NIST AI 100-5

Draft for Public Comment

A Plan for Global Engagement on AI Standards

This publication is available free of charge from:

[DOI link TK]

April 2024

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST AI 100-5

Draft for public comment

A Plan for Global Engagement on AI Standards

This publication is available free of charge from:

[DOI link TK]

April 2024

Comments are especially requested on:

1. Prioritization of topics for standardization work, including additional topics.
2. Prioritization of activities and actions.
3. Important ideas that are missing.
4. Completeness of standards activities listed in Appendix B.

Comments on NIST AI 100-5 may be sent electronically to NIST-AI-100-5@nist.gov with "NIST AI 100-5, A Plan for Global Engagement on AI Standards" in the subject line. Comments may also be submitted via www.regulations.gov: enter NIST-2024-0001 in the search field, click on the "Comment Now!" icon, complete the required fields, including "NIST AI 100-5, A Plan for Global Engagement on AI Standards" in the subject field, and enter or attach your comments. Comments containing information in response to this notice must be received on or before **June 2, 2024, at 11:59 PM Eastern Time.**

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.

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.

Table of Contents

1. Executive Summary

TK

2. Introduction

As the world leader in Artificial Intelligence (AI), the United States recognizes the importance of advancing global 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 a coordinated effort to work with key international allies and partners and with standards developing organizations to 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 global AI standards 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 actors 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 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. ISO/IEC[1] Guide 2:2004 Standardization and related activities—General vocabulary[2] 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.

The plan is guided by principles set out in the National Institute of Standards and Technology (NIST) AI Risk Management Framework[3] (AI RMF) and U.S. Government National Standards Strategy for Critical and Emerging Technology[4] (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

[1] ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission)
[2] https://www.iso.org/standard/39976.html
[3] https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf
[4] https://www.nist.gov/standardsgov/usg-nss

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.[5]

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 typically is needed. That includes producing guidelines that might form the basis for a standard and building consensus 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. Activities related to all of these stages are in scope for this plan.

[5] https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

3. Objectives for Engagement on AI Standards

Standards-related engagement activities are most effective when they are aimed at achieving clear, specific objectives. The actions laid out in this plan are designed to further the outcomes below.

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. They can achieve that purpose if they are widely accepted and implemented. As in other technological domains, while some AI standards will be required by government regulations, their effectiveness generally will depend on organizations to voluntarily adopt those standards – which they will do only if they find the relevant standards **implementable and useful**.

New standards typically are based on novel discoveries and technical insights from scientific research and innovation. The more grounded a standard is in the underpinning science, 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 attempts to get ahead of the underpinning science may be built on less rigorous technical foundations, may prove unhelpful, or even counterproductive or 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 more timely fashion. Where there are gaps in foundational understanding (see Section 3.4), new research can fill those gaps so that implementable and useful standards can be developed.

To achieve international adoption, a standard needs to be clear, implementable, viewed as unlikely to inhibit innovation, perceived as neutral (i.e., without favoring specific nations or organizations) 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.

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 around the world. Standards are most likely to achieve this if they are:

- **Context-sensitive**, providing 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 human needs, interactions, and values; and
- **Sensitive 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, 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), which they can reflect in their standards development activities and in their discussions about technical standards in international policy fora.

Human rights commitments with respect 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 Rights[6] and joining human rights treaties. Many public and private actors have endorsed instruments such as the UN Guiding Principles on Business and Human Rights.[7] Some SDOs also have indicated a desire to align their work with the broader context of international human rights law and norms. For example, the IEEE's Ethically Aligned Design[8] 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 Responsibility[9] 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. (See UN Human Rights Council resolution A/HRC/RES/53/29.[10])

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.

[6] https://www.un.org/en/about-us/universal-declaration-of-human-rights
[7] https://www.ohchr.org/sites/default/files/documents/publications/guidingprinciplesbusinesshr_en.pdf
[8] https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf
[9] https://www.iso.org/iso-26000-social-responsibility.html
[10] https://undocs.org/A/HRC/RES/53/29

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 on AI standards needs around the globe). 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

Low- and middle-income countries particularly stand to benefit from AI innovations through applications such as identifying better agricultural practices or strengthening health systems. These countries can also be disproportionately vulnerable to certain risks, such as employment shocks or AI-enabled cybercrime in areas where expanding network access takes priority over security. Without meaningful participation by representatives from these countries, AI standards may not fully reflect such concerns.

