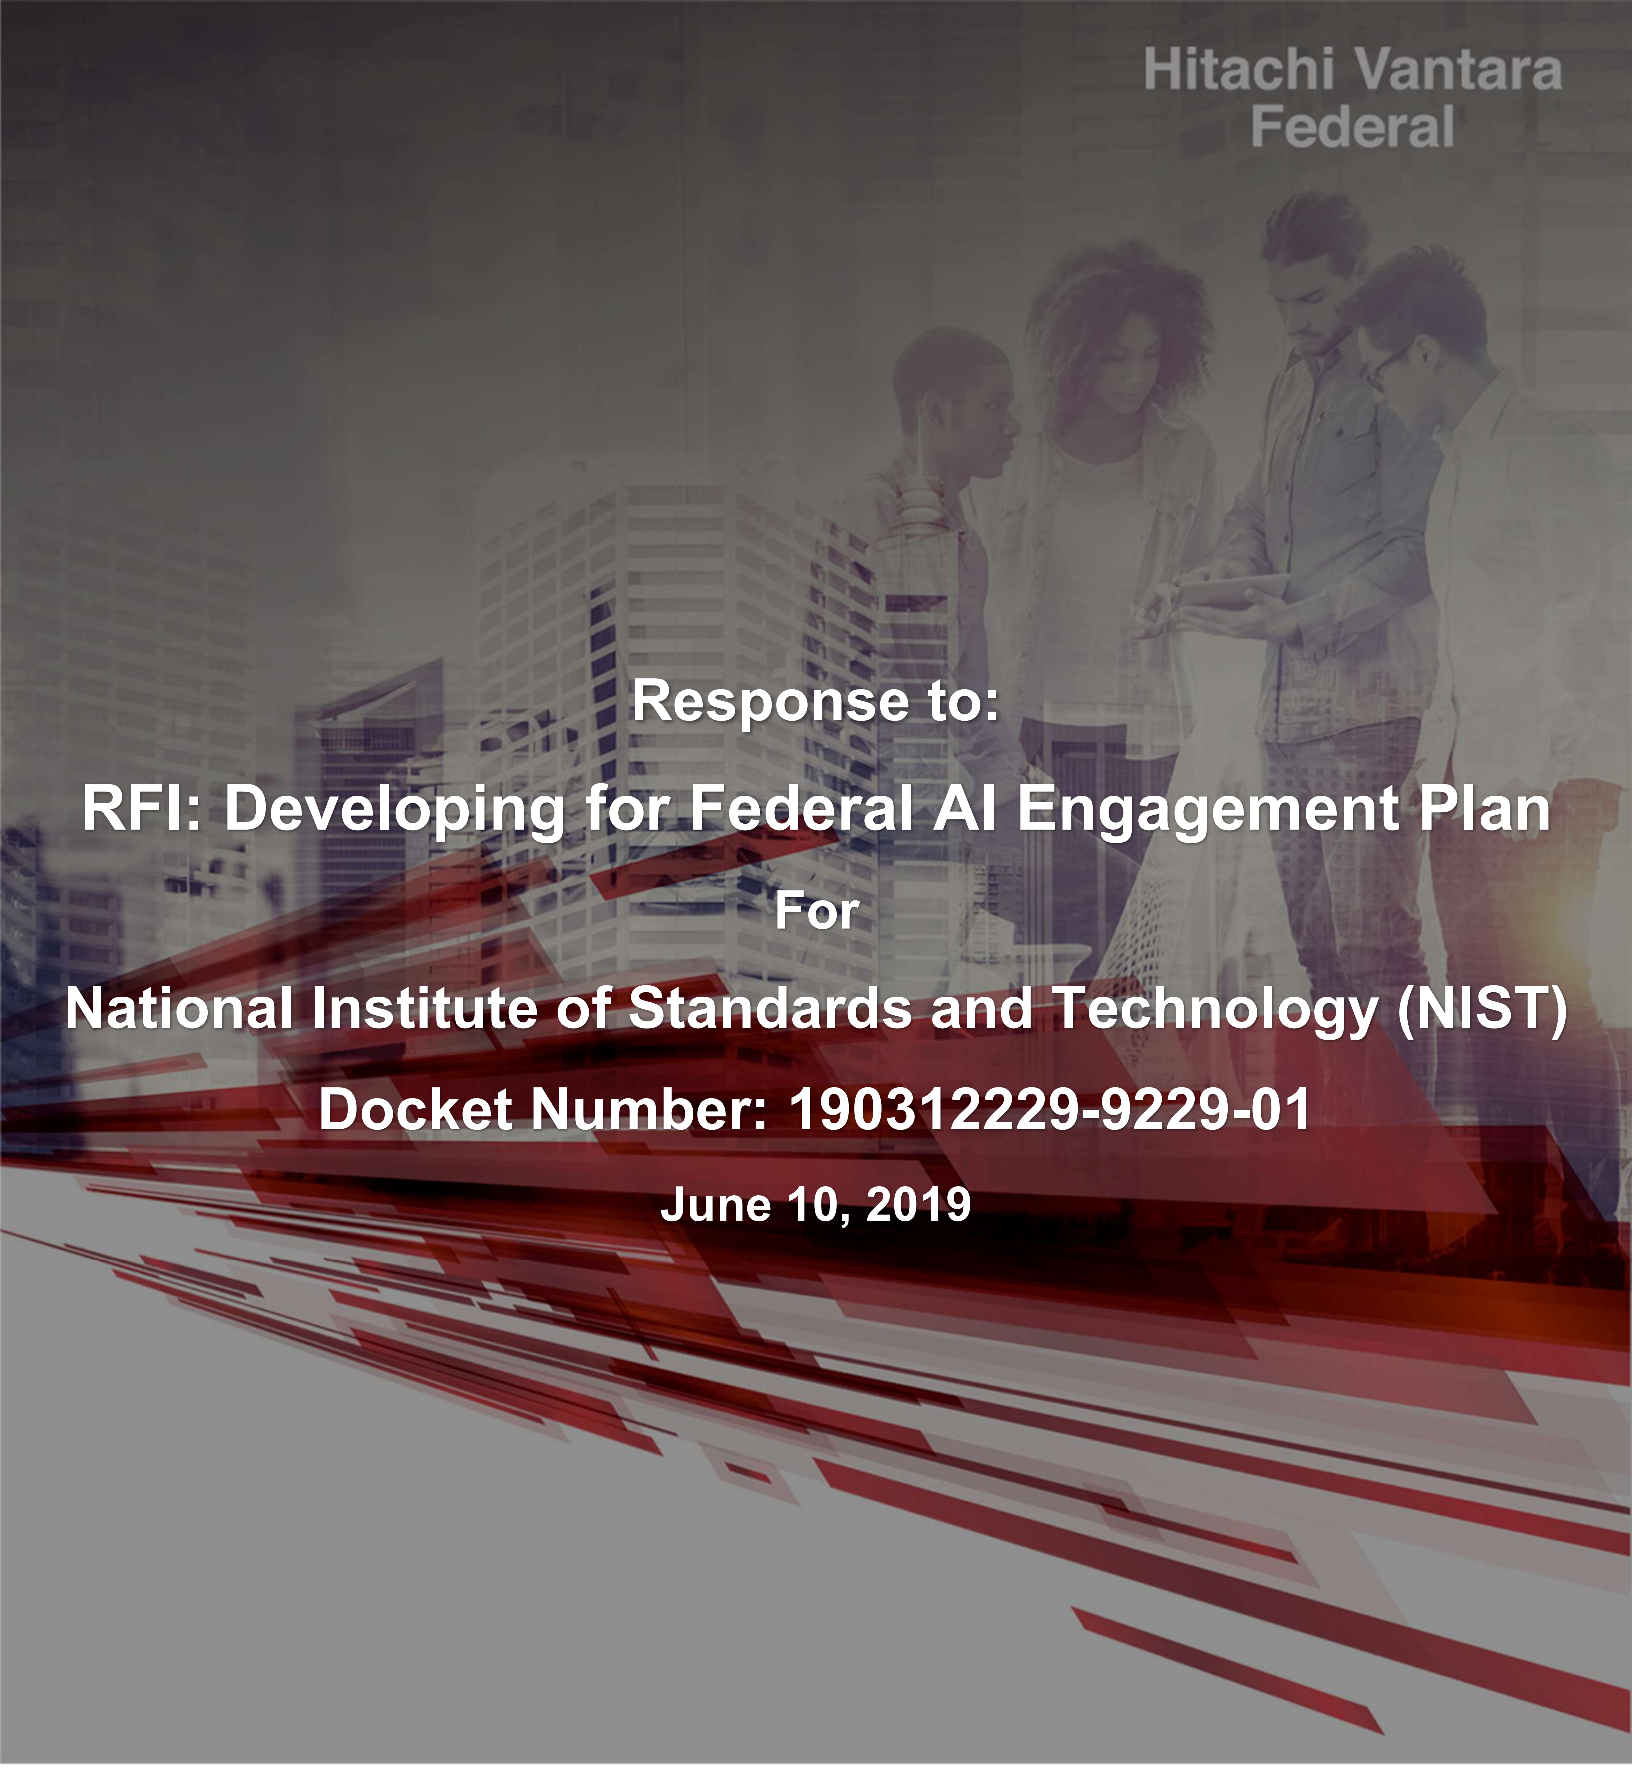

Hitachi Vantara
Federal

Response to:

RFI: Developing for Federal AI Engagement Plan

For

National Institute of Standards and Technology (NIST)

Docket Number: 190312229-9229-01

June 10, 2019

Table of Contents

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

May 31, 2019

Subject: *RFI: Developing a Federal AI Standards Engagement Plan*
Docket Number: 19031229-9229-01
Document Number: 2019-08818

Dear Sir or Madam,

Hitachi Vantara Federal (HVF) is pleased to respond with information pertaining to the National Institute for Standards and Technology's (NIST) Request for Information (RFI) for RFI: Developing a Federal AI Standards Engagement Plan, Docket Number 19031229-9229-01. As you will note in our response to each of the RFI questions, HVF understands the valuable partnership between Government and Industry to establish U.S. Government standards that are required through Executive Order #13859 – Maintaining American Leadership in Artificial Intelligence. We look forward to meeting with the appropriate technical teams and working groups to further discuss HVF's involvement and assistance with setting Artificial Intelligence standards that would support NIST's objectives.

Should you have questions regarding our response, please contact Mr. Gary Hix, Director of Engineering, Federal at 703-772-6971 or via written email at gary.hix@hitachivantarafederal.com.

