AI-in-Gov Council

https://cec.gmu.edu/AI-in-Gov-Council

A Public-Private Partnership between the George Mason University College of Engineering and Computing and Public Sector Technology Providers

George Mason University | Brillient Corporation | Amazon Web Services | 4A Consulting | 22nd Century Technologies | InterFuze | Unissant | CoAspire

Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO), National Science Foundation (NSF).
Attn: Faisal D'Souza, NCO 2415 Eisenhower Avenue Alexandria, VA 22314, USA

Email: ostp-ai-rfi@nitrd.gov

Subject: Response to RFI on the Development of an Artificial Intelligence (AI) Action Plan

Dear Mr. D'Souza,

On behalf of the AI-in Gov Council, a public-private partnership uniting academia, public sector technology providers, and government CXOs from local, state, and federal organizations, we appreciate the opportunity to contribute to the development of the AI Action Plan. Led by George Mason University Vice President and Chief AI Officer, Dr. Amarda Shehu, and Brillient Corporation's Chief Digital Officer and Senior Vice President, Mr. Richard Jacik, our Council is committed to fostering AI-driven innovation that enhances efficiency, security, and transparency in government. Drawing from our experience working with government agencies we present recommendations to sustain and enhance America's AI leadership while ensuring the responsible development and deployment of AI in the public sector.

In compliance with RFI requirements, we stipulate that ***This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.***

Our recommendations are grouped into three sections each targeting a different action area.

The first section, **Policy Reforms to Accelerate AI Innovation**, outlines near-term opportunities to ***remove barriers and encourage AI innovation by government (and its contractors) in delivering its mission.*** We propose actionable low- or no-cost reforms to contracting frameworks that would enable industry to better support government missions by deploying AI innovations. These recommendations focus on streamlining governance processes, reducing unnecessary regulatory burdens, and empowering private sector partners to implement tailored safeguards aligned with project objectives.

The second section, **Private/Public Opportunities in AI Development and Deployment,** outlines a ***holistic framework of practice*** for intelligent hardware/software ecosystems that could benefit from policy guidance and private/public partnerships for concurrently innovative and safe AI practices. We outline key policy elements that would benefit from clear guidance and strengthened public-private partnerships. This framework balances innovation with responsible AI development, ensuring safety, fairness, and accountability.

The third section, **Strategic Investment Areas to Strengthen U.S. AI Leadership** describes a number of ***targeted areas for investment by government*** that would serve as a catalyst for AI advancements in both the public and private sectors. These targeted investments will enable the U.S. to maintain a competitive edge in AI development and deployment, and foster innovation and national leadership in the field.

Section I: Policy Reforms to Accelerate AI Innovation

I-A: Establishing a Clear and Actionable Definition of AI

AI should be more precisely defined, narrowing its scope to systems that critically align with a specific, meaningful definition predicated on the unpredictability of computing outcome. Current definitions are arbitrary, are overly broad, or lack sufficient clarity to be mapped to emergent technologies in a straightforward manner thus creating ambiguity in regulation and policy. Many AI-driven technologies have become components of mainstream software architectures, such as NLP (natural language processing), and domain-specific rule bases (expert systems) and don't warrant additional rigor above current architectural reviews and control assessments. Similarly, all longitudinally driven, data dependent systems are model-based to some degree yet derive predictions with full transparency, predictability, and repeatability.

Each computable component that is swept up and included in overly broad criteria requires additional attention, review, and cost; even though the specific AI approach engenders no additional mission, societal, or computational risk. Narrowing the definition to address only those components with a high risk profile will immediately reduce the level of inspection, planning, and decision-making effort and cost while still providing deployment protection.

A precise and actionable definition of AI is required, ensuring that research, policy, and industry innovation are grounded in clear principles rather than ambiguous classifications. A well-defined scope will drive more targeted advancements, enable effective regulation (as opposed to over-regulation), and foster trust in AI technologies, ultimately accelerating meaningful innovation.

I-B: Safe Harbor Liability Protections to Encourage Responsible AI Innovation

The government should provide a degree of legal protection for contractors and innovators against liabilities arising from the unintended consequences of AI/ML system deployments, provided those deployments are conducted in good faith and adhere to established guidelines, ethical standards, and best practices. Given the inherent unpredictability of current AI/ML systems—especially in complex, real-world environments—contractors should not be held solely responsible for unforeseen errors, biases, or system failures beyond their reasonable control. For government-sponsored projects, this can be accomplished with minor acquisition reforms, while more universal liability protection will require legislative support.

