M-21-06: Guidance for Regulation of Artificial Intelligence Applications

November 17th, 2020

Summary:

M-21-06: Guidance for Regulation of Artificial Intelligence Applications was issued by the Office of Management and Budget (OMB) as a requirement of Executive Order 13859. The memorandum "sets out policy considerations that should guide, to the extent permitted by law, regulatory and non-regulatory approaches to AI applications developed and deployed outside of the Federal government."

Principles for the Stewardship of AI Applications

The memorandum lists off the following principles for agencies to consider:

Public Trust in AI: Since the continued adoption and acceptance of AI will depend significantly on public trust and validation, the government's regulatory and non-regulatory approaches to AI should contribute to public trust in AI by promoting reliable, robust, and trustworthy AI applications. For example, an appropriate regulatory approach that reduces accidents can increase public trust and thereby support the development of industries powered by AI.

Public Participation: Agencies must provide ample opportunities for the public to provide information and participate in all stages of the rulemaking process, to the extent feasible and consistent with legal requirements (including legal constraints on participation to, for example, protect national security and address imminent threats or respond to emergencies). Agencies are also encouraged, to the extent practicable, to inform the public and promote awareness and widespread availability of voluntary frameworks or standards and the creation of other informative documents.

Scientific Integrity and Information Quality: Agencies should hold information, whether produced by the government or acquired by the government from third parties, that is likely to have a clear and substantial influence on important public policy or private sector decisions (including those made by consumers) to a high standard of quality and transparency.

Risk Assessment and Management: Agencies should be transparent about their evaluations of risk and re-evaluate their assumptions and conclusions at appropriate intervals so as to foster accountability. Correspondingly, the magnitude and nature of the consequences should an AI tool fail, or for that matter succeed, can help inform the level and type of regulatory effort that is appropriate to identify and mitigate risks.

Benefits and Costs: Agencies should, when consistent with law, carefully consider the full societal costs, benefits, and distributional effects when considering regulations related to the development and deployment of AI applications. Such consideration will include the potential benefits and costs of employing AI, when compared to the systems AI has been designed to complement or replace; whether implementing AI will change the type of errors created by the system; and comparison to the degree of risk tolerated in other existing systems.

Flexibility: Targeted agency conformity assessment schemes, to protect health and safety, privacy, and other values, will be essential to a successful, and flexible, performance-based approach. To advance American innovation, agencies should keep in mind international uses of AI, ensuring that American companies are not disadvantaged by the United States' regulatory regime.

Fairness and Non-Discrimination: Agencies should consider in a transparent manner the impacts that AI applications may have on discrimination. When considering regulations or non-regulatory approaches related to AI applications, agencies should consider, in accordance with law, issues of fairness and nondiscrimination with respect to outcomes and decisions produced by the AI application at issue, as well as whether the AI application at issue may reduce levels of unlawful, unfair, or otherwise unintended discrimination as compared to existing processes.

Disclosure and Transparency: Agencies should carefully consider the sufficiency of existing or evolving legal, policy, and regulatory environments before contemplating additional measures for disclosure and transparency. What constitutes appropriate disclosure and transparency is context-specific, depending on assessments of potential harms (including those resulting from the exploitation of disclosed information), the magnitude of those harms, the technical state of the art, and the potential benefits of the AI application.

Safety and Security: Agencies should promote the development of AI systems that are safe, secure, and operate as intended, and encourage the consideration of safety and security issues throughout the AI design, development, deployment, and operation process. Agencies should pay particular attention to the controls in place to ensure the confidentiality, integrity, and availability of the information processed, stored, and transmitted by AI systems. Agencies should also consider methods for providing systemic resilience, and for preventing bad actors from exploiting AI systems, including cybersecurity risks posed by AI operation, and adversarial use of AI against a regulated entity.

Interagency Coordination: Agencies should coordinate with each other to share experiences to ensure consistency and predictability of AI-related policies that advance American innovation and adoption of AI, while appropriately protecting privacy, civil liberties, national security, and American values and allowing sector- and application-specific approaches.

Non-Regulatory Approaches to AI

Sector-Specific Policy Guidance or Frameworks: Agencies should consider using any existing statutory authority to issue non-regulatory policy statements, guidance, or testing and deployment frameworks, as a means of encouraging AI innovation in that sector. Agencies should provide guidance where a lack of regulatory clarity may impede innovation. This may also include work done in collaboration with industry, such as development of playbooks and voluntary incentive frameworks.

Pilot Programs and Experiments: Agencies should consider using any authority under existing law or regulation to grant waivers, deviations, and exemptions from regulations, or to allow pilot programs that provide safe harbors for specific AI applications. Such programs may also include events such as hackathons, tech sprints, challenges, and other types of piloting programs.

Voluntary Consensus Standards: Agencies should consider relying on private-sector conformity assessment programs, credentialing, and related activities, before proposing either regulations or compliance programs. Whenever relying on work done by private sector or other stakeholders or collaborating with them, agencies must ensure that their actions do not contribute to entrenchment by market incumbents or erect barriers to entry.

Voluntary Frameworks: Agencies should consider how to promote, leverage, or develop datasets, tools, frameworks, credentialing, and guidelines to accelerate understanding, innovation, and trust in AI. Agencies should carefully consider what existing products can be used for particular AI needs and work with stakeholders if any gaps are identified to update existing datasets, tools, frameworks, credentialing, and guidelines or develop new ones, which could include risk management frameworks.