



NIST Trustworthy and Responsible AI
NIST AI 100-5

A Plan for Global Engagement on
AI Standards

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.AI.100-5>

NIST Trustworthy and Responsible AI
NIST AI 100-5

**A Plan for Global Engagement on
AI Standards**

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.AI.100-5>

July 2024



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

About AI at NIST: The National Institute of Standards and Technology (NIST) develops measurements, technology, tools, and standards to advance reliable, safe, transparent, explainable, privacy-enhanced, and fair artificial intelligence (AI) so that its full commercial and societal benefits can be realized without harm to people or the planet. NIST, which has conducted both fundamental and applied work on AI for more than a decade, is also helping to fulfill the 2023 Executive Order on Safe, Secure, and Trustworthy AI. NIST established the U.S. AI Safety Institute and the companion AI Safety Institute Consortium to continue the efforts set in motion by the E.O. to build the science necessary for safe, secure, and trustworthy development and use of AI.

About this document: In accordance with Section 11(b) of Executive Order 14110 on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, this plan has been developed by the Department of Commerce in coordination with the Department of State and agencies across the U.S. Government. In December 2023, NIST released a Request for Information on selected tasks related to EO 14110. More than 65 comments addressing AI standards were received. Multistakeholder listening sessions covering multiple sectors were held with representatives of federal and non-U.S. governments, businesses, academia, and civil society, which provided further input and comments. These inputs were reviewed and combined with insights from across NIST, other agencies in the Department of Commerce, the Department of State, United States Agency for International Development, and other departments and agencies into an initial public draft, released in April 2024 for public comment. This document has been revised from the initial public draft based on the 57 comments received and based on a discussion about the draft with the AI Safety Institute Consortium.

NIST Technical Series Policies

[Copyright, Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

Publication History

Approved by the NIST Editorial Review Board on 07-25-2024

Contact Information

ai-inquiries@nist.gov

National Institute of Standards and Technology
Attn: NIST AI Innovation Lab, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8900) Gaithersburg, MD 20899-8900

Additional Information

Additional information about this publication and other NIST AI publications are available at <https://airc.nist.gov/Home>.

Disclaimer: Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Table of Contents

1. Executive Summary	1
2. Introduction	3
3. Desired Outcomes from Engagement on AI Standards.....	5
3.1. Scientifically sound AI standards that are accessible and amenable to adoption.....	5
3.2. AI standards that reflect the needs and inputs of diverse global stakeholders	6
3.3. AI standards that are developed in a process that is open, transparent, and consensus-driven.	8
3.4. International relationships that are strengthened by engagement on AI standards	9
4. Priority Topics for Standardization Work.....	9
4.1. Urgently needed and ready for standardization	10
4.2. Needed, but requiring more scientific work or maturity before standardization	12
4.3. Needed, but requiring significant foundational work.....	13
5. Recommended Global Engagement Activities	14
5.1. Prioritize engagement in standards work, including research and related technical activities..	14
5.2. Facilitate diverse multistakeholder engagement in AI standards development and adoption ..	16
5.3. Promote global alignment on AI standards approaches.....	18
Appendix A. Standards in Relation to AI.....	20
A.1. What are standards and why are they important?	20
A.2. How are standards developed?.....	20
A.3. How do AI standards differ from other technical standards?.....	22
Appendix B. The Current Landscape of AI Standardization	23
B.1. Horizontal standards: SDOs and topics.....	23
B.2. Participation in AI standards development	34

1. Executive Summary

Recognizing the importance of technical standards in shaping development and use of Artificial Intelligence (AI), the President's October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (EO 14110) calls for *"a coordinated effort...to drive the development and implementation of AI-related consensus standards, cooperation and coordination, and information sharing"* internationally. Specifically, the EO tasks the Secretary of Commerce to *"establish a plan for global engagement on promoting and developing AI standards... guided by principles set out in the NIST AI Risk Management Framework and United States Government National Standards Strategy for Critical and Emerging Technology"* (NSSCET). This plan, prepared with broad public and private sector input, fulfills the EO's mandate. As called for by the EO, within 180 days from publication of this document, NIST will report to the President on priority U.S. government actions undertaken pursuant to the plan.

The scope of the plan is deliberately broad, on several dimensions:

- U.S. standards stakeholders engage with **many kinds of interlocutors** around the globe on standards-related work, including standards development organizations (SDOs), industry, academia, civil society, and foreign governments.
- In the United States, the Federal government is just one among many participants in a dynamic, private sector-led standards ecosystem. Recognizing that U.S. global leadership on AI standards hinges on engagement by stakeholders from across this ecosystem, the plan lays out objectives, topical priorities, and actions that can be taken up not just by the Federal government but by the **full array of U.S. stakeholders in AI standards**.
- The plan addresses the **full lifecycle of standards-related activities**. This includes the foundational technical work that is typically necessary before a formal standard can be developed, the collaborative process of developing a consensus standard, and the development of complementary tools to help with implementing the standard.
- The plan addresses **AI-related standards of all scopes**, both "horizontal" (applicable across sectors) and "vertical" (designed for the needs of a particular sector).

The engagement activities identified in this plan are aimed at furthering four sets of outcomes:

- **Scientifically sound AI standards that are accessible and amenable to adoption.** Standards are most likely to meet these criteria when they are grounded in underpinning science and technical experience, clear, implementable, understood as neutral and unlikely to impede innovation, and globally accessible in a timely fashion. "Horizontal" (non-sector-specific) AI standards can further promote adoption by serving sectoral needs and minimizing the need for per-sector adaptation.
- **AI standards that reflect the needs and inputs of diverse global stakeholders.** Standards are most likely to achieve this if they are context-sensitive, performance-based, human-centered, and responsive to societal considerations. Though views on societal considerations vary, existing international instruments can provide a starting point for finding consensus, and standards participants can reflect their existing commitments in their standards work. AI standards are

most likely to reflect stakeholders' needs when they are based on inputs from participants from many backgrounds and geographic areas who are empowered to contribute meaningfully to standards development. Human-centered design approaches for standards may be helpful.

- **AI standards that are developed in a process that is open, transparent, and consensus-driven.** The United States continues to support standards efforts that are voluntary and market-driven—with the Federal government participating as one among many stakeholders—so that standards can best achieve consensus and serve stakeholder needs. It is also well-established that standards are most technically sound, independent, and responsive to needs when developed through open, transparent, consensus-driven processes.
- **International relationships that are strengthened by engagement on AI standards.** These engagement activities can strengthen relationships between participating experts and contribute to broader cross-border connections between companies, governments, and other stakeholders.

Engagement on AI standards will often need to focus on specific AI-related topics that standards could address. The plan defines three tiers of topics based on a qualitative assessment of need, impact of global involvement on the standards, urgency, and maturity of technical foundations. The tiers are:

1. **Urgently needed and ready for standardization.** These topics include terminology and taxonomy; testing, evaluation, verification, and validation (TEVV) methods and metrics, including cross-cutting TEVV methods and TEVV for bias; mechanisms for enhancing awareness and transparency about the origins of digital content, including for synthetic content; risk-based management of AI systems; security and privacy; transparency among AI actors about system and data characteristics; incident response and recovery plans; and training-data practices. Some of these topics have sub-topics that need additional study, perhaps on an accelerated timeline, before being ready for standards.
2. **Needed, but requiring more scientific work or maturity before standardization.** These topics include energy consumption of AI models; conformity assessment with other standards; and testing and evaluation datasets.
3. **Needed, but requiring significant foundational work.** These topics include techniques for interpretability and explainability and configuring human-AI interactions for effective decision-making and operations.

The plan recommends a variety of engagement activities relevant to all stakeholders, along with some specific high-priority ways for the U.S. government to implement the broader recommendations. Recommended activities include:

1. **Prioritize engagement in standards work, including pre-standardization research and related technical activities.** This includes bolstering foundational research on priority topics; facilitating timely development of standards by participating in standards development; encouraging alignment between sectoral practices and standards via thoughtful design and use of horizontal standards; developing and sharing implementation tools; and exploring processes for tighter feedback loops with potential standards adopters. Priority U.S. government implementation actions focus on strategic guiding of research efforts, agency participation in standards work,

information exchange within and beyond the government, and international collaborations with a standards component.

2. Facilitate diverse multistakeholder participation in AI standards development and adoption.

Domestic capacity-building engagements could include regularly convening AI standards stakeholders; developing and disseminating information such as standards training and handbooks for AI stakeholders; and devoting organizational resources to standards participation. *Global capacity-building engagements* could include broadening global access to frameworks and standards; increasing resources to support diverse participation in AI standards development; bringing education about AI standards to the settings where global AI experts gather; and building a global scientific network of AI standards experts. Priority U.S. government implementation actions focus on internal capacity-building, resourcing, and education, making relevant U.S. government documents broadly accessible, foreign assistance programming, and prioritizing work with countries at varied stages of development.

3. Promote global alignment on AI standards approaches, seeking common understandings of the role standards can and should play in the broader AI ecosystem. Engagements could include advocating for a standards ecosystem driven by multistakeholder involvement and global consensus; arranging exchanges among experts from different countries on global needs, priorities for, and experiences with AI standards; continuing to seek alignment between varied frameworks, but focusing on standardization where possible; and convening discussions on the relationship between standards and open-source software. Priority U.S. government actions focus on working AI standards issues into diplomatic meetings, communications, and outputs, including via interagency collaboration.

