation and testing. In doing so, they can ehance assessment and documentation of gaps between predicted and achieved AI systems' outcomes. To address differing needs, in addition to developing cross-sector tools for accountability and auditing, sector-specific tools can aid in focusing on the risks and impacts associated with particular sectors and applications.

F. WHAT ARE OTHER IMPORTANT CONSIDERATIONS?

Like several other pioneering areas of science and technology, the development of AI raises a host of legal, ethical, and societal issues that create real and perceived challenges for developers, policy makers, and users—including the general public. These are matters appropriate for consideration in the policy realm to be applied in the development and deployment of AI technologies and systems. Standards are one tool for implementing or informing policies and principles related to such issues.

Public input on this Federal engagement plan has highlighted the importance of establishing aspirational principles and goals in developing AI standards—along with the associated need to be mindful of the current state of the practice and its limitations. In this arena, standards flow from principles, and a first step toward standardization will be reaching broad consensus on a core set of AI principles. These kinds of principles are being forged by multiple organizations, including the Organisation for Economic

---

[31]SRI International and USC Information Sciences Institute, "Cybersecurity Experimentation of the Future (CEF): Catalyzing a New Generation of Experimental Cybersecurity Research", Final Report, July 31, 2015.

Cooperation and Development (OECD), whose member countries (including the United States) recently adopted such principles.[32]

While stakeholders in the development of this plan expressed broad agreement that societal and ethical considerations must factor into AI standards, it is not clear how that should be done and whether there is yet sufficient scientific and technical basis to develop those standards provisions. Moreover, legal, societal, and ethical considerations should be considered by specialists trained in law and ethics. Three areas where there is some consensus are:

- The degree to which ethical considerations might be incorporated into standards should be tied tightly to the type, likelihood, degree, and consequence of risk to humans,

- Privacy risks are different depending on the use case, the type of data involved, the societal and cultural context, and many other factors. Privacy considerations should be included in any standards governing the collection, processing, sharing, storage, and disposal of personal information, and

- Standards should facilitate AI systems that function in a robust, secure and safe way throughout their life cycles.

Legal, ethical, and societal considerations also can come into play as developers and policy makers consider whether and how to factor in the management of risk to individuals, communities, and society at large. Some standards and standards-related tools aim to provide guidance for evaluating risks, which can be used by developers and policy makers in considering how to manage those risks. Ultimately, it is up to system owners and users to determine what risks they are willing to accept, mitigate, or avoid within existing regulations and policies.

The degree of potential risk presented by particular AI technologies and systems will help to drive decision making about the need for specific AI standards and standards-related tools.

---

[32]https://www.oecd.org/going-digital/ai/principles/.

2. U.S. GOVERNMENT AI STANDARDS PRIORITIES

A. PRIORITIES FOR FEDERAL GOVERNMENT INVOLVEMENT

Which standards development efforts merit federal engagement?

In deciding which standards efforts merit strong Federal government involvement, U.S. government agencies should prioritize AI standards efforts that are:

- *Consensus-based*, where decision-making is based upon clearly established terms or agreements that are understood by all involved parties, and decisions are reached on general agreement.

- *Inclusive and accessible*, to encourage input reflecting diverse and balanced communities of users, developers, vendors, and experts. Stakeholders should include representatives from diverse technical disciplines as well as experts and practioners from non-traditional disciplines of special importance to AI such as ethicists, economists, legal professionals, and policy makers: essentially, accommodating all desiring a "seat at the table."

- *Multi-path*, developed through traditional and novel standards-setting approaches and organizations that best meet the needs of developers and users in the marketplace as well as society at large. (See "Nimble, Multi-path Standards Development" on Page 19.)

- *Open and transparent*, operating in a manner that: provides opportunity for participation by all directly- and materially- affected; has well-established and readily accessible operating rules, procedures, and policies that provide certainty about decision making processes; allows timely feedback for further consideration of the standard; and ensures prompt availability of the standard upon adoption.[33]

- *Result in globally relevant and non-discriminatory standards*, where standards avoid becoming non-tariff trade barriers or locking in particular technologies or products.

This plan intentionally does not single out any individual SDO for priority attention in terms of AI standards development. Federal policy defines goals and calls for each agency to make decisions about participation in the development and use of standards based on its specific needs. (See "Maximizing Use of the Voluntary Consensus Standards Process" on Page 18.) Moreover, especially in fast-moving areas of technology such as AI, new standards initiatives are launched by existing—and new—organizations aiming to develop standards. This section of the plan outlines priorities for government engagement in AI standards and related tools. Section 3 makes recommendations for how the Federal government can engage effectively to meet U.S. needs in order to retain its leadership in AI.

Which standards characteristics are important?

There are a variety of characteristics related to AI standards that deserve priority Federal government consideration, including:

---

[33]Some Federal agencies operate with specific restrictions on information sharing (e.g., due to the classification of information) that limits their full participation in AI standards-related activities. These agencies still may be able to contribute in a meaningful way to standards development, and they should be expected to consider prioritizing use of such standards when they are produced.

- *Innovation-oriented* to keep pace with rapid technological changes, including maximum flexibility, technology- and platform-neutral, and a preference for performance-based (versus prescriptive) requirements to accommodate varied approaches in meeting the standard.

- *Applicable across sectors* (horizontal) to allow for wide-scale deployments in multiple areas of industry, government, and society.

- *Focused on particular sectors and applications* (vertical), especially where there are specific risks and impacts.

- *Clearly stated scope and intended use* to allow users to decide whether an AI standard that is intended for one application is appropriate for other applications based on the data or algorithms used, or if the level of risk is deemed acceptable.

- *Addresses the need to monitor and manage AI systems* throughout the entire product lifecycle.

---

[34]A-119, Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities, is avail-able at https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/circulars/A119/revised_circular_a-119_as_of_1_22.pdf.

- *Reflective of the early state of development and understanding of AI technologies, risk, and societal implications* so that standards initiatives appropriately represent the state of AI technological feasibility and understanding.

- *Regularly updated* to reflect the rapid pace of change in AI technology and to avoid locking out new developments and knowledge about technological as well as social impacts.

- *Effective in measuring and evaluating AI system performance* to assist in determining degree of risk, deciding on fit-for-purpose and readiness, considering conformance, and monitoring effectiveness.

- *Human-centered* to ensure that human interactions and values—including abilities, disabilities, diversity—are considered during AI data collection, model development, testing, and deployment.

- *Harmonized and using clear language* to define AI-related terms and concepts and to promote interoperability and facilitate uniformity across the board in vocabulary and nomenclature.

- *Sensitive to ethical considerations*, identifying and minimizing bias, and incorporating provisions that protect privacy and reflect the broader community's notions of acceptability.

