28 May 2019

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

Via email: ai_standards@nist.gov

Re: NIST *RFI: Developing a Federal AI Standards Engagement Plan* (Docket Number: 190312229-9229-01)

IEEE-USA is pleased to submit the following information in response to the above-captioned request for information, published in 84 FR 18490 (1 May 2019). We commend NIST for its effort to create a plan for federal engagement in developing the necessary standards and tools to support reliable, robust and trustworthy systems that use AI technologies.

IEEE-USA represents approximately 200,000 engineers, scientists, and allied professionals living and working in the US. Our members work in the new AI-related industries, developing and working with the emerging technologies used in artificial intelligence systems. This expertise provides us with a unique perspective on the benefits of these technologies.

The IEEE Standards Association (IEEE-SA), a globally recognized standards-setting body within IEEE, develops consensus standards through an open process that engages industry and brings together a broad stakeholder community. IEEE standards set specifications and best practices based on current scientific and technological knowledge. IEEE-SA has a portfolio of over 1,250 active standards and over 650 standards under development.

We are pleased to provide input on the following areas, as requested:

1. Current status and plans regarding the availability, use and development of AI technical standards and tools in support of reliable, robust, and trustworthy systems that use AI technologies.
2. Needs and challenges regarding the existence, availability, use, and development of AI standards and tools.
3. The current and potential future role of Federal agencies regarding the existence, availability, use, and development of AI technical standards and tools to meet the nation's needs.

1. Current status and plans regarding the availability, use, and development of AI technical standards and tools in support of reliable, robust, and trustworthy systems that use AI technologies

IEEE is working to engage and grow communities to develop technical standards that address AI and AI systems. We develop these technical standards through an open and transparent process that welcomes all stakeholders. IEEE standards developers work together to codify technical practices that can be applied across sectors; and that developers and regulators can use to support interoperability, safety and the public interest.

IEEE has standards working groups engaged in developing global technical standards for identifying and mitigating algorithmic bias, for documenting and protecting users interests in personal data, for evaluating

reliability of online messaging, for protecting children's personal information online, for assuring the safety of autonomous and intelligent systems, and many other relevant areas.

Noted below for reference are the IEEE AI relevant standards projects in development.

Standard No.	Title	Description	Stakeholders
IEEE P2801	Recommended Practice for the Quality Management of Datasets for Medical Artificial Intelligence	This recommended practice identifies best practices for establishing a quality management system for datasets used for artificial intelligence medical device. It covers a full cycle of dataset management, including items such as but not limited to data collection, transfer, utilization, storage, maintenance and update. This practice recommends a list of critical factors that influence the quality of datasets, such as but not limited to: data sources, data quality, annotation, privacy protection, personnel qualification/training/evaluation, tools, equipment, environment, process control and documentation.	Medical device industry, regulators, medical community and academia.
IEEE P2802	Standard for the Performance and Safety Evaluation of Artificial Intelligence Based Medical Device: Terminology	This standard establishes terminology used in artificial intelligence medical devices, including definitions of fundamental concepts and methodology, that describe the safety, effectiveness, risks and quality management of artificial intelligence medical device. It provides definitions using the following forms, such as but not limited to literal description, equations, tables, figures and legends. The standard also establishes a vocabulary for the development of future standards for artificial intelligence medical device.	Medical device industry, regulators, medical community and academia.

