RFI AI Commission

Action

Request for information.

Summary

The Artificial Intelligence (AI) Commission on Competition, Inclusion, and Innovation is a non-partisan, multistakeholder, expert body established by the U.S. Chamber of Commerce Technology Engagement Center (C_TEC). The Commission's charge is to study and provide policy recommendations for establishing responsible leadership and a thriving workforce in an AI-driven economy. The Commission will host a series of public hearings across the country and internationally over the next year to guide the expected policy recommendations.

The Commission seeks public comment on the following Request for Information (RFI) on a rolling basis. The Commission seeks comment around fairness and ethical concerns around AI and global competitiveness. This RFI will help inform the Commission's work in developing strong bipartisan recommendations.

Addresses

Comments may be submitted by any of the following methods:

- *Email*: Comments in electronic form can be sent to AICommission@uschamber.com any of the following formats: HTML; ASCII; Word; RTF; or PDF. Please submit comments only and include your name and your organization's name. Also, please indicate if you are ok with your comments being made public in the email.

Supplemental Information

There is an ongoing discussion around the specific responsibility of the government in regulating artificial intelligence versus the duties of private entities that develop and consume AI. Many questions are still being debated around the optimal regulatory approach to providing the necessary transparency, explainability, and accountability of artificial intelligence that will not stifle innovation. For these reasons, the Commission on Competition, Inclusion, and Innovation is asking for comments on these critical issues to seek further clarity to develop bipartisan and durable policy recommendations.

Request for Information

The following list covers the major areas about which the Commission seeks information. The listed areas are not intended to limit the topics that may be addressed.

Regulatory Frameworks: Effective AI regulation will ideally provide guardrails that build public trust and prevent societal harm; however, there is active debate on which philosophical approach is most appropriate. There appears to be a current consensus among regulatory authorities in the U.S. and E.U. toward risk-based approaches to AI regulation that, for example, recommends a lighter legal regime for AI applications with negligible risk, heavier auditing for applications with moderate to high risk, and complete banning of applications with unacceptable risk. The risk assessments are typically based on the AI application itself (e.g., facial recognition) or the context of use (e.g., healthcare). On the other hand, many civil society organizations criticize this approach for not addressing broad and grave risks to fundamental rights. They have called for an approach that puts the burden of proof on the entity wanting to develop or deploy the AI system to demonstrate that it does not violate the rights of a specific group or society at large, rather than their own operational risk.

- What regulatory frameworks other than a risk-based approach should the United States look at?

- What approach should the U.S. pursue? Please explain.

- Is a risk-based approach necessarily incompatible with a human-rights approach?

- How should "high risk" be defined? What factors should be considered? Is it linked to human rights violations? Is it linked to the individual's health, safety, and fundamental rights?

- How important is it for the U.S. to mirror regulatory standards of the E.U. and other key markets in order to support compliance developers operating in both markets?

- Are there any AI applications that should be banned entirely, whether from a risk or human rights-based perspective?

- If AI applications should be banned, how should such a ban be determined?

- Are there AI application risks that are more appropriately addressed by non-regulatory approaches? What existing authorities would be best positioned to deploy such frameworks, if any? Explain.

Transparency and Oversight Audits: Regulators or certified third parties would like to have a consistent set of standards and frameworks against which to measure the impact of AI systems to ensure they are safe, reliable, and equitable. Many proposed regulations like the draft E.U. AI Act require developers of high-risk AI systems to perform both pre-deployment conformity assessments and post-market monitoring analyses to demonstrate that their systems are in compliance. Others advocate for additional process-based assessments, as is done in the fraud and safety industries (e.g., OSHA's Process Safety Management standard. Process-based approaches emphasize adhering to and reporting on an agreed set of best practices to mitigate and communicate known risks, incidents, and near misses. However, there are many different methods for inspecting or auditing algorithms— each requiring different kinds of access and ways to measure outcome-based concepts such as fairness and bias, with the most appropriate approach depending on the use case.

- What metrics are appropriate in determining how one assesses the variations in the result?

- What pre-deployment and post market assessments should be considered? How can we make such requirements flexible enough to account for diverse use cases and applications?

- Outcome assessments can often look at the impact on the population as a whole and overlook larger impacts on vulnerable groups, such as individuals who are ethnic minorities, elderly, disabled, impoverished, digitally illiterate, etc. How can we ensure that regulation adequately protects the vulnerable?

- What, if any, liability protections should be provided to companies whose technology is misused (unintended use cases) in ways that are harmful to society? What diligence is required on the side of the developer to prevent unintended use cases?

- Is bias in AI systems truly eliminable?

- Should the government require companies to be transparent about the use of AI, and should the government be the clearinghouse for such private information?

