September 14, 2021

Dr. James Olthoff
Director
National Institute of Standards and Technology
Re: Artificial Intelligence Risk Management Framework (Docket Number: 210726–0151)

Dear Director Olthoff,

As a leading global professional services company, Accenture provides a broad range of services and solutions in strategy and consulting, technology, interactive, and operations, that span all industries. We combine artificial intelligence (AI) with deep industry and analytics expertise to help our clients embrace these emerging, intelligent technologies confidently and responsibly.

At Accenture Labs and Accenture Federal Services, we incubate new concepts and apply the latest technologies to design and deliver breakthrough solutions for business, government, and society. In addition, Accenture's Applied Intelligence practice delivers AI applications at scale and our Responsible AI practice focuses on the ethical, transparent, and accountable use of AI technologies in a manner consistent with user expectations, organizational values, and societal laws and norms.

Thank you for the opportunity to provide input on the development of the National Institute of Standards and Technology's Artificial Intelligence Risk Management Framework. We look forward to further participation in the development of the framework.

Sincerely,

Paul Daugherty
Group Chief Executive – Technology & Chief Technology Officer
Accenture

**Accenture**
**Comments to National Institute of Standards and Technology**
**Artificial Intelligence Risk Management Framework**

***1. The greatest challenges in improving how AI actors manage AI-related risks—where "manage" means identify, assess, prioritize, respond to, or communicate those risks;***

When helping our clients manage AI-related risks, Accenture has encountered a wide range of challenges. Some challenges result from resource constraints and incentives that make investing in AI risk management difficult, while the rapid pace of technological evolution has made it difficult for skills, regulations, and professional credentialling to keep up. Some of the greatest challenges in improving the management of AI-related risk are:

**Not an Enterprise Priority**

- Incentives to rapidly adopt AI and achieve the promise of algorithms and compete with competitors outpace incentives to manage risk
- Speed at which AI methodologies are advancing, making it hard to keep skills and processes up to date
- Lack of awareness among an organizations' senior leaders
- Lack of skills, human resources, and funding focused on risk management
- Overconfidence in performance of the AI systems without proper accountability/audit mechanisms.

**Regulatory Immaturity**

- Growing patchwork of local, state, federal, and international regulations and standards
- Absence of defined governance processes (especially outside of the financial services industry)
- Lack of technical expertise and accredited qualifications

**Lack of Tools and Frameworks**

- Lower maturity of tools that support risk measure, impact, and mitigation that does not keep pace with tools that allow automated application of AI capability for the citizen user
- Lack of multi- and interdisciplinary engineering framework, which results in fragile and unquantifiable measures
- Anthropomorphism and attribution fallacy
- Under specification in modeling
- Continual changes of AI systems and machine learning models once deployed in production environments

AI systems should not be implemented in a technology vacuum. AI systems should be integrated into an overall ecosystem that brings together cross-functional teams across an organization – in addition to C-Suite accountability and leadership, key players should also be the chief information officer, chief information security officer, human resources, legal, risk, ethics, compliance and quality.

***2. How organizations currently define and manage characteristics of AI trustworthiness and whether there are important characteristics which should be considered in the Framework besides: Accuracy, explainability and interpretability, reliability, privacy, robustness, safety, security (resilience), and mitigation of harmful bias, or harmful outcomes from misuse of the AI;***

To the extent that AI is an integral part of a system, and in anticipation of broader use and scale of AI, there are several additional characteristics of AI trustworthiness that NIST should consider:

> **Interoperability**: Design systems with the ability to interact and work with other systems. When systems are interoperable, users gain confidence in the technology, and the cooperative marketplace becomes a better arbiter for trustworthiness and risk.
>
> **Auditability/Traceability**: Ability to track and monitor the performance of AI through all stages of the AI lifecycle. Organizations can achieve this by taking into consideration the following: ensuring the provenance and validity of training data; automatically logging all events; ensuring that systems can detect drift and quickly resolve such issues.
>
> **Soundness**: Ensure proper context, data quality and model validity during development and deployment of AI. Organizations can achieve this by taking into consideration the following: developing technical assessments and processes to address data completeness and representativeness; developing risk assessments and procedures to address data limitations; discussing with data owners and developers the intention of using data and AI; conducting regular performance tests to ensure that AI functionality does not drift.
>
> **Sustainability**: Associate the needs of end users, society, or the environment into the development of AI. Organizations can achieve this by taking into consideration the following: incorporating input from relevant internal and external sources to diversify AI development; conducting trials that involve end users from differing societal circumstances; developing AI with minimal impact to the environment; educating all stakeholders on how AI could provide solutions to existing societal and the environmental problems.

