MERCATUS CENTER
George Mason University

PUBLIC INTEREST COMMENT

OMB'S AI GUIDANCE EMBODIES WISE TECH GOVERNANCE

ADAM THIERER
Senior Research Fellow, Mercatus Center at George Mason University

Request for Comments on a Draft Memorandum to the Heads of Executive Departments and Agencies, "Guidance for Regulation of Artificial Intelligence Applications"
Agency: Office of Management and Budget
Comment Period Opens: January 13, 2020
Comment Period Closes: March 13, 2020
Comment Submitted: March 13, 2020
Document No. 2020-00261

The Office of Management and Budget (OMB) has requested comments on a draft memorandum intended to provide "Guidance for Regulation of Artificial Intelligence Applications" (herein after the *AI Guidance*) as federal agencies consider regulatory and nonregulatory approaches for artificial intelligence (AI) technologies. The *AI Guidance* also asks agencies to consider ways to reduce barriers to the development and adoption of AI technologies.

The Mercatus Center at George Mason University is dedicated to bridging the gap between academic ideas and real-world problems and to advancing knowledge about the effects of regulation on society. This comment does not represent the views of any particular party or special interest group but is designed to assist OMB in creating a policy environment that will facilitate increased innovation, competition, and access to technology to the benefit of the public.

I appreciate the opportunity to comment on this guidance document. The *AI Guidance* represents a balanced approach to the governance of AI-based applications because it is rooted in humility, flexibility, and forbearance, which are the touchstones of wise emerging tech governance.[1] Specifically, we applaud the *AI Guidance's* admonition to agencies to "avoid regulatory or non-regulatory actions that needlessly hamper AI innovation and growth" and I applaud the requirement that "agencies should assess the effect of the potential regulation on AI innovation and growth" before deciding to regulate AI applications.[2]

Federal agencies have already examined applications of AI, notably autonomous and connected vehicles, and have so far elected to use "soft law" methods of governance, such as

1. Adam Thierer, Andrea Castillo O'Sullivan, and Raymond Russell, "Artificial Intelligence and Public Policy" (Mercatus Research, Mercatus Center at George Mason University, Arlington, VA, August 2017), 5.
2. Thierer, O'Sullivan, and Russell, "Artificial Intelligence and Public Policy," 2.

For more information, contact
Mercatus Outreach, 703-993-4930, mercatusoutreach@mercatus.gmu.edu
Mercatus Center at George Mason University
3434 Washington Blvd., 4th Floor, Arlington, VA 22201

The ideas presented in this document do not represent official positions of the Mercatus Center or George Mason University.

nonbinding guidance documents. While guidance documents and other modes of soft law allow for the regulatory flexibility and forbearance sought by the *AI Guidance*, they also raise concerns about accountability and transparency. Executive Order 13891, which will require each federal agency to create a webpage for the publication of its active guidance documents, is a good start. Yet, general purpose technologies like AI will surely have applications that fall under the jurisdiction of most federal agencies.[3] A single, cross-agency portal for guidance documents on driverless cars and other AI applications would allow for better coordination and traceability of soft law activities in the federal government.

A LIGHT-TOUCH APPROACH ENCOURAGES EXPERIMENTATION

The *AI Guidance* sets forth a vision for technological governance that is consistent with the principle of permissionless innovation. Permissionless innovation refers to the idea that "experimentation with new technologies and business models should generally be permitted by default. Unless a compelling case can be made that a new invention will bring serious harm to society, innovation should be allowed to continue unabated and problems, if they develop at all, can be addressed later."[4] The United States made permissionless innovation the basis of internet policy beginning in the early 1990s, and the principle drove the modern digital revolution.[5] US-based information technology companies became household names across the globe as a result. If policymakers wish to replicate America's success with the internet and e-commerce, they need to adopt a similar light-touch approach for the governance of AI technologies.[6] Policymakers can accomplish that by ensuring that the policy defaults toward AI applications are close to permissionless innovation instead of the precautionary principle.[7] The precautionary principle is the idea that innovations should be curtailed or disallowed until their developers can prove that they will not cause any harm to individuals, groups, specific entities, cultural norms, or various existing laws, norms, or traditions.

Thanks to sensible, bipartisan decisions made by Congress and the Clinton administration in the 1990s, the United States generally rejected a precautionary principle approach for digital computing and the internet.[8] Adoption of the precautionary principle as the policy default and its application through technocratic, top-down regulatory edicts tend to suffocate technological experimentation and stifle entrepreneurial opportunities and economic growth.[9] It is vital, therefore, that public policy regarding emerging technologies such as AI applications continue this policy tradition and generally favor a permissionless innovation approach over a precautionary-principle-based regulatory system.

---

3. Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics" (NBER Working Paper No. 24001, National Bureau of Economic Research, Cambridge, MA, November 2017).
4. Adam Thierer, *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom* (Arlington, VA: Mercatus Center at George Mason University, 2016).
5. Adam Thierer, "Embracing a Culture of Permissionless Innovation" (essay, Reviving Economic Growth: A Cato Online Forum, Washington, DC, November 17, 2014).
6. Adam Thierer and Andrea Castillo O'Sullivan, "Preparing for the Future of Artificial Intelligence" (Public Interest Comment, Mercatus Center at George Mason University, Arlington, VA, July 22, 2016).
7. Andrea Castillo O'Sullivan and Adam Thierer, "Counterpoint: Regulators Should Allow the Greatest Space for AI Innovation," *Communications of the ACM* 61, no. 12 (2018): 33–35.
8. Adam Thierer, "What 20 Years of Internet Law Teaches Us about Innovation Policy," *FedSoc Blog*, May 12, 2016.
9. Adam Thierer, "How Many Lives Are Lost Due to the Precautionary Principle?," *The Bridge*, October 31, 2019.

The *AI Guidance* wisely rejects such an approach as America's policy default for AI: "Agencies must avoid a precautionary approach that holds AI systems to such an impossibly high standard that society cannot enjoy their benefits," the memorandum notes