Sincerely,

Gary Hix
Director of Engineering, Federal
703-772-6971
gary.hix@hitachivantarafederal.com

A. HITACHI VANTARA FEDERAL CORPORATE INFORMATION

1. Description Data	
a. Company Name	Hitachi Vantara Federal
b. Mailing Address and Website	11921 Freedom Drive, Suite 900, Reston, VA, 20190, USA www.hitachivantarafederal.com
c. Telephone Number	703 787-2900
d. Physical Location	11921 Freedom Drive, Suite 900, Reston, VA, 20190, USA
e. Applicable Business Category under U.S. Government Social-Economic Programs and Preference	Large Business
f. Data Universal Numbering System (DUNS) Number	07-875-6944
g. Commercial and Government Entity (CAGE) Code and/or Tax Identification Number	6V5K1
h. Business Size Standard and Applicable NAICS code(s)	334111, 334112, 423430, 511210, 811212

2. Point of Contact	
a. Name	Gary Hix
b. Title	Director of Engineering, Federal
c. Responsibility	- Pre-Sales Solution Development - Product Development in Support of Mission Programs
d. Telephone Number	703-772-6971
e. Email Address	gary.hix@HitachiVantaraFederal.com

1.1 Company Experience/Capabilities

As a wholly-owned subsidiary of Hitachi Vantara, Hitachi Vantara Federal (HVF) helps data-driven government leaders find and use the value in their data, to innovate intelligently and reach outcomes that matter. We combine technology, intellectual property and industry knowledge to deliver data-managing solutions to help the U.S. Government innovate and lower the costs of operations. As a mission partner, HVF elevates the end-user’s innovation advantage by combining deep information technology (IT), operational technology (OT) and domain expertise. HVF has been providing data management solutions and data resiliency since 2013. HVF is subject to Foreign Ownership, Control or Influence (FOCI), which is mitigated under a Proxy Agreement (PROXY) with the Defense Security Service (DSS). HVF operates under a Department of Defense Top Secret facility clearance (FCL) with all Key Management cleared at the level of the FCL. Hitachi Vantara has been a leader in IT for over 50 years—bringing IT applications, analytics, content, cloud, and infrastructure solutions to market that have transformed the way enterprises do business. Hitachi Vantara gives customers a powerful, collaborative partner in data.

1.2 Executive Summary RFI: Developing a Federal AI Standards Engagement Plan

There is a need for development and publishing of AI standards. In most organizations, the development of generic AI standards is still under design, and development will follow. Different groups working in AI within the organization might have guidelines that are expected to be followed by the team, but typically there are no formal standards. As the standards get developed there would be a need for both sector specific and cross-sector standards.

From a technical perspective there are several components to think about. At a high-level, think of AI as an algorithm ingesting data and outcome being used in an application:

1. In terms of data, there are existing standards (such as GDPR, CCPA, HIPPA, etc.). However, there might be a need to review them considering the application that uses the data.

2. In terms of algorithms, there are different variants of open source implementations (such as Tensorflow) but they are not yet standardized. Many people use these as black boxes and there is need to have "standardized" implementation of popular AI libraries. Since the field is rapidly evolving, there will be open source, but it will be good to have standard implementations that can be relied on. These standardized implementations will be used across sectors.

3. The most challenging is to define standards for applications, where more amorphous attributes need to be defined. Many such concepts are being discussed in the community such as "Explainable AI", "Safe AI", "Ethical AI", "Privacy Preserving AI", "Fair AI", etc. These standardizations will be both cross sector and sector specific. For example, applications that use data about individuals will need adhere to standards of "Fair AI", "Privacy Preserving AI", etc. A good example in this category would be stipulation regarding similar accuracy across different races for facial recognition. On the other hand, applications that use environmental data and are deployed directly in the physical world such as autonomous driving will need to be adhere to standards such as "Safe AI." From this perspective:

   - There is a need to identify and define properties that are desirable for an AI based application to have

   - There could be standardized data sets for some key cross sector tasks, such as face recognition, and all algorithms would need to be trained on. The data sets will need to ensure that they are representative (for example, racially representative) and remove bias.

The United States Government needs to, as one of the world’s largest spenders on Research and Development, establish standards in the field of AI. This would not only ensure continued leadership, but also that the field develops in a socially responsible way. As AI driven applications become increasingly integrated into our life, for continued public confidence and acceptance, the public and private sector needs to work together to establish and adhere to standards.

B. RESPONSES TO REQUEST FOR INFORMATION QUESTIONS & RESPONSES

I. AI Technical Standards and Related Tools Development: Status and Plans – Questions and Responses - 1 - 8

1. AI Technical Standards and Tools have been developed, and the developed organization, including the aspects of AI these stands and tools address, and whether they address sector-specific needs or are cross-sector in nature;

Currently, there are two main working groups on International Standards relating to Artificial Intelligence (AI). What neither working group or standards body focuses on is “responsible deployment” of those AI tools or safety specification. The name and mission, or scope, of each organization are provided below:

- Institute of Electrical and Electronics Engineers (IEEE) – the IEEE is mainly concerned with AI Ethical Concerns. The IEEE Global Initiative’s mission is, “To ensure every stakeholder involved in the design and development of autonomous and intelligent systems is educated, trained, and empowered to prioritize ethical considerations so that these technologies are advanced for the benefit of humanity.”
- International Organization for Standardization (ISO) – the ISO’s Technical Committee on Information Technology for Learning, Education, and Training’s scope is:
    - Standardization in the field of information technologies for learning, education, and training to support individuals, groups, or organizations, and to enable interoperability and reusability of resources and tool

As AI has continued to grow and evolve in recent years, Hitachi and the global AI community have seen many changes. For example, AI technical methods and processes have become more mainstream, which has also led AI to be used for activities such as more efficient information espionage.

AI methods and processes that are cross-sector in nature, opposed to sector specific, have become the collective leading practice in the AI marketplace. These cross-sector AI applications address basic intellectual abilities such as computer vision, facial recognition, speech recognition, and natural language processing.

Conversely, many sector specific use cases, code bases, and applications are evolving to cross-sector as entities commercialize AI services as part of a broader service enhancement or stand-alone products. Examples of sector-specific, or stand-alone products and tools, would be predictive maintenance, risk and asset evaluation, and stock market analysis.

An example of Hitachi’s Ai Development would be specific to Hitachi Vantara. Hitachi Vantara Labs (formerly Pentaho Labs) has created and updates an experimental plugin for Pentaho Data Integration (PDI) called “Plugin Machine Intelligence” (PMI) as an opensource and freely available to use plugin via the Pentaho Marketplace.