Such indemnification would encourage innovation, strengthen public-private collaboration, and promote the responsible adoption of AI technologies without excessive legal risk. However, safeguards should be in place to prevent reckless development and ensure accountability for negligence, fraud, or willful misconduct.

I-C: Streamlining Government Procurement and Risk Mitigation for Efficient Adoption of AI Innovations

Inconsistent acquisition standards and subsequent project delivery approaches across government organizational units are key inhibitors to the re-application of innovative solutions created for one government organization to peer organizations. This results in an additional taxpayer burden to support the cost of delivering newly developed solutions when transfer solutions or product/service solutions should easily suffice.

Procurement requirements for narrowed AI systems should be standardized across government to promote sharing, re-use, and cost-saving collaboration. Contractors should be required to submit AI/ML risk mitigation plans that are evaluated at source selection time. Minimal ongoing surveillance should be needed.

We propose adopting a framework similar to Organizational Conflict of Interest (OCI) mitigation plans. Under this framework, a contractor's AI/ML risk mitigation proposal, if deemed reasonable, would be accepted without imposing excessive additional layers of governance. This would streamline the approval process while ensuring accountability and effective risk management.

The focus should be on proactive accountability, ensuring that risk is addressed at the outset rather than burdening innovation with excessive, reactive oversight. A more streamlined approval process will enable faster AI deployment, reduce bureaucratic slowdowns, and encourage agile and cost-effective innovation in the public sector.

Section II. Private/Public Opportunities in AI Development and Deployment

II-A: Key Industry Considerations for Secure, Accountable, and Trusted AI Solutions in Government

As AI technologies evolve, industry participants developing AI foundation models must ensure adherence to baseline security requirements to safeguard against potential threats and vulnerabilities. Security and compliance standards are crucial for ensuring that AI solutions are resilient, scalable, and aligned with evolving regulatory frameworks. Equally important is prioritizing confidentiality and data privacy, especially in sensitive government applications, where privacy-preserving techniques, such as federated learning and differential privacy, should be implemented to protect citizens' personal data and avoid breaches.

Additionally, AI providers must incorporate robust algorithmic governance and accountability into their solutions, ensuring that their systems include mechanisms for auditability and risk mitigation, along with provisions for redress if issues arise. In procurement decisions, government entities should prioritize AI solutions that demonstrate effective risk mitigation strategies, fostering public trust by ensuring the technology is ethically designed, transparent, and rigorously tested.

II-I: Establishing Robust AI Governance: Ensuring Compliance, Trust, and Ethical Accountability

Establishing a standardized framework for AI governance is critical to ensure that AI technologies are deployed in ways that are both compliant with regulations and aligned with ethical principles. This framework should focus on risk mitigation, transparency, accountability, and fostering trust in AI.

It is essential that the governance structure not only provides clear guidelines for the responsible development and deployment of AI but also includes mechanisms for continuous oversight and accountability. Industry contributions are vital in supporting government efforts by helping define key AI-related terms, best practices, and standards that can guide effective AI governance. Given the rapid pace of technological advancements and the evolving nature of ethical considerations, governance frameworks must be regularly updated to address new challenges and opportunities. By doing so, we can ensure that AI technologies are used responsibly and that they maintain the trust of the public and all stakeholders involved.

Establishing stronger cross-agency coordination among different regulatory bodies will create more consistent oversight across sectors while preventing regulatory gaps. A centralized AI governance function should facilitate interagency collaboration, align policies, and minimize regulatory fragmentation.

Finally, the modernizing legacy government AI systems presents an opportunity to leverage industry expertise and emerging AI architectures to replace outdated or orphaned systems with secure, adaptable and high-performance solutions, ensuring that government applications remain relevant and effective in a rapidly advancing technological landscape.

II-B: Developing Tailored, Risk-Managed AI Solutions for Public Sector Applications

Implementing risk-based approaches tailored to the unique needs of public sector applications is crucial for ensuring that AI solutions are not only effective but also accountable. By focusing on the specific risks associated with government use cases, AI systems can be designed to minimize potential harm while maximizing operational impact.

