MP190421

June 10, 2019

Response of The MITRE Corporation to the National Institute of Standards and Technology (NIST) Request for Information on Artificial Intelligence Standards

For additional information about this response, please contact:
Duane Blackburn, S&T Policy Analyst
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

dblackburn@mitre.org
(434) 964-5023

Table of Contents

List of Figures

INTRODUCTION

The MITRE Corporation is pleased to provide this response to the National Institute of Science and Technology's (NIST) Request for Information (RFI) on Artificial Intelligence (AI) Standards. MITRE regards the effort described in the President's Executive Order (EO13859) to "plan for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies"[1] as vital to the economic and national security of the United States.

We view AI as one of the most important enablers of the emerging global information technology (IT) ecosystem, which our nation must help shape. Characteristics of this ecosystem include near-ubiquitous fifth generation (5G) networks offering speedy and direct connectivity to IT and Internet-of-Things (IoT) devices alike. Internet Protocol version 6 (IPv6) will provide a practically unlimited set of IP addresses for these devices for both the private and public sectors, while thousands of low-earth-orbit satellites will extend connectivity globally and in all terrains. Getting the most benefit from the new, complex, interconnected networks will require AI-enabled analytics, which will mediate resources for government, business, and national infrastructure. The first instantiations of "smart city" technology, enabled by AI, are already in use.

AI is becoming the technology from which other technologies derive increasing value. Dependence on AI will only grow in a world that increasingly defines enterprises by their "information intensity."[2] Overall, effective implementation of AI, as described by the Executive Order, will serve the dual imperatives of U.S. economic security – through increased economic innovation and competitiveness – and national security – through more effective and interoperable systems supporting the nation's defense.

The AI standards challenge is complex. How can standards guide the effective development of a discipline that changes rapidly? How can we use standards to influence the application of AI in ways we have not yet anticipated? The creation of a "plan for Federal engagement" must allow for the dynamic nature of the AI discipline and its application. We offer our views and recommendations accordingly.

The question of tools is equally complex, though replete with opportunity. Tools, such as MITRE's ATT&CK threat framework, can help shape the development of AI standards by providing a common understanding of the cybersecurity threats AI systems face. Other tools can improve the interoperability of AI-based applications. Tools might also be used to drive a common approach to the characterization and management of information subject to AI-based processing.

MITRE serves as the manager of the National Cybersecurity Federally Funded Research and Development Center, sponsored by NIST's National Cybersecurity Center of Excellence (NCCoE). As a result, our experience in working with the NCCoE has informed our understanding of NIST's mission. MITRE's technical and domain expertise, however, extends well beyond cybersecurity into AI, machine learning, computer science, data analytics, and other fields. Our views and recommendations reflect our broad understanding of the effects of IT on the economic and national security of our country and the role we have played in support of the development national strategies and communities of interest. In addition, our work with the NCCoE gives MITRE broad exposure to the private sector, both to the

[1] Executive Order 13859 of February 11, 2019 - Maintaining American Leadership in Artificial Intelligence.
[2] See: https://profesores.virtual.uniandes.edu.co/~isis1404/dokuwiki/lib/exe/fetch.php?media=bibliografia:10_how_information_gives_you_competitive_advantage.pdf

companies developing and deploying new, complex IT infrastructures and systems, and to the companies developing the technologies necessary to defend those infrastructures and systems.

This exposure to the private sector reinforces our view that government cannot, and should not, attempt to develop AI standards on its own. The preponderance of technology innovation and deployment takes place in our nation's private sector. As a result, the insights necessary to drive AI development and application must encompass the resources of the private sector, even as the Government brings to this process the broader public interest.

Our views regarding AI are also informed by our work operating six other FFRDCs, including those for transportation, health care, aviation, national security, homeland security, and other missions. Each of these domains is certain to adopt aspects of AI technology, and MITRE is being asked to help define the future state of the information systems that support these domains, as well as the strategies needed to achieve this future state. For example, MITRE's work in support of the future of civil aviation is likely to encompass increasingly broad application of AI. Our insights are shaped to a great extent by this exposure to a wide range of national IT imperatives.

Overall, we have crafted insights and recommendations that reflect our understanding of the need for a public-private partnership approach to the development of AI standards, our insights into a broad range of potential missions to which AI can be applied, our experience in working with both the Government and the broad private sector, and the national economic and security imperatives that an effective AI standards strategy can help meet. We look forward to discussing with you our views and recommendations, and offer to provide additional assistance as the American AI Initiative evolves.