B. LEVELS OF U.S. GOVERNMENT ENGAGEMENT IN AI STANDARDS

AI standards needs are expansive and challenging, and it is widely acknowledged that serious work on AI-specific standards has only recently begun in earnest. U.S. engagement in establishing AI standards is critical; AI standards developed without the appropriate level and type of involvement may exclude or disadvantage U.S.-based companies in the marketplace as well as U.S. government agencies. Furthermore, due to the foundational nature of standards, the lack of U.S. stakeholder engagement in the development of AI standards can degrade the innovativeness and competitiveness of the U.S. in the long term. Possible levels and types of Federal involvement in the standards development process can be grouped into four categories ranked from least-to-most engaged:

NIMBLE, MULTI-PATH STANDARDS DEVELOPMENT

There is a need for stability (provided by the arena of formal standards bodies), coping with rapid change (provided by consortia and alliances), specific intellectual property and marketing environments, and the need for robust community involvement (provided by Open Source). To tackle the vast emerging standardization needs for AI and AI systems, the groups within each arena need to work together more effectively to create standards of the highest quality, through open systems and open standardization processes that effectively contribute to the public good.[35]

Standardization processes must be sufficiently nimble to effectively address the development and commercial application of rapidly evolving technologies such as AI, and they must be open to addressing ethically aligned design concepts from the onset.[36]

---

[35]Text from IEEE response to the NIST RFI. Available at www.nist.gov/sites/default/files/documents/2019/05/28/nist-ai-rfi-ieee_001.pdf.
[36]Same as above.

- *Monitoring:* Following either a specific standards effort or broader programs and evolving standards being produced by SDOs to address unique needs or interests.

- *Participating:* Commenting on and providing meaningful contributions to strategically important standards, including potentially serving as an observer on a committee.

- *Influencing:* Developing a deeper understanding of, and relationships with, the key players—working directly with industry and international players and exerting influence through formal and informal discussions and by providing expertise.

- *Leading:* Leading standards efforts by convening or administering consensus groups, serving as standards project editor or in similar technical leadership roles, or acting as the liaison representative between standards groups. Leadership also can be exercised by serving on the Board of Directors or in other executive positions of an SDO.[37]

Each of these categories of engagement requires having qualified U.S. government participants (Federal employees or contractors) function in these capacities based on their expertise, relationships, and knowledge of specific standards development processes and best practices.

C. PRACTICAL STEPS FOR AGENCY ENGAGEMENT IN AI STANDARDS

1. Identify how AI technologies can be used to further the agency's mission—for example, research, technology development, procurement, or regulation.

2. Know existing statutes, policies and resources relating to participation in the development of, and use of standards (e.g., OMB Circular A-119, Interagency Committee on Standards Policy.[38])

3. Conduct a landscape scan and gap analysis to identify standards and related tools that exist or need to be developed.

4. Use appropriate standards if they exist.

5. Engage in the development of standards if they do not exist or do not fully meet agency requirements.

    i. Coordinate with other Federal agencies that may have similar needs.

    ii. Identify the appropriate level of engagement: see section 2(B).

    iii. Follow guidance on where and how to engage: see section 2(A).

    iv. Identify, train, and enable staff to participate in standards development.

---

[37]See OMB Circular A-119: Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities, pp 27-28, for the policy regarding participation on boards of standards development organizations.

[38]Standards.gov (https://www.standards.gov).

AGENCIES DETERMINING THEIR AI STANDARDS NEEDS

Federal agencies contributing to the development of standards leading to reliable, robust, and trustworthy AI must first understand and articulate the anticipated AI role in agency operations and its regulations and regulated entities. Agencies also should have and share their visions of how AI will affect stakeholders and communities.

A Federal agency cannot optimize its standards activities if it has not yet determined its needs for standards and specific requirements, let alone participate in an effective leadership capacity to meet these needs. Several Federal departments and agencies are ahead of the curve in examining the use and impact of AI and strategies for considering the role of AI standards.

The Department of Transportation report, Preparing for the Future of Transportation: Automated Vehicles 3.0 (AV 3.0)[39] provides a vision for using AI and its potential impact. Voluntary consensus standards are mentioned throughout the report as a strategy for supporting Automated Driving Systems and Automated Vehicle development.

The Food and Drug Administration report, Proposed Regulatory Framework for Modification to Artificial Intelligence/Machine Learning (AI/ML)-based Software as a Medical Device (SaMD), leans forward in "considering a...product lifecycle-based regulatory framework for technologies that would allow for modifications to be made from real-world learning and adaptation, while still ensuring that the safety and effectiveness...is maintained."[40]

Both of these agencies articulate an understanding of the impact of AI and propose a path ahead. The white paper "ai: using standards to mitigate risk,"[41] published jointly through the Department of Homeland Security and the Office of the Director of National Intelligence, serves to "start a dialogue on creating standards that will reduce the risk from use, misuse, and exploitation of AI".[pg. 3] These examples highlight the urgent need for every agency to comprehend and appreciate the impact of AI on their missions.

---

[39]https://www.transportation.gov/av/3/preparing-future-transportation-automated-vehicles-3.

[40]Abstracted from webpage: https://www.fda.gov/medical-devices/software-medical-device-samd/40artificial-intelligence-and-machine-learning-software-medical-device.

[41]AI: Using Standards to Mitigate Risks, produced by the Public-Private Analytic Exchange Program, 2018. https://www.dhs.gov/sites/default/files/publications/2018_AEP_Artificial_Intelligence.pdf.

3. RECOMMENDED FEDERAL GOVERNMENT STANDARDS ACTIONS TO ADVANCE U.S. AI LEADERSHIP

In addition to the guidance provided regarding priorities and levels of engagement called for in the previous section of this plan, ***the Federal government should commit to deeper, consistent, long-term engagement in AI standards development activities to help the United States to speed the pace of reliable, robust, and trustworthy AI technology development.*** Specifically, the Federal government[42] should:

1. Bolster AI standards-related *knowledge, leadership, and coordination* among Federal agencies to maximize effectiveness and efficiency.

- The National Science and Technology Council (NSTC) Machine Learning/Artificial Intelligence (ML/AI) Subcommittee should designate a Standards Coordinator with responsibility to gather and share AI standards-related needs, strategies, roadmaps, terminology, use cases, and best practices in support of reliable, robust, and of trustworthy AI in government operations, including:

  – planned and ongoing standards approaches and engagement activities, including a robust feedback loop with SDOs,

  – specific horizontal or vertical areas for prioritization,

  – requirements for input into proposed standards activities, and

  – analyses of whether ongoing standards activities meet Federal government needs and whether additional guidance is appropriate.

- Make maximum use of existing standards that are broadly adopted by industry sectors and can be used or evolved within the new context of AI solutions.

- Reinforce the importance of agencies' adherence to Federal policies for standards and related tools; for example, data access and quality. *Suggested lead: Office of Management and Budget*

- Be flexible in selecting AI standards for use in regulatory or procurement actions to meet agency objectives. Flexibility is required to adapt to the rapid pace of AI technology developments and standards and to reflect the increasing understanding about trustworthiness, accessibility, and human-centered implications of AI.

- Grow a cadre of Federal staff with the relevant skills and training. Develop and provide a clear career development and promotion path that values and encourages participation in and expertise in AI standards and standards development. *Suggested lead: Department of Commerce[43], Office of Personnel Management.*

---

[42]Except where specific agencies are noted, all agencies with AI-related needs and activities should consider their possible contribution to implementing each recommendation.

[43]As the U.S. national measurement institute and leader in research supporting technology standardization NIST leads Department of Commerce efforts to implement recommendations of this plan.