2. Introduction

As a leader in Artificial Intelligence (AI), the United States recognizes the importance of advancing technical standards for safe, secure, and trustworthy AI development and use. Toward that goal, this document establishes a plan for global engagement on promoting and developing AI standards. The plan calls for outreach to and engagement with international stakeholders and standards developing organizations to help drive the development and implementation of AI-related consensus standards, cooperation and coordination, and information sharing.

This plan furthers the policies and principles in the Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (EO 14110), which instructs the Federal government to “promote responsible AI safety and security principles and actions with other nations, including our competitors, while leading key global conversations and collaborations to ensure that AI benefits the whole world, rather than exacerbating inequities, threatening human rights, and causing other harms.” By advancing AI standards globally with these goals in mind, the U.S. government seeks to assist both the private and public sectors to seize the benefits of AI while managing risks to people domestically and across the globe.

Standards play a crucial role in the development and adoption of new and emerging technologies. They are especially important in the field of AI, where policymakers and regulators in the United States and

abroad are looking to the standards ecosystem to guide AI actors¹ on how to implement high-level principles and policies. This plan, developed in accordance with Section 11(b) of the EO, highlights how engagement by U.S. stakeholders, including the U.S. government, on technical standards for AI technologies can enhance global cooperation, coordination, and alignment.

For the purpose of this plan, “technical standards” refer to “documentary” standards.² ISO/IEC³ Guide 2:2004 Standardization and related activities—General vocabulary⁴ 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.⁵ Standards are typically adopted and implemented on a voluntary basis, although they can support implementation of specifications outlined in policies and regulations.

The plan is guided by principles set out in the National Institute of Standards and Technology (NIST) AI Risk Management Framework⁶ (AI RMF) and U.S. Government National Standards Strategy for Critical and Emerging Technology⁷ (NSSCET). The NIST AI RMF, released in January 2023, is a framework to better manage risks to individuals, organizations, and society associated with AI. It is intended for voluntary use to improve the ability of organizations to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems. The framework was developed through a consensus-driven, open, transparent, and collaborative process with the private and public sectors.

The NSSCET recognizes the importance of standards to enable technology that is safe, universal, and interoperable. That strategy renews the United States’ rules-based approach to standards development. It also emphasizes the Federal government’s support for international standards for critical and emerging technologies, which will help accelerate standards efforts led by the private sector to facilitate global markets, contribute to interoperability, and promote U.S. competitiveness and innovation. AI is one of those technologies.

This plan also expands on the priorities outlined in the Plan for Federal Engagement in AI Standards and Related Tools.⁸

¹ For the purposes of this document, an AI actor is any person or organization performing any of the AI actor tasks defined in Appendix A of the AI Risk Management Framework:

https://airc.nist.gov/AI_RM_F_Knowledge_Base/AI_RM_F/Appendices/Appendix_A

² This is in contrast to other types of standards such as measurement standards (which are often physical), standard reference materials, and standard reference data.

³ ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission)

⁴ <https://www.iso.org/standard/39976.html>

⁵ Open-source software (OSS) has a nuanced relationship with standards. The concepts are not the same, but there can be commonalities and overlaps, and they are often subject to the same policies. See also Section 5.3

⁶ <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

⁷ <https://www.nist.gov/standardsgov/usg-nss>

⁸ https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

This plan addresses activities before, during, and after the creation of a formal standard. Before a standard can be developed, a foundational body of scientific and technical work is typically needed. That includes producing guidelines that might form the basis for a standard and building consensus within SDOs around other informative documents such as technical reports. The standards development process draws from this foundational material to establish consensus on the rules, guidelines, or characteristics that make up the standard. Once a standard is finalized, complementary standards-related tools are often needed to help with implementation. These include datasets, benchmarks, reference implementations, implementation guidance, verification and validation tools, and conformity assessment procedures (some of which may themselves be standards). Activities related to all these stages are in scope for this plan.⁹

3. Desired Outcomes from Engagement on AI Standards

Standards-related engagement activities are most effective when they aim for clear, specific, achievable objectives. The actions laid out in this plan are designed to further the outcomes below with respect to AI standards.

3.1. Scientifically sound AI standards that are accessible and amenable to adoption

A central purpose of standards and related tools is to facilitate safety, interoperability, and competition, including by lowering barriers to trade. Their ability to achieve that purpose depends on being widely accepted and implemented. As in other technological domains, while some AI standards may be required by government regulations, most standards will generally have to be adopted voluntarily to be effective—which organizations will do if they find the relevant standards **implementable and useful**.

New standards typically are based on technical insights and novel discoveries from scientific research and innovation. The more grounded a standard is in the underpinning science and in stakeholders’ real-world experiences, the more implementable and useful it will be for the global AI community, and the greater its chances of international adoption. Conversely, a standard that would attempt to get ahead of the underpinning science and engineering may be built on less rigorous technical foundations; it may prove unhelpful, counterproductive, or even technically incoherent. The same holds true for related tools.

Accordingly, where a science-backed body of work exists, AI standards can be developed in a timelier fashion. Where there are gaps in foundational understanding (see Section 4.3), new research can fill those gaps so that implementable and useful standards can be developed.

To achieve widespread adoption, standards need to be clear, implementable, viewed as unlikely to unduly inhibit innovation or raise intellectual property (IP) concerns, understood to be neutral (i.e.,

⁹ In other words, this plan broadly addresses “standards-related activities” as defined by the Department of Commerce, including any “action taken for the purpose of developing, promulgating, revising, amending, reissuing, interpreting, implementing or otherwise maintaining or applying...a standard.”

without favoring specific nations or technology solutions), and accessible in a timely fashion to potential users across the globe.

One particularly important adoption-related issue is sectoral adoption or adaptation of horizontal standards (those intended to be used across many applications and industries). Horizontal standards may directly serve the needs of a given sector, but sector-specific practices, clarifications, and adjustments will also often be needed. In such cases, horizontal standards will be most amenable to adoption and implementation if they serve many or most sectoral needs, minimize necessary adaptation, and provide for interoperability across sectors.

Facilitating implementation of AI standards may require creating and maintaining additional standards-related tools such as datasets, benchmarks, reference implementations, implementation guidance, verification and validation tools, and conformity assessment procedures. As with the AI standards themselves, these additional standards, procedures and tools are most useful and amenable to adoption when developed and made available through inclusive, fair, and transparent processes.

3.2. AI standards that reflect the needs and inputs of diverse global stakeholders

AI standards will be most useful if they respond to the needs of a diversity of potential users and other stakeholders around the world. Standards are most likely to achieve this if they are:

- **Context-sensitive**, providing the necessary flexibility to enable adoption by small, medium, and large entities in their own contexts of use;
- **Performance-based**, providing flexibility by focusing on outcomes rather than prescribing specific ways of achieving those outcomes;
- **Human-centered**, accounting for people's needs and how they interact with the system; and
- **Responsive to societal considerations** that may arise from the design, development, deployment, or use of the technologies.

Views of what societal considerations should be reflected in AI standards are likely to vary across international contexts and stakeholders. However, attempts to find common ground on commonly accepted societal considerations can be anchored in bilateral, multilateral, regional, and global agreements. This includes international human rights instruments, particularly those that articulate governments' duties to protect people's rights and private actors' responsibilities to respect people's rights. Participants in standards development activities often represent organizations and governments that have expressed human rights commitments (see text box). Other commonly discussed societal considerations regarding AI include concerns around privacy and IP rights. The policy and legal issues surrounding these topics will be context-specific, but, as above, some international instruments address these issues, and many SDOs and standards development participants have expressed relevant commitments or policies. Stakeholders can help standards be more responsive to such considerations by reflecting any prior commitments or legal obligations in their standards development activities and in their discussions about technical standards in international policy fora.

Human rights commitments as they relate to technical standards

Participants in standards development activities include representatives from many governments and organizations that have expressed commitments to human rights. Governments have expressed these commitments by signing the United Nations (UN) Universal Declaration of Human Rights¹⁰ and joining human rights treaties. Many public and private actors have endorsed instruments such as the UN Guiding Principles on Business and Human Rights.¹¹ Some SDOs have indicated a desire to align their work with the broader context of international human rights framework. For example, the IEEE's Ethically Aligned Design¹² vision for autonomous and intelligent systems states that these systems should not infringe on human rights as its first principle. Similarly, ISO 26000: Guidance on Social Responsibility¹³ includes respect for human rights as a principle and emphasizes the role of human rights due diligence. Alongside many partner governments, the U.S. Government remains committed to protecting human rights in all its activities, including standards-setting for emerging technologies such as AI, as reflected in UN Human Rights Council Resolution 53/29.¹⁴ The U.S. Department of State, in close coordination with NIST and the U.S. Agency for International Development (USAID), has developed a "Risk Management Profile for AI and Human Rights." Informed by extensive multi-stakeholder consultations, the Profile demonstrates how to apply NIST's AI Risk Management Framework in concert with the international human rights framework.

AI standards are more likely to reflect stakeholders' needs if they are based on inputs from participants with diverse backgrounds and expertise. Especially given that AI standards so frequently involve *sociotechnical* phenomena—that is, interactions between technical systems and people (see Appendix A.3)—it is helpful for AI standards development to draw on insights from a broad set of multi-disciplinary stakeholders including enterprises of various sizes, governments, civil society, and academics.