Additionally, NIST should consider exploring the inherent contradictions and subsequent tradeoffs between their characteristics when developing the AI Risk Management Framework. For example, principles such as privacy and explainability can be contradictory when personal data is involved. An overemphasis on the explainability of an AI system can lead to diminishing privacy. If explainability regulations required detailed information regarding the training dataset being used, the type of machine learning algorithm, or other information, the regulations could make attacks on machine learning models, such as model extraction, much easier and more successful. In the absence of such details, attackers are left to less successful black-box attacks. How organizations weigh these various tradeoffs should be a subject of discussion for NIST as it develops the framework.

***3. How organizations currently define and manage principles of AI trustworthiness and whether there are important principles which should be considered in the Framework besides:***

In this request for information, NIST appears to be making a distinction between so-called "characteristics" and "principles" of AI trustworthiness. NIST should provide more details on this distinction in future requests and workshops. For example, Accenture believes that "privacy" – currently classified as a characteristic – is fundamental to successfully develop and deploy AI. In that context, could or should "privacy" be considered a principle, or is still just a characteristic?

Setting aside questions of classification, organizations can employ the following approaches to manage the principles of AI trustworthiness identified by NIST:

> **Accountability**: Organize governance structure for management of processes, roles, and responsibilities. Trustworthy systems should have clear and transparent frameworks by which a stakeholder can rely on due process-like mechanisms to hold the organizational system accountable for good and bad outcomes. Organizations can achieve this by taking into consideration the following: setting up a centralized board to adjudicate key decisions; ensuring that employees fully understand their roles and responsibilities; establishing a transparent chain of command to designate authority; and developing a feedback loop to oversee and respond to issues from end users.
>
> **Fairness**: Mitigate bias throughout an AI lifecycle. Organizations can achieve this by taking into consideration the following: checking for equality among subgroups in the data; encouraging diverse backgrounds and views among developers; developing rigorous procedures to process data; employing methods to monitor AI for bias after deployment; establishing clear end goal targets to aim towards; examining the origins of how data was gathered; ensuring that all data and AI processes can be repeated.
>
> **Transparency**: Engender trust and clarity with all stakeholders of data and AI processes. Organizations can achieve this by taking into consideration the following: producing documentation that clearly explains AI methodology for data and algorithms used; taking into account the opinions and desires of end users; engaging with relevant stakeholders to ensure AI meets production level quality; providing training to end users for them to understand the full capabilities of AI.

Too often ethical concerns are restricted just to the technical aspects of a model. In practice a biased model can be used fairly, and an unbiased one can be used unethically. Technically, it may even be unethical *not* to use a biased model if the alternative is to use a lower performing approach. Guidelines are necessary, therefore, not just to evaluate the fairness of a model, but also the fairness of the application that uses it, the business process that uses that application, and the overall context into which the application and business process are being placed.

---

***4. The extent to which AI risks are incorporated into different organizations' overarching enterprise risk management—including, but not limited to, the management of risks related to cybersecurity, privacy, and safety;***

In considering how to incorporate AI risks into organizations' enterprise risk management, NIST should carefully consider the scope of AI risk. There are risks that are of particular concern with regards to AI, such as bias and model drift. However, many risks that may be associated with AI may come more generally from digital transformation and not be intrinsically tied to AI. NIST should consider how more general risks and AI-specific risks are addressed in the Risk Management Framework.