2. **Promote focused research to advance and accelerate broader exploration and understanding of how aspects of trustworthiness can be practically incorporated within standards and standards-related tools.**

  - Plan, support, and conduct research and evaluation that underlies technically sound, fit-for-purpose standards and related tools for reliable, robust, and trustworthy AI. *Suggested lead: Department of Commerce, National Science Foundation and research funding agencies.*

  - Develop metrics and data sets to assess reliability, robustness, and other trustworthy attributes of AI systems, focusing on approaches that are readily understandable, available, and can be put on a path to standardization. *Suggested lead: Department of Commerce and research funding agencies.*

  - Conduct research to inform standardization of risk management strategies including identifying, monitoring and mitigating risks. *Suggested lead: research funding agencies.*

  - Identify research needs, requirements and approaches that help advance scientific breakthroughs for reliable, robust, and trustworthy AI to give us confidence in AI technologies, and cultivate trust in the design, development, and use of AI. *Suggested lead: Department of Commerce and research funding agencies.*

3. **Support and expand public-private partnerships to develop and use AI standards and related tools to advance reliable, robust, and trustworthy AI.**

  - Strategically increase participation in the development of technical AI standards in targeted venues. Participation may include a variety of engagement options ranging from monitoring to leading—especially at the early stage of standards development where major decisions can be made about scoping and leadership. In making decisions about involvement in standards development, consider the priorities and guidelines cited in Section 2(A) and (B) and SDO activities cited in Appendix II.

  - Advance non-traditional collaborative models for standards development, such as open source efforts and Federal open data initiatives.

  - Increase data discoverability and access to the Federal government data that is authorized for public use to enable more widespread training and use of AI technologies particularly as they relate to standards development.

  - Spur benchmarking efforts to assess the reliability, robustness, and trustworthiness of AI systems. Ensure that these benchmarks are widely available, result in best practices, and improve AI evaluations and methods for verification and validation. *Suggested lead: Department of Commerce.*

  - Foster collaborative environments to promote creative problem solving through AI challenge problems and testbeds to advance standards development.

  - Facilitate research and collaboration across scientific disciplines to increase understanding of how societal and ethical considerations relate to, and can advance the development and use of standards for reliable, robust, and trustworthy AI technologies. *Suggested lead: National Science Foundation, Department of Commerce.*

4. Strategically engage internationally to advance AI standards for U.S. economic and national security needs.

- Champion U.S. AI standards priorities in AI standards development activities around the world.

- Accelerate the exchange of information between Federal officials and counterparts in like-minded countries through partnering on development of AI standards and related tools. *Suggested lead: Department of Commerce, Department of State, Department of Justice.*

- Track and understand AI standards development strategies and initiatives of foreign governments and entities. *Suggested lead: Department of Commerce, Department of State, Department of Justice.*

APPENDIX I: DEFINITIONS

ANSI INCITS 172-2002 (R2007) Information Technology—American National Standard Dictionary of Information Technology (ANSDIT) (Revision and Redesignation Of ANSI X3.172-1996)

artificial intelligence (AI):
(1) A branch of <u>computer science</u> devoted to developing <u>data processing systems</u> that performs <u>functions</u> normally associated with human intelligence, such as <u>reasoning</u>, <u>learning</u>, and self-improvement. (2) The capability of a device to perform functions that are normally associated with human intelligence such as reasoning, learning, and self-improvement.

ISO/IEC 3WD 22989 Information Technology—Artificial Intelligence—Artificial Intelligence Concepts and Terminology

artificial intelligence:
capability of a system to acquire, process, and apply knowledge
Note 1 to entry: knowledge are facts, information, and skills acquired through experience or education

AI system:
technical system that uses artificial intelligence to solve problems

APPENDIX II: AI STANDARDS

Noting that standards development is an ongoing effort with new projects and new technical focus areas being added regularly, any listing of standards bodies and associated AI standards development activities is current only at the time the list was developed. The following are examples of activities provided by stakeholder input from the NIST Request for Information, the NIST AI Standards Workshop, contacts and discussions with members of federal agencies involved with artificial intelligence and related topics, and comments on a draft of this plan.

**International Organization for Standardization (ISO)
and the International Electrotechnical Commission (IEC)**

ISO/IEC JTC 1/SC 42 Artificial Intelligence was established in October 2017 to develop AI standards that can be applied across applications and industries. SC 42 is chartered to work on Information Technology standards, with current work items focused on topics such as updated AI terminology, interoperable framework for AI systems, AI lifecycle, big data, AI trustworthiness (e.g., robustness, unbiased, and risk managed), use cases, and computational approaches.

Published Standards under ISO/IEC JTC 1/SC 42 Artificial Intelligence

ISO/IEC 20546:2019 Information technology — Big data — Overview and vocabulary

ISO/IEC TR 20547-2:2018 Information technology — Big data reference architecture – Part 2: Use cases and derived requirements

ISO/IEC TR 20547-5:2018 Information technology — Big data reference architecture – Part 5: Standards roadmap

Standards under development by ISO/IEC JTC 1/SC 42 Artificial Intelligence

ISO/IEC AWI TR 20547-1: Information technology — Big data reference architecture — Part 1: Framework and application process

ISO/IEC DIS 20547-3: Information technology — Big data reference architecture — Part 3: Reference architecture

ISO/IEC WD 22989: Artificial Intelligence Concepts and Terminology

ISO/IEC WD 23053: Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML)

ISO/IEC NP TR 24027: Information technology — Artificial Intelligence (AI) — Bias in AI systems and AI aided decision making

ISO/IEC NP TR 24028: Information technology — Artificial Intelligence (AI) — Overview of trustworthiness in Artificial Intelligence

ISO/IEC NP TR 24029-1: Artificial Intelligence (AI) — Assessment of the robustness of neural networks — Part 1: Overview

ISO/IEC NP TR 24030: Information technology — Artificial Intelligence (AI) — Use cases

ISO/IEC NP 38507: Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations

ISO/IEC NP 23894: Information technology — Artificial Intelligence — Risk Management

Other ISO/IEC JTC 1 and ISO Standards Activities

AI-related cross sector and sector specific standards have been and are being developed in several committees and subcommittees.

JTC 1 SC 7: Software and systems engineering

JTC 1 SC 17: Cards and security devices for personal identification

JTC 1 SC 22: Programming languages, their environments and system software interfaces

JTC 1 SC 24: Computer graphics, image processing and environmental data representation

JTC 1 SC 27: Information Security, cybersecurity and privacy protection

JTC 1 SC 28: Office equipment

JTC 1 SC 29: Coding of audio, picture, multimedia and hypermedia information

JTC1 SC 32: Data management and interchange

Examples:

ISO/IEC 24707:2018 Information technology -- Common Logic (CL) -- A framework for a family of logic-based languages

ISO/IEC DIS 21838-2 Information technology -- Top-level ontologies (TLO) -- Part 2: Basic Formal Ontology (BFO)

JTC 1 SC 36: Information technology for learning, education and training

JTC 1 SC 37: Biometrics

JTC 1 SC 40: IT Service Management and IT Governance

JTC 1 SC 41: Internet of Things and related technologies

ISO TC 184: Automation systems and integration

ISO TC 199: Safety of machinery

ISO TC 299: Robotics

Institute of Electrical and Electronics Engineers (IEEE)

Standards under development by IEEE

Starting in 2016, the *IEEE P7000™* series of standards projects addresses specific issues at the intersection of technological and ethical considerations for AI.