Similarly, the needs of stakeholders from countries and regions around the world may not be reflected if a standard is not developed with adequate geographic representation (see text box below). Standards developers can address global needs by bringing geographically diverse stakeholders to the table and remaining sensitive to their concerns and views.

AI standards needs around the globe

In some cases, low- and middle-income countries particularly stand to benefit from AI innovations, which can accelerate progress on sustainable development through applications such as identifying better agricultural practices or strengthening health systems. At the same time, these countries can also be disproportionately vulnerable to certain risks, such as labor market disruptions or AI-enabled cybercrime. Countries around the world may also vary regarding which needs among the UN Sustainable Development Goals (SDGs)¹⁵ are most acute for them, and AI stands to have significant impact on many SDGs. Without meaningful participation by representatives from diverse countries, including low- and middle-income countries, AI standards may not fully reflect such priorities and concerns.

¹⁰ <https://www.un.org/en/about-us/universal-declaration-of-human-rights>

¹¹ https://www.ohchr.org/sites/default/files/documents/publications/guidingprinciplesbusinesshr_en.pdf

¹² https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf

¹³ <https://www.iso.org/iso-26000-social-responsibility.html>

¹⁴ <https://undocs.org/A/HRC/RES/53/29>

¹⁵ <https://sdgs.un.org/goals>

Stakeholders from all backgrounds and regions will be better equipped to influence standards if they have the necessary knowledge about both AI technologies and standardization processes. They may also need to be prepared to communicate and seek mutual understanding of conceptual frameworks, areas of expertise, and field-specific expectations.

One way to maximize AI standards' value could be to develop such standards following a human-centered design approach, where stakeholder needs are analyzed at the outset of a project and then guide the work. This approach can be particularly useful for AI standards development as AI requires an understanding of risks, impacts, and potential harms with multiple AI actors working together to manage those risks to achieve trustworthy AI. Such an approach could also provide a basis for assessing how successfully standards are meeting various stakeholders' needs once completed.

3.3. AI standards that are developed in a process that is open, transparent, and consensus-driven

In the United States, documentary technical standards are overwhelmingly developed through open, consensus-driven, private sector-led processes within domestic and international SDOs. As articulated in the NSSCET, the United States supports standards efforts that are voluntary and market-driven. The Federal Government engages primarily through foundational research, coordination, education, and participation in standards development processes as one of many stakeholders. Retaining this model for AI standards, with standards development **led by the private sector and featuring diverse representation from industry, civil society, government, and academia**, will help ensure that the standards meet the needs of those who will need to apply them or will be affected by them and that the standards reflect broad consensus.

It is well-established that standards development is best done through an **open, transparent, consensus-driven process**.¹⁶ This helps ensure that the resulting standards are technically sound, independent, and responsive to broadly shared market and societal needs—all characteristics that are as important for AI standards as for other areas.

Governments that desire to promote or require standards can best facilitate both technical interoperability and regulatory alignment by using consensus-driven standards. Where **international standards**¹⁷ are available, using them to the maximum extent possible reduces market friction and incompatibility and promotes efficiencies for buyers and sellers alike. Use of international standards as a

¹⁶ As noted in the NSSCET, the six principles that traditionally govern the international standards development process are transparency, openness, impartiality and consensus, effectiveness and relevance, coherence, and a commitment to participation by low- and middle-income countries.

¹⁷ The World Trade Organization (WTO) defines "international standard" as "a standard adopted by a governmental or non-governmental body whose membership is open to all [WTO] Members, one of whose recognized activities is in the field of standardization." (https://www.wto.org/english/docs_e/legal_e/21-psi_e.htm#fnt-2)

means to facilitate trade is encouraged in the World Trade Organization Technical Barriers to Trade Agreement.^{18,19}

3.4. International relationships that are strengthened by engagement on AI standards

Global engagement activities, such as active participation in standards bodies, fora, and bilateral expert exchanges, can strengthen relationships between the experts who will need to come to consensus through the standards development process. These relationships can facilitate information flow among SDO participants even outside of formal engagements and make it easier to identify common views and approaches.

In addition, engagement activities contribute to broader cross-border connections between companies, governments, and other stakeholders. For example, as creators of different frameworks of guidelines compare them with each other, they may build relationships that form the foundation of future business collaborations or diplomatic exchanges.

4. Priority Topics for Standardization Work

This plan defines three tiers of topics for engagement in international AI standardization work. The tiers are based on the degree to which:

- Experts and stakeholders have identified a well-defined, widely recognized need for international AI consensus standards;
- Global involvement can substantially enhance the speed, quality, relevance, or adoption of the resulting standards;
- Delivering standards in a timely fashion would significantly enhance the impact of those standards, including by providing a prerequisite foundation for further practices and standards; and
- Foundational scientific work exists or can be enhanced to develop technically robust standards that meet identified needs.

The tiers have been formulated by applying these criteria qualitatively. As the standardization community's needs and collective knowledge change over time, urgency and readiness may need to be reassessed. The community would benefit from further thinking—already underway in some fora, including NIST—on how to rigorously evaluate a topic's readiness for standardization, and such thinking can inform ongoing prioritization.

¹⁸ https://www.wto.org/english/docs_e/legal_e/17-tbt_e.htm

¹⁹ When the U.S. government identifies standards are needed but international standards are not available, the U.S. government seeks to work with the private sector to develop an approach for addressing the gap.

Some topics have multiple subtopics that are at different levels of standardization readiness. For ease of elaboration, each topic is addressed only in one tier but with discussion of which subtopics are more or less mature.

4.1. Urgently needed and ready for standardization

Top-priority topics are those where global stakeholders have identified a pressing need for a standard, including based on likely marketplace impact; accelerating the work would offer significant payoff; and there exists a reasonable scientific underpinning. These topics are urgent in the sense that certain foundational standards can either immediately increase the trustworthiness of AI systems or be the basis for developing further practices and standards that facilitate the responsible adoption of AI and sector-specific use cases. The payoff may come from producing a consensus standard based on existing foundational scientific work, if that is already feasible, or from bringing the community closer to agreeing on a highly-impactful future standard that would help to advance innovation, trustworthiness, and market acceptance.

For the topics listed below, the available scientific work provides a strong basis for developing standards on at least one sub-topic. In some cases, there are additional sub-topics where further research needs to be conducted before the issue can be fully addressed by standards. Although these less mature sub-topics would thus fit well as priority areas for accelerated study in the next section, for ease of elaboration they have nonetheless been included in this section under the broader topical headings. In many cases, ongoing research and updates will be required even for more mature sub-topics as technologies progress.

Topics meeting these criteria include:

- **Terminology and taxonomy.** Existing standards on AI concepts and terminology (e.g., ISO/IEC 22989:2022²⁰) provide a critical starting point, but further clarity and alignment on terminology are needed. This is particularly true for terms related to developments in AI since the existing standards were developed, including (but not limited to) concepts related to generative AI models, model fine-tuning, AI red-teaming, and model openness. Such terms and concepts underlie many other standards, policy discussions, and regulations, so technical consensus on the terminology would quickly yield wide-ranging benefits. Multiple projects outside of SDOs (e.g., academic papers,²¹ the U.S.-European Union (EU) Trade and Technology Council,²² and the work of the Organisation for Economic Co-operation and Development [OECD]) provide extensive thinking to draw upon for standardizing such terms, and updates to some existing terminology standards are already underway (e.g., ISO/IEC 22989 AMD1).
- **Measurement methods and metrics.** Shared testing, evaluation, verification, and validation (TEVV) practices for AI models and systems would open the way for more rigorous discussions about capabilities, limitations, risks, benefits, appropriate or inappropriate use, AI assurance,

²⁰ <https://www.iso.org/standard/74296.html>

²¹ <https://crfm.stanford.edu/assets/report.pdf>

²² <https://digital-strategy.ec.europa.eu/en/library/eu-us-terminology-and-taxonomy-artificial-intelligence>

and more. In particular, TEVV standards would define the performance metrics for performance-based standards, which in turn allow defining what constitutes an effective risk mitigation. Completed and ongoing foundational research in this space (e.g., NIST AI 200-2²³, ISO/IEC AWI TS 17847) provides a rich platform for standardization work, particularly on cross-cutting TEVV methods (e.g., demonstrating construct validity) and appropriate TEVV practices for bias and fairness risks, where much foundational work (e.g., NIST SP 1270²⁴) has taken place. Further work may be needed before it is possible to standardize some testing and evaluation protocols (e.g., AI red-teaming and evaluating impacts on people in realistic settings), some types of measurements (e.g., measuring effectiveness and robustness of interventions to mitigate risks), and what TEVV approaches are most effective for some types of risks (e.g., risks to safety or information integrity), all of which would merit accelerated study.