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 to users 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 completed standards are meeting various stakeholders' needs.

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

In the United States, documentary technical standards are overwhelmingly developed through open, consensus, private sector-led processes within domestic and international standards developing organizations (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 largely by industry but also civil society, government, and academia,** will help ensure that the standards meet the needs of those who will need to apply them and that they reflect broad consensus.

It is well-established that standards development is best done through an **open, transparent, consensus-driven process**.[11] This helps ensure that the resulting standards are technically sound, independent, and

[11] 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.

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** 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 means to facilitate trade is encouraged in the World Trade Organization Technical Barriers to Trade Agreement.[12]

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

Global engagement activities, such as active participation in standards bodies, forums, 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.

[12] https://www.wto.org/english/docs_e/legal_e/17-tbt_e.htm

4. Priority Topics for Standardization Work

This plan defines three areas for engagement in international AI standardization work, based on the degree to which:

- Experts and stakeholders have identified a need for international AI consensus standards;
- Global involvement can substantially enhance the speed, quality, relevance, or adoption of the resulting standards;
- Delivering timely standards would significantly enhance the impact of those standards, including trustworthiness and the acceptance of new technology solutions in international markets; and
- Foundational scientific work exists or can be enhanced to develop technically robust standards that meet identified needs.

Within each area, several priorities are identified for standardization and/or accelerated study.

4.1. Urgently needed and ready for standardization

Top-priority topics are those where stakeholders have identified a pressing need for a standard, 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 be the basis for other standards to be built upon in order to 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 some of the topics listed below, the available scientific basis may be sufficient for standards development; in other cases, additional research needs to be conducted.

Topics meeting these criteria include:

- **Terminology and taxonomy**. Existing standards on AI concepts and terminology (e.g., ISO/IEC 22989:2022[13]) provide a critical starting point, but further clarity and alignment on terminology is needed, particularly on terms related to recent developments in AI. For example, consensus is needed around terms and concepts related to foundation models, model fine-tuning, AI red-teaming, open models, and synthetic content. 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,[14] the U.S.-European Union (EU) Trade and Technology Council,[15] and U.S. AI Safety Institute (AISI) Consortium) provide extensive thinking to draw upon for standardizing such terms.

[13] https://www.iso.org/standard/74296.html
[14] https://crfm.stanford.edu/assets/report.pdf
[15] https://digital-strategy.ec.europa.eu/en/library/eu-us-terminology-and-taxonomy-artificial-intelligence

- **Measurement and mitigations for risks and safety issues**. For example, for bias and equity issues, some standards-related efforts on related topics (e.g., NIST SP 1270[16]) have taken place. More work is needed to adapt to rapid changes in AI technology, such as measurement methods and metrics for the effectiveness of mitigation methods for harmful biases, security, safety, and other risks.

- **Testing, evaluation, verification, and validation (TEVV)**. Shared TEVV practices for AI systems would open the way for more rigorous discussions about risks, benefits, appropriate or inappropriate use, and more. Completed and ongoing foundational research in this space offers a platform for standardization on some AI TEVV topics. However, the technical community lacks clarity and consensus around how different testing methodologies interact, what kinds of testing are best for what systems, and which AI actors should be involved. Other notable gaps include protocols for AI red-teaming and approaches to measure interactions with and impacts on humans. Some methods are mature enough for standardization, while others merit accelerated study.

- **Mechanisms for enhancing awareness and transparency about the origins of digital content**, particularly of whether content is authentic or synthetic (i.e., AI-generated or AI-modified), as well as greater context of the origins and history and context. An example of mechanisms that may be mature enough for standardization is metadata recording (a technique for provenance data tracking.) Other mechanisms, such as watermarking, and synthetic content detection merit accelerated study 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) provide an important basis for risk based management of AI systems. However, more work is needed to adopt or revise those documents to account for changes in the technology as well as risks for specifical applications, contexts, or industry verticals.

- **Security.** 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 has 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 need for taxonomy and terminology of attacks and mitigation. There is a foundation of technical work to draw from on these topics (e.g., NIST AI 100-2 E2023[17]).

- **Transparency among AI actors about system and data characteristics.** System deployers and users often need information from designers and developers about training data, performance testing results, areas of intended or intended use, and the like. These needs and mechanisms for

[16] https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence
[17] https://csrc.nist.gov/pubs/ai/100/2/e2023/final

filling them (e.g., model cards and data cards) have been well-studied, laying the groundwork for standardization.