P7000 - Model Process for Addressing Ethical Concerns During System Design

P7001 - Transparency of Autonomous Systems

P7002 - Data Privacy Process

P7003 - Algorithmic Bias Considerations

P7004 - Standard for Child and Student Data Governance

P7005 - Standard for Transparent Employer Data Governance

P7006 - Standard for Personal Data Artificial Intelligence (AI) Agent

P7007 - Ontological Standard for Ethically Driven Robotics and Automation Systems

P7008 - Standard for Ethically Driven Nudging for Robotic, Intelligent and Autonomous Systems

P7009 - Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems

P7010 - Wellbeing Metrics Standard for Ethical Artificial Intelligence and Autonomous Systems

P7011 - Standard for the Process of Identifying and Rating the Trustworthiness of News Sources

P7012 - Standard for Machine Readable Personal Privacy Terms

P7013 - Inclusion and Application Standards for Automated Facial Analysis Technology

P7014 - Standard for Ethical considerations in Emulated Empathy in Autonomous and Intelligent Systems

Other IEEE AI Standards

P2049.1 - Standard for Human Augmentation: Taxonomy and Definitions

P2049.2 - Standard for Human Augmentation: Privacy and Security

P2049.3 - Standard for Human Augmentation: Identity

P2049.4 - Standard for Human Augmentation: Methodologies and Processes for Ethical Considerations

P2801 Recommended Practice for the Quality Management of Datasets for Medical Artificial Intelligence

P2802 Standard for the Performance and Safety Evaluation of Artificial Intelligence Based Medical Device: Terminology

P2089 - Standard for Age Appropriate Digital Services Framework - Based on the 5Rights Principles for Children

P2817 - Guide for Verification of Autonomous Systems

P3333.1.3 - Standard for the Deep Learning-Based Assessment of Visual Experience Based on Human Factors

P3652.1 Guide for Architectural Framework and Application of Federated Machine Learning

ASTM Technical Committees

Several ASTM technical committees are developing standards to support the reliable, robust, and trust-worthy systems that use AI.

ASTM Committee F15 on Consumer Products was formed in 1973 and maintains over 100 standards. Subcommittee F15.75 on Connected Products is working on a standard to provide guidance for consumer Internet of Things (IoT) as it relates to connected product hazards. It will apply to consumer products connected to the internet that need testing and evaluation of software to prevent cybersecurity vulnerabilities and software weaknesses that could compromise safety related performance of the product and create a safety hazard. Consumer IoT product (CIP) means a physical object that transmits or receives data remotely through a network, other than a mobile phone or personal computer, primarily intended for consumer use remotely through a network. Examples of these types of products include baby monitors, wearable health trackers, and connected appliances. Consumer IoT standards will be intended to apply in conjunction with product specific standard requirements to address the overall system safety of a connected end product.

ASTM Committee F45 on Driverless Automatic Guided Industrial Vehicles was formed in 2014. This Committee addresses issues related to performance standards and guidance materials for 'automatic'- (e.g., automatic guided vehicles) through 'autonomous'- (e.g., mobile robots) unmanned ground vehicles (A-UGVs) with industrial applications. A-UGV applications include, but are not limited to: indoor warehouse, manufacturing, and medical facilities and outdoor security and shipyards. It also works closely with industrial vehicle safety standards organizations.

ASTM Committee F38 on Unmanned Aircraft Systems was formed in 2003 and maintains over 15 standards. This Committee addresses issues related to design, performance, quality acceptance tests, and safety monitoring for unmanned air vehicle systems. F38 is working on standards to assist unmanned aircraft in detection and avoidance and containing complex functions sometimes referred to as "autonomous."

ASTM Committee F42 on Additive Manufacturing (AM) Technologies was formed in 2009 and maintains over 22 standards. This committee addresses standards related to the process of creating three-dimensional objects by the successive addition of material – whether plastic, metal, ceramic, or composite. Artificial intelligence, machine learning (ML), and deep learning (DL) are used in the

selection of AM materials and the development of AM devices/systems to find the best combinations of processing routes to obtain required properties or functionalities. Such technologies help rapidly suggest candidate materials for AM or predict functionalities of devices/systems based on multiple AM design parameters. Such digital, smart AM frameworks operate by reducing the huge design space needed for materials, guiding processes, and facilitating integration of complex data from design, processing, characterization, and simulation. In addition, AI/ML/DL for AM are intimately connected with other data-intensive activities such as AM data management/databases with respect to the data FAIR (findable, accessible, interoperable, and reusable) principles, as well as data-driven areas such as integrated computational materials engineering (ICME) and the Materials Genome Initiative (MGI) to identify structure-property-processing-performance relationships.

ASTM Administrative Committee 377, Autonomy Design and Operations in Aviation (AC377) was formed in 2017 through the collective actions of four of ASTM's aviation-related technical committees. The purpose of the administrative committee is to help harmonize standards development efforts related to autonomy/AI in the aviation community.

In June 2019, ASTM published a technical report entitled "Autonomy Design and Operations in Aviation: Terminology and Requirements Framework" to serve as a guide for terminology and requirements for increasingly autonomous and complex aviation systems. The committee is now working on technical guidance regarding specific aspects of AI/autonomy to assistant technical committees with their standards development efforts.

The Consumer Technology Association (CTA) is currently developing three standards:

Definitions and Characteristics of Artificial Intelligence (under development)
Scope: This standard defines terms related to artificial intelligence and associated technologies.

Definitions and Characteristics of Artificial Intelligence in Health Care
Scope: This standard defines terms related to artificial intelligence and associated technologies in health care.

The Use of Artificial Intelligence in Health Care: Trustworthiness
Scope: Artificial Intelligence (AI) is quickly becoming a pervasive tool in the health care industry. This standard explores the impact of the trustworthiness of AI in health care through the lens of the end user (e.g., physician, consumer, professional and family caregiver). Additionally, the standard will identify the unique challenges and opportunities for AI in the health care sector.

The International Telecommunication Union Telecommunication Standardization Sector (ITU-T) is investigating possible standardization work for AI in the following focus groups:

Focus Group on Machine Learning for Future Networks including 5G

Focus Group on Artificial Intelligence for Health

Focus Group on Environmental Efficiency for Artificial Intelligence and other Emerging Technologies

The **Object Management Group (OMG)** cross-sector AI-related specifications under development include:

Application Programming Interfaces for Knowledge Platforms (API4KP)

Robotics Service Ontology (RoSO)

OMG's sector-specific AI-related specifications under development include:

A retail specification for digital receipts that embodies an ontology defining not only the receipts themselves, but also content related to jurisdiction-specific taxation.

A joint effort between OMG's Retail and Robotics Task Forces to create a standard for point-of-sale/point-of-service (POS) robotic interfaces for the 2020 Olympics specifically, but which will be broadly applicable to POS robotic services.

Society of Automotive Engineering International (SAE International)

SAE J 3016-2018, *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*

SAE CRB 1-2016 (SAE CRB1-2016) *Managing The Development Of Artificial Intelligence Software (Stabilized: May 2016)*

U.S. Department of Transportation

Examples of both horizontal cross sector and vertical sector-specific standards for AI systems are found in the Department of Transportation report, *Preparing for the Future of Transportation: Automated Vehicles 3.0* (AV 3.0). Voluntary consensus standards are mentioned throughout this report as a strategy for supporting Automated Driving Systems and Automated Vehicle development. Appendix C, "Voluntary Technical Standards for Automation," lists numerous AI-relevant horizontal and vertical standards in the functional areas of: Definitions and Architecture; Data; Design; Maintenance and Inspections; Functional/Performance; Protocols (Communications); Security; and Testing/Test Target.