- **Mechanisms for enhancing awareness and transparency about the origins of digital content**, particularly of whether content is AI-generated or AI-modified, as well as broader information on content's origins, history, and context. An example of a type of mechanism that may be mature enough for standardization is metadata recording (a technique for provenance data tracking), for which standardization work is already underway. Other sub-topics such as watermarking and synthetic content detection merit accelerated study, including assessments of efficacy, across modalities to help address widespread and pressing concerns about the societal impacts of synthetic content.
- **Risk-based management of AI systems.** Existing frameworks (NIST AI RMF) and standards (ISO/IEC 23894:2023, ISO/IEC 42001) provide an important basis for risk-based management of AI systems. Such process-oriented frameworks are especially important in the absence of TEVV metrics that would support performance-based standards. However, more work is needed to adopt or revise those documents to account for changes in the technology as well as risks for specific applications, contexts, or industry verticals. The existing documents and stakeholders' experiences with them offer a strong basis for this work. More foundational work is needed to establish how organizations can best determine and apply their risk tolerances.
- **Security and privacy.** While many traditional cybersecurity practices apply naturally to AI systems, AI technologies also introduce a variety of new security issues. The latter category of distinct risks encompasses adversarial machine learning attacks, which include risks to the integrity of AI algorithms and data and the confidentiality of data that have been used to train an AI system (often a privacy issue). A related issue is when and how various privacy-enhancing technologies (PETs) can be used to improve privacy and security. Standards are needed for many of these topics, including mitigations against AI-specific threats, and there is a foundation of technical work to draw from (e.g., NIST AI 100-2 E2023²⁵).
- **Transparency among AI actors about system and data characteristics.** System deployers and users often need information from designers and developers about training data, performance

²³ <https://doi.org/10.6028/NIST.AI.200-2>, to be published.

²⁴ <https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence>

²⁵ <https://csrc.nist.gov/pubs/ai/100/2/e2023/final>

testing results, areas of intended or unintended use, AI systems' supply chains (the underlying software, data, and model components), and the like. These needs and mechanisms for filling them such as model cards, data cards, software bills of materials, and other disclosure mechanisms, have been well-studied, laying the groundwork for standardization. It may also be helpful for users and deployers to have standardized ways to share information upstream to designers and developers about usage patterns and issues that have been observed in deployment (including security vulnerabilities that have been discovered). This sub-topic would require further research before standardization.

- **Training data practices.** Challenges in managing training data for AI systems include data quality maintenance (especially in very large datasets), preprocessing technique selection, dataset change management, efficient use of limited data, managing diverse data formats, and identification of data intended to be permitted for or excluded from training use. Similar issues have been studied and standardized extensively in standards bodies (e.g., the ISO/IEC SC 42 working group on data for AI systems), and work is already underway on fleshing out frameworks further and adapting them to recent types of datasets and systems.
- **Incident response and recovery plans.** Some organizations already implement such plans in their own ways, and other fields may offer informative insights, particularly cybersecurity but also non-computational fields such as human rights and healthcare. Organizations such as the OECD²⁶ have developed insights on key AI-specific questions such as what constitutes an AI incident and how such incidents would be identified and shared, offering some basis for standardization. One subtopic that may need more work is how different kinds of incidents can be addressed via plans, policies, and procedures, which may include proactive baseline mitigations as well as responsive controls after a risk has been demonstrated. Further work will also be needed to align horizontal AI practices with existing sectoral standards and policies.

4.2. Needed, but requiring more scientific work or maturity before standardization

This grouping encompasses topics where there is a clear need for standardization, but more work is needed on most or all sub-topics before a standard can be developed.

For some topics, the path to standardization is longer due to a lack of foundational understanding about metrics, methods, or other critical components of a potential standard. These topics include:

- **Measuring resource consumption of AI models.** As some kinds of AI models have become both more compute-intensive and more widely used, concerns about cost and environmental impacts have grown in tandem. Though research has explored measurement methods and metrics for measuring energy usage, standardized approaches remain an important technical gap, and more foundational work seems necessary before standardization work can begin in earnest.

²⁶ https://www.oecd-ilibrary.org/science-and-technology/defining-ai-incidents-and-related-terms_d1a8d965-en.

In other cases, there is a need for tools for implementing standards, but these tools would be difficult to develop before the base standards exist. Topics with payoffs that are more distant for this reason include:

- **Conformity Assessment.**²⁷ Conformity assessment and compliance procedures can provide confidence that the specifications in a given standard have been met, but they depend on having first defined the standardized practices with which to assess conformity. As more standards mature to a point where this is possible, conformity assessment may rise in priority.²⁸ Developers of the underlying standards can also address conformity assessment needs by designing the standards to be amenable to conformity assessment.
- **Testing and evaluation datasets.** To implement testing and evaluation protocols, it is often necessary to have agreed-upon datasets before applying those protocols. Those datasets may also need to be subject to standard practices for data integrity, data quality assessment, change management, and data formats. Moreover, settling on standard datasets would depend on having reached consensus on what and how to test and evaluate (see TEVV, above).

4.3. Needed, but requiring significant foundational work

This priority consists of topics where standards would be helpful, but significant foundational work (e.g., foundational research and development) remains to be done. Examples include:

- **Techniques for interpretability and explainability.** There is ongoing research on how to better help users, affected individuals, and other stakeholders make sense of AI system outputs (e.g., NISTIR 8367,²⁹ Gunning et al. 2021³⁰). Existing research has proposed many techniques for explainability—providing information about how an AI system makes its decisions and why it behaved as it did. However, establishing empirically to what extent such techniques are useful for what purposes remains a significant gap. Techniques for interpretability, or enabling humans to understand how to act on system output, are also needed. Discussion around interpretability and explainability standards should consider the extent to which testing and transparency may yield benefits similar to those achieved with these techniques.
- **Human-AI configuration.** Interactions between humans and AI systems are generally designed with the goal of producing effective decision-making and operations. Configuring AI systems, people, organizations, and their relationships to achieve this will rely on a number of measurement methods and metrics, including for performance, bias, and trust. Metrics and potential standards in this area will be important for training, testing, and evaluation of human-AI teaming before wide-scale integration into critical operations.

²⁷ <https://www.iso.org/conformity-assessment.html>

²⁸ One example is ISO/IEC 42001 – Management System, for which conformity assessment standards are already being developed.

²⁹ <https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8367.pdf>

³⁰ <https://onlinelibrary.wiley.com/toc/26895595/2021/2/4>

- **Measurement of AI hardware performance and resource use.** Many kinds of specialized hardware, such as neural processing units (NPUs) and tensor processing units (TPUs), have been developed to accelerate AI-related computations. Companies and researchers are also experimenting with novel architectures that may not even share the same technical building blocks as traditional chips, and some kinds of hardware allow trading off characteristics such as precision and speed. Such diversity may make it difficult to compare performance and resource use (e.g., energy efficiency) across hardware platforms, which in turn would make it difficult to determine when a given form of hardware is an improvement and how it relates to existing procedures and policies. Foundational work may be needed to identify procedures, metrics, and benchmarks that would enable appropriate comparisons.

An additional research need, beyond the development of specific standards, is assessing the effectiveness of standards. Considered within the context of explosive growth in global trade, standards impact trillions of dollars of trade—with benefits and costs well beyond their economic implications. Nevertheless, research assessing the effectiveness of standards focuses primarily on specific examples of their use. (One NIST study³¹ estimated a \$250 billion economic impact just from the development of its Advanced Encryption Standard over a 20-year period.) With the emergence and forecasted explosive growth of AI technologies, AI and standards stakeholders would benefit from a more explicit and quantitative estimate and understanding of the effectiveness of AI standards—and economic impact is only one way to assess that effectiveness.

5. Recommended Global Engagement Activities

EO 14110 directs the Department of Commerce to “establish a plan for global engagement on promoting and developing AI standards.” In this case, “engagement” includes a wide variety of ways U.S. standards stakeholders can interact with current and potential stakeholders across the globe.

In recognition that AI presents global issues that require global solutions, and that AI standards, like other standards, require investment and engagement across society, the core recommendations below are scoped more broadly than U.S. government activity. These are recommendations for U.S. stakeholders writ large, and many will depend on private-sector leadership and joint efforts from the global AI and standards communities. Specific suggestions for how the U.S. government could implement these recommendations are included in text boxes, with suggested lead departments and agencies in parentheses.

5.1. Prioritize engagement in standards work, including research and related technical activities

By advancing research that can underpin standards, working on standards projects, and developing tools that facilitate adoption, AI actors and relevant stakeholders can contribute directly to standardization

³¹ <https://www.nist.gov/news-events/news/2018/09/nists-encryption-standard-has-minimum-250-billion-economic-benefit>

and lead by example. They can take the following actions to increase their direct involvement in standardization activities on AI and maximize its effectiveness:

- **Bolster foundational (pre-standardization) research on the priority topics listed above** by increasing investment in and focusing on relevant research, emphasizing international collaboration whenever appropriate and possible. When developing practices, frameworks, and guidelines that may ultimately inform standards, seek out and incorporate perspectives from as wide a range of stakeholders as is feasible.
- **Facilitate timely development of science-backed consensus-based, voluntary standards** by participating, contributing to, influencing, or leading standards development efforts. To maximize timeliness, consider the full range of SDOs, product types, and existing mechanisms for shortening development processes. Promote international cooperation whenever appropriate and possible, and prioritize SDOs and projects whose products are likely to meet the objectives described in Section 3. Work toward both horizontal and vertical standards as appropriate.
- **Encourage alignment between sectoral practices and standards** by striving to develop international horizontal standards that are as amenable as possible to adoption or adaptation across sectors, including by maximizing their incorporation of, or reference to, global documents (e.g., terminology lexicons and taxonomies). When developing vertical standards, build on and back-reference applicable horizontal standards.
- **Develop and widely share tools to assist with implementing standards and guidelines.** Make them as accessible as possible, including to potential users (organizations and nations) who may have fewer resources available to maintain awareness of standards and develop their own implementation tools.
- **Explore processes for potential standards adopters to share insights via direct feedback loops.** Human-centered design principles may offer inspiration for maximizing the usefulness, buy-in, and likelihood of adoption for AI standards. For example, potential adopters outside of SDO committees might be interested in helping to articulate a standard's goals, determine whether the standard achieves the goals, track adoption, and share experiences with implementation.