AI TECHNICAL STANDARDS AND RELATED TOOLS DEVELOPMENT: STATUS AND PLANS

MITRE has provided responses to the questions for which we believe we have either an in-depth understanding or specialized knowledge that can aid NIST in develoing a Federal AI Standards Engagement Plan. Our responses to those specific questions follow.

Throughout this document, the term "standards" uses the current definition within FIPS 201, i.e., a standard is "a published statement on a topic specifying characteristics, usually measurable, that must be satisfied or achieved," in this specific case, for the goal of ensuring a level of assurance in AI to reflect an aspect such as accuracy, robustness, or transparency.

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

MITRE, through its relationships with industry, the Institute of Electrical and Electronics Engineers (IEEE), the Association for Computing Machinery (ACM), the National Institute of Standards and Technology (NIST), the Intelligence Advanced Research Project Agency (IARPA), and the Defense Advanced Research Projects Agency (DARPA), and through our own internal research efforts, recognizes an overarching concern about aspects of AI, such as standards, tools, ethics, and transparency. This concern focuses on various efforts concentrated on ethics, transparency, and the "support of reliable, robust, and trustworthy systems," highlighted as a specific need in the *Executive Order (EO) on Maintaining American Leadership in Artificial Intelligence*.

The technical report by Peter Cihon from the University of Oxford, Standards for AI Governance: International Standards to Enable Global Coordination in AI Research & Development, states that,

"There are two existing international standards bodies that are currently developing AI standards. First is a joint effort between International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC)… The second international standards body that is notable in developing AI standards is the IEEE Standards Association… A third international standards body may become increasingly relevant for AI in the future: the International Telecommunications Union (ITU). The ITU has historically played a role in standards for information and communications technologies, particularly in telecommunications. It has a Focus Group on Machine Learning for Future Networks that falls within this telecommunications remit."[3] Our analysis of the materials from these bodies shows that IEEE standards have been more focused on ethics, privacy, accuracy, transparency, and bias, with only two such standards focused on reliability, robustness, and trustworthiness: Fail-safe Design for AI Systems (IEEE P7009), and Process of Identifying and Rating the Trustworthiness of News Sources (IEEE P7011). By comparison, the Joint Committee, ISO/IEC JTC 1 Standards Committee on Artificial Intelligence (SC 42), has committed to developing standards more tightly coupled with the security of AI in anticipation of developing three publications focused on "robustness of neural networks, bias in AI systems, and overview of trustworthiness in AI".

Companies that provide AI as a service appear to be the primary drivers of sector-specific AI efforts, including Amazon (Web Services)[4], Google (Cloud)[5], Microsoft (Azure)[6], and IBM (MWatson)[7]. Other companies that cater to large development communities, like Facebook and Apple, are engaging in activities aimed at securing AI. These companies provide open-source libraries of attack types and defense methods, which are intended to help AI developers and practitioners assess multiple concerns, such as the susceptibility of their applications to adversarial manipulations or vulnerabilities to privacy concerns.

Open source libraries of these attack types and defense methods are very valuable for making robust and secure AI; they are of cross-sector interest and directly support the Executive Order. Given the Executive Order's interest in training "the next generation of American AI researchers," developing standards for AI tools will be valuable in building AI training programs.

As a result of our involvement with government research programs, MITRE is aware that the Government is standing up organizations to address the research gaps not covered by industry. Figure 1 highlights these organizations and shows their instantiation dates.

---

[3] Cihon, Peter. "Standards for AI Governance: International Standards to Enable Global Coordination in AI Research & Development," University of Oxford: https://www.fhi.ox.ac.uk/wp-content/uploads/Standards_-FHI-Technical-Report.pdf

[4] Amazon Web Services, Inc. (2018.) *Machine learning on AWS [Online]*. Available: https://aws.amazon.com/machine-learning/.

[5] Google LLC. (2018). *Cloud machine learning engine* [Online]. Available: https://cloud.google.com/ml-engine/.

[6] Microsoft. (2018.) *Azure machine learning* [Online]. Available: https://azure.microsoft.com/en-us/overview/machine-learning/.

[7] IBM. (2018). *IBM Watson* [Online]. Available: https://www.ibm.com/watson/.