4.2. Needed, but requiring more scientific work before standardization

This grouping encompasses venues of work where there is a clear need for standardization, but more work is needed before a standard can be developed or the payoffs from standardization may be more distant.

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:

- **Energy consumption of AI models.** As AI models have become both more compute-intensive and more widely used, concerns about 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.
- **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. Plans, policies, and procedures may include proactive baseline mitigations as well as responsive controls after a risk has been demonstrated. Significant work on areas such as terminology (e.g., what constitute incident, mechanisms to report, etc.) remains to align practices and arrive at shared approaches.

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.**[18] Conformity assessment and compliance procedures can provide confidence that the specifications in a given standards have been met, but they depend on having first defined the standardized practices with which to assess conformity.
- **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 and data quality assessment. Moreover, settling on standard datasets would depend on having reached consensus on what and how to test and evaluate (see TEVV, above).
- **Channels for upstream reporting.** In addition to forms of providing transparency from designers and developers to deployers and users, it would also be helpful for users and deployers to have standardized ways to share information back to designers and developers about usage patterns and issues that have been observed in deployment, which would require further research.

[18] https://www.iso.org/conformity-assessment.html

Before this type of reporting channel can be standardized, the content of what should be communicated and how would need to be agreed.

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,[19] Gunning et al. 2021[20]). Existing research has proposed many techniques for explainability—providing information about how an AI system makes its decisions. 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 that can lead to effective decision-making and operations rely on a number of measures 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 (global) integration into critical operations.

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 study[21] 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, the community 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.

[19] https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8367.pdf
[20] https://onlinelibrary.wiley.com/toc/26895595/2021/2/4
[21] https://www.nist.gov/news-events/news/2018/09/nists-encryption-standard-has-minimum-250-billion-economic-benefit

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 international partners.

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; 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.

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

By continuing and advancing research that can underpin standards and developing tools which facilitate adoption, AI actors and relevant stakeholders can contribute directly to standardization and lead by example. They can take the following actions to increase and maximize the effectiveness of their direct involvement in standardization activities on AI:

- **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.
- **Facilitate development of science-backed consensus-based, voluntary standards** by participating, contributing to, influencing, or leading standards development efforts, promoting international cooperation whenever appropriate and possible.
- **Encourage horizontal standards that are applicable across sectors** by maximizing their incorporation of, or reference to, global standards (including terminology, taxonomies, and crosswalks) and by striving to develop horizontal standards that are as amenable as possible to adoption or adaptation across sectors.
- **Develop and widely share tools to assist with implementing standards and guidelines,** making them as accessible as possible, including to potential users (organizations and nations) that are less well resourced.

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

- Identify and allocate resources to priority AI work related to standards projects that align with agency missions and encourage participation by agency experts.
- Consult with private sector and civil society organizations about AI standards-related priorities and views – including participation in SDO projects.
- Share priorities and views with other agencies, including sector-specific agencies, and 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.
- Work on standards development projects jointly with other governments around the globe (see Section 5.3).
- Leverage opportunities to align and collaborate on standards such as Joint Committee Meetings, AI working groups, public-private partnerships, U.S. AISI and AISI Consortium engagements, and multilateral mechanisms such as the Quadrilateral Security Dialogue, Global Partnership on AI, and NIST's bilateral efforts with the Canada, EU, Japan, Singapore, the United Kingdom (UK), and other nations.

5.2. Facilitate diverse multistakeholder engagement in AI standards development

Many potential contributors to the development of AI standards and related tools could benefit from more extensive and meaningful engagement with current participants both domestically and more broadly. Special attention should be given to drawing in stakeholders from all regions and backgrounds, particularly who have historically been less well represented in standards development processes. Considering the risks and potential harms related to AI, along with the enormous benefits, it is critical that these voices be part of the standards development process, and that both calls for and leads to building greater capacity.

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

- **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.
- Convene periodic meetings of government AI standards experts that include private sector and civil society.
- 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.
- Aligned with OMB A-119,[22] emphasize the need to integrate standards requirements into government AI acquisition processes.

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 counties.
- **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 forums 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. This network could be called upon for standards specific work, knowledge about potential impacts of standards, and possibly scientific input on global AI issues as they emerge.