World Wide Web (W3C)

The W3C Web Ontology Language (OWL) is a Semantic Web language designed to represent rich and complex knowledge about things, groups of things, and relations between things. OWL is a computational logic-based language such that knowledge expressed in OWL can be exploited by computer programs, e.g., to verify the consistency of that knowledge or to make implicit knowledge explicit. OWL documents, known as ontologies, can be published in the World Wide Web and may refer to or be referred from other OWL ontologies. OWL is part of the W3C's Semantic Web technology stack.

The W3C Semantic Web Standards

The term "Semantic Web" refers to W3C's vision of the Web of linked data. Semantic Web technologies enable people to create data stores on the Web, build vocabularies, and write rules for handling data.

Standards include: RDF, OWL, SPARQL, RDFa, JSONLD, SKOS, RDFS, GRDDL, POWDER, PROV, RIF, SAWSDL, RDB2RDF.

APPENDIX III: AI-RELATED TOOLS

The following are examples of activities provided by stakeholder input from the NIST Request for Information, the NIST AI Standards Workshop, contacts and discussions with members of federal agencies involved with artificial intelligence and related topics, and comments on a draft of this plan.

Data Repositories (e.g., for algorithm training)

Data repositories are essential for training and evaluating AI models. Some examples of public AI datasets include:

CIFAR-10 dataset (Canadian Institute for Advanced Research) is a collection of images that are commonly used to train machine learning and computer vision algorithms.

COCO is a large-scale object detection, segmentation, and captioning dataset.

Data.gov maintains a catalog of over 250,000 open sourced government datasets in topics such as agriculture, climate, consumer, ecosystems, education, energy, finance, health, local government, manufacturing, maritime, ocean, public safety and science.

ImageNet project is a large visual database designed for use in visual object recognition software research. More than 14 million images have been hand-annotated by the project to indicate what objects are pictured and in at least one million of the images, bounding boxes are also provided.

MNIST dataset of handwritten digits has a training set of 60,000 examples, and a test set of 10,000 examples. It is a subset of a larger set available from NIST.

Google AI periodically releases data of interest to a wide range of computer science disciplines for researchers to use.

Kaggle provides an extensive collection of thousands of datasets from a wide variety of different sources and domains.

OpenML is a data set repository that links data to algorithms to teach machines to learn better.

Pascal VOC data sets provides standardized image data sets for object class recognition, a common set of tools for accessing the data sets and annotations, and enables evaluation and comparison of different methods.

Registry of Open Data on AWS provides an extensive collection of thousands of datasets from a wide variety of different sources and domains.

UC Irvine Machine Learning Repository maintains and hosts a diverse collection of now over 400 datasets ranging between the 1980's to 2019.

Evaluations and Benchmarks

NIST TREC efforts today extend to more sophisticated AI tasks including complex question answering, incident management, and news summarization, as well as to industry specific challenges. NIST TREC has also expanded into modalities beyond text, such as with the NIST TRECVID evaluations for tasks related to digital video, NIST Multimedia Event Detection (MED), and NIST Multimedia Event Recounting (MER). These evaluations are important for driving fundamental advancements in the accuracy of AI technologies on a growing field of tasks using data modalities such as images, video, speech, and text.

The IIC Deep Learning Facilities Testbed is intended to optimize diagnosis, maintenance, and repair of monitored assets; increase energy efficiency by adjusting power-consuming services, and improve visitor experience relative to wait times and ambient climate control.

The MLPerf effort aims to build a common set of benchmarks that enables the machine learning (ML) field to measure system performance for both training and inference from mobile devices to cloud services.

AI-Matrix is an AI benchmark suite aiming at measuring the performance of AI hardware platforms and software frameworks. This deep learning benchmark suite currently consists of three types of workloads: layer-based benchmark, macro benchmark, and micro benchmark.

AIIA DNN is a benchmark to objectively reflect the current state of AI accelerator capabilities and all metrics are designed to provide an objective comparison dimension.

AnTuTu is a benchmarking tool for Android smartphones and tablets for checking device performance.

DeepBench is intended to benchmark operations that are important to deep learning on different hardware platforms.

Fathom provide reference workloads for modern deep learning.

Metrics

In its TREC evaluations, NIST has helped to establish important metrics for the AI field - precision vs. recall, mean average precision, and false alarm vs. miss rate.

Industry has also played a prominent role in the development of metrics, such as in the case of the BLEU metric created by IBM Research, which has achieved wide use for evaluating natural language-related AI tasks.

Industry- and academia-driven evaluations are also using metrics such as top-1 and top-5 accuracy for evaluating classification results and intersection-over-union to measure localization in object detection.

CLEVER (Cross-Lipschitz Extreme Value for nEtwork Robustness) is a metric for measuring the robustness of deep neural networks.

Open Source Software

There are a number of open source AI frameworks available such as the following:

Caffe was developed by Berkeley AI Research (BAIR) and by community contributors.

Keras is a deep-learning library that sits atop TensorFlow and Theano.

Machine Learning in R (mlr) provides a generic, object- oriented, and extensible framework for classi-fication, regression, survival analysis and clustering for the R language. It provides a unified interface to more than 160 basic learners and includes meta-algorithms and model selection techniques to improve and extend the functionality of basic learners with, e.g., hyperparameter tuning, feature selection, and ensemble construction.

Markov Decision Process Toolbox for Python implements the classes and functions needed to perform backwards induction, linear programming, policy iteration, q-learning, value iteration and their variations.

MxNet is an open-source deep learning software framework, used to train, and deploy deep neural networks.

PyTorch is an open-source machine learning library. It has an extensive ecosystem of tools and libraries that extend PyTorch and support development in computer vision, NLP and more

Scikit-learn is a software machine learning library for the Python programming language.

Swi-Prolog provides a comprehensive Prolog environment. Prolog is a declarative logic programming language with a variety of applications in developing AI systems.

TensorFlow is an end-to-end open source platform for machine learning. It has a comprehensive, flexible ecosystem of tools, libraries and community resources that lets researchers push the state-of-the-art in ML and developers easily build and deploy ML powered applications.

Theano is a Python library and optimizing compiler for manipulating and evaluating mathematical expressions, especially matrix-valued ones.

Torch is an open-source machine learning library, a scientific computing framework, and a script language based on the LUA programming language.

Other AI open source software projects include:

Acumos AI is a platform and open source framework that makes it easy to build, share, and deploy AI apps. Acumos standardizes the infrastructure stack and components required to run an out-of-the-box general AI environment.

Adversarial Robustness Toolbox (ART) IBM has released this open source toolbox. ART implements state-of-the-art attacks and defenses, including adversarial training and data poisoning detection, as well as multiple metrics for robustness.

AI Fairness 360 (AIF360) IBM has released this open source. AIF360 implements more than ten bias mitigation algorithms and seventy state-of-the-art metrics related to fairness in a common software framework. The AIF360 toolbox is industry sector neutral, and thus, can be applied to a wide range of problem domains.