Figure 1. Landscape of Government Initiatives in AI Privacy, Security, and Assurance. *Beginning with the AI Security Initiative in 2017, Government organizations continue to address AI through research studies and other activities.*

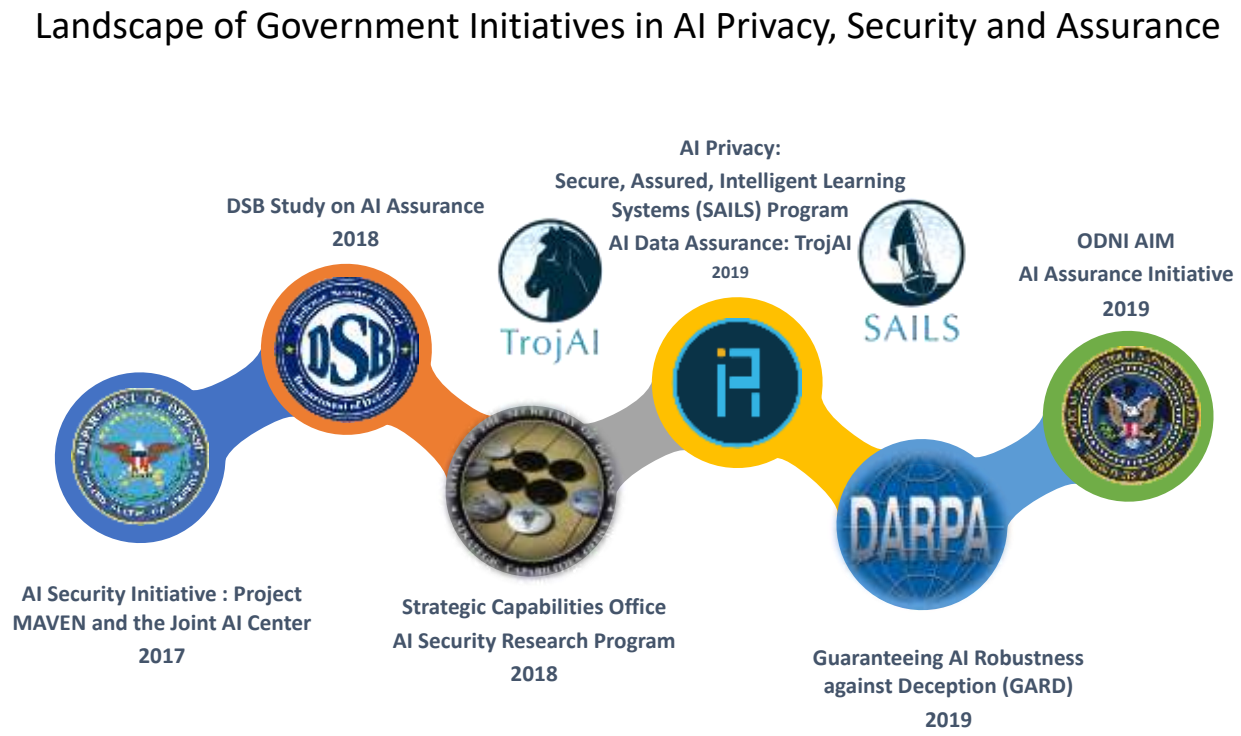

MITRE recognizes that AI technical standards and tools developed by academia, industry, and government that address sector-specific needs may also be leveraged in a cross-sector approach.

2.0 Reliable sources of information about the availability and use of AI technical standards and tools

The response to Question 1 identifies and discusses the presence of reliable sources for tool development. While reliable sources for standards are scarce, this deficiency does not reflect communities' lack of desire for reliable sources. For example, IEEE has stood up engagement opportunities, such as Dependable and Secure Machine Learnning,[8] but these opportunities have not yet evolved to a point where they have produced and promoted their information about standards and tools development.

Other information sources are also being adopted, but they are extremely new developments. For example, "The Recommendation on Artificial Intelligence (AI)," the first intergovernmental standard on AI, was adopted by the Organisation for Economic Cooperation and Development (OECD) Council at the Ministerial level on May 22, 2019, on the proposal of the Committee on Digital Economy Policy (CDEP). The recommendation aims to foster innovation and trust in AI by promoting the responsible

[8] See: *https://dependablesecureml.github.io/index.html*

stewardship of trustworthy AI while ensuring respect for human rights and democratic values. The recommendations complement existing OECD standards in areas such as privacy, digital security risk management, and responsible business conduct. They also focus on AI-specific issues and set a standard that is implementable and sufficiently flexible to stand the test of time in this rapidly evolving field." [9]

MITRE recognizes that there are a few souces of information for AI standards and tools, but the newness of these types of developments has created only a few reliable sources that are not yet widely publicized.

