May 31, 2019

Elham Tabassi
National Institute of Standards and Technology
100 Bureau Drive, Stop 2000
Gaithersburg, MD 20899

Via email to: ai_standards@nisg.gov

Re: Developing a Federal Artificial Intelligence Standards Engagement Plan

Dear Ms. Tabassi:

BSA | The Software Alliance appreciates the opportunity to provide feedback in response to the National Institute of Standards and Technology's (NIST's) Request for Information regarding the development of a Federal Artificial Intelligence Standards Engagement Plan.[1] BSA is the leading advocate for the global software industry. [2] Our members are at the forefront of software-enabled innovation that is fueling global economic growth and advancing the development and deployment of Artificial Intelligence (AI).

As global leaders in the development of cutting-edge technologies, BSA members recognize the important role that technical standards and benchmarks can play in promoting trust and confidence in new technologies by establishing baseline measures for quality assurance, facilitating interoperability, promoting best practices, and enabling collaboration. Standards can help unlock marketplace efficiencies by establishing a common framework of understanding between developers and consumers of technologies. BSA therefore strongly supports NIST's effort to develop a Standards Engagement Plan to support the development of "reliable, robust, and trustworthy" AI systems and promote the "creation of new AI-

[1] 84 Fed. Reg. 18490 (May 1, 2019) [hereinafter "RFI"].

[2] BSA's members include: Adobe, Akamai, Apple, Autodesk, Bentley Systems, Box, Cadence, CNC/Mastercam, DataStax, DocuSign, IBM, Informatica, Intel, MathWorks, Microsoft, Okta, Oracle, PTC, Salesforce, ServiceNow, Siemens PLM Software, Sitecore, Slack, Splunk, Symantec, Trend Micro, Trimble Solutions Corporation, Twilio, and Workday.

related industries and the adoption of AI by today's industries."[3] We offer below two recommendations for advancing these objectives.

Heighten Current Level of Engagement with International Standards Development Organizations

The Standards Engagement Plan should prioritize robust US participation with the range of international standards development organizations that are currently developing AI standards. In addition to promoting trust, confidence, and marketplace efficiencies, international standards have the added benefit of mitigating the risks that can accompany country-specific standards. The proliferation of national standards can undermine global commerce and stunt the development of technology in two related ways. First, it can give rise to a patchwork of inconsistent national standards that act as an unintentional barrier to international trade, making it more costly for companies to develop and sell their AI-related products and services to the global marketplace. Second, national standards can also serve as overt barriers to trade when they are manipulated to "create unfair advantages for national firms, including with respect to participation by foreign firms."[4]

A Standards Engagement Plan that prioritizes US leadership in the development of international standards can serve as an important safeguard against these risks. Pursuant to the World Trade Organization's Agreement on Technical Barriers to Trade, countries that are considering the adoption of technical regulations must ensure that they are consistent with any existing (or imminent) international standards.[5] Robust engagement with international standards bodies can therefore help promote a globally harmonized approach to the governance of AI and prevent the use of domestic standards as a tool for protectionism.

Both the International Organization for Standards (ISO) and the Institute of Electrical and Electronics Engineers (IEEE) are currently developing an ambitious suite of technical and process-based standards pertaining to the development of AI. The ISO's Standards

---

[3] RFI at 18491.

[4] *See* United States Trade Representative, 2019 Special 301 Report at page 17, available at https://ustr.gov/sites/default/files/2019_Special_301_Report.pdf.

[5] World Trade Organization, Technical Barriers to Trade Agreement, Article 2.4 ("Where technical regulations are required and relevant international standards exist or their completion is imminent, Members shall use them, or the relevant parts of them, as a basis for their technical regulations except when such international standards or relevant parts would be an ineffective or inappropriate means for the fulfilment of the legitimate objectives pursued, for instance because of fundamental climatic or geographical factors or fundamental technological problems.").

Committee on Artificial Intelligence[6] has completed work on three standards and has 11 additional standards currently under development.[7] Among other things, the ISO work is focused on establishing "foundational" standards (e.g., developing a common AI terminology framework) and pursuing standards to promote AI "trustworthiness" (e.g., bias, explainability, security, and privacy). IEEE's work is taking place through the Global Initiative on Ethics of Autonomous and Intelligent Systems,[8] with individual working groups that are exploring more than a dozen standards on a range of issues, including the transparency of autonomous systems,[9] algorithmic bias considerations,[10] and automated facial analysis technology.[11]

Because these international standards could shape the technological and regulatory landscape for the future development of AI, ensuring robust US government involvement in their development should form the centerpiece of the Standards Engagement Plan. Indeed, other leading AI nations have already signaled their intent to help shape the development of these standards,[12] and have established formal working groups to coordinate engagement on international and national standards.[13] To ensure that US interests are adequately represented, the Standards Engagement Plan should include a roadmap for how the US government plans to engage with ISO and IEEE. NIST should also consider the establishment of formal mechanisms (e.g., working groups) by which industry can receive updates and provide inputs to help inform US government participation in international standards development processes.