High-priority implementation actions specific to the U.S. government

- Identify and allocate resources to priority AI work related to horizontal and vertical standards projects that align with agency missions and encourage participation by agency experts, in line with existing policies on SDO participation.³² (NIST, Department of Homeland Security [DHS], U.S. Food and Drug Administration [FDA])
- Promote international collaboration on pre-standardization research. (National Science Foundation [NSF], NIST)
- Consult with private sector and civil society organizations about AI standards-related priorities and views. This could include views on SDO projects to prioritize for participation, potential U.S. contributions and proposals, and finalized standards to promote adoption of. (NIST)
- Share priorities and views on AI standards development and adoption between agencies, including sector-specific agencies. Identify intersections between standards work and AI policy as well as ways to optimize interagency collaborations and coordination to improve effectiveness and efficiency. Utilize current interagency mechanisms, especially the AI Standards Coordination Working Group. (NIST, DHS)
- Work on standards development projects jointly with other governments around the globe (see Section 5.3). (NIST, State)
- Strategically direct research efforts toward priority topics. (NSF, Department of Energy, Department of Defense, Department of Health and Human Services, National Institutes of Health, NIST)
- Leverage opportunities to align and collaborate on standards such as Joint Committee Meetings, AI working groups, public-private partnerships. (NIST, State, International Trade Administration [ITA])

5.2. Facilitate diverse multistakeholder engagement in AI standards development and adoption

Many stakeholders domestically and abroad could benefit from more extensive engagement with standards processes. Special attention should be given to drawing in stakeholders from all regions and backgrounds, particularly those who have historically been less well represented in standards development processes. Considering the breadth and scale of potential harms and benefits related to AI, it is critical that these voices be part of the standards development process. For standards to have the desired impact, it is also essential that diverse global stakeholders are aware of the standards and able to implement them. These needs call for building greater capacity and awareness among potential contributors to standards development and potential adopters.

5.2.1. Domestic capacity-building

- **Regularly convene stakeholders on AI standards.** As AI standards-related activities, including research, increase, so too do the opportunities for expanding training and the exchange of

³² E.g., OMB Circular A-119 (<https://www.whitehouse.gov/wp-content/uploads/2017/11/Circular-119-1.pdf>) and M-12-08 (https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/memoranda/2012/m-12-08_1.pdf).

information on AI standardization and discussion of AI standards issues. When groups convene on AI standards matters, they have a platform for encouraging robust information exchange about the substance of AI standards and the process of their development among subject matter experts (SMEs) in the private sector, academia, and civil society. Some such SMEs may have knowledge of AI but less experience with standards. Pre-meeting tutorials, ancillary discussions outside formal standards development sessions, and information exchange about adoption can aid in creating a more informed, more diverse, and more capable AI standards community.

- **Develop and disseminate information, including training and handbooks on standards development, participation, and adoption, for AI stakeholders.** Building on existing material, prepare and promote materials to help those from small- and medium-sized companies, academia, and civil society to understand how, where, and when they can provide their input to be most effective—including mechanisms for contributing to and improving U.S. inputs on international standards. In addition, host workshops and engagement exercises to facilitate AI standards literacy, including in partnership with the Small Business Administration (SBA) and other federal agencies.
- **Support standards participation with organizational resources.** Prioritize AI standardization staffing needs in organizational decisions about budgets, training programs, and staff incentives. Provide materials that articulate the value of standards participation and use them to make the case for prioritizing and incentivizing participation in standards work.

High-priority implementation actions specific to the U.S. government

- Increase agencies' capacity for standards participation, including when making resourcing decisions and setting staff work expectations and developing incentives. (NIST, DHS, FDA)
- Convene periodic meetings for knowledge sharing between government AI standards experts, including some open to private sector and civil society. (NIST)
- Educate U.S. government staff on the importance and benefits of participating in standards activities, including clarifying policies on committee participation and leadership as a U.S. government representative. (NIST)
- Promote participation in standards activities by recipients of federal R&D funding. (NSF, NIST)

5.2.2. Global capacity-building

- **Broaden global access to frameworks and standards.** Translate higher-priority AI standards-related documents into multiple languages. For standards that are not freely available, explore mechanisms for increasing access, particularly for potential users in developing countries.
- **Increase resources to support diverse participation in AI standards development.** Provide or fund training on participation for international stakeholders, particularly non-traditional standards participants such as those from small- or medium-sized entities, academia, and civil society and particularly those from low- and middle-income countries.
- **Bring education about AI standards to the settings where AI experts gather.** In particular, look for ways to raise awareness about standards work at AI conferences (e.g., via an AI standards

“roadshow”). These conferences bring together large groups of academics and industry practitioners, many of whom have little awareness of the standards ecosystem but much AI-related expertise across a variety of domains to contribute. Online fora where AI experts congregate virtually also are fruitful avenues for education and raising awareness.

- **Build a global scientific network of AI standards experts.** Collaboration on standards development could be facilitated by a scientific network of AI standards experts across the globe, ideally including strong representation from low- and middle-income countries. This network could be called upon for standards-specific work, guidance on implementing standards, knowledge about potential impacts of standards, and possibly scientific input on global AI issues as they emerge.

High-priority implementation actions specific to the U.S. government

- Translate key U.S. government documents, AI standards, and related resources into multiple languages. (NIST, State)
- Incorporate private sector participation or bilateral private sector exchanges into existing government-to-government engagements such as technology dialogues. (State, ITA)
- Leverage foreign assistance funds and other diplomatic programming, in collaboration with civil society and the private sector, to arrange training and support for SDO participation by stakeholders in partner countries. (State, USAID, U.S. Trade and Development Agency, ITA)
- Expand resources for government bodies that facilitate standards development. (Commerce, DHS, FDA)
- Prioritize countries for engagement that are in varied stages of development. (State, ITA, USAID)

5.3. Promote global alignment on AI standards approaches

The standards ecosystem provides the greatest value when parties around the world that develop, use, or are affected by standards and guidelines are aligned on those documents’ role and how they fit into the broader AI ecosystem. This includes widespread recognition of standards’ potential to minimize trade barriers by facilitating compatible practices. Stakeholders can work toward alignment on the role of standards through the following activities:

- **Advocate for a standards ecosystem driven by global stakeholder involvement and consensus.** In a coordinated fashion, push for standards-setting activity to take place in open, transparent, consensus-driven venues. Where applicable, prefer developing international standards over developing domestic or regional ones, seek to align any domestic standards with international standards, and advocate for others to do the same. Seek collaboration on standards that achieve broad consensus to avoid competition between potentially incompatible standards.
- **Arrange bilateral and multilateral exchanges among experts from different countries.** These exchanges would cover public and private sector AI standards needs and how they are using existing standards and guidelines. Interactions such as these would promote greater understanding between standards developers and users, including government representatives, about global needs, priorities, and experiences. Expert-to-expert exchanges can be leveraged to encourage contributions from low- and middle-income countries and strengthen mutual understanding of the benefits and limitations of standardization.

- **Continue seeking to maximize alignment between frameworks and their points of intersection but focus on standardization where possible.** While “crosswalks” between AI standards and frameworks,³³ including the NIST AI RMF, are helpful, international consensus standards have advantages over crosswalks. They tend to be more efficient, durable, and internationally acceptable than multiple frameworks and crosswalks. That said, international consensus standards also are typically much slower-moving. The fast pace of AI and dearth of international standards work on some key AI topics leads to multiple national and regional approaches. Where possible, global collaboration efforts would be most productive if focused on identifying shared ideas and taking them into the standardization process on a faster timescale. In the meantime, crosswalks will continue to add value.
- **Explore the relationship between standards and open-source software (OSS) in the context of AI systems.** OSS development is a significant venue for technical collaboration and sometimes consensus-building, including for AI. However, OSS development processes differ significantly from those of traditional documentary standards. OSS packages, including those that underpin or incorporate AI systems and models, can also implement standards. More work is needed to understand the relationship between standards and OSS in the context of AI systems.

High-priority implementation actions specific to the U.S. government

- Work with allies and partners to articulate shared principles for AI standards in multilateral diplomatic outputs. (State, Commerce)
- Build standards discussions, prioritizing international standards development and adoption, into bilateral engagements on AI policy and bilateral or multilateral collaborations on scientific research, including international partnerships formed with the U.S. AI Safety Institute within NIST/DOC. Also incorporate discussions with the local private sector (e.g., via online meetings). (State, NSF, ITA)
- Leverage or refresh existing diplomatic engagements on AI standards to promote deep exchanges between technical experts. (State, NIST)
- Expand on successful examples of coordination of U.S. government agencies on international standards engagement, such as the coordination between the Department of Commerce’s International Trade Administration and NIST via standards attachés and the Department of State and NIST on translating, standards training, and professional exchange programs. (State, Commerce, USAID)
- Strengthen communication about domestic progress on foundational technical work underlying and supporting AI standards via diplomatic channels. (State)

³³ <https://www.nist.gov/itl/ai-risk-management-framework/crosswalks-nist-artificial-intelligence-risk-management-framework>

Appendix A. Standards in Relation to AI

A.1. What are standards and why are they important?