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

Today, AI communities are challenged by the inconsistent use of terminology, and no standard definition exists for what these various terms mean. To help drive communities to a common approach to terminology, MITRE is providing direct support to NIST as they develop a foundation of common terminology and taxonomy of attacks and defense, with respect to adversarial machine learning. We must identify and understand what reliable definitions of these standards reflect before reliable sources of standards can be developed.

While the "definition" stage is considered critical, it is not in and of itself sufficient for standards and tools development. The second challenge involves developing a framework that can be adopted acoss communities that captures data, processes, and risks. This framework would illustrate the various attributes and models relevant to the development of AI and identify various concerns that include, but are not limited to, risk, ethics, privacy, and transparency. For example, the development of a risk model for AI would provide a framework to better understand potential threats, vulnerabilities, and risks to systems that are dependent on machine learning. This threat model should be a function of the likelihood of a given counter-AI vulnerability and the severity of the resulting impact.

Consequently, MITRE recommends that, first, we must define and use terminology uniformly. This first step supports the creating of a framework. Once these two developments (i.e., terminology and framework) are met, we can develop reliable standards and tools. .

5.0 Any supporting roadmaps or similar documents about plans for developing AI technical standards and tools

There is little publicly available information on sector-specific AI roadmap development. We believe, however, that organizations have their own roadmaps for standards and tools development to aid the evolution of their AI products. Cross-sector organizations have established working groups to focus on these needs, but MITRE has not found any timelines concerning dates for development or release.

[9] See: https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

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

Given the nascent development of standards, it is hard to determine if sector-specific or cross-sector organizations are meeting the need for standards development. Both, however, are playing vital roles in pioneering development efforts.

NIST is developing foundations for measuring the trustworthiness of AI algorithms, systems, and data. NIST's efforts include relevant projects (for example, a project that involves developing metrics to evaluate the security of control systems).[10] Another NIST AI program closely related to developing metrics and best practices is the NIST Foundry for Living Measurement Systems. This program involves engineering cells for safe and reliable use in dynamic and unpredictable environments, tenets both relevant to AI. Other standards organizations, such as the Internet Engineering Task Force (IETF), are placing AI and machine learning (ML) in their proposals for future working groups[11]. These standards developments, created by cross-sector organizations, demonstrate their usefulness across mulitple sectors.

Industry sectors have been playing an important role as well by investigating how to evaluate and standardize AI-enabled systems in the automobile[12] [13] [14] and health-care[15] [16] [17] industries. IBM proposed that Supplier's Declarations of Conformity (SDoCs) can constitute a standardized way to add transparency and accountability about the security and safety of AI systems and services[18]. The authors argue that providers of AI services could increasingly adopt SDoCs to remain competitive in the market, and, as a result, SDoCs have the potential to bootstrap a new era of trusted AI. The first two developments may at first seem more sector-specific, but in fact all three, as they evolve, may lend themselves to cross-sector adoption.

The technology industry has also approached the problem of secure AI by deviating from the concept of traditional standards and instead trying to develop libraries that can help AI practitioners and

---

[10] D.I. Urbina et al., "Survey and new directions for physics-based attack detection in control systems," NIST, Gaithersburg, MD, NIST GCR 16-010, 2016.
[11] Edgewall Software. (2018). *IRTF Wiki* {Online]. Available: https://trac.ietf.org/trac/irtf/wiki.
[12] D.J. Fagnant and K. Kockelman, "Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations," *Transp. Res. Part Policy Pract.*, vol. 77, pp. 167-181, 2015.
[13] N. Kaira and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?," *Transp. Res. Part Policy Pract.*, vol. 94. pp. 182-193, 2016.
[14] SAE International. (2016, Sep. 20). *U.S. DoT chooses SAE J3016 for vehicle-autonomy policy guidance – SAE International* [Online]. Available: http://articles.sae.org/15021/.
[15] U.S. Food and Drug Administration. (2018, Apr. 11). *FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems* [Online]. Available: https://www.fda.gov/newsevents/newsroom/pressannouncements/ucm604357.htm.
[16] S. Gottlieb, U.S. Food and Drug Administration. (2018, Apr. 26). *Transforming FDA's approach to digital health* [Online]. Available: https://www.fda.gov/newsevents/speeches/ucm605697.htm.
[17] D.S. Char, N.H. Shah, and D. Magnus, "Implementing machine learning in health care – Addressing ethical challenges, "*N. Engl. J. Med.*, vol. 378, no.11, p. 981, 2018.
[18] M. Hind et al., "Increasing trust in AI services through supplier's declarations of conformity," ArXiv Prepr. ArXiv1808.07261, 2018.

application designers to evaluate the resilience of their algorithms and systems [19] [20]. However, this evaluation approach goes beyond efforts to secure AI[21]. For example, organizations such as OpenAI have developed full systems for comparing reinforcement learning algorithms[22] and for benchmarking competitions for AI systems.[23] DeepMind,[24] Google,[25] and Microsoft[26] [27] have developed comparable simulation environments.