IEEE P3652.1	Guide for Architectural Framework and Application of Federated Machine Learning	Federated learning defines a machine-learning framework that allows for constructing a collective model from data distributed across data owners. This guide provides a blueprint for data usage and model building across organizations--while meeting applicable privacy, security and regulatory requirements. It defines the architectural framework and application guidelines for federated machine learning, including: 1) description and definition of federated learning; 2) the types of federated learning, and the application scenarios to which each type applies; 3) performance evaluation of federated learning; and 4) associated regulatory requirements.	Organizations that federate data to build AI models; users of those models; and regulatory agencies that are concerned with privacy and security. Impacted industry sectors include, but are not limited to--finance, marketing, education, research, healthcare and telecommunications.
IEEE P7000	Standard for Model Process for Addressing Ethical Concerns During System Design	This standard outlines an approach for identifying and analyzing potential ethical issues in a system or software program from the onset of the effort. The values-based system design methods address ethical considerations at each stage of development to help avoid negative unintended consequences, while increasing innovation.	Innovation managers, engineers and technologists involved in product or systems life cycles, development, operations, maintenance, and disposal; end users; suppliers; acquirers; regulatory bodies; the public-at-large.
IEEE P7001	Standards for Transparency of Autonomous Systems	This standard describes measurable, testable levels of transparency, so that autonomous systems can be objectively assessed and levels of compliance determined. A key concern over autonomous systems (AS) is that their operation must be transparent to a wide range of stakeholders, for different reasons. For designers, the standard will provide a guide for self-assessing transparency during development, and suggest mechanisms for improving transparency (e.g., the	Users; certification; regulation or accident investigation agencies; expert professionals; and society-at-large, in addition to autonomous systems designers.

		need for secure storage of sensor and internal state data, comparable to a flight data recorder or black box).	
IEEE P7002	Standard for Data Privacy Process	This standard specifies how to manage privacy issues for systems or software that collect personal data. It will do so by defining requirements that cover corporate data collection policies and quality assurance. It also includes a use case and data model for organizations developing applications involving personal information. The standard will help designers, by providing ways to identify and measure privacy controls in their systems, utilizing privacy impact assessments.	Employees, customers, or members of the public whose personal data is used in creating or manufacturing products, services and systems generated by organizations using this process. Organization based stakeholders—including senior leaders, designers, engineers, producers, suppliers and marketers who design products, services and systems in a way to honor end-user privacy and personal data. Stakeholders include local, national, and international governing bodies responsible for the transfer, storage, (potential) sale, or dissemination of said personal data.
P7003	Standard for Algorithmic Bias Considerations	This standard describes specific methodologies to help users certify how they worked to address and eliminate issues of negative bias in creating their algorithms--where "negative bias" infers the usage of overly subjective or uniformed data sets, or information known to be inconsistent with legislation concerning certain protected characteristics; or with instances of bias against groups not necessarily protected explicitly by legislation, but otherwise diminishing stakeholder or user well-being, and for which there are good reasons to be considered inappropriate.	Programmers, manufacturers, researchers, or others creating an algorithm, along with any actors defined as end-users of said algorithm, and any non-users affected by the use of the algorithm, including but not limited to, customers, citizens and website visitors.

P7004	Standard for Child and Student Data Governance	The standard defines specific methodologies to help users certify how they approach accessing, collecting, storing, utilizing, sharing, and destroying child and student data. The standard provides specific metrics and conformance criteria regarding these types of uses from trusted global partners; and it demonstrates how vendors and educational institutions can meet them.	Vendors, adult students, children, their parents / caregivers / guardians, educators, and any organization wishing to access a child's data-- whether in overt form (in the form of an educational based technology, for example); or hidden (algorithms accessing child data for advertising, marketing, email / transfer); or any other purposes.
P7005	Standard for Transparent Employer Data Governance	The standard defines specific methodologies to help employers to certify how they approach accessing, collecting, storing, utilizing, sharing and destroying employee data. The standard provides specific metrics and conformance criteria regarding these types of uses from trusted global partners, and how vendors and employers can meet them.	Stakeholders within the value chain of an organization, including but not limited to shareholders, C-Suite level management, Managers, HR, CSR, and all other staff, their trade unions, shop stewards and representatives. Central stakeholders, including employees-- whether full time or part time--to provide education, training and ongoing support--to ensure they have the tools and knowledge to protect and utilize their data to their own best advantage--while also ensuring trusted information exchange with their employers.
P7006	Standard for Personal Data Artificial Intelligence (AI) Agent	This standard describes the technical elements required to create and grant access to a personalized Artificial Intelligence (AI) composed of inputs, learning, ethics, rules and values controlled by individuals.	Individuals whose information is utilized for any data interaction, but also academics, engineers, programmers, marketers or technologists of any kind wishing to utilize said data.
P7007	Ontological Standard for Ethically Driven Robotics and	The standard establishes a set of ontologies with different abstraction levels that contain the concepts, definitions and axioms necessary to	Manufacturers, service and solution providers, equipment suppliers in the robotics and users.