Apache Jena is an open source Semantic Web framework for Java. CaffeOnSpark brings deep learning to Hadoop and Spark clusters. By combining salient features from deep learning framework Caffe and big-data frameworks Apache Spark and Apache Hadoop, CaffeOnSpark enables distributed deep learning on a cluster of GPU and CPU servers.

CleverHans is a python library that provides tools used in benchmarking machine learning systems' vulnerability to the adversarial examples it produces. It has recently been absorbed into the TensorFlow framework.

CROWN is a verification framework to certify robustness of neural networks with general activation functions including but not limited to ReLU, tanh, sigmoid, arctan, etc.

Deep Learning Benchmark Suite has been developed by HPE, in conjunction with Hewlett Packard Labs. It is an open source performance benchmark suite for comparing Deep Learning frameworks, models, and compute platforms.

Explainable Artificial Intelligence (XAI) program by DARPA has the goal of developing a toolkit library consisting of machine learning and human-computer interface software modules that could be used to develop future explainable AI systems.

Flora-2 is an advanced object-oriented knowledge representation and reasoning system.

H2O.ai is open-source software for big-data analysis. H2O allows users to fit thousands of potential models as part of discovering patterns in data.

Hierarchical Data Format 5 (HDF5) from the HDF Group, is a standard representation of scientific data sets, together with metadata, and is used in particular for the interchange of training data sets used in machine learning.

Plugin Machine Intelligence (PMI) project is a plugin for the Pentaho Kettle engine that provides access to supervised machine learning algorithms from various underlying "engines".

Neural Network Exchange Format (NNEF), developed by the Khronos Group, "reduces machine learning deployment fragmentation by enabling a rich mix of neural network training tools and infer-ence engines to be used by applications across a diverse range of devices and platforms."

Open Neural Network eXchange (onnx) is an open-source, community-driven effort to allow devel-opers to more easily move between machine learning frameworks. The initiative was launched by Facebook and Microsoft and was subsequently supported by IBM, Huawei, Intel, AMD, ARM and Qualcomm.

OpenAI Gym/Universe is a reinforcement learning toolkit, a wide range of environments, and an online scoreboard for developing and comparing reinforcement learning algorithms.

Pellet is an open-source Java based OWL 2 reasoner. It can be used in conjunction with both Jena and OWL API libraries; it can also be included in other applications.

Protégé is an open-source platform that provides a suite of tools to construct domain models and knowledge-based applications with ontologies.

EXECUTIVE ORDER ON MAINTAINING AMERICAN LEADERSHIP IN ARTIFICIAL INTELLIGENCE

Emphasizing the importance of artificial intelligence (AI) to the future of the U.S. economy and national security, on February 11, 2019, the President issued an Executive Order (EO 13859)[44] directing Federal agencies to take a variety of steps designed to ensure that the nation maintains its leadership position in AI.

Among its objectives, the EO aims to "Ensure that technical standards minimize vulnerability to attacks from malicious actors and reflect Federal priorities for innovation, public trust, and public confidence in systems that use AI technologies; and develop international standards to promote and protect those priorities."

The order directs the Secretary of Commerce, through the National Institute of Standards and Technology (NIST), to issue "a plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies." That plan is to be completed within 180 days of the EO—by August 10, 2019.

The EO specifies:

(i) Consistent with OMB Circular A-119, this plan shall include:

(A) Federal priority needs for standardization of AI systems development and deployment;

(B) identification of standards development entities in which Federal agencies should seek membership with the goal of establishing or supporting United States technical leadership roles; and

(C) opportunities for and challenges to United States leadership in standardization related to AI technologies.

(ii) This plan shall be developed in consultation with the Select Committee, as needed, and in consultation with the private sector, academia, non governmental entities, and other stakeholders, as appropriate."

THE PROCESS NIST USED TO DEVELOP THIS PLAN

NIST reached out widely to solicit input for the AI standards engagement plan that is the basis of this document. That outreach and consultation included:

- Publication of a Request for Information in the Federal Register that attracted 97 comments, including recommendations regarding AI standards priorities and the appropriate Federal role for

---

[44]Maintaining American Leadership in Artificial Intelligence https://www.govinfo.gov/content/pkg/FR-2019-02-14/pdf/2019-02544.pdf.

engaging in the standards development process. See Appendix V for the text of the Request for Information.

- Contacts and discussions with members of the White House Select Committee on Artificial Intelligence and other Federal agencies involved with artificial intelligence and related topics, especially through the National Science and Technology Council (NSTC) Machine Learning/Artificial Intelligence (ML/AI) Subcommittee and the Networking and Information Technology Research and Development (NITRD) Program.

- Contacts and discussions with members of the Interagency Committee on Standards Policy[45].

- A workshop on a Federal AI standards engagement strategy that attracted more than 400 representatives (about 215 in person and 210 webcast participants) from the private and public sectors, including standards developing organizations and other non-profit organizations, companies, academia, Federal agencies, and others. See Appendix VI for the workshop agenda.

- Private sector and Federal agencies' review and comment on a draft version of this AI standards Federal engagement plan.

This plan for Federal engagement in AI standards is *one component* of the overall Federal strategy for AI called for by the AI executive order.

---

[45]https://www.nist.gov/standardsgov/what-we-do/federal-policy-standards/interagency-committee-standards-policy-icsp.

APPENDIX V: REQUEST FOR INFORMATION

DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
Docket Number: [190312229-9229-01]
Artificial Intelligence Standards
Agency: National Institute of Standards and Technology, U.S. Department of Commerce.
Action: Notice; Request for Information (RFI)

SUMMARY: The February 11, 2019, Executive Order on Maintaining American Leadership in Artificial Intelligence (AI) directs the National Institute of Standards and Technology (NIST) to create a plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies (Plan). This notice requests information to help NIST understand the current state, plans, challenges, and opportunities regarding the development and availability of AI technical standards and related tools, as well as priority areas for federal involvement in AI standards-related activities. To assist in developing the Plan, NIST will consult with Federal agencies, the private sector, academia, non governmental entities, and other stakeholders with interest in and expertise relating to AI.

DATES: Comments in response to this notice must be received on or before May 31, 2019 at 5:00 pm Eastern Time.

ADDRESSES: Written comments in response to this RFI may be submitted by mail to AI-Standards, National Institute of Standards and Technology, 100 Bureau Drive, Stop 2000, Gaithersburg, MD 20899. Online submissions in electronic form may be sent to ai_standards@nist.gov. Submissions may be in any of the following formats: HTML, ASCII, Word, RTF, or PDF. Please cite "RFI: Developing a Federal AI Standards Engagement Plan" in all correspondence. All relevant comments received by the deadline will be posted at https://www.nist.gov/topics/artificial-intelligence/ai-standards and regulations.gov without change or redaction, so commenters should not include information they do not wish to be posted (e.g., personal or confidential business information). Comments that contain profanity, vulgarity, threats, or other inappropriate language or content will not be posted or considered.

FOR FURTHER INFORMATION CONTACT: For questions about this RFI contact: Elham Tabassi, NIST, MS 8900, 100 Bureau Drive, Gaithersburg, MD 20899, telephone (301) 975-5292, e-mail elham.tabassi@nist.gov. Please direct media inquiries to NIST's Public Affairs Office at (301) 975-NIST.