MITRE finds that many developments are occurring across sector-specific organizations and cross-sector organizations. All developments, though, are meeting various needs and no one organization is dominating in meeting the needs of AI technical standards.

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

The AI community must make significant progress toward gaining a deep understanding about how to secure AI systems through standardization. MITRE is not alone in this belief, as multiple authors have confirmed. [28] [29] [30] The same AI technologies widely adopted by various legitimate industries are also being used by adversaries to extend the capabilities of human operators, increase the difficulty of attack attribution, and reduce command and control communication with compromised systems. Therefore, potential threats are likely to increase rapidly in step with evolving technology.

The AI community must have standards on the underlying technologies and their dependencies to ensure the overall accuracy, security, privacy, and other aspects of AI are satisfactory for use in the U.S. Some of these standards exist, but standards and guidance still need to be provided on how to apply these existing practices to new concepts such as AI. Also, an AI framework needs maturing and community involvement. The development of a framework would helpl industries understand what concerns are specific to their sector and ultimately allow for interoperability across industries.

---

[19] V. Zantedeschi, M.I. Nicolae, and A. Rawat, "Efficient defenses against adversarial attacks," *CoRR*, vol. abs/1707.06728, 2017.
[20] N. Papernot et al., "Cleverhans v.1.0.0: an adversarial machine learning library, "ArXiv Prepr. ArXiv161000768, 2016.
[21] A. Mojsilovic. (2018, Aug. 22). *Factsheets for AI services* [Online]. Available: https://www.ibm.com/blogs/research/2018/08/factsheets-ai/.
[22] OpenAI. (n.d.). *OpenAI Gym* [Online]. Available: https://gym.openai.com/read-only.html.
[23] OpenAI. (2018, Jul. 18). *OpenAI five benchmark* [Online]. Available: https://blog.openai.com/openai-five-banchmark/.
[24] C. Beattie et al., "Deepmind lab" ArXiv161203801 Cs, Dec. 2016.
[25] J. Wexler. (2018 Sep. 11). *The what-if tool: Code-free probing of machine learning models* [Online]. Available: https://ai.googleblog.com/2018/09/the-what-if-tool-code-free-probing-of.html.
[26] S. Shah 35 al., "AirSim: High-fidelity visual and physical simulation for autonomous vehicles, *Field and Service Robotics*, pp. 621-635, 2017.
[27] M. Johnson et al., "The Malmo platform for artificial intelligence experimentation," in *25th International Joint Conference on Artificial Intelligence*, New York, NY, 2016, pp. 4246-4247.
[28] O. Suciu et al., "When does machine learning FAIL? Generalized transferability for evasion and poisoning attacks," in *27th USENIX Security Symposium*, Baltimore, MD, 2018, pp. 1299-1316.
[29] V. Wang et al., "With Great Training Comes Great Vulnerability: Practical Attacks against Transfer Learning," in 27th USENIX Security Symposium, Baltimore, MD 2018, pp. 1281-1297.
[30] M. Brundage et al., "The malicious use of artificial intelligence: Forecasting, prevention, and mitigation, "ArXiv Prepr. ArXiv180207228, 2018.

DEFINING AND ACHIEVING U.S. AI TECHNICAL STANDARDS LEADERSHIP

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

The President's Executive Order (EO) has mandated an entirely new level of imperative action for establishing technical standards supporting the pursuit of AI, and NIST must respond to that challenge. That same sense of urgency purported by the EO leading to broad adoption by U.S. interests must become mainstream in order to realize new, resilient, and durable AI technical standards. Underscoring this urgency is the global competition for building AI-enabled complex infrastructures and the possible ubiquity of foreign hardware and software as the building blocks of these infrastructures. The use of AI to mediate resources across complex systems is likely to come from China, through its "made in China 2025" program, attempts self-sufficiency and global market dominance in the AI domain.