The main goal of PMI is to make the flow of Machine Learning, including Deep Learning, easier to use and deploy by combining the extract-transform-load (ETL) and data integration tool, PDI, and Machine Learning into one powerful data science toolset.

This is accomplished by using popular industry standard Machine Learning libraries from popular Machine Learning development environments. Currently, the Machine Learning libraries integrated into PMI are,

- Python Scikit-learn
- R – Machine Learning in R (MLR)
- Apache Spark Machine Learning library – MLlib
- Waikato Environment for Knowledge Analysis – Weka Machine Learning
- Deep Learning for java – DL4j

These libraries are integrated and used in a consistent and uniform way, using the same graphical user interface for all supported Machine Learning algorithms. This makes supporting Machine Learning easier and more consistent. Also, the performance metrics used to measure the accuracy of the algorithms are calculated within the PMI framework to be uniform across all algorithms. Therefore, libraries are consistent when comparing performance in any combination of library and algorithm.

The current scheme supported in PMI is for Supervised Machine Learning.

2. Reliable Sources of Information about the availability and use of AI Technical Standards and Tools;

Traditionally, the United States is hands-off and will let the capital markets and enacted legislation drive the appropriate changes. However, these changes only come after a significant negative event such as the 2008 financial collapse. A more proactive approach would be to take a Leadership position regarding AI standards that results in more guiding, opposed to compliance and enforcement. The latter would naturally come from federal agencies as required (e.g., consumer data privacy and FTC/FCC regulation).

There are multiple communities for AI standards and tools which are tailored to their respective communities. Examples include the IEEE Computer Society, which provides advance perspectives in mathematics and computation models. Conversely, Hitachi Ltd, Google, Microsoft, IBM, and IBM Research have published standards to use for their commercial AI services with business partners and system integrators.

Here is a list of blogs from the Hitachi Vantara Community – Pentaho Community web site:

- PMI announcement - https://community.hitachivantara.com/community/products-and-solutions/pentaho/blog/2018/03/06/operationalizing-machine-learning
- Machine Learning Model Management - https://community.hitachivantara.com/community/products-and-solutions/pentaho/blog/2018/03/06/4-steps-to-machine-learning-model-management
- Deep Learning - https://community.hitachivantara.com/community/products-and-solutions/pentaho/blog/2018/11/14/pmi-v14-with-deep-learning-is-here
- Artificial Intelligence with Pentaho - https://community.hitachivantara.com/community/products-and-solutions/pentaho/blog/2018/10/16/deep-learning-coming-to-pentaho

Lastly, here are the leading conferences in the field of AI:

- Institute of Modern Language Research (IMLR)
- International Conference on Learning Representations (ICLR)
- Neural Information Processing Systems (NIPS)
- International Conference on Machine Learning (ICML)
- Association of Computational Linguistics (ACI)
- Empirical Methods in Natural Language Processing
- Genetic and Evolutionary Computation Conference (GECCO)
- Special Interest Group on Computer GRAPHics and Interactive Techniques (SIGGRAPH)
- Computer Vision and Pattern Recognition (CVPR)

3. Reliable Sources of Information about the availability and use of AI Technical Standards and Tools;

This is a similar question to one that was raised over a decade ago in the medical field regarding genetic engineering. The following framework should be considered:

- Leveraging Academy & Government to set the appropriate behavior thresholds and legal consequences.
- Establishing a community of interest with the major stakeholders who develop and sell AI capabilities to provide education, solicit support, and provide self-policing of practices.
- Partnering with the broader international community (e.g., EU Commission on AI) to provide reasonable consistency in standards, technology and technology availability.

The broader challenge this group needs to address is the growing “black hat” community, which has its own standards and is in the business of developing and selling algorithms. Many times these “black hat” communities purposes are not in society’s best interest.

Additionally, the digital divide will continue to grow as technology moves faster than government’s ability to keep up. The workflow for AI use cases implementation can be broken down into 6 main parts according to CRoss-Industry Standard Process for Data Mining (CRISP-DM) methodology:

- Business Understanding
- Data Understanding
- Data Preparation
- Data Modelling
- Evaluation
- Deployment

The set of tools and standards should be based on the stages for this methodology. One tool can cover several stages of CRISP-DM (e.g. data preparation and deployment), however, one set of tools and standards might be required for business and data understanding. This part is highly important in setting up AI projects. At this stage one can also classify to which subfields of AI this use case can be related to. That can further determine a set of tools needed for modelling. For instance, AI models and tools can be different for speech recognition and predictive maintenance. It is important to highlight that business/use case and data understanding are highly important. The approach of looking into a data and trying to find some interesting (“knowledge discovery”) insights and correlations proved to be inefficient by Hitachi and other organizations. Therefore, standards and tools should follow CRISP-DM.

At Hitachi Vantara Labs, we are watching the continued development of eXplainable Artificial Intelligence (XAI), sometimes called Transparent AI, to add model forensics to the Plug-In Machine Intelligence (PMI) framework. XAI is a DARPA program and Hitachi R&D is heavily studying this capability. The XAI program would allow regulator and compliance auditors to analyze the decisions of Machine Learning model and the primary influences that impact their outcomes.

Once again, standards establish “guard rails” in place, which provides direction and protection from going too far off the path. However, once the guard rails are up, they require ongoing maintenance by the respective federal organization as they do business.

4. Reliable Sources of Information about the availability and use of AI Technical Standards and Tools;

AI is a rapidly developing field, and while some applications are cross-sector, others can be more niche and sector-specific. For example, computer vision is a very broad field. It includes object recognition for self-driving cars, facial recognition, poses, and actions recognition. Another example is natural language processing, which includes text interpretation and creation of assistants which will be able to understand human languages without need for specific trigger words or commands, as well as to speak human language.

The two areas are not sector specific because they are focusing on basic intellectual abilities – vision and speech interpretation. There are other areas, such as recommendation and predictive maintenance, which are more in demand in retail, e-commerce sector and in manufacturing respectively. As the AI field continues to develop, there will be more niche applications. So far, in addition to the majority of universities conducting research in the AI and Machine learning areas, main organizations actively developing AI tools and techniques are:

- Google
- Facebook
- OpenAI
- Tesla

Furthermore, society at large should realize that AI will eliminate job functions just like the industrial and digital revolutions, only faster. In addition to the many companies that are designing and proto-typing autonomous cars, AI's impact on the traditional taxis made them digitally obsolete in less than a decade, after being mainstream for more than 80 years. Assuming that AI keeps progressing at the rate it has in recent years; driverless car fleets could make services like Uber and Lyft obsolete in even less time.

The IEEE and ISO Approach to AI Standards

These changes in the world that have been driven by AI, while convenient for the most part, could have unintended consequences. The need for the United States to develop AI standards, and maintain is technology position, is a "must-have" to ensure all major AI companies operate to the same criteria. Currently, the leading standards organizations are both cross-sector in nature and are:

- Institute of Electrical Engineers and Electronics Engineers (IEEE)
- International Organization for Standardization (ISO)

IEEE's critical question has traditionally been, "What do we wish to prioritize in the creation of technology in the algorithmic age?" Simply acting, and then later thinking about the potentially negative sides of what we have done, may take humanity to places where nobody wants to go, at least consciously. In the case of AI, it is only by defining the deep ethical considerations we wish to address as a society before we create technology that we can best align with people's values who use it and avoid negative unintended consequences.

To help in this process of societal definition, the IEEE Standards Association (IEEE-SA) launched The IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems (AI/AS) in April of 2016. This was done for two key reasons:

- To move beyond both the fear and the uncritical admiration regarding autonomous and intelligent technologies.
- To show that aligning technology with ethical values will help advance innovation with these new tools, while diminishing fear in the process.

In an effort to pragmatically address specific ethical issues in AI/AS, the IEEE Global Initiative on AI Ethics was tasked with two primary deliverables. The first was the creation of *Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems*. Given the sensitive and complex nature of the matter, we chose an iterative approach. The first version was created by over 100 AI/Ethics thought leaders from the IEEE Global Initiative and contains over 80 pragmatic Issues and Candidate Recommendations for technologists to use in their work today to create a positive future. It was published as an explicit open call for opinions and feedback via our submission guidelines process, to help inform the creation of the second version.

Along with creating and evolving Ethically Aligned Design, members of the IEEE Global Initiative are encouraged as their second deliverable to recommend standardization projects to IEEE-SA based on their work. Here are titles for each of these standardization projects, and more information is available via the links included. Along with the projects listed below, the IEEE Global Initiative recently submitted three more standardization ideas for consideration:

- IEEE P7000™: Model Process for Addressing Ethical Concerns During System Design (Working Group already in process)
- IEEE P7001™: Transparency of Autonomous Systems (Working Group already in process)
- IEEE P7002™: Data Privacy Process (Working Group already in process)
- IEEE P7003™: Algorithmic Bias Considerations (Project has been approved as a Working Group. More information will be available on The IEEE Global Initiative's website soon.)

From what has been shared, IEEE P7000™ is the first Standard in the history of IEEE that is directly focused on the implementation of applied ethical methodologies to technology. To be clear, this is not to infer that engineers and technologists have not always focused on prioritizing sound ethical practices in the creation of their work. Likewise, IEEE has had a professional code of ethics guiding its work and membership for decades.

But as the purpose of the IEEE Global Initiative states, our goal is to ensure every technologist is educated, trained, and empowered to explicitly prioritize ethical considerations in the design and development of autonomous and intelligent systems. By this we mean that along with a code of ethics providing direction for member behavior, technologists in the algorithmic era need to use methodologies that provide more rigorous due diligence regarding the values of stakeholders and end users than they may be using today. Examples of these methodologies along these lines include Value Sensitive Design and Responsible Research and Innovation (RRI). As the RRI Tools website notes, this means, "involving society in science and innovation 'very upstream' in the processes of R&I to align its outcomes with the values of society."

By creating the IEEE P7000™ series of Standards based on Ethically Aligned Design, our goal is to institutionalize the rigors of this "upstream analysis" to further aid the scientists and engineers involved in the creation of the intelligent, autonomous, and other emerging technologies driving our human future.

The launch of the IEEE Global Initiative and the subsequent development of the IEEE P7000™ family of Standards are pushing the boundaries of the art of consensus that builds into key facets of the AI/AS ecosystem. These activities contribute to the IEEE TechEthics™ program, which is a broader effort being launched at IEEE to foster an open, broad, and inclusive conversation about ethics in technology. It is because of these efforts of the entire organization to prioritize ethics that we can collectively create a societal standard for our future that truly advances technology for humanity, and for a healthy and innovative future.

5. Any supporting roadmaps or similar documents about plans for developing AI Technical Standards and Tools;

AI is a tool. Its primary use today is to help humans deal with information volumes that exceed our capacity to process. The result is that we have the ability crash harder and faster when it goes wrong. A prime example of this are the Stock Markets and all the automated trading that occurs each day. Over the past 10 years, Wall Street has had to create multiple new gates to keep automation from burying the market during negative events.

Most companies innovating AI capabilities for themselves or for others (as system integrators) have roadmaps and view this information highly confidential. These roadmaps provide a competitive edge in competing for AI-related sales and services, and in protecting Intellectual Property as it relates to future product development.

Hitachi Vantara Labs is looking to add additional libraries to our existing framework like Keras/Tensorflow, as well as potentially expanding to include an Unsupervised Machine Learning framework.

6. Whether the need for AI Technical Standards and related Tools is being met in a timely way by organizations;

In the AI field, there is a widely used term called "AI-Readiness." AI-Readiness simply characterizes if the organization is capable of incorporating AI practices into its daily workflow. Currently, only few companies call themselves AI-ready. The heart of the AI-first organization should be AI algorithms and they should be incorporated in every business process. In addition, the focus in the organization should be on to the performance and accuracy of AI-models.

Some of organizations may start to employ AI in some specific and usually very narrow tasks, but there is still a very long way ahead to become AI-ready company. One of the concerns on this path is AI-maturity, and there is much trial and error in the AI field. Sometimes companies deploy models into production that are not yet optimal. It can be an informed decision because AI needs to gather data and, in this case, the company is prepared for the temporary performance drawbacks. But usually it is because organizations are too eager to start deploying AI and models that are not production ready. In this case, the company can end up with a more significant performance draw back and even worse – irresponsible usage of AI can lead to legal consequences.

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

Both will occur and both require standards. For example, automated piloting of large mechanical vehicles will always require specific standards more associated with the application of AI. Clinical trials may be prescriptive as to the acceptable AI used to ensure consistent measurable outcomes over time. Web site Search Engine Optimizers (SEOs) will look to more cross-sector standards related to specific AI algorithms. Currently all the above rely on the market of ideas and self-regulation to determine what is acceptable or not.

Additionally, cross sector standards can occur for “high-level” AI task such as data preparation, models’ management, typical architectures of AI solutions, etc. However, sector specific AI technical standards can take place depending on the application of AI techniques. Such application of AI in health care will require a different set of algorithm evaluation metrics and procedures compared to recommendation engines on websites like Amazon. In addition, AI applications in transportation and industry (predictive maintenance, processes optimization) largely rely on domain specific knowledge. Therefore, technical AI standards might differ for those sectors-specific applications.

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

Key Performance Indicators (KPIs) for both sector-specific and cross-sector applications should be related to performance, not technology. The means that a different set of KPIs would be required to look at AI usage. This approach must be balanced with protection of intellectual property rights as well. However, it is recommended that as much transparency as possible be placed on the AI used by any government agency.

Transparency in AI code should be a top priority as well. The reasons for this priority directly influence the other KPI's: accuracy, security, privacy, robustness and performance:

- Accuracy – without public review it is likely that the AI used will have bias built in. It is a common problem, and without large scale review it is very hard to detect, much less remove, that bias. If the AI contains bias it is likely that the results will be skewed, and we will miss the key value trying to be achieved.
- Security – The security of any government used system is tested in any way possible on a constant basis. The logic that public review allows an attacker to find security flaws is faulty due to the laws of human nature that good actors outweigh bad ones by a significant margin. Giving good actors access to the code will allow for a much more robust security assessment to be made. Those same people will also suggest how to remedy any security flaws found.
- Privacy – public review allows for the privacy challenges to be discussed and mitigated before millions are spent on a project that will have to be scrapped after judicial review.
- Robustness and Performance – both will be tested, evaluated and improved through public review, as that is the first and foremost value from such a process.

Lastly, one of the more pressing areas that needs to be addressed will be in the regulations and compliance of AI models. In one responder's opinion, this is a segment, lineage, that needs to be added to the current concept of MLMM as described in the blog in Question #2.

- Monitoring accuracies in a uniform way across all algorithms and implementations in order to compare them in a non-biased way should be standardized. There are a number of standard metrics used to measure a model's accuracy in use today. These are a list of the evaluation metrics output from all algorithm implementations in the Plug-In Machine Intelligence (PMI) Framework:

    - Scheme name
    - Scheme options
    - Evaluation mode
    - Unclassified instances
    - Correctly classified instances
    - Incorrectly classified instances
    - Percent correct
    - Percent incorrect
    - Mean absolute error
    - Root mean squared error
    - Relative absolute error
    - Root relative squared error
    - Total number of instances
    - Kappa statistics
    - *class* TP rate
    - *class* FP rate
    - *class* Precision

- *class* Recall
- *class* F-measure
- *class* MCC
- *class* ROC area
- *class* PRC area

II. Defining and Achieving U.S. AI Technical Standards Leadership – Questions and Responses – 9 - 12

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

The US should desire to be the world leader in both invention and innovation, as both are core to our DNA as a nation. Not having well defined standards in place does not foster either invention nor innovation and leaves the US trailing behind the EU, not to mention potentially antagonistic governments like China, Iran or Russia.

Additionally, AI is at or near its peak in popularity. That creates confusion when AI projects are being implemented as well as mistakes in overall AI strategy. Development of AI technical standards and related tools is important for moving from being popular, to being productive, and having successful implementations. Effectiveness in AI technical standards development can be achieved by focusing on specific areas of AI workflow (data preparation, models deployment etc.) and accounting for specific AI subcategories (image recognition, autonomous driving and other).

Another important aspect is ethics in AI solutions, as AI algorithms should not discriminate based on gender, national origin, race and other factors. Development of AI technical standards can effectively address points raised above. The need for AI technical standards and related tools development is urgent, as it will accelerate AI adoption across many sectors benefiting society, such as transportation, retail, health care, climate monitoring, agriculture, etc.

Currently, areas with bigger pressure of regulatory requirements are developing their own regulations regarding AI. For example, compliance and risk assessment departments in many banks are using machine learning techniques for better credit risk scoring and fraud detection. Application of these techniques require a lot of paperwork from the bank side, though, to justify the use of AI. Having a single technical standard for AI being deployed would let the organization use it more freely and in a more controllable way. At the same time, very big areas of AI usage are now not properly regulated or regulated, because certain standards were not built for AI, like facial recognition, self-driving cars, speech recognition and speech synthesis.

There are clearly identifiable leaders in each field of the AI applications. These companies should be in collaboration with federal agencies to create “AI Technical Standards” for specific industries.

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

The AI community is a global community. Trying to define where the US is a leader or follower is not as simple as a clear report. Most companies, regardless of HQ location, utilize a global workforce who creates and influences the innovation of the corporation. It is probably more effective to look at who the leaders in the world are, based on who owns the IP. Those leaders would be: United States, Israel, and Japan. China and Russia are less well defined but should be considered in the top 5 leaders in AI study, with China also being a leader in thought and execution. China is ahead of the United States in defining their AI standards and have already deployed large scale AI projects in the public sector. It would be problematic to start an AI race with China or Russia, as those nations have very different views of privacy, security and liberty than the US.

Comparatively, the European Union is approximately 5+ years ahead of the United States in Data Privacy and Protection standards. These standards set the foundation for more advanced topics around how data is processed and used. With California leading the US (CCPA - California Consumer Protection Act), it is anticipated 35 states will have enforceable laws creating a patchwork of varying requirements. Private sector estimates do not expect to have a federal law in place until at least 2022. With the lack of Federal thought-leadership, many US companies are relocating businesses between states, or outside of the US. Likewise, major companies leading the US around AI (e.g., Google, Facebook, etc.) have been specifically targeted by the EU for failing to govern data since the US favors private sector self-regulation.

In an effort to remain competitive with the European Union, and the larger Global Community, there have been other entities established in an effort to foster standardization across Artificial Intelligence (AI). One example of this is the Linux Foundation Artificial Intelligence group (LFAI). LFAI is an umbrella foundation of the Linux Foundation that supports open source innovation in artificial intelligence (AI), machine learning (ML), and deep learning (DL). From LF AI’s webpage, it the About section, it reads “LF AI was created to support open source AI, ML, and DL, and to create a sustainable open source AI ecosystem that makes it easy to create AI products and services using open source technologies. We foster collaboration under a neutral environment with an open governance in support of the harmonization and acceleration of open source technical projects.”

Lastly, the OPEN Data Act was passed in Jan 2019 which is a huge step in enabling access to unclassified government data for the citizens. Data.gov, code.gov, search.gov-akin websites enable the objectives of this Act. Transparency in government spending has been enabled via usaspending.gov. The mentioned sites are actively addressing providing access to the data. Similar platforms for the knowledge sharing of the developed solutions, specifically analytics and AI-based solutions, is missing.

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

Every year, the landscape of the AI technologies is growing extremely fast and becoming increasingly complex. More and more start-ups and firms continue to appear on the AI arena. It can be quite challenging for the organizations to understand what technology or AI-platform to use. Problems with choosing the right technology, as well as keeping up with the pace of development in the AI-field, can be a serious obstacle on the AI-adoption path.

One of the possible paths of standardization of the whole AI field can be separation of this big area into sub-areas, and further development of the standards for each of the smaller pieces. For example, one of the essential steps for the AI model development is data blending and data preparation. It is possible to develop a standard for data cleaning and preparation based on best practices known from both academia and the industry. Another part of AI deployment process is AI model management, which also can be standardized based on best practices.

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

The US should incorporate a good mix of Government, Private and Public groups (those that represent the interests of the citizens). This will ensure that what is proposed has good representation of the needs of the different stakeholders, and the public groups could provide perspective to help prevent future legal challenges. Additionally, the United States should provide investment and support for the development of AI models. This should also include promoting and protecting general use licensing, while reforming and updating current patent laws.

As part of maintaining technical standards and development, appropriate public and private sector incentives should also be implemented to drive adoption and improve US capabilities in:

- Health Sciences and Welfare – improve operational use of medical facilities to drive down healthcare cost (Medicare, Medicaid)
- Transportation – improve operational efficiencies (e.g., autonomous trucking)
- Education – Continue to refine and extend STEM research in K-12 to enable our next generation to co-exist and be responsible for utilizing AI capabilities/services
- Workforce Transition – Assist the workforce in something the transition from occupations that become obsolete from AI technology

Investing in AI talent and creating opportunities is vital. There are many US universities that provide great education in STEM and raise AI talent. However, current immigration policy and strict work visa conditions can hamper AI talent retention in US and therefore leadership in AI technical standards development. If these policies do not improve, other places in the world with softer immigration policy can become more attractive for international AI talent and can dominate development of AI technical standards.

III. Prioritizing Federal Government Engagement in AI Standardization – Questions and Responses – 13 - 18

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

Many, if not all, of the federal AI standards and tools are applicable to a larger US, and ROW (rest of world) audience. The Federal government needs the same standards with regards to ethical and responsibility challenges, security procedures and test criteria as the rest of the world. AI is evolving too quickly to place technical requirements in the standards, and so the whole world is on a level playing field regarding need.

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

The United States should take a leading role in the development, adoption, and creation of these standards to declare to the world:

- AI is a critical tool in the protection of the people's safety.
- AI has the capability for gross misuse, and the US government will not entertain the use of AI for bad actors.
- We can provide for the common safety and protect liberty at the same time, and the US is representing the citizens of the US and the world in providing standards for both security and liberty and guidelines to determine what the right level is of each.

Additionally, there are many Departments and Agencies (D/As) within the United States Government that could benefit from using AI for their benefit, as well as State and Local Governments. Examples are included but not limited to:

- Department of Transportation: Not only should standards be developed and implemented as it relates to Autonomous Vehicles, but AI could be used to reduce manufacturers costs compiling information on vehicle safety and crash standards.

- The U.S. Department of Health and Human Services: AI could be used to make operational improvements to health care, resulting in reduced cost to not only the Government, but the consumer.
- The U.S. Department of Labor, specifically Occupational Safety and Health Administration (OSHA): For high risk areas and occupations, AI could be used to reduce to the risk to human life by providing results to test criteria.
- The U.S. Food and Drug Administration: Improvements recommended by or implemented through AI could assist in automation and improve the oversight of food handling and management.
- Data Privacy & Governance issues that are different for each agency: Define what and how data may be used to enable AI capabilities as well as in the use (direct or 3rd party resell) of data generated from AI algorithms. For example, the U.S. Department of Health and Human Services will have different Data Privacy laws and issues than many other agencies, due to the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule.

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

The Federal Government should focus on core parts of AI workflow first, such as task definition, data preparation (accessing data sources, data processing, data blending, data cleansing), feature engineering, models selection, management, and deployment. These steps are at the core of most of AI applications regardless of sectors.

Additionally, a Risk/Benefit assessment should be performed to determine:

- Where is the Federal Government most at risk (e.g., behind in specific standards, sector automation, etc.)?
- Do private sector focus areas align with Federal Government needs?
- Where are the greatest benefit drivers to the Federal Government to drive or establish AI standards, tools and capabilities to improve government capabilities and the constituents that serve them

Some of the sectors, like banking regulation and compliance, already have their own standards of AI usage, based on sector-specific best practices. Adopting currently existing standards could be a start for broader standards development for related sectors.

16. 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

Accounting standards, such as GAAP, are private sector led and adopted by the Federal government. The established standards universally apply to all industries, organizations, commercial activities. Further, the established standards serve a wide variety of derived purposes both commercial, regulatory and informational. e.g. capital funding, legal and regulatory adherence, tax obligation, economic reporting, industry surveys, market research, etc.

The focus here is on usage, transparency and outcomes. The standards seek to identify:

- appropriate use of accounting techniques in both context and manner of employment
- documentation and information required to judge adherence to the standard
- reasonable and acceptable outcomes for techniques employed

AI requires a similar set of standards that are focus on application regardless of industry or purpose.

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

In the telecommunications arena, the United States Government plays an active role in the standards arena. Examples of this would include, but not be limited to, the following:

1. Standards that originated from the National Telephony Security Working Group (NTSWG). The NTSWG became that Telephone Security Group and the purpose of the standard is as follows:

    - This standard was prepared by the Telephone Security Group (TSG). The charter members of the TSG are: Department of the Air Force, Department of the Army. Central Intelligence Agency. Defense Intelligence Agency Department of Energy, Federal Bureau of Investigation, Department of the Navy. National Security Agency, US Secret Service, and Department of State. The TSG is the primary technical and policy resource in the US Intelligence Community for all aspects of the TSCM (technical surveillance countermeasures) program involving telephone systems. The TSG standards contain guidance for providing on-hook security to telephone systems in areas where sensitive government information is discussed. Implementation of TSG standards neither prevents the application of more stringent requirements nor satisfies the requirements of other security programs such as TEMPEST. COMSEC, or OPSEC. This standard establishes requirements for planning, installing, maintaining, and managing a computerized telephone system (CTS). The requirements established in this standard are necessary in order to achieve on-hook audio security for computerized telephones located in sensitive discussion areas. For a CTS conforming to this standard, all protected on-hook telephones will be completely isolated from all transmission media and wires that are physically unprotected. This standard requires that the isolation for most telephones be achieved in the CTS itself.

- The TSG-6 Approved Telephones, which allow for unclassified communications meet Federal standards regarding telecommunications inside:
    - Protecting US Embassies and Consulates
    - Protecting Intelligence Community Facilities Domestic and Abroad
    - Technical Threat Assessment and Technical Surveillance Countermeasures

2. Joint Interoperability Test Command (JITC) - With the vision of being experts in configuration and certification, the Department of Defense's Joint Interoperability Certifier and only non-Service Operational Test Agency for Information Technology (IT)/National Security Systems (NSS) is one of multiple JITC programs. JITC provides risk-based Test, Evaluation & Certification services, tools, and environments to ensure Joint Warfighting IT capabilities are interoperable and support mission needs. While JITC has many programs, the program mainly used by telephony equipment vendors is:

- Joint Interoperability Certifier – Test and Evaluation of the interoperability of IT and NSS is essential to reduce the risks faced by warfighters in the field. JITC is constantly reviewing processes to ensure it is performing tests as efficiently as possible into today's austere environment. JITC's ability to provide outstanding support to DISA and the Warfighter is characterized by several unique features. Although each Service has its own test organizations, JITC has sole responsibility for certifying joint and combined interoperability of all DoD IT and NSS

The National Institute of Standards and Technology, following each of the above models, could perform system security, operational, and configuration evaluations and testing at each AI vendor's site. For an overall interoperability testing model, perhaps adding another JITC Interoperability Testing Program could be used for both Commercial and Government uses with a cost sharing approach to maintain the testing environment.

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

AI technical standards should be taught at universities (as well as open universities platforms e.g. Courses). This could be accomplished through not only Educational, but Research and Development, Grants to all post-high school education levels. By including AI technical standards into curriculum, one can achieve a high level of AI technical standards adoption.

There is a growing need in AI/data science specialists in US and beyond. Many people use online education to gain skills needed for AI professionals. Many people choose AI/computer science as their degrees in university education. If AI technical standards are taught at universities, they will likely to be adopted in practice.

Several metrics can be developed to access how useful AI technical standards in practice and if anything must be changed or adjusted in general or for specific sectors. The metrics can involve: percentage of time spent on data preparation tasks in AI projects, business outcomes of AI solutions (“developed AI model(s) increased company revenue/improved safety/increased transportation reliability by xx%”) , number of AI projects and their impacts implemented per year etc.