SUPPLEMENTARY INFORMATION:

Genesis of the Plan for Federal Engagement in Artificial Intelligence Standards

The Executive Order (EO) on AI[46] states that "[c]ontinued American leadership in AI is of paramount

---

[46]https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/.

importance to maintaining the economic and national security of the United States and to shaping the global evolution of AI in a manner consistent with our Nation's values, policies, and priorities." Accordingly, Section 1 of the EO calls for a coordinated Federal Government strategy, the American AI Initiative, and notes that the U.S. must drive development of appropriate AI technical standards in order to enable the creation of new AI-related industries and the adoption of AI by today's industries. This can be achieved through the work and partnership of industry, academia, and government.

Section 1(d) of the EO states that the U.S. must foster public trust and confidence in AI technologies and protect civil liberties, privacy, and American values in their application in order to fully realize the potential of AI technologies for the American people.

Section 2(d) of the EO directs Federal agencies to ensure that technical standards minimize vulnerability to attacks from malicious actors and reflect Federal priorities for innovation, public trust, and public confidence, and to develop international standards to promote and protect those priorities.

Section 6(d) of the EO directs the Secretary of Commerce, acting through the Director of NIST, to issue a Plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies. It further directs NIST to lead the development of the Plan with participation from relevant agencies, as determined by the Secretary of Commerce.

Approach for Developing this Plan

NIST will develop the Plan in a manner that fulfills the objectives of the EO and is consistent with relevant provisions of the Office of Management and Budget (OMB) Circular A-119, "Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities," and NIST's mission to promote U.S. innovation and industrial competitiveness. NIST has a special interest in advancing the development and use of standards relied upon by all sectors of the economy and society, recognizing that the vast majority of standards are developed through a voluntary process led by the private sector.

NIST will be informed through an open process that will include this RFI and other opportunities, such as a public workshop, to provide input. NIST expects to develop a draft Plan on which it will seek comment from the public and Federal agencies. Information about this effort, including ways to provide input, and future steps, will be available at https://www.nist.gov/topics/artificial-intelligence/ai-standards.

Goals of this Request for Information

Timely and fit-for-purpose AI technical standards—whether developed by national or international organizations—will play a crucial role in the development and deployment of AI technologies, and will be essential in building trust and confidence about AI technologies and for achieving economies of scale.

NIST seeks to understand the:

- Current status and plans regarding the availability, use, and development of AI technical standards and tools in support of reliable, robust, and trustworthy systems that use AI technologies;
- Needs and challenges regarding the existence, availability, use, and development of AI standards and tools; and

- The current and potential future role of Federal agencies regarding the existence, availability, use, and development of AI technical standards and tools in order to meet the nation's needs.

For purposes of this Plan[47], AI technologies and systems are considered to be comprised of software and/or hardware that can learn to solve complex problems, make predictions or solve tasks that require human-like sensing (such as vision, speech, and touch), perception, cognition, planning, learning, communication, or physical action. Examples are wide-ranging and expanding rapidly. They include, but are not limited to, AI assistants, computer vision systems, automated vehicles, unmanned aerial systems, voice-mail transcriptions, advanced game-playing software, facial recognition systems, as well as application of AI in both Information Technology (IT) and Operational Technology (OT).

Responding to This Request for Information

The scope of this RFI includes AI technical standards and related tools regardless of origin or use.[48] Respondents may define "standards" as they desire, indicating clearly what they mean when using the term. AI technical standards and related tools should include those necessary or helpful to reduce barriers to the safe testing and deployment of AI and to support reliable, robust, and trustworthy systems that use AI technologies.

Respondents may define tools as broadly or as narrowly as they wish. They should indicate clearly what they mean when using specific terms (e.g., practices, datasets, guidelines). An illustrative, non-exclusive list of standards-related tools includes:

- Test tools (e.g., executable test code) for conformance testing, performance testing, stress testing, interoperability testing, and other purposes;

- Use cases;

- Reference data and datasets;

- Reference implementations; and

- Training programs

---

[47]This RFI is intended to be broadly directed to any and all technologies that might be considered AI by the US Government and other interested parties. AI systems have been defined in different ways, and this RFI is directed to any information that might fall within any of these definitions. See, for example, section 238(g) of the John S. McCain National Defense Authorization Act, 2019 (P.L. 115-232), in which AI is defined to include the following:

    (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets;

    (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action;

    (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks;

    (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task; and

    (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

[48]OMB Circular A-119 defines standards broadly to include: (1) Common and repeated use of rules, conditions, guidelines or characteristics for products or related processes and production methods, and related management systems practices; and (2) The definition of terms; classification of components; delineation of procedures; specification of dimensions, materials, performance, designs, or operations; measurement of quality and quantity in describing materials, processes, products, systems, services, or practices; test methods and sampling procedures; or descriptions of fit and measurements of size or strength.

Where this RFI uses the term "organizations," it refers to private, public, and non-profit bodies, and includes both national and international organizations. If desired, commenters may provide information about: the type, size, and location of their organization(s); and whether their organization develops AI technology and related tools; uses or potentially uses AI technology and related tools; and/or participates in the development of AI standards or related tools. Provision of such information is optional and will not affect NIST's full consideration of the comment.

Comments containing references—including specific standards and related tools—studies, research, and other empirical data that are not widely published (e.g., available on the Internet) should include paper or electronic copies of those materials, unless they are restricted due to copyright or are otherwise propri-etary. In those cases, NIST encourages respondents to provide clear descriptions and designations of those references. Do not include in comments or otherwise submit any information deemed to be propri-etary, private, or in any way confidential, as all comments relevant to this RFI topic area that are received by the deadline will be made available publicly at https://www.nist.gov/topics/artificial-intelligence/ai-standards and regulations.gov.

The following list of topics covers the major areas about which NIST seeks information. This list is not intended to limit the topics that may be addressed by respondents, who may provide information about any topic which would inform the development of the Plan. Possible topics, subdivided by area, are:

AI Technical Standards and Related Tools Development: Status and Plans

1. AI technical standards and tools that have been developed, and the developing organization, including the aspects of AI these standards and tools address, and whether they address sector-specific needs or are cross-sector in nature;

2. Reliable sources of information about the availability and use of AI technical standards and tools;

3. The needs for AI technical standards and related tools. How those needs should be determined, and challenges in identifying and developing those standards and tools;

4. AI technical standards and related tools that are being developed, and the developing organi-zation, including the aspects of AI these standards and tools address, and whether they address sector-specific needs or are cross sector in nature;

5. Any supporting roadmaps or similar documents about plans for developing AI technical stan-dards and tools;

6. Whether the need for AI technical standards and related tools is being met in a timely way by organizations; and

7. Whether sector-specific AI technical standards needs are being addressed by sector-specific organizations, or whether those who need AI standards will rely on cross-sector standards which are intended to be useful across multiple sectors.

8. Technical standards and guidance that are needed to establish and advance trustworthy aspects (e.g., accuracy, transparency, security, privacy, and robustness) of AI technologies.