Leadership responsibility for this vast, mutli-disciplined domain, which traverses the entirety of IT and many other practical and functional areas, will require a unification of forces and influences that takes precedence over individual government agency roles or industry markets. Instead, the movement must capitalize on the synergy of AI to converge and align national interests into a coordinated technical transformation. NIST and its technical peers must optimize and streamline the effort surrounding how standards are formulated, developed, and accepted so those standards are timely, suitable, and useful.

NIST must push to shorten, optimize, and consolidate the usual process for proposal, analysis, discussion, and community comment of typical technical standards development. The scope of AI is vast, while at the same time interthreaded across nearly every technical influence and domain. Thus, AI standards development will require a specialized and highly optimized leadership model that pulls together such strengths as the Government's expectations for performance, industry's ability to innovate, and the nation's willingness to adopt and operationalize, all at once. NIST will have to redesign and recast many traditional practices that have steered public/private partnerships in the past, and reconcile the motivations of both sides of these partnerships to achieve the maximum level of progress, while mitigating all of the destabilizing forces that exist within these partnerships. The leadership exhibited to develop these partnerships must be grounded in its objective purpose and then expeditiously held to that path.

A great urgency also reflects the international AI landscape. The competition for global technical dominance has a direct influence on U.S. national and international standing, and the U.S. is admittedly well behind other countries in both its adoption of AI and its investments in it. U.S. leadership must recognize this and establish phased planning, which can account for that definciency, undertake corrective actions to equalize U.S. influence in AI, and propel U.S. interests ahead. As our international peers and political rivals establish and harness efficiencies in and through AI, the U.S. Government and its interests must find ways to rapidly prioritize, and then operationalize, AI-centric approaches to realize the greatest, and fastest, positive effect.

NIST must prioritize focus areas for AI implementation for maximum effect, along with economy of effort. Not all U.S. interests are under direct Government influence, but establishing AI technical standards, and the tools through which those standards are achieved and reinforced, brings with it collective improvement. This is most significant in the context of national security and defense interests, where the contention among different nations' technical advances can be played out in warfighting, where any efficiency realized through AI may lead to a strategic or tactical advantage.

National security imperatives are only part of the proposition for AI technical standards. As the EO suggests, nearly everything that AI can influence or impact can be considered of national interest. Industry's innovations are rapidly evolving and all U.S. interests can benefit from these innovations, but leadership must oversee and manage it for collective effect through a leadership model that can account for, and be considerate of, all of the AI use cases bearing on broad AI implementation, and close the gap between innovation and time to implementation.

NIST, working through the National Science and Technology Council (NSTC), could undertake an approach to unify stakeholder interests and actions that consists of several important steps. These include: develop a national AI stakeholder community; identify key AI standards challenges; allocate those challenges broadly against groups of stakeholders; and put in place an AI standards collaboration and information sharing network. A prior NSTC-led national-level standards effort and supporting policy[31], which guided development and mandated consistent federal adoption of biometric standards to support post-9/11 screening advances, can provide lessons-learned and a starting point for developing this collaborative approach.

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

Leadership for achieving technical standards development has not been as focused, or as motivated, as the exploding AI field requires. AI technical standards have been largely localized in their creation and adoption, and have come about through functional needs instead of unified AI standards leadership. Early adopters in Government and industry have been seeking standards but are not finding them readily available, particularly in the emergent areas of cutting-edge AI. Consequently, early adopters have improvised specific standards to suit their own propositions, hoping that at a later date there can be a for technical convergence with the larger AI community. So far, that convergence has been voluntary.

The same leadership necessary to achieve effective AI standards must fuse the interests of both the public and private sectors to capitalize on the the innovation of private industry and the receptivity of Government to transform the technology. To date, this fusion has not been consistent because the leadership has not been identified or sufficiently empowered. NIST must establish a role in the technical dialogue between industry and Government and promote an persistent, attractive narrative for cooperation and complementary effort. Industry and Government each bring strengths to the table that are required for an effectiveness solution. This suggests a hybridized model for leadership in AI that can objectively combine the interests of all stakeholders into a future technical vision, and pursue standards to achieve this end. To promote the atmosphere of shared success between public and private interests, the leadership model must include the development of more effectively focused messaging and education that lays out the basic expectations of AI.