[22] https://www.whitehouse.gov/wp-content/uploads/2017/11/Circular-119-1.pdf

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

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

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 what role those documents should serve and how they should fit into the broader AI ecosystem. Stakeholders can work toward that goal through the following activities:

- **Encourage a standards ecosystem driven by multistakeholder involvement and global consensus.** Push for standards-setting activity to take place in multistakeholder consensus-driven venues. Prefer international standards over domestic or regional ones, seek to align any domestic standards with international standards, and advocate for others to do the same.
- **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,[23] 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 AI leads to multiple national and regional approaches. Where possible, global collaboration efforts would be most productive if focused on identifying shared ideas and taking

[23] https://www.nist.gov/itl/ai-risk-management-framework/crosswalks-nist-artificial-intelligence-risk-management-framework

them into the standardization process on a faster timescale. In the meantime, crosswalks will continue to add value.

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.
- Build standards discussions into bilateral engagements on AI policy and bilateral or multilateral collaborations on scientific research, including international partnerships formed with the U.S. AISI. Also incorporate discussions with the local private sector (e.g., via online meetings).
- Leverage or refresh existing diplomatic engagements on AI standards to promote deep exchanges between technical experts.
- 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 translations and standards training.
- Strengthen communications about domestic progress on foundational technical work underlying and supporting AI standards via diplomatic channels.

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 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."

Standards are typically adopted and implemented on a voluntary basis, although they can support implementation of specifications outlined in policies and regulations. 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.

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. 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, outlines the contribution of private-sector led standards development to overall competition and innovation in the U.S. economy.

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 ensuing phase following standardization is about adoption: potential users of a standard may need significant additional tools to be able to adopt 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. AI technologies are so fast-moving that existing standardization processes may well struggle to keep up.

Most SDOs do not track the impact of their standards once completed. 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 Organisation for Economic Co-operation and Development 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 prevalence[24] of such frameworks of principles, AI standards take on extra societal significance beyond their usual role in facilitating trade and technological innovation.

[24] 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 state of standardization on various AI topics is shown in the table, and the subsections below elaborate further on individual SDOs and their projects.

Topic	Availability of standards
AI nomenclature and terminology	TK (based on input)
Data capture, processing, protection, privacy, confidentiality, handling, and analysis	TK (based on input)
Trustworthiness, verification, and assurance of AI systems	TK (based on input)
AI risk management	TK (based on input)

B.1.1. ISO/IEC JTC 1

ISO/IEC JTC 1 SC 42 Artificial Intelligence is a subcommittee (SC) of the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) Joint Technical Committee (JTC) 1. The purpose of this subcommittee is to develop technical standards and guidelines for AI and its associated technologies.[25] The subcommittee focuses largely on horizontal foundational standards.

Most of the mature documents[26] 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, ML computing devices, and a review of AI algorithms and system characteristics. While these documents represent a consensus of conceptual thought, few appear to have led directly to operationalizable standards.

Relatively few of SC 42's standards projects (5 of 24, as of March 2024) have been measurement-focused. Measurement topics covered are neural network robustness, data quality, classification performance, benchmarking quality characteristics, and evaluation metrics for AI use cases and

[25] https://jtc1info.org/wp-content/uploads/2023/06/01_01_Overview_ISO_IEC_AI_for_ISO_IEC_AI_Workshop_0623.pdf

[26] Documents considered here as "mature" include the ISO/IEC stages of Draft International Standard (DIS), Final Draft International Standard (FDIS), and Publication Stage. See https://www.iso.org/stages-and-resources-for-standards-development.html.

applications. 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.