Accenture believes that organizations should integrate their AI risk management into existing processes where they already exist. For example, well-developed risk and compliance processes already exist within the financial services industry, and AI risk management should be integrated into these pre-existing structures. As AI is embedded into every part of the business, lines of business using AI need to customize and adopt their own risk management but base guidance from a central team that may be part of security, legal, responsible business, and data and AI center of excellence.

However, many organizations in fields with less robust risk and compliance processes will need to create new risk management structures and processes. To best assist organizations that may find themselves building out new risk and compliance processes for AI risk, NIST should provide examples of frameworks that can be used to develop and implement AI risk management processes. Further, examples of different maturity models for AI risk management may be helpful (e.g., per industry or business function).

---

***5. Standards, frameworks, models, methodologies, tools, guidelines and best practices, and principles to identify, assess, prioritize, mitigate, or communicate AI risk and whether any currently meet the minimum attributes described above;***

The nature of AI presents new and unique challenges that businesses must address (and address early) to ensure continued success. However, universal standards and safeguards do not yet exist to assure outcomes and comfort stakeholders. Accenture often finds that businesses have a limited understanding of the contents of the data used to train AI, the methods and algorithms that make it work, and the potential societal impacts. Therefore, it's critical that organizations develop and implement a plan to enable the creation of fair and transparent AI systems as well as mitigating the impact of unintended consequences and malicious use.

One framework that Accenture uses to manage AI risk and ensure that AI is used responsibly is our Responsible AI framework, which ensures the ethical, transparent, and accountable use of AI and leads to consistency with user expectations, organizational values, and societal laws and norms.

Accenture's Responsible AI Framework is comprised of four pillars,

1. **Principles & Governance**: Refine capabilities into guiding principles and establish structure for processes, roles, and responsibilities to enact governance

2. **Risk, Policy & Control**: Manage risks, policies, and controls to comply with data and AI ethics regulations

3. **Technology & Enablers:** Provide tools and techniques to support the implementation of refined capabilities

4. **Culture & Training:** Educate employees to enable them to adopt refined capabilities into their day-to-day operations

AI is increasingly being developed and deployed to critical processes (e.g. healthcare, employment, judicial, policing, etc.) where it could pose risks to safety, privacy, and human rights. Therefore, it is important to evaluate the level of risk posed by AI and its intended application to determine an appropriate course of action for mitigating any existing or potential risks. Accenture recommends conducting a "risk triage" to determine whether the risk posed by a particular use of AI is low or high. Organization can determine the category of risk by asking difficult questions, such as, "What is the size of the potentially impacted audience?" and "How long will the impacts of the use case affect end users?" Organizations need to put in place different processes to deal with low- and high-level risk.

Moreover, as AI is developed across data cooperatives where multiple lines of business or even partners work together to supply data and create algorithms, standards that include common measures and metrics are essential to enable working together. As an example, Accenture is sponsoring UC Berkeley Professor Marc van der Laan's research on targeted learning, which incorporates machine learning (ML) methods for causal inference. The goal is to support multiple parties whose experiments can run off shared data across distributed environments.

---

***7. AI risk management standards, frameworks, models, methodologies, tools, guidelines and best practices, principles, and practices which NIST should consider to ensure that the AI RMF aligns with and supports other efforts;***

While countries around the world will undoubtedly pursue their own ways of regulating AI, NIST should strive to harmonize definitions of key terms with those already published so that the global AI community is speaking the same language. A common lexicon will give organizations and society more confidence and promote greater alignment of standards, frameworks, models, etc.

Legislative bodies, regulatory agencies, standards bodies, and others have started developing their own definitions related to AI, the proliferation of which unnecessarily complicates operations for global organizations. For example, Article 3 of the European Commission's proposed Artificial Intelligence Act includes numerous definitions of AI-related terms, including 'artificial intelligence system' and 'training data.' In the United Kingdom, the Information Commissioner's Office has published its own definition of artificial intelligence. In the United

States, Congress defined artificial intelligence in the John S. McCain National Defense Authorization Act. The Institute of Electrical and Electronics Engineers (IEEE) has also weighed in with its own definitions. Ensuring some degree of harmony in the terms used to talk about AI is one way that NIST's Risk Management Framework can promote innovation and adoption of AI.