As technical leadership in AI is already manifesting throughout industry, the momentum to define and drive standards most suitable to industry is already evident. Early adopters in many sectors have evolved their own working technical standards to fortify and optimize their AI-driven commercial technologies. However, the Government is not efficiently aligned with those innovations. Government's technical leadership is not widely recognized or organized, causing individual federal agencies to solve their own problems and follow their own strategic plans in isolation. It is essential for government to understand

[31] https://www.nist.gov/sites/default/files/documents/2017/04/12/nstc_policy_bio_standards.pdf

prospective technical futures and be mindful of the differences between industry's priorities and the requirements of Government transformation.

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

Effectiveness leadership for technical AI standards must channel the competitive innovation of the private sector and adapt it to Government. The AI community must rally an atmosphere of innovation, optimization, and transformative change, not only in Government agencies, but also across and among them, interlaced with the momentum of the private sector. It is critical to leverage a systemic view of the Government's as-is state and the understanding of the circumstances and drivers around it that are suitable to AI transformation, and leadership must account for this. Private industry has already achieved a level of excellence in creating market environments compatible with AI development and integrating their results into operational infrastructure. Government can harness this through forceful and dynamic engagement.

In meeting the President's AI challenge, Government must not slow the pace of commercial development, but rather accelerate Government interests through cooperative alignment with industry. MITRE strongly asserts that this alignment must be facilitated and engineered to achieve maximum gain in the shortest space of time. Standards that support AI technical innovation useful to Government must be responsive to the same opportunities to transform that industry sees, and they must be closely examined to discover the advantages for satisfying government use cases. This is a test of dedication to a future state for U.S. interests, and while assurance of success is often qualified, movement toward an improved operational state through AI is highly likely.

Because there is no consolidated authority over all U.S. AI interests, the lack of centralization has prompted industry to pursue its own agenda. This state of affairs presents a significant challenge to U.S. interests in both the public and private sectors. The U.S. must initiate a joint effort, aligned and synchronized through fully informed, insightful, and dedicated leadership to gain ground on, compete with, and surpass other parties in the international AI landscape. Leadership must create and adopt a new model and new rules, including influence and input on joint ventures, identifying areas of urgent technical need, managing peripheral interests, developing a phased plan for the pursuit of AI within critical infrastructure, and recognizing all direct and indirect drivers affecting AI. The desired outcome—compatible uniformity across the shifting AI landscape—will require close cooperation and risk sharing in working relationships.

Leadership is required to manage this dynamic interaction, which can be provided by NIST and the NSTC. Once enabled by the support of government, industry, and academic members of a national AI standards community, the U.S. can achieve the desired outcome.

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

The idea that AI is an attainable national priority must be reinforced for both Government and industry. They must work together to ground this goal in firm technical standards to facilitate U.S. interest-wide transformative change. While the President's EO sets this in motion, adoption and dedication at the working level is often a function of circumstance. NIST must drive AI standards development with an intentional acceleration of purpose: standards must come faster, and there must be a shared sense of urgency. Maintaining and expanding the advantage in developing and driving technical standards must

be based on the perceived importance to the national interest. This leverages a notion of leadership more aligned with industry's experience than the Government's. How this can be devised rests on any number of influences (authority, budget, short- and long-term goals, programmatic expectations, and other factors), which have not yet been articulated. There is no substitute for this work, and the drive must be continually and actively reinforced. Industry must recognize the positive incentives for advocating standards, and Government must enable this through a system of acceptance and adoption which is timely to its own demands.

There must be broad cooperative alignment, and incentives may be needed to entice industry to consider pursuing community-wide technical standards. Leadership must strengthen and communicate stronger value propositions and compelling arguments for undertaking collective efforts as a reinforcing exercise to continuously focus the national agenda. NIST should state this proposition at the point of commencement to the national AI agenda to keep technical standards flowing. Incentivizing industry, academia, and small business interests to support AI technical standards through federal grant vehicles may aid in building momentum toward standards recognition and adoption at higher levels of activity and interest.

PRIORITIZING FEDERAL GOVERNMENT ENGAGEMENT IN AI STANDARDIZATION

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

The uniqueness of the Federal Government's use cases for embracing AI is compelling because of the high potential for dramatic change and improvement to static processes beneficial to Government efficiencies. This change will ride on technical standards that provide consistency in how AI is implemented and leveraged, as well as enabling the associated tools necessary for use in the AI environment. The complementary nexus between these two elements, standards and tools, must be reproducible to meet future needs as priorities change and trends emerge.