- If the government is the right clearinghouse, how can the government conduct effective auditing across diverse use cases without unnecessarily revealing proprietary information about a training data set or exactly how a model works?

Explainability and Human-in-the-Loop as safety mechanisms: To protect against potentially harmful AI, many proposed regulations have suggested that "high risk" AI systems be designed to allow for oversight by humans who will be tasked with working alongside the system to help prevent or minimize risks. Often expressed as the "human-in-the-loop" solution, this approach requires creating explanations of AI behavior that humans can understand in order to detect bias or bad recommendations. However, bringing explainability is a challenging technical issue that generally involves a tradeoff between accuracy and interpretability: the most accurate

models are usually the most difficult to understand (e.g., deep neural networks) while the most easily interpreted are usually less accurate (e.g., linear regression). Some users would like explanations that reflect the way in which human decision-makers provide explanations, others note this simplification reduces accuracy. Moreover, too detailed explanations risk the potential to disclose commercially sensitive material or allow a system to be gamed. Finally, many fear humans are likely to easily misinterpret explanations or over (or under) estimate the AI's capabilities and trustworthiness.

- What should explainability metrics be considered for federal regulation? Should explainability requirements vary based on application or certain types of risk?

- What does "meaningful human oversight" mean and how should it be defined in federal regulation? When is it most needed and appropriate?

- What mechanisms are needed to address or protect employees (in both public and private settings) who might decide to override or otherwise challenge an automated decision?

- Given the challenges around explainability and intellectual property, what regulatory requirements would offer the most protection without hindering innovation?

- Does requiring a human-in-the-loop somehow shield producers from liability from faulty systems?

- How can we create regulations that empower individuals to act responsibly but do not hold them unfairly accountable for actions they lack the power or knowledge to own?

Transparency Requirements and Oversight of Government Use of AI: It is essential that any regulation provide adequate oversight of the government's use of privately-developed AI systems, given the important impact government decisions can have on human rights. While other regulatory proposals, such as the draft E.U. AI Act, include some transparency requirements to help the public understand how governments use AI, human rights organizations argue that more specific data is needed for meaningful transparency, such as the

names of specific government agencies using specific systems, dates of service, and what the system is being used for.

- What information should government agencies be required to disclose about AI systems they are using and for what purpose? What exemptions, if any, should be made for sensitive use cases?

- What responsibility should the government—versus the private entity that developed the AI—have for the impact of the application on society? What standards should be met before a government agency is allowed to take on a privately-developed system for use? What kinds of impact assessments should be required?

- Should there be an independent federal oversight body that is responsible for conducting regulatory inspections on government use of AI? If so, should a new body be created or are there current agencies that are well-positioned to conduct this work? How should such a body receive public concerns or respond to requests for review?

- Should citizens have distinct rights to appeal decisions facilitated by an automated system?

Fostering Innovation: Many fear that AI regulations will hinder innovation, particularly for AI applications that are deemed high risk but also may be critical to U.S. competitiveness vis à vis authoritarian regimes that are rapidly pursuing AI innovation.

- Please define and explain what United States competitiveness looks like when it comes to AI?

- What specific goals and metrics do you believe are relevant in determining global competitiveness?

- How can the U.S. help foster the development and retention of local expertise with AI R&D?

- Are there relevant innovations in intellectual property regulations that might help the U.S. maintain an edge over near-peer competitors?

- How should the U.S. posture itself on issues related to data localization and data privacy? And what are the innovation implications of its posture on these dimensions?

- How can the U.S. support R&D, data-sharing, testing, experimentation, regulatory sandboxes, and investment in high-risk AI areas?

- How can the U.S. use tax incentives and other tools to encourage companies to build or buy high-risk AI systems that adhere to regulations?

- If we want to be globally competitive, artificial intelligence and the productivity it can precipitate must be a focal point of both workforce preparation and production processes. What can the public sector do to keep the U.S. on top in this sector while minimizing collateral labor market damage?

- How can the U.S. reduce uncertainty around legislative implementation once rules are established? Should the government work together with industry bodies to translate regulatory goals into practical steps for compliance? What bodies are best positioned for this?

- How can the U.S. accelerate the development and adoption of AI quality management systems that legislation may require?

- On our current trajectory, do we know about the likely impact of AI on labor market needs?

- Should the U.S. advocate for a unified international approach to AI governance to circumvent barriers to trade and protect against the worst use cases and certain military applications?

- As artificial intelligence technologies continue to move at the speed of business in their development, complexity, and gain of function, how do we encourage businesses that consume such technology to use it to augment human labor inputs and productivity rather than replace them? To what degree should this be a concern and, if it is one, what policies can government adopt to ensure alignment of incentives for a pro-worker AI policy?