Additionally, NIST should note the increasing number of countries that are incorporating the concept of risk triaging into their frameworks, including both Canada and the European Commission. By distinguishing between high- and low-risk uses of AI, regulators are able to focus their limited resources on the uses of AI that could have the most significant impacts on individuals' lives.

---

***8. How organizations take into account benefits and issues related to inclusiveness in AI design, development, use and evaluation—and how AI design and development may be carried out in a way that reduces or manages the risk of potential negative impact on individuals, groups, and society.***

In a recently published paper, "AI for Disability Inclusion: Enabling Change with Advanced Technology," Accenture partnered with Disability:IN and the American Association of People with Disabilities (AAPD) to examine how AI – when developed and used responsibly and ethically – can help organizations be more inclusive of people with disabilities. In that report, Accenture looked at the benefits of inclusive design.

Inclusive design considers the needs of all users as a product or service is being developed, from start to finish. Organizations that design for diversity of all users and edge cases (e.g. individuals with disabilities) will create better solutions and experiences for all users. Effectively incorporating inclusive design into the way organizations create products and services requires adopting a new culture. Organizations need to take charge, consciously integrating accessibility, diversity, and inclusion into their ways of working.

Ideally, accessibility (the "what") and inclusive design (the "how") work together to make experiences that are not only compliant with standards, but truly usable and open to all. And the "why": Together, inclusive design and accessibility can allow teams to build more diverse products, attract talent, ensure product teams are representative of the diversity of the customers we serve, and help identify new opportunities for revenue growth.

In order to reduce the risk of potential negative impacts on individuals, groups, and society, Accenture recommends a set of guiding principles called R(AI)S:

> **Responsible**: means adopting and scaling AI responsibly and ethically; innovating with purpose; and placing a premium on compliance, accountability, transparency and explainability.

**Accessible**: indicates that all AI endeavors put a premium on accessibility and that includes features and functionality of the tool itself, the capabilities of vendors involved, the experience of all people using it.

**Inclusive**: means taking action for fairness in/with AI: Use inclusive design approaches that incorporate the lived experience of all persons, and debiasing techniques to create a culture of equality and inclusion.

**Secure**: at one level means ensuring that using AI will not put data privacy at risk. However, it also recognizes security in the sense that individuals should not need to ask that their particular needs be considered in the design of AI (thus potentially depriving them of the autonomy and privacy that would otherwise be protected when using the technology).

---

***9. The appropriateness of the attributes NIST has developed for the AI Risk Management Framework. (See above, "AI RMF Development and Attributes");***

Accenture believes that the attributes that NIST has developed for the AI Risk Management Framework are appropriate and commends NIST for its thoughtful approach. Accenture supports the use of a risk-based standards and regulatory approach to AI that accounts for the varying magnitude and nature of consequences when considering risk mitigation. One of the key advantages of a risk-based approach is that it minimizes constraints on innovation, while also applying sufficient controls where needed.

Additionally, Accenture supports the development of voluntary consensus standards because they can help create and safeguard trust at the heart of AI-driven systems and business models and permit the flexibility for innovation, allowing codes to develop with the technology. Examples include, the IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems and its nine pipeline standards on Ethically Aligned Design.

---

***10. Effective ways to structure the Framework to achieve the desired goals, including, but not limited to, integrating AI risk management processes with organizational processes for developing products and services for better outcomes in terms of trustworthiness and management of AI risks. Respondents are asked to identify any current models which would be effective. These could include—but are not limited to—the NIST Cybersecurity Framework or Privacy Framework, which focus on outcomes, functions, categories and subcategories and also offer options for developing profiles reflecting current and desired approaches as well as tiers to describe degree of framework implementation; and***

Accenture believes that the NIST Cybersecurity Framework is generally effective because it brought a shared understanding of the vernacular and flexible approaches to think about risk and risk management.