---

[6] *See* ISO/IEC JTC 1/SC 42 at https://www.iso.org/committee/6794475.html

[7] https://www.iso.org/committee/6794475/x/catalogue/p/0/u/1/w/0/d/0

[8] https://ethicsinaction.ieee.org/

[9] https://standards.ieee.org/project/7001.html

[10] https://standards.ieee.org/project/7003.html

[11] https://standards.ieee.org/project/7013.html

[12] Jeffrey Ding, Paul Triolo, and Samm Sacks, *Chinese Interests Take a Big Seat at the AI Governance Table – Government and Industry Team to Shape Emerging AI Standards-Setting Process*, New America (June, 2018), available at https://www.newamerica.org/cybersecurity-initiative/digichina/blog/chinese-interests-take-big-seat-ai-governance-table/

[13] *See* Peter Cihon, *Standards for AI Governance: International Standards to Enable Global Coordination in AI Research & Development*, University of Oxford – Future of Humanity Institute (April 2019), available at https://www.fhi.ox.ac.uk/wp-content/uploads/Standards_-FHI-Technical-Report.pdf.

Convene a Process to Develop an AI Lifecycle Risk Management Framework

As NIST develops the Standards Engagement Plan, it should consider the full range of potential tools for promoting trustworthy AI. In addition to technical standards and benchmarks, NIST has considerable expertise in guiding the development of risk management frameworks that help enterprises of all sizes design, operate, and use technologies with greater confidence and trust. While individual standards can help address specific challenges, NIST should draw from its experience in developing the Cybersecurity Framework, and its ongoing effort to develop a Privacy Framework, and consider developing a similar tool to help stakeholders manage the spectrum of risks that could undermine the trustworthiness of an AI system.

Like efforts to secure networks and personal data, ensuring that AI systems are trustworthy requires a lifecycle approach to risk management. Issues that may impact the trustworthiness of an AI system can arise during multiple stages of the AI system lifecycle, including when an AI system is being designed, when its training datasets are constructed, when its models are defined and trained, when it is tested, and after it has been deployed. Individual standards and benchmarks will play an important role in mitigating specific risks that may arise during discrete phases of the AI lifecycle. But, NIST can also foster the development of "trustworthy" AI by developing more holistic tools for identifying and mitigating risks through the various stages of the AI lifecycle.

To that end, NIST should consider convening a multistakeholder process for the purpose of developing an "AI Lifecycle Risk Management Framework." Like the process that led to Cybersecurity Framework, development of an AI Lifecycle Risk Management Framework would enable stakeholders to identify a voluntary, consensus-based set of standards, guidelines, best practices, methodologies, procedures, and processes to cost-effectively mitigate risks that AI systems may pose. By organizing such a framework around the specific phases of the AI lifecycle, NIST could help stakeholders (including designers, deployers, and users of AI systems) identify the range of existing standards, system architectures, governance processes, technical tools, and best practices that can be employed for the purposes of mitigating specific risks and promoting trustworthiness.

Like the Cybersecurity Framework, an AI Lifecycle Risk Management Framework could also help facilitate communication across the AI supply chain. By establishing a common set of base definitions for the conceptual underpinnings of trustworthy AI (e.g., fairness, explainability, robustness, and transparency), a Risk Management Framework will help AI stakeholders more seamlessly communicate about potential risks and available mitigation measures. Because the risks implicated by any particular AI system are entirely context-specific, the Risk Management Framework will need to account for the fact that the appropriate mechanisms for promoting trustworthiness will vary depending on the nature of the particular use case. Thus, it is important not to define the underlying concepts in an overly-prescriptive manner. NIST should instead seek to establish a common frame of reference by which specific risks can be identified and communicated.

* * * * *

We strongly support the NIST's effort to develop an AI Standards Engagement Plan and appreciate this opportunity to provide our perspective.

Sincerely,

Christian Troncoso
Director, Policy

20 F Street, NW, Suite 800
Washington, DC 20001

P 202-872-5500
W bsa.org