Defining and Achieving U.S. AI Technical Standards Leadership

9. The urgency of the U.S. need for AI technical standards and related tools, and what U.S. effectiveness and leadership in AI technical standards development should look like;

10. Where the U.S. currently is effective and/or leads in AI technical standards development, and where it is lagging;

11. Specific opportunities for, and challenges to, U.S. effectiveness and leadership in standardization related to AI technologies; and

12. How the U.S. can achieve and maintain effectiveness and leadership in AI technical standards development.

Prioritizing Federal Government Engagement in AI Standardization

13. The unique needs of the Federal government and individual agencies for AI technical standards and related tools, and whether they are important for broader portions of the U.S. economy and society, or strictly for Federal applications;

14. The type and degree of Federal agencies' current and needed involvement in AI technical standards to address the needs of the Federal government;

15. How the Federal government should prioritize its engagement in the development of AI technical standards and tools that have broad, cross-sectoral application versus sector- or application-specific standards and tools;

16. The adequacy of the Federal government's current approach for government engagement in standards development,[49] which emphasizes private sector leadership, and, more specifically, the appropriate role and activities for the Federal government to ensure the desired and timely development of AI standards for Federal and non-governmental uses;

17. Examples of Federal involvement in the standards arena (e.g., via its role in communications, participation, and use) that could serve as models for the Plan, and why they are appropriate approaches; and

18. What actions, if any, the Federal government should take to help ensure that desired AI technical standards are useful and incorporated into practice.

Kevin A. Kimball,
Chief of Staff

---

[49]See the National Technology Transfer and Advancement Act, https://www.nist.gov/standardsgov/national-technology-transfer-and-advancement-act-1995, and OMB Circular A-119, https://www.whitehouse.gov/wp-content/uploads/2017/11/Circular-119-1.pdf.

DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
Docket Number: [190312229-9229-01]
Artificial Intelligence Standards
AGENCY: National Institute of Standards and Technology, U.S. Department of Commerce.
ACTION: Notice; extension of comment period.

SUMMARY: The National Institute of Standards and Technology (NIST) extends the period for submitting written comments on the request for information (RFI) entitled "Artificial Intelligence Standards," published on May 1, 2019. The public comment period was originally scheduled to close on May 31, 2019; the public comment period is extended to now close on June 10, 2019. NIST is taking this action to provide additional time to submit comments because multiple interested parties have expressed difficulty in submitting comments by the original deadline and have asked for an extension.

DATES: Comments must be received on or before June 10, 2019 at 5:00 pm Eastern Time.

ADDRESSES: Written comments in response to this RFI may be submitted by mail to AI-Standards, National Institute of Standards and Technology, 100 Bureau Drive, Stop 2000, Gaithersburg, MD 20899. Online submissions in electronic form may be sent to ai_standards@nist.gov. Submissions may be in any of the following formats: HTML, ASCII, Word, RTF, or PDF. Please cite "RFI: Developing a Federal AI Standards Engagement Plan" in all correspondence. All relevant comments received by the deadline will be posted at https://www.nist.gov/topics/artificial-intelligence/ai-standards and regulations.gov without change or redaction, so commenters should not include information they do not wish to be posted (e.g., personal or confidential business information). Comments that contain profanity, vulgarity, threats, or other inappropriate language or content will not be posted or considered.

FOR FURTHER INFORMATION CONTACT: For questions about this RFI contact: Elham Tabassi, NIST, MS 8900, 100 Bureau Drive, Gaithersburg, MD 20899, telephone (301) 975-5292, e-mail elham.tabassi@nist.gov. Please direct media inquiries to NIST's Public Affairs Office at (301) 975-NIST.

SUPPLEMENTARY INFORMATION: On May 1, 2019, NIST published a notice and RFI in the Federal Register (84 FR 18490), about Artificial Intelligence Standards. The notice requested public comments on or before May 31, 2019. Multiple interested parties have expressed difficulty in submitting comments by the original deadline, and have asked for an extension. In light of these requests, NIST extends the period for submitting public comments to June 10, 2019. Previously submitted comments do not need to be resubmitted.

Kevin A. Kimball,
Chief of Staff

APPENDIX VI: WORKSHOP AGENDA

FEDERAL ENGAGEMENT IN ARTIFICIAL INTELLIGENCE STANDARDS WORKSHOP

National Institute of Standards and Technology

100 Bureau Drive, Gaithersburg, MD 20899

May 30, 2019 – Final Agenda	
9:00 AM	**Welcome and Overview of Logistics** – Elham Tabassi, Acting Chief of Staff, NIST Information Technology Laboratory (Green Auditorium)
9:20 AM	**Opening Remarks** – Walter G. Copan, NIST Director and Under Secretary of Commerce for Standards and Technology
9:30 AM	**Panel Session – What's Next in Standards Setting for AI** Panel presented by the Center for Data Innovation, https://www.datainnovation.org/ This panel will explore the many facets of AI standards and federal engagement in standards development, specifically. Introduction of panelists by Chuck Romine, Director, NIST Information Technology Laboratory. Panel Discussion: • Jason Matusow, General Manager, Corporate Standards Group, Microsoft • Joshua New, Senior Policy Analyst, Center for Data Innovation (Moderator) • Lynne Parker, Assistant Director for Artificial Intelligence, White House Office of Science and Technology Policy • Anthony Robbins, Vice President, North America Public Sector, Nvidia
10:45 AM	*Transition to Working Session #1 (multiple breakouts, locations varied)*
11:00 AM	**Working Session #1 – What's Out there Already? What's in the Works?** This breakout session will review the current status and plans for, and identify needs and challenges regarding, the availability, use, and development of AI technical standards and tools to support reliable, robust, and trustworthy systems that use AI technologies. Among topics to be addressed: sector specific vs. cross-sector standards, available inventories of AI standards, and AI standards roadmaps.
11:55 AM	*Return to Green Auditorium for Plenary*
12:00 PM	**Plenary Session—Insights from Working Session #1** All-attendee readout/discussion
12:45 PM	*Lunch – NIST Cafeteria (on your own)*

2:00 PM	**Panel Session – What AI Standards are Needed by Federal Government Agencies?** This session will offer examples of the widely varying needs of federal agencies for AI standards to carry out their missions including, but not limited to, R&D, national security, economic development, and oversight. Panel Discussion: • Dan Chenok, Executive Director, Center for The Business of Government, IBM Global Business Services (Moderator) • Rob High, IBM Fellow, Vice President and Chief Technology Officer, IBM Cloud and Cognitive Software • Timothy A. Klein, Director Technology Policy and Outreach, Office of the Assistant Secretary for Research and Technology, U.S. Department of Transportation • Bakul Patel, Director of Division of Digital Health, Food and Drug Administration, U.S. Department of Health and Human Services • Jon White, Deputy National Coordinator for Health Information Technology, U.S. Department of Health and Human Services
3:00 PM	*Transition to Working Session #2 (multiple breakouts, locations varied)*
3:05 PM	**Working Session #2 – How Should the Federal Government Engage?** This breakout session will review the current and potential future engagement of Federal agencies in the development and use of AI technical standards and tools in order to meet the nation's needs. Participants will discuss the unique needs of the federal government and individual agencies for AI technical standards and related tools, and how federal agencies' priorities can be aligned with national needs and priorities.
4:00 PM	*Break*
4:15 PM	**Plenary – Insights from Working Session #2,** Green Auditorium All-attendee readout/discussion
5:00 PM	*Wrap up + Adjourn*