The Government's increased involvement in AI must also consider both the very specific needs and perceived use cases at the individual agency level and the need to create more generic operational use cases across the broader Government superstructure and beyond. The assembled leadership must realize that individual agencies have their own views of, and demand for, harnessing AI in the shadow of mandated policies or government-wide initiatives, which are often less considerate of agency-specific success. Compounding this problem are budgetary considerations, technical resource refresh initiatives, IT roadmaps, expanding mission areas, specific constraints on privacy of data and data protection, data retention, and fluctuations in available workforce staffing plans. Leadership must consider and reconcile these enormous variables at the local agency level, since doing so contributes to the momentum of adopting AI throughout the Government. While viewed in the context of individual mission areas and responsibilites, AI transcends these boundaries and suggests that the way we define traditional agency mandates will change as AI is accepted and implemented, altering the boundaries between agencies.

Government's AI adoption has massive influence outside of its own scope of interest through the reinvention of Government itself, and, empowered with AI, this influence will be transformative to society, culture, and the end-user experience through the sheer breadth of the AI undertaking.

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

Government itself must define the scope, scale, and context of where and how AI can be most beneficially implemented for two reasons: it validates the Government's commitment to harness AI by embedding itself in its own solution path, and it makes Government beholden to the consequences of decision-making that reach beyond its own logical boundaries, which are swiftly dissolving. Surrendering individual agency autonomy to a larger authority can blind agencies to their own opportunity to transform, and removes them from being seen as stakeholders by the larger leadership model which AI requires. That transformation will be built on technical standards, and agencies are their own stakeholders. Being prepared to invest the time and resources necessary to embrace the AI opportunity is a display of commitment to the outcome at all levels of scale.

To that end, Government must seek to create a fortress of uniformity for AI beneficial to its own interests, by requiring standards be pursued and accomplished in a timely manner to avoid stalling implementation and innovation. AI requires participation at all levels of Government, so that each layer understands the signficiance and relevance of AI to both the agency-level mission and the larger Government need.

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

The prioritization of AI technical standards should be based on a blended approach that considers both cross-sectional application and sector/application-specific standards all at once. This is because AI solutions have considerable dependencies on peripheral and fringe processes and resources that typically cannot be isolated from those mechanisms dependent on them. Pursuing AI is a holistic exercise.

AI technical standards could also be prioritized using those technical use cases most readily achievable across the broadest possible resource landscape, coupling both cross-sectional and application-specific standards together. Depending on how the Government's AI-receptive technical landscape is perceived and accounted for, those processes and data resources most significant to the technical core of Government are likely the most influential, and subject to greater impact through optimization. This should not, however, marginalize the need to understand and consider peripheral processes and technical conditions that surround those core use cases as being essential to core innovations.

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

Current AI technical standards have been largely driven by industry's own pursuit of innovation. Government's investment of interest is clear, but it is often considered as an afterthought to industry's actions to achieve standards. Industry's motivations go beyond simple Government use cases. Moving forward, Government must adopt a much more aggressive role in brokering the recognition of technical standards and driving those standards through formalization. Government/industry partnerships are essential to ensure that technical standards are suitable, appropriate, germane, and adaptable to broad technical use cases, both in Government and within the strategic direction of industry.

The actual process of formal standards establishment is not currently scalable or timely to the level of demand commanded by the AI agenda. One element of technical standards achievement that can be influenced by the government is the speed at which those standards are reviewed, researched, processed, and fomalized. This will have a direct bearing on how industry sees government's commitment to AI. Because the prospects for success in infusing AI into government rest on industry innovation, and that innovation is grounded in standards, leadership will need to develop entirely new mechanisms for realizing technical standards more rapidly. This includes the redesign of traditional standards workflow management, technical research, due diligence, quality assurance, and the reconciling of community input to provide for faster determination of needed standards, at the volume and pace necessary to keep them relevant to real-time advancements. This will not be a comfortable circumstance for either government or industry. Success requires improvements to throughtput of standards evaluation.

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

The EO establishes a bold vision for the harnessing of AI, and involves many aspects of industry and technical stakeholders in what is really a Government-mandated transformation. This challenges the status quo of both industry and Government through a coupling of interests that has traditionally been divergent. Government's enlistment of industry is for mutual and collective benefit, not just for the present, but for the future, and that partnership is undeniable. As industry's investment in the effort to realize Government-centric AI grows, so does its interest in the sustainability and resiliency of those capabilities it has enabled. This will offer powerful assurances that as technical standards are crafted and fielded, they are applied and adopted. The strongest and most potent action the Government can take is to assemble its industry partnerships, establish its roadmap to transformation, and account for its level of effort to ensure that the roadmap is followed, at the individual agency level, within other U.S. interests, and by the Government as a whole.