	Automation Systems	establish ethically driven methodologies for the design of Robots and Automation Systems.	
P7008	Standard for Ethically Driven Nudging for Robotic, Intelligent and Autonomous Systems	"Nudges" as exhibited by robotic, intelligent or autonomous systems are defined as overt or hidden suggestions or manipulations designed to influence the behavior or emotions of a user. This standard establishes a delineation of typical nudges (currently in use, or that could be created). It contains concepts, functions and benefits necessary to establish and ensure ethically driven methodologies for designing the robotic, intelligent and autonomous systems that incorporate them.	Manufacturers, service and solution providers, robotics equipment suppliers and users.
P7009	Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems	This standard establishes a practical, technical baseline of specific methodologies and tools for developing, implementing and using effective fail-safe mechanisms in autonomous and semi-autonomous systems. The standard includes (but is not limited to): clear procedures for measuring, testing and certifying a system's ability to fail safely--on a scale from weak to strong--and instructions for improving unsatisfactory performance.	Technology companies, engineers, developers, researchers, and other agents creating autonomous and semi-autonomous systems. This standard also includes, but is not limited to, regulators and the society-at-large who are directly and indirectly affected by these systems.
P7010	Standard for Well-being Metrics for Ethical Artificial Intelligence and Autonomous Systems	This standard establishes well-being metrics relating to human factors directly affected by intelligent and autonomous systems, and it establishes a baseline for the types of objective and subjective data these systems should analyze and include (in their programming and functioning) to proactively increase human well-being.	Manufacturers, service and solution providers, programmers, engineers, technologists
P7011	Standard for the Process of	This standard provides semi-autonomous processes using	News Media, Internet Search, Social Media,

	Identifying and Rating the Trustworthiness of News Sources	standards to create and maintain news purveyor ratings for purposes of public awareness. It standardizes processes to identify and rate the factual accuracy of news stories to produce a rating of online news purveyors and the online portion of multimedia news purveyors. This process will be used to produce truthfulness scorecards through multi-faceted and multi-sourced approaches.	Online Advertising, software engineers, open source community, general consumers. website developers
P7012	Standard for Machine Readable Personal Privacy Terms	The standard identifies/addresses the manner in which personal privacy terms are proffered and how they can be read and agreed to by machines.	Individuals who interact as first parties with others in the networked world and entities--mostly companies-- that agree as second parties to individuals' terms.
P7013	Standard for Benchmarking Accuracy, Increasing Transparency, and Governing Use of Automated Facial Analysis Technology	The standard provides phenotypic and demographic definitions that technologists and auditors can use to assess the diversity of face data used for training and benchmarking algorithmic performance, establishes accuracy reporting and data diversity protocols/rubrics for automated facial analysis, and outlines a rating system to determine contexts in which automated facial analysis technology should not be used.	Engineers, technologists, and researchers who develop automated facial analysis technology; procurement officers, auditors, and other decision-makers who assess the suitability of using these technologies in a given context; and the public-at-large.

2. Needs and challenges regarding the existence, availability, use and development of AI standards and tools.

The forms of standardization serve purposes, especially in information and communications technology (ICT), including AI and AI systems. There is a need for stability (provided by the arena of formal standards bodies), coping with rapid change (provided by consortia and alliances), specific intellectual property and marketing environments, and the need for robust community involvement (provided by Open Source). To tackle the vast emerging standardization needs for AI and AI systems, the groups within each arena need to more effectively work together to create standards of the highest quality, through open systems and open standardization processes that effectively contribute to the public good.

Commercial stakes, as well as those related to cultural values such as privacy and ethics, are bringing to the forefront a new era of standardization which needs to be global and inclusive. Standardization processes must be sufficiently nimble to effectively address the development and commercial application of rapidly evolving technologies such as AI, and they must be open to addressing ethically aligned design concepts from the onset.

Developing AI standards needs to be inclusive of diverse communities of experts and users, including economists, ethicists, legal professionals, philosophers, educators, policy-makers, regulators, and community representatives, in addition to technologists and scientists. With participants from fields not traditionally directly engaged in technical standards development, there will be a need for training and mentoring programs to hasten understanding of, and engagement in, the process of developing standards.