In this plan, “standards” and “technical standards” both refer to documentary standards, defined by International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) Guide 2:2004 Standardization and related activities—General vocabulary 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.” Standards can be developed in many types of organizations that span a wide range of formality, structure, subject matter, and approach.

Widespread use of standards can facilitate technological advancement and adoption by providing common foundations from which to build. They can make products and services more interoperable, avoid technical barriers to trade, and facilitate an efficient marketplace. Standards can also make products and services safer and more trustworthy by establishing well-vetted consensus practices. In 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 reliably and safely.

For some technologies and domains, including AI, standards are important not just for technical interoperability but also for regulatory interoperability. Standards define shared concepts, metrics, and practices that governments can refer to and build on as they develop policies and regulations. If different jurisdictions can standardize on the same building blocks, then even if regulatory environments are not fully aligned, it is at least easier for market participants to move smoothly between markets. Global cooperation and coordination on AI standards will be critical for defining a consistent or at least interoperable set of “rules of the road.”

As noted in the introduction, standards are typically adopted and implemented on a voluntary basis. Voluntary compliance and conformity regimes can bring significant benefits. First, they can adapt more easily and quickly as technology changes or new and better practices emerge. Voluntary standards, particularly those that are performance- and outcome-based, can also be far more flexible; because they do not depend on compulsory compliance mechanisms, they can leave more freedom to adopters to account for their own contexts. This flexibility can advance innovation and help standardized practices be as widely applicable as possible.

A.2. How are standards developed?

The U.S. standards system differs significantly from the government-driven standards systems in many other countries and regions. Hundreds of standards developing organizations (SDOs)—most of which do not develop AI standards—are domiciled within the United States. These organizations provide the infrastructure for the preparation of standards documents. U.S. government personnel participate in SDO activities along with representatives from industry, academia, and other organizations and consumers. It is important to emphasize that these SDOs are primarily private-sector organizations, and

that the Federal government is simply one of many stakeholders and participants. The American National Standards Institute (ANSI) United States Standards Strategy, elaborated through a private-public partnership in 2005 and updated by ANSI every five years, outlines the contribution of private-sector led standards development to overall competition and innovation in the U.S. economy. The Strategy sets a strategic vision to support U.S. competitiveness, innovation, health and safety, and global trade, guiding how the United States develops and uses standards, and participates in the international standards development process.

In many other standards systems, the government plays a larger role in standards development-related activities. In such cases, these governments have more leverage to use standards as tools for competition, innovation policy, and geopolitical influence. While U.S. Government agencies possess certain responsibilities related to standards, such as in the use of standards in regulation, procurement, or other activities, there is a much greater reliance in the United States than in the European Union or China on obtaining input from industry groups, consumers, and other interested parties in making decisions related to the technical content of standards and on allowing the private sector to drive standards development.

By contrast, other governments have instituted top-down standards systems, which may involve governmental direction to stakeholders to develop particular standards, the provision of funding to national delegations, and hosting meetings.

The formal process of developing a standard tends to be relatively long, and the full process of standardization extends significantly further, both before and after formal development, review, and approval. Before a standard can even be proposed, there is often a need for significant foundational scientific work, such as technical research and pilot experiments, to explore what rules, guidelines, characteristics, or activities ought to be standardized. The standards development process itself builds on that foundational work, incorporating additional views and the need to establish consensus. The phase following standardization is about adoption: potential users of a standard may need significant additional tools to be able to implement it, including datasets, benchmarks, reference implementations, implementation guidance, verification and validation tools, and conformity assessment procedures.

To be useful, standards need to be timely. If standards development is attempted before foundational work has yielded a critical mass of technical understanding, the resulting standard may prove ill-founded or even counterproductive. Voluntary standards developed in this manner will likely fail to be adopted, and, if they are adopted (or mandated), they can impede innovation while providing little or no countervailing benefit. However, a standard is not useful if it arrives after the technologies have already moved on. Standards can also fail to gain market acceptance if they are produced late enough that market incumbents have built up infrastructure and market power, which can also hinder innovation. All technologies are so fast-moving that existing standardization processes may well struggle to keep up. The tradeoff between timeliness and rigor can sometimes be addressed by pursuing products such as guidelines and technical reports, which can be developed faster but are less prescriptive and typically represent less thorough consensus. These products could be the basis of future standards.

Most SDOs do not track the impact of their standards once completed, though many would like to be able to. SDOs may be able to track downloads or sales of standards documents, and national standards

bodies may arrange with the SDO to publish a standard as a national standard, in which case the SDO would be aware of the standard's national adoption. However, these are at best loose proxies for how extensively standards are being implemented and how well they are meeting users' needs.

Broadly, AI standards can address *horizontal* (cross-sector) or *vertical* (sector-specific) needs. Horizontal AI 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 AI standardization efforts and vice versa.

A.3. How do AI standards differ from other technical standards?

Unlike in some other technical fields such as communications technologies, where inter-system technical compatibility is vital, AI technologies often do not depend on standardized interfaces and protocols to work. Accordingly, standards in AI have tended to serve more of a “trailing edge” function. As AI stakeholders consider technologies that are already gaining traction, standards help them to:

- Converge on foundational concepts and terminology, essential for interoperability of technical approaches and evaluation methodologies as well as productive policy conversations;
- Set norms for governance and accountability processes (e.g., for risk management and trustworthiness), which raises the bar for developers' and deployers' practices and helps AI actors, especially lower-resourced ones, innovate with confidence; and
- Measure and evaluate their systems in comparable ways, facilitating confidence by developers, deployers, users, and affected parties in the usefulness and trustworthiness of AI systems.

Many of these areas of standardization must account for or directly address interactions between AI systems and people and institutions. In other words, AI systems and their impacts are inherently *sociotechnical*, hinging on complex interactions between AI systems and humans. The standards addressing these systems, such as for institutional governance practices or processes for measuring impact, are therefore often sociotechnical as well, addressing these interactions head-on.

Because AI standards are generally more detailed than the high-level AI policy principles discussed in multilateral settings such as the OECD or the G7, they can provide actionable guidance for developers, project managers, senior leaders, and other hands-on AI actors on how to implement high-level principles. Given the prevalence³⁴ of such frameworks of principles, AI standards take on extra societal significance beyond their usual role in facilitating trade and technological innovation.

This is not to say that standards have no role in defining interfaces and protocols for interactions with or between AI systems. It remains possible that standards in this direction (e.g., for standardized dataset formats, model access APIs, or large language model plugins reusable across models) will yet be identified by standards stakeholders as important needs.

³⁴ <https://cyber.harvard.edu/publication/2020/principled-ai>

Appendix B. The Current Landscape of AI Standardization

To paint the backdrop for this plan’s objectives and engagement actions, this section briefly overviews SDO efforts to date on AI.

B.1. Horizontal standards: SDOs and topics

Several SDOs have been particularly active in developing horizontal (i.e., sector-independent) AI standards. The subsections below elaborate on individual SDOs and their projects. This list is based on public comments, and may not be fully exhaustive.

A.1.1. ISO/IEC JTC 1

ISO/IEC JTC 1 SC 42 Artificial Intelligence is a subcommittee (SC) of the ISO/IEC Joint Technical Committee (JTC) 1. The purpose of this subcommittee is to develop technical standards and guidelines for AI and its associated technologies.³⁵ The subcommittee focuses largely on horizontal standards, with a particular emphasis on core concepts and practices.

Many of the documents produced by SC 42 focus on topics around concepts and governance. Topics include a management system standard, impact assessment, the data lifecycle, AI systems software quality, requirements for audit and certification, and risk management guidance. The committee has also produced some pre-standardization work in the form of Technical Reports (TRs), which provide general overview and discussion. TR topical areas include functional safety, ethical and societal concerns, machine learning (ML) computing devices, and a review of AI algorithms and system characteristics. These documents represent a consensus of conceptual thought and inform future standardization work, though not all have led directly to operationalizable standards.

Relatively few of SC 42’s standards projects (3 of 28 published projects and 4 of 33 in-progress projects, as of June 2024) have been measurement-focused. Measurement topics covered are classification performance, neural network robustness, data quality, evaluating natural language processing systems, and benchmarking quality characteristics. None address monitoring and measuring societal outcomes and impacts of deployed AI systems.

Other subcommittees of ISO/IEC JTC 1 have also produced a few AI-focused work items, such as SC 27 on cybersecurity and SC 7 on software engineering.