---

***11. How the Framework could be developed to advance the recruitment, hiring, development, and retention of a knowledgeable and skilled workforce necessary to perform AI-related functions within organizations.***

Ensuring that a workforce has the requisite knowledge and skills to perform AI-related functions is a challenge faced by many organizations. According to research from Accenture and the MIT Chief Data Office and Information Quality Symposium (MIT CDOIQ), the top challenge faced by chief data officers (CDOs) in implementing their organization's data vision is the lack of talent to operationalize the vision.[1] Addressing these workforce issues starts at the top. When Accenture works with clients, we advise them that there are five steps that CEOs should take to ensure that they set the right priorities and improve their organizations' ability to use AI in transformative ways:

1. **Develop personal expertise and leadership in AI**: CEOs should immerse themselves in the details of how AI works, not merely engage in broad conversations about its potential value to an organization.

2. **Craft AI strategies that support the organization's core goals**: CEOs must develop effective strategies to capture, store, and process the mountains of data that fuel AI.

3. **Elevate the role of *Chief AI Officer***: Even CEOs who are knowledgeable about AI need experts to advise them; elevating the position of 'Chief AI Officer' creates accountability and accelerates transformation.

4. **Keep humans in charge of who decides**: The reason a business deploys AI is to help better achieve its goals, not to cede control of the company over to machines. AI works best when humans and machines are working together, which means CEOs need to strike a balance between computer-led and human-led decision making, ultimately leaving humans in charge of who decides.

5. **Build a data-fluent culture across the workforce**: Ensuring that the most value is gained from AI requires the full commitment and talent of an entire organization. Companies need to demystify AI for all of their workers and

Scaling AI effectively for the long term will require the professionalization of the industry, and NIST can help advance this goal by including recommendations and best practices for the professionalization of organizations' AI workforces in the Risk Management Framework. Accenture explored the need to professionalize the AI workforce, and how organizations can do so, in a recent report, "Scaling enterprise AI for business value." Stakeholders – from practitioners to leaders across the private and public sector – must come together to distinguish clear roles and responsibilities for AI practitioners; demand the right level of education and training for practitioners; define processes for developing, deploying and managing AI; and democratize AI literacy across the enterprise. By formalizing AI as a trade with a shared set of norms and principles, companies will be poised to achieve more value from AI.

[1] *Accenture Research, MIT CDOIQ Survey, July 2021*

Real value can only be realized when trained AI practitioners are working hand in hand with the business to accomplish their organization's goals, and those interdisciplinary teams are guided by standards, rules, and processes. By following these steps to standardize professionals and processes, organizations can better set themselves up to scale AI and, in so doing, make the most of this quickly evolving technology.

1. **Distinguish clear AI roles**: A hallmark of a professionalized industry or trade is that practitioners understand the individual roles that contribute to a final product. Multidisciplinary teams of diverse perspectives, skills, and approaches must work together to innovate and deliver AI products or services. The mix and the ratio of roles is going to depend on the use cases pursued at the time and will vary from project to project. Establishing a blueprint for how teams should operate will help this process become more turnkey over time. But one thing remains true across all projects – organizations need to establish ownership and expectations from the start.

2. **Demand education and AI training**: It's important for organizations to establish education and training requirements for their AI practitioners. To establish an effective professionalized workforce, it's up to companies to assess which skills they need, the qualifications of their talent, and match appropriate skills to roles, and identify skills gaps. To enable a consistent approach to upskilling, companies should create clear career paths for their AI practitioners with prerequisite coursework, training, and experiences.

3. **Define AI processes**: While some argue that formalized processes and governance could stifle innovation, Accenture's research has shown the opposite. In professionalized industries, there's a standard approach to testing and benchmarking during the creation (or optimization) of products and services. Similarly, whether a company is making smart devices or building a data science model to improve the online retail experience, establishing systems and processes to support the development of the AI product or solution allows people to innovate in a predictable and efficient way.

4. **Democratize AI literacy across the organization**: While there's certainly growing interest from leaders to invest in AI technologies, true professionalization will result in (and rely on) AI literacy across an entire organization. Organizations owe it to their employees and to their bottom lines to provide upskill their talent on AI. To start, companies should define the minimum level of AI knowledge they require from their employees. Helping the entire workforce understand what AI is, how it impacts their jobs and how it benefits the company are part of building confidence in AI and driving adoption and usage.

---