Public sector technology providers should be incentivized to develop AI applications that are *specifically geared towards solving agency-specific challenges*. These solutions must be crafted in alignment with stringent security protocols and ethical standards to maintain public trust and ensure compliance with regulatory requirements. Furthermore, standardized risk assessment frameworks should be established to evaluate AI deployments proactively, reducing implementation delays and regulatory uncertainties. Encouraging the development of such tailored, risk-aware systems will enable the government to harness the full potential of AI while safeguarding against unintended consequences.

II-C: Enhancing AI Development with Accessible, High-Quality Public Data Sets

Funding initiatives aimed at creating, maintaining, and making government-curated datasets accessible are critical for advancing the effectiveness of AI model training and evaluation. These public datasets provide a foundation for developing AI systems that are robust, transparent, and widely applicable across various industries.

By ensuring that these datasets are accessible, machine-readable, and structured consistent with standard metadata frameworks, the government can encourage innovation and improve the accuracy of AI models. Collaboration with industry partners is essential to ensure that the datasets remain representative, well-documented, unbiased, and relevant to real-world applications. Such partnerships can also facilitate ongoing dataset validation and bias mitigation. In turn, these publicly available resources will drive progress in AI research and deployment, fostering greater collaboration between the public and private sectors and encourage evidence-based policymaking.

II-D: Centralized AI Repository for Public Sector Collaboration and Transparency

While current efforts, such as the 2024 Executive Order 13690, encourage agencies to publish their AI use cases, establishing a centralized AI repository would significantly enhance collaboration, standardization, and transparency across federal, state, and local agencies. A federated hub would provide a unified platform where agencies can share insights, lessons learned, and best practices related to AI deployments, fostering more efficient and effective government operations.

The repository could also serve as a valuable resource for ensuring consistency and reducing duplication of effort in AI projects across the public sector. To maintain the repository's relevance and quality, industry stakeholders should be engaged in coordinating governance, policy compliance, and rigorous testing. This collaboration would ensure that the repository remains current, functional, and beneficial for all users, promoting continuous improvement and innovation in AI adoption within the public sector.

II-E: Promoting Open-Source AI Development for Innovation and Efficiency in Government

Open-source frameworks are crucial in fostering innovation and enabling cost-effective AI adoption within government agencies. By embracing open-source models, agencies can leverage a vast pool of shared resources, tools, and expertise, accelerating the development and deployment of AI technologies. Open-source solutions also offer the advantage of greater transparency, community-driven enhancements, and distributed validation allowing for more robust testing, validation, and improvement of AI systems.

To maximize these benefits, federal AI policy should incentivize open-source initiatives and encourage industry participation in AI projects, ensuring that government agencies can access the latest advancements without the high costs of proprietary solutions. Policy must, however, also address security concerns and protect intellectual property rights, balancing open collaboration with the need for privacy and intellectual property protections.

II-F: Ensuring AI Security and Safety: Strengthening Cybersecurity and Data Privacy Measures

Strengthening AI security measures is essential to safeguarding government systems against adversarial attacks, data breaches, and other potential vulnerabilities that could compromise the integrity of AI-driven solutions. As AI technologies become increasingly integrated into critical public services, ensuring the protection of sensitive data and the security of AI models must be prioritized at both the infrastructure and application layers.

Establishing comprehensive AI assurance frameworks would play a key role in facilitating the development of standardized methodologies for evaluating and ensuring the safety, reliability, and robustness of AI systems. These frameworks would provide a structured approach to threat profiling, anomaly detection, and risk prediction implementing mitigation strategies, and monitoring performance to ensure that AI systems operate securely and in compliance with data privacy regulations. By prioritizing cybersecurity and data privacy, the government can foster trust in AI technologies while simultaneously minimizing the risks associated with their deployment.

II-G: Ensuring AI Accountability: Validation, Transparency, and Stakeholder Engagement

To ensure the responsible deployment and use of AI in public sector applications, a framework for routine validation and testing of AI models should be mandated. This process must focus on assessing key aspects such as bias, security risks, and overall efficacy to ensure that AI systems are fair, reliable, and safe for public use.

Standardizing AI labels and encouraging metadata standards across components of AI models can provide greater clarity, enabling users to better understand the intended use, risks, and limitations of these systems. Additionally, implementing structured engagement models (including robust stakeholder feedback mechanisms) will allow for ongoing improvements to AI services, ensuring that they meet the evolving needs of citizens while addressing any concerns that may arise. By establishing these practices, the government can foster greater transparency and accountability in AI applications, mitigate unintended consequences, promote trust and enhance the overall effectiveness of AI-driven solutions.