Project Identifier	Project Title
ISO/IEC 38507	Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations
ISO/IEC AWI TS 29119-11	Information technology — Artificial intelligence — Testing for AI systems — Part 11:
ISO/IEC 22989	Information technology — Artificial intelligence — Artificial intelligence concepts and terminology
ISO/IEC 23053	Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML)
ISO/IEC CD 42001.2	Information Technology — Artificial intelligence — Management system
ISO/IEC CD 42001.2	Information Technology — Artificial intelligence — Management system
ISO/IEC CD 5259-1, ISO/IEC AWI 5259-2, ISO/IEC CD 5259-3, ISO/IEC CD 5259-4	Artificial intelligence — Data quality for analytics and machine learning (ML) Parts 1-4
ISO/IEC FDIS 24668	Information technology — Artificial intelligence — Process management framework for big data analytics
ISO/IEC TR 24027	Information technology — Artificial intelligence (AI) — Bias in AI systems and AI aided decision making
ISO/IEC DIS 24029-2	Artificial intelligence (AI) — Assessment of the robustness of neural networks — Part 2: Methodology for the use of formal methods
ISO/IEC FDIS 23894	Information technology — Artificial intelligence — Guidance on risk management
ISO/IEC TR 24368	Information technology — Artificial intelligence — Overview of ethical and societal concerns
ISO/IEC DTR 5469	Artificial intelligence — Functional safety and AI systems
ISO/IEC DIS 25059	Software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Quality model for AI systems
ISO/IEC AWI TS 6254	Information technology - Artificial intelligence - Objectives and approaches for explainability of ML models and AI systems
ISO/IEC TR 24030	Information technology - Artificial intelligence (AI) - Use cases
ISO/IEC DIS 5338	Information technology - Artificial intelligence - AI system life cycle processes
ISO/IEC CD 5339	Information Technology - Artificial Intelligence - Guidelines for AI applications
ISO/IEC TR 24372	Information technology - Artificial intelligence (AI) - Overview of computational approaches for AI systems
ISO/IEC PRF TS 4213	Information technology - Artificial intelligence - Assessment of machine learning classification performance

ISO/IEC CD 5392	Information technology - Artificial intelligence - Reference architecture of knowledge engineering
ISO/IEC PWI 18966	Artificial intelligence - Oversight of AI systems (Governance/Management)
ISO/IEC PWI TS 42108	AI Operational Design Domain
ISO/IEC NP IS 25059	AI Software Quality
ISO/IEC WD IS 24029-3	NN Robustness
ISO/IEC WD TR 42106	AI Benchmarking
ISO/IEC WD TS 42105	AI Oversight
ISO/IEC WD TS 22443	AI Ethical & Societal Concerns
ISO/IEC CD IS 12792	AI Transparency
ISO/IEC FD TS 8200	AI Controllability
ISO/IEC FD TS 12791	AI Bias
ISO/IEC TS 25058:(pending)	AI Software Quality

B.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.[27] CENELEC plays a similar role for electrotechnical standardization.

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 EU legislation, policies, principles, and values.[28]

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.

To date, CEN/CLC JTC 21 has not published any standards of its own, although it has adopted some ISO/IEC standards. Its current work program includes technical reports on conformity assessment, risk management checklists, environmental sustainability, and natural language processing. The standardization request from the Commission, which is expected to drive future work, includes

[27] https://www.cencenelec.eu/about-cen/

[28] https://www.cencenelec.eu/areas-of-work/cen-cenelec-topics/artificial-intelligence/

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 are all slated to be delivered by January 2025.

B.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."[29]

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.[30]

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 P 2894	IEEE Draft Guide for an Architectural Framework for Explainable Artificial Intelligence
IEEE P 2976	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
IEEE P2986	Recommended Practice for Privacy and Security for Federated Machine Learning
IEEE P2975	Standard for Industrial Artificial Intelligence (AI) Data Attributes

B.1.4. ITU

The International Telecommunication Union (ITU) is the United Nations specialized agency for information and communication technologies (ICTs).[31] The Study Groups of ITU's Telecommunication Standardization Sector (ITU-T) assemble experts from around the world to develop international

[29] https://standards.ieee.org/about/
[30] https://sagroups.ieee.org/ai-sc/
[31] https://www.itu.int/en/about/Pages/default.aspx

standards known as ITU-T Recommendations which act as defining elements in the global infrastructure of ICTs.[32]

Though its mandate historically was limited to telecommunications standards, ITU has initiated many projects on AI.

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
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

[32] https://www.itu.int/en/ITU-T/about/Pages/default.aspx

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
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

B.1.5. Sectoral standards: SDOs and topics

In industries that are coming to rely heavily on AI, sector-specific standards projects have also begun to emerge. SAE International, a global association of engineers and related technical experts in the aerospace, automotive and commercial vehicle industries,[33] is developing standards products on foundational concepts and certification processes related to AI in aeronautical systems. EUROCAE, a European non-profit that develops standards for European civil aviation, also has a working group on AI. The Consumer Technology Association has published a standard on characteristics of AI in healthcare, among other AI topics. 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.[34]

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.

[33] https://www.sae.org/

[34] https://x9.org/aistudygroup/