SC 42’s published projects as of June 2024 include:

Project Identifier	Project Title
ISO/IEC 38507:2022	Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations

³⁵ https://jtc1info.org/wp-content/uploads/2023/06/01_01_Overview_ISO_IEC_AI_for_ISO_IEC_AI_Workshop_0623.pdf

Project Identifier	Project Title
ISO/IEC 22989:2022	Information technology — Artificial intelligence — Artificial intelligence concepts and terminology
ISO/IEC 23053:2022	Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML)
ISO/IEC 42001:2023	Information Technology — Artificial intelligence — Management system
ISO/IEC 24668:2022	Information technology — Artificial intelligence — Process management framework for big data analytics
ISO/IEC TR 24027:2021	Information technology — Artificial intelligence (AI) — Bias in AI systems and AI aided decision making
ISO/IEC 24029-2:2023	Artificial intelligence (AI) — Assessment of the robustness of neural networks — Part 2: Methodology for the use of formal methods
ISO/IEC 23894:2023	Information technology — Artificial intelligence — Guidance on risk management
ISO/IEC TR 24368:2022	Information technology — Artificial intelligence — Overview of ethical and societal concerns
ISO/IEC TR 5469:2024	Artificial intelligence — Functional safety and AI systems
ISO/IEC 25059:2023	Software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Quality model for AI systems
ISO/IEC TR 24030:2024	Information technology - Artificial intelligence (AI) - Use cases
ISO/IEC 5338:2023	Information technology - Artificial intelligence - AI system life cycle processes
ISO/IEC 5339:2024	Information Technology - Artificial Intelligence - Guidelines for AI applications
ISO/IEC TR 24372:2021	Information technology - Artificial intelligence (AI) - Overview of computational approaches for AI systems
ISO/IEC TS 4213:2022	Information technology - Artificial intelligence - Assessment of machine learning classification performance
ISO/IEC 5392:2024	Information technology - Artificial intelligence - Reference architecture of knowledge engineering
ISO/IEC TS 8200:2024	Information technology — Artificial intelligence — Controllability of automated artificial intelligence systems

Project Identifier	Project Title
ISO/IEC TS 25058:2024	Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Guidance for quality evaluation of artificial intelligence (AI) systems
ISO/IEC 8183:2023	Information technology — Artificial intelligence — Data life cycle framework
ISO/IEC TR 17903:2024	Information technology — Artificial intelligence — Overview of machine learning computing devices
ISO/IEC 20546:2019	Information technology — Big data — Overview and vocabulary
ISO/IEC TR 20547-1:2020	Information technology — Big data reference architecture — Part 1: Framework and application process
ISO/IEC TR 20547-2:2018	Information technology — Big data reference architecture — Part 2: Use cases and derived requirements
ISO/IEC 20547-3:2020	Information technology — Big data reference architecture — Part 3: Reference architecture
ISO/IEC TR 20547-5:2018	Information technology — Big data reference architecture — Part 5: Standards roadmap
ISO/IEC TR 24028:2020	Information technology — Artificial intelligence — Overview of trustworthiness in artificial intelligence
ISO/IEC TR 24029-1:2021	Artificial Intelligence (AI) — Assessment of the robustness of neural networks — Part 1: Overview

SC 42 has the following additional projects under development as of June 2024:

Project Identifier	Project Title
ISO/IEC 5259-1	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 1: Overview, terminology, and examples
ISO/IEC FDIS 5259-2	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 2: Data quality measures
ISO/IEC 5259-3	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 3: Data quality management requirements and guidelines
ISO/IEC 5259-4	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 4: Data quality process framework
ISO/IEC DIS 5259-5	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 5: Data quality governance framework

Project Identifier	Project Title
ISO/IEC CD TR 5259-6	Artificial intelligence — Data quality for analytics and machine learning (ML) — Part 6: Visualization framework for data quality
ISO/IEC CD TS 6254	Information technology — Artificial intelligence — Objectives and approaches for explainability and interpretability of ML models and AI systems
ISO/IEC DTS 12791.2	Information technology — Artificial intelligence — Treatment of unwanted bias in classification and regression machine learning tasks
ISO/IEC DIS 12792	Information technology — Artificial intelligence — Transparency taxonomy of AI systems
ISO/IEC AWI TS 17847	Information technology — Artificial intelligence — Verification and validation analysis of AI systems
ISO/IEC AWI TR 18988	Artificial intelligence — Application of AI technologies in health informatics
ISO/IEC CD TR 20226	Information technology — Artificial intelligence — Environmental sustainability aspects of AI systems
ISO/IEC CD TR 21221	Information technology – Artificial intelligence – Beneficial AI systems
ISO/IEC AWI TS 22440-1	Artificial intelligence — Functional safety and AI systems — Part 1: Requirements
ISO/IEC AWI TS 22440-2	Artificial intelligence — Functional safety and AI systems — Part 2: Guidance
ISO/IEC AWI TS 22440-3	Artificial intelligence — Functional safety and AI systems — Part 3: Examples of application
ISO/IEC AWI TS 22443	Information technology — Artificial intelligence — Guidance on addressing societal concerns and ethical considerations
ISO/IEC 22989:2022/AWI Amd 1	Information technology — Artificial intelligence — Artificial intelligence concepts and terminology — Amendment 1
ISO/IEC 23053:2022/AWI Amd 1	Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML) — Amendment 1
ISO/IEC AWI TR 23281	Artificial intelligence — Overview of AI tasks and functionalities related to natural language processing
ISO/IEC AWI 23282	Artificial Intelligence — Evaluation methods for accurate natural language processing systems

Project Identifier	Project Title
ISO/IEC AWI 24029-3	Artificial intelligence (AI) — Assessment of the robustness of neural networks — Part 3: Methodology for the use of statistical methods
ISO/IEC AWI 24970	Artificial intelligence — AI system logging
ISO/IEC AWI 25029	Artificial intelligence — AI-enhanced nudging
ISO/IEC AWI 25059	Software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Quality model for AI systems
ISO/IEC AWI TS 29119-11	Software and systems engineering — Software testing — Part 11: Testing of AI systems
ISO/IEC DIS 42005	Information technology — Artificial intelligence — AI system impact assessment
ISO/IEC DIS 42006	Information technology — Artificial intelligence — Requirements for bodies providing audit and certification of artificial intelligence management systems
ISO/IEC AWI 42102	Information technology — Artificial intelligence — Taxonomy of AI system methods and capabilities
ISO/IEC AWI TR 42103	Information technology — Artificial intelligence — Overview of synthetic data in the context of AI systems
ISO/IEC AWI 42105	Information technology — Artificial intelligence — Guidance for human oversight of AI systems
ISO/IEC AWI TR 42106	Information technology — Artificial intelligence — Overview of differentiated benchmarking of AI system quality characteristics
ISO/IEC AWI TR 42109	Information technology — Artificial intelligence — Use cases of human-machine teaming

Beyond SC 42, JTC1 is developing the following standards from subcommittee 27 on cybersecurity:

Project Identifier	Project Title
ISO/IEC CD 27090	Cybersecurity — Artificial Intelligence — Guidance for addressing security threats and failures in artificial intelligence systems
ISO/IEC WD 27091.2	Cybersecurity and Privacy — Artificial Intelligence — Privacy protection

A.1.2. CEN-CENELEC

CEN, the European Committee for Standardization, is an association that brings together the National Standardization Bodies of 34 European countries. CEN provides a platform for the development of European Standards and other technical documents in relation to various kinds of products, materials, services, and processes.³⁶ CENELEC, the European Committee for Electrotechnical Standardization, plays a similar role for electrotechnical standardization. The openness of CEN and CENELEC processes to participation by U.S. stakeholders and other stakeholders from outside Europe is sometimes limited.

In 2020, CEN and CENELEC established a new JTC 21 “Artificial Intelligence.” CEN-CLC/JTC 21 identifies and adopts international standards already available or under development from other organizations like ISO/IEC JTC 1 and its subcommittees, such as SC 42. Furthermore, CEN-CLC/JTC 21 focuses on producing standardization deliverables that address European market and societal needs, as well as underpinning European Union (EU) legislation, policies, principles, and values.³⁷

CEN/CLC JTC 21 was formed partly in response to the European Commission white paper that initiated the creation of the EU AI Act. The committee has accepted a standardization request from the Commission to fulfill the standardization needs of the AI Act, which will drive much of its work in the coming months. The committee is expected to produce “harmonized standards” (standards developed for the purpose of being referenced by regulation). These standards will be voluntary, but, nonetheless, will have legal implications: Referenced EU harmonized standards carry a presumption of conformity, making compliance with these standards the recommended but not the only method to meet regulatory requirements. Per a 2016 ruling from the European Court of Justice, such standards form part of EU law, as they have legal effects.

As of June 2024, CEN/CLC JTC 21 has not published any standards of its own, although it has adopted some ISO/IEC standards. The standardization request from the Commission, which is expected to drive future work, includes standards for risk management systems, dataset quality and governance, record keeping, transparency, human oversight, accuracy specifications, robustness specifications, cybersecurity specifications, quality management systems, and conformity assessment, which were all requested for delivery by April 2025. Projects under development that are not shared with ISO/IEC include:

Project Identifier	Project Title
FprCEN/CLC/TR 18115	Data governance and quality for AI within the European context
prCEN/CLC/TR 17894	Artificial Intelligence Conformity Assessment
prCEN/CLC/TR XXX	Impact assessment in the context of the EU Fundamental Rights
prCEN/CLC/TR XXX	AI Risks - Check List for AI Risks Management

³⁶ <https://www.cencenelec.eu/about-cen/>

³⁷ <https://www.cencenelec.eu/areas-of-work/cen-cenelec-topics/artificial-intelligence/>

Project Identifier	Project Title
prCEN/CLC/TR XXX	Environmentally sustainable Artificial Intelligence
prCEN/CLC/TR XXXX	Artificial Intelligence - Overview of AI tasks and functionalities related to natural language processing
prEN XXX	AI trustworthiness framework
prEN XXX	AI tasks and evaluation methods of computer vision systems
prEN XXX	AI Risk Management
prEN XXXXX	AI-enhanced nudging

A.1.3. IEEE

“IEEE Standards Association (IEEE SA) is a consensus building organization that nurtures, develops, and advances global technologies, through IEEE. It brings together a broad range of individuals and organizations from a wide range of technical and geographic points of origin to facilitate standards development and standards-related collaboration.”³⁸

Starting in 2016, the IEEE P7000 series of standards projects addresses specific issues at the intersection of technological and ethical considerations for AI. The AI Standards Committee is responsible for standards that enable the governance and practice of AI as related to computational approaches to machine learning, algorithms, and related data usage.³⁹

Other topics addressed by IEEE’s AI standards include organizational governance, explainable AI, federated learning, autonomous system verification, and technical details such as data attributes and formats.