II-H: Navigating Intellectual Property Challenges in AI Development

As AI technologies advance, intellectual property (IP) rights become increasingly complex. The rapid development of AI models and systems raises questions about ownership, patentability, and the protection of innovations. Traditional IP frameworks do not fully address the unique challenges posed by AI, such as the difficulty in determining authorship when models generate novel solutions or the potential for AI systems to infringe upon existing patents without clear human oversight.

As described above, common-sense liability limitations for innovators can serve to encourage experimentation, discovery, invention, and innovation in AI and emergent technologies. Deploying these technologies for both the public and private good must be balanced with IP protection, especially when infringement is accidental or unforeseeable.

Moreover, as AI is integrated into collaborative environments, the issue of shared ownership and IP distribution between developers, organizations, and contributors must be carefully considered. Striking the right balance between incentivizing innovation and ensuring open access to AI advancements is critical, as overly restrictive IP policies could stifle progress. In contrast, overly permissive policies might will undermine private-sector investment. Government and industry partnerships should explore a clear and adaptable framework for IP in AI development, which is necessary to foster innovation while protecting the rights of all stakeholders.

II-J: Building AI Literacy and Workforce Development: Preparing Government and Society for the Future

Expanding AI training programs for government employees is essential to ensure that public sector workers are well-equipped to adopt and effectively use AI technologies while remaining aware of potential risks. These programs should focus not only on the technical aspects of AI but also on fostering an understanding of its ethical implications, biases, and socioeconomic impacts.

To further promote AI literacy, targeted investments should support state grants to align AI literacy programs with local K-12 education and workforce reskilling initiatives. By integrating AI education into early schooling and offering reskilling opportunities for workers, we can create a future-ready workforce equipped with the knowledge and skills necessary to thrive in an increasingly AI-driven economy. These efforts will empower individuals to engage with AI technologies responsibly, ensuring that current and future generations can contribute to, and benefit from, the AI-powered world.

Section III: Strategic Investment Areas to Strengthen U.S. AI Leadership

III-A: Model Immutability Service

We propose a critical initiative to enhance AI accountability and transparency, the creation of a ***Model Registry*** service, akin to the National Archives and Records Administration (NARA) for AI models. This service would ensure that every iteration of a trained AI model is preserved in permanent storage, allowing for auditable tracking, historical benchmarking, review, and immutability.

This registry would serve as a centralized repository where all government agencies and contractors could upload and securely access model versions in an environment guaranteeing model integrity and immutability. This recommendation aims to re-target industry away from fragmented, inconsistent, and incompatible decentralized model version management and focus more on innovation, while maintaining the ability to track and verify models for accountability and compliance purposes. This registry could also be accessed for audit and analysis by academics and 3rd parties. Such a service would both enhance government efficiency and ensure that AI systems are rigorously tested, reviewed, and understood; enabling explainability, reproducibility, and trust in AI applications within the public sector.

III-B: Advancements in Quantum Computing for AI

Quantum computing holds the potential to significantly enhance AI capabilities, especially in high-dimensional optimization, probabilistic modeling, and complex problem-solving. By leveraging the unique properties of quantum mechanics, quantum computers can process vast amounts of data simultaneously, enabling AI systems to solve problems that were previously computationally infeasible.

These advancements can revolutionize sectors such as drug discovery, climate modeling, and cybersecurity, where massive data sets and intricate simulations are common. Moreover, quantum-enhanced AI can accelerate federated learning and privacy-preserving AI techniques, reinforcing national security and advancing scientific frontiers.

For the U.S. to maintain its leadership in AI, investment in quantum computing research, development, and secure cloud-based integration into AI technologies is crucial. Establishing a national quantum computing initiative with dedicated AI-aligned research hubs can position the U.S. at the forefront of both fields, driving future innovation in AI.

III-C: AI Agents that can Reason

AI systems capable of reasoning, goal-setting and iterative self-directed actions— presents an exciting frontier in AI development. Research activity in neuro-symbolic AI and LLM-modulo frameworks is increasing. Incentivizing research that hybridizes foundation models with structured reasoning, cognitive architectures, and probabilistic logic frameworks is critical for enhanced decision-making, increased efficiency, and personalized services, particularly in mission-critical domains like defense, healthcare, and government service delivery.