3. The current and potential future role of Federal agencies regarding the existence, availability, use, and development of AI technical standards and tools in order to meet the nation's needs.

Standards are critical components of our nation's technology infrastructure, and vital to promoting innovation, industry, trade and commerce; many are also crucial to the safety of Americans. As AI technologies become more prevalent, the federal government's role could not be more fundamental for protecting its citizens. In an attempt to preempt the possible negative effects of their misuse, emerging AI technologies require a collective response for rapidly implementing these innovative technologies.

Because no single organization, public or private, controls the U.S. standards development system, the federal government should play an active role in helping to guide the process and ensure it meets the American public's needs.

Due to its ability to leverage its purchasing and convening powers, the federal government is indispensable in promoting the development and adoption of trustworthy AI standards. Federal government technology acquisition decisions can help establish agreed-upon standards for safety, security, transparency and algorithmic bias.

Federal agencies should designate individuals and teams of highly qualified employees representing their agencies to contribute to standards development. By collaborating with standards development organizations like the IEEE Standards Association, these interagency workgroups can improve standards quality and process inclusiveness. For example, the SBA should invite the participation of small and medium enterprises not normally included in these discussions.

Interagency working groups should track the latest state-of-the-art technology, and AI and autonomous systems development, to ensure that the resulting standards and internal agency processes match fast AI development and its applications. Such tracking can be achieved by consulting with academic and industry consortia (like IEEE Future Directions), or USPTO semantics libraries, to provide indicators of the latest technology development. USPTO's PatentsView, for example, is a helpful tool to examine inventor-patenting dynamics over time.

We also recommend that NIST internally prioritize the pooling of data, methods, and models that contribute to developing, verifying, and validating benchmarks for AI systems performance and effectiveness.

IEEE provides the following comments – organized under the six subtopics listed in the RFI:

Subtopic 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.

AI-enabled systems can provide federal judges and law enforcement officials with powerful tools. Those same tools, however, can impinge significantly on the privacy, liberty, and even the lives of citizens whose actions come under the purview of the federal law enforcement and judicial systems. It is essential, therefore, that standards be developed and adopted that allow these agencies to realize the benefits offered by AI-enabled systems, while at the same time protecting core societal values. Consequently, the federal government ought to involve civil rights and civil society groups in developing standards.

Subtopic 14: The type and degree of Federal agencies' current and needed involvement in AI technical standards to address the needs of the Federal government.

The Federal Judicial Center (FJC) provides federal judges with educational materials and model orders regarding the use of AI-enabled systems in legal proceedings. Such materials are essential to ensuring the safe, effective and consistent use of AI-enabled systems in legal proceedings. Such efforts should be supported and expanded.

Subtopic 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 (through NIST) should prioritize the development of standards for AI-enabled products and services that deployed the vital societal functions of **medicine and law.**

Also, to address fears that emerging AI technologies threaten the availability and creation of jobs, the federal government could employ "workforce modeling" to monitor the need for particular skills related to AI technologies. Workforce modeling could help the private sector create the appropriate and much-needed jobs in industry and automation, and close the skills gap with innovation training that incorporates technology and offers micro credentialing to retrain the current workforce. This effort would necessitate strong coordination and data sharing with the Department of Labor Bureau of Statistics and Bureau of Census.

Subtopic 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.*

The federal government should **support and expand initiatives that will provide comprehensive, all-inclusive input to standards development efforts**. Even in the case of standards developed under private sector leadership, **the federal government has a crucial role to play in providing fundamental input to the success of those initiatives**. In the case of AI-enabled systems, this input includes on-going benchmarking exercises designed to assess the effectiveness of AI-enabled systems; publicly available datasets for the development and testing of AI-enabled systems; and standard, readily understandable, metrics for gauging the effectiveness of AI-enabled systems.

NIST should consider the input that state agencies might provide and collaborate with state governments to provide constant communication regarding the efforts of agencies across the United States to roll out AI technologies.