Project Identifier	Project Title
IEEE P2863	Recommended Practice for Organizational Governance of Artificial Intelligence
IEEE P2894	IEEE Draft Guide for an Architectural Framework for Explainable Artificial Intelligence
IEEE P2976	Standard for XAI - eXplainable Artificial Intelligence - for Achieving Clarity and Interoperability of AI Systems Design
IEEE P3123	Standard for Artificial Intelligence and Machine Learning (AI/ML) Terminology and Data Formats
IEEE P2817	IEEE Draft Standards Project Guide for Verification of Autonomous Systems

³⁸ <https://standards.ieee.org/about/>

³⁹ <https://sagroups.ieee.org/ai-sc/>

Project Identifier	Project Title
IEEE P2986	Recommended Practice for Privacy and Security for Federated Machine Learning
IEEE P2975	Standard for Industrial Artificial Intelligence (AI) Data Attributes

A.1.4. ITU

The International Telecommunication Union (ITU) is the United Nations specialized agency for information and communication technologies (ICTs).⁴⁰ The Study Groups of ITU's Telecommunication Standardization Sector (ITU-T) assemble experts from around the world to develop international standards known as ITU-T Recommendations that act as defining elements in the global infrastructure of ICTs.⁴¹

ITU-T's work on AI standards relates primarily to use of AI for increasing efficiency of telecommunication/ICT networks and systems. ITU also engages in activities related to information sharing and capacity building on AI for sustainable development as related to telecommunications/ICTs. ITU collaborates with other United Nations agencies and stakeholders to support realization of the benefits of AI use cases for sustainable development, such as through its AI for Good platform, which includes information exchanges and workshops related to AI standards.

Project Identifier	Project Title
F.ADT4MM	Requirements and framework of AI-based detection technologies for 5G multimedia messages
F.ACIP-GA	Technical specifications for AI cloud platform: general architecture
F.ACIP-MD	Technical specification for AI cloud platform: AI model development
F.AI-CPP	Technical specification for AI cloud platform: performance
F.AI-DMPC	Technical framework for deep neural network model partition and collaborative execution
F.AI-FASD	Framework for audio structuralizing based on deep neural network
F.AI-ILICSS	Technical requirements and evaluation methods of intelligent levels of intelligent customer service system
F.AI-ISD	Requirements for intelligent surface-defect detection service in industrial production line

⁴⁰ <https://www.itu.int/en/about/Pages/default.aspx>

⁴¹ <https://www.itu.int/en/ITU-T/about/Pages/default.aspx>

Project Identifier	Project Title
F.AI-MKGDS	Requirements for the construction of multimedia knowledge graph database structure based on artificial intelligence
F.AI-MVSLWS (ex F.AI-VDSLWS)	Requirements for artificial intelligence based machine vision service in smart logistics warehouse system
F.AI-RSRSreqs	Requirements for real-time super-resolution service based on artificial intelligence
F.AI-SF	Requirements for smart factory based on artificial intelligence
F.FDIS	Requirements and framework for feature-based distributed intelligent systems
F.FML-TS-FR	Requirement and framework of trustworthy federated machine learning based service
F.ML-ICSMIReqs	Requirements and framework for intelligent crowd sensing multimedia interaction based on deep learning
F.REAIOCR	Requirements and evaluation methods for AI-based optical character recognition service
F.SCAI	Requirements for smart class based on artificial intelligence
F.TCEF-FML	Trusted contribution evaluation framework on federated machine learning services
Y.3181 (ex Y.ML-IMT2020-SANDBOX)	Architectural framework for Machine Learning Sandbox in future networks including IMT-2020
Y.3182 (ex Y.ML-IMT2020-E2E-MGMT)	Machine learning based end-to-end multi-domain network slice management and orchestration
Y.CNAO	Requirements and functional framework for Customer-oriented Network Quality Auto Optimization with Artificial Intelligence
Y.IMT2020-DJLML	Requirements and framework for distributed joint learning to enable machine learning in future networks including IMT-2020
Y.IMT2020-AINDO-req-frame	Requirements and framework for AI-based network design optimization in future networks including IMT-2020
Y.ML-IMT2020-VNS	Framework for network slicing management enabled by machine learning including input from verticals

Project Identifier	Project Title
Y.ML- IMT2020- MLFO	Requirements and architecture for machine learning function orchestrator
Q.AIS-SRA	Signalling requirements and architecture to support AI based vertical services in future network, IMT2020 and beyond

A.1.5. Consumer Technology Association

The Consumer Technology (CTA) oversees the development of standards by experts from across the consumer technology industry, including manufacturers, service providers, regulators and other industry leaders. These 1,500 tech pioneers — from engineers to doctors and scientists — help produce specifications that define how products work and the ways consumers interact with them.⁴²

In addition to its sector-specific standards (see below), CTA has developed several horizontal AI standards as of June 2024:

Project Identifier	Project Title
CTA-2089-A	Definitions and Characteristics of Artificial Intelligence
CTA-5203	Cybersecurity Threats and Security Controls for Machine Learning Based Systems
ANSI/CTA-2096	Guidelines for Developing Trustworthy Artificial Intelligence Systems
CTA-5200	What is Artificial Intelligence?

A.1.6. Non-AI-specific horizontal standards highly relevant to AI

Several standards that are not AI-specific are nonetheless highly relevant to AI and are sometimes applied to AI systems. These include:

Project Identifier	Project Title
IEEE 7000-2021	IEEE Standard Model Process for Addressing Ethical Concerns during System Design
ISO/IEC Guide 51:2014	Safety aspects — Guidelines for their inclusion in standards

⁴² <https://shop.cta.tech/collections/standards/artificial-intelligence>

A.1.7. Sectoral standards: SDOs and topics

In industries that are coming to rely heavily on AI, sector-specific standards projects have also begun to emerge:

- For **AI in aeronautical systems**, SAE International, a global association of engineers and related technical experts in the aerospace, automotive and commercial vehicle industries,⁴³ is developing standards products on foundational concepts and certification processes. EUROCAE, a European non-profit that develops standards for European civil aviation, also has a working group on AI.
- For **AI in ground vehicle applications**, SAE also has standards in progress.
- For **AI in healthcare** (where regulatory requirements can make standards especially critical, as in the recently finalized [HTI-1 rule](#)⁴⁴), the Consumer Technology Association has published a variety of standards on topics such as foundational definitions and characteristics, bias, and data management and is working on more. ISO and IEC have also published healthcare-related standards addressing topics such as medical electrical equipment employing autonomy, risk management, and quality management systems for AI/ML in medical devices. IEEE has several standards related to AI in healthcare. The Association for the Advancement of Medical Instrumentation has published a technical report on risk management in AI for medical devices. Finally, Integrating the Healthcare Enterprise has published “profile” standards that specify the application of base standards to health care use cases.
- For **AI in finance**, X9, an ANSI-accredited developer of financial services standards, has started an AI study group aiming to identify areas where standards are or could be needed to safeguard financial, infrastructure and user data,⁴⁵ and two standing working groups assigned to work on AI issues.
- For **AI in biotechnology**, American Type Culture Collection has begun work on standards and is looking to develop authenticated reference data for use in training AI models.

Other SDOs that have indicated that they have sectoral AI work underway or under consideration include the Robotics Industries Association, CSA Group, Alliance for Telecommunications Solutions, National Information Standards Organization, National Council for Prescription Drug Programs, American Society for Nondestructive Testing, and the Instrument Society of America.

⁴³ <https://www.sae.org/>

⁴⁴ <https://www.healthit.gov/topic/laws-regulation-and-policy/health-data-technology-and-interoperability-certification-program>

⁴⁵ <https://x9.org/aistudygroup/>

B.2. Participation in AI standards development

During consultations, some parties noted that the majority of participants in AI standards bodies are from industry. Large, well-resourced technology companies were cited as the participants most aware of and active in standards development, while relatively few SMEs have been participating. Startups may be aware of standards-setting work, but they do not always have the resources to effectively participate.

Many commenters also noted that civil society and academia have historically not been well-represented in standards development work, including on AI. Some commenters attributed this to confusion about what standards are, what they can and cannot do, and when and how they are developed. It was also suggested that these entities tend not to recognize how standards development might contribute to their goals, and that they find procedures for participating opaque.

Low- and middle-income countries seem to be particularly missing from AI standards, as reported with great concern by numerous commenters.