However, the deployments of autonomous AI systems introduce challenges related to risk management, performance, and ethical concerns. To fully harness this technology's potential, a balanced approach is necessary—one that includes policy-aligned governance mechanisms, red teaming evaluations, and robust oversight mechanisms.

Human-in-the-loop (HITL) controls must be integrated to ensure that AI decision-making aligns with ethical standards and societal values, preventing unintended consequences and ensuring accountability in the adaptive decision-making process. By establishing clear guidelines and ethical considerations, we can ensure that autonomous AI benefits the public sector while minimizing risks.

III-D: Zero Trust Architecture for AI Security

As AI systems become more integrated into government infrastructure, security and resilience become an increasing concern. Adopting a Zero Trust Architecture (ZTA) paradigm for AI systems ensures that security protocols are applied consistently at every level of the AI stack. Zero Trust is based on the principle that no entity, inside or outside the network, should be trusted by default and is particularly important for AI systems in government, where sensitive data and critical services are at stake.

Implementing Zero Trust will require granular access controls, continuous authentication, rigorous monitoring, and real-time threat detection to safeguard AI models from cyber threats and malicious actors. Integrating AI-driven security analytics with Zero Trust policies can further enhance real-time anomaly detection and automated threat response. By incorporating Zero Trust principles, government agencies can better protect AI systems from attacks, ensuring that AI deployments are adaptive, tamper-resistant, secure, reliable, and resilient to evolving security challenges.

III-E: Securing AI-Driven Infrastructure: Strategies for Protecting Critical Systems and Services

As AI becomes integral to critical infrastructure sectors like energy, transportation, and telecommunications, securing these AI-driven systems is essential to protect against cyber threats and vulnerabilities. AI can enhance infrastructure security through autonomous threat detection, real-time monitoring, predictive maintenance, and anomaly detection, but it also introduces risks such as adversarial attacks on machine learning models.

A comprehensive approach to AI security is needed to mitigate these risks, including robust encryption techniques, federated learning architectures, zero-day threat detection, adversarial defense strategies, and continuous monitoring. Cross-sector collaboration between government, industry, and academia is crucial to develop best practices, set security standards, and ensure that AI systems in critical infrastructure are resilient, transparent, and accountable. By establishing proactive security frameworks, essential services can be safeguarded, public trust can be maintained, and AI vulnerabilities are mitigated.

III-F: Advancing Innovation through AI Ethics and Accountability Infrastructure

As AI becomes more integrated into government operations, it is crucial to establish a framework that fosters ethical use while encouraging innovation. The goal is to ensure responsible AI deployment without impeding technological progress through overly restrictive regulations.

This can be achieved by creating infrastructure that facilitates ethics and accountability, such as advisory committees and cross-sector AI ethics boards, which support AI development in alignment with ethical guidelines and legal standards. The focus should be on enabling ongoing evaluation and transparency, such as independent audits of AI models to ensure explainability, compliance with fairness metrics, accountability, and transparency, rather than imposing restrictive measures.

By addressing public concerns regarding algorithmic bias and ensuring that AI technologies align with societal values, we can build trust and advance innovation in a manner that upholds ethical responsibility while maximizing AI's transformative potential.

Conclusion

The AI-in-Gov Council strongly supports the collaborative efforts of OSTP and NITRD to shape policies that advance AI leadership and foster innovation. We appreciate your consideration of these recommendations and welcome the opportunity to discuss them further. We look forward to continued collaboration between government and industry to develop innovative and impactful AI solutions.

Sincerely,

Amarda Shehu
Academic Chair, Vice President and Chief AI Officer, George Mason University

Richard Jacik
Industry Chair, Chief Digital Officer and Senior Vice President, Brillient Corporation

Dominic Delmolino
Vice President, Worldwide Public Sector Technology & Innovation at Amazon Web Services

Ramesh Ponnada
Chief Executive Officer, 4A Consulting

Anil Sharma
Chief Executive Officer, 22nd Century Technologies

Daniel Gade
Owner & Chief Executive Officer, Interfuze

Manish Malhotra
Founder and Executive Chairman, Unissant

Matt Fisher
Vice President of Engineering of CoAspire