Additionally, the **federal government should ensure that all interested parties have access to both the public and private aspects of standards development process to** safeguard the privacy interests of the American public. In addition to the scientists and engineers who develop AI systems and technology, and the standards development organizations that actually develop the standards, interested parties could include small business advocates, health care providers, educators and civil rights organizations. For example, some AI systems are unique in that their use generates ethical concerns; for example, the potential to violate individual privacy via the possible use of irrelevant personal data for making financial or healthcare decisions. **This collaborative approach will be extremely important for ensuring that AI technologies do not negatively affect our culture and society.**

Subtopic 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.*

Since 1992, NIST has sponsored a research initiative, the Text Retrieval Conference (TREC), focused in conducting annual evaluations of the effectiveness of advanced technologies at carrying out a range of different information retrieval tasks. While TREC itself is not a standards development initiative, the results of its research have been crucial to the informed adoption of AI-enabled systems to information

retrieval tasks (e.g., the results of evaluations conducted within TREC's Legal Track provided the empirical data courts needed to give approval to the use of AI-enabled systems in legal discovery). These research results could be central to developing standards and standards-related tools in this field. Such research initiatives should be supported and expanded, in the interest of ensuring well-informed and effective standards.

Subtopic 18*: What actions, if any, the Federal government should take to help ensure that desired AI technical standards are useful and incorporated into practice.*
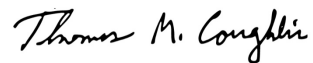
1) To ensure that any standards that are developed can serve the objective of trustworthy adoption of AI, these developed standards should fit a coherent, well-articulated framework of principles--which, if duly operationalized through standards and certifications--would meet the objective of trustworthy adoption.

2) To ensure a well-informed process in the particularly complex area of developing standards for AI-enabled products and services, the federal government (e.g., through NIST) should develop an awareness of and engage with national and international endeavors that have substantial standards-development expertise, and who already have mature AI standards initiatives under way.

3) In contributing to standards development, the federal government (e.g., through NIST) should recognize that most AI-enabled products and services remain socio-technical systems, involving technologies and human intervention. Hence, in certain circumstances, in particular those where the life, liberty, or health of humans are at stake, standards for AI operators should be considered (for example those in the healthcare system who are tasked with operating and interpreting diagnostic AI; or those in the legal system who operate in contexts such as recidivism, risk-assessment, or fact-finding in civil and criminal proceedings.

4) In contributing to standards development, the federal government should recognize the importance of ensuring that any developed standards are broad in scope, covering all human agents involved in the design, development, procurement, deployment, operation and validation of effectiveness of the AI-enabled systems in question.

5) In contributing to standards development, the federal government should recognize the importance of ensuring that any standards developed have broad public support. To that end, federal agencies should facilitate dialogue among all stakeholders in the safe and effective use of AI-enabled systems (those engaged in designing, developing, procuring, deploying, operating, and validating the technology's effectiveness; those who may be immediately affected by the technology results; those who may be indirectly affected by the technology results, including the general public; and those with specialized expertise in ethics, politics and the law).

6) NIST should prioritize spending and investment in developing annual evaluations to enable the ongoing assessment by all stakeholders of the extent to which AI-enabled products and services (including innovative products and services) meet or exceed the performance levels established by the standards. (Along the lines of the groundbreaking TREC Legal Track studies, but with the primary purpose of serving as ongoing annual evaluations targeted in their design and reporting to a broad swath of stakeholders in society.)

7) The federal government should facilitate creating data sets that can be used both for the purpose of developing standards, and for the purpose of evaluating the adherence of AI-enabled systems to any standards developed. In assisting in creating such data sets, governments and administrative agencies must take potentially competing societal values, such as the protection of personal data, into consideration--and arrive at solutions that maintain those values, while enabling the creation of usable, real-world data sets.

8) Federal government agencies should support educational initiatives designed to create greater awareness among all stakeholders of the potential benefits and risks of adopting A/IS in the legal system, and of the ways of mitigating such risks (including the development and adoption of

technical standards and related tools).

9) In the interest of encouraging the adoption of any standards developed, federal and state agencies should be encouraged to develop procurement policies favoring products and services consistent with NIST Standards, or other similar standards in industry.

We thank NIST for consulting interested stakeholders with expertise in developing technology standards and welcome any further discussions with the agency on these matters.

Respectfully submitted,

Thomas M. Coughlin
2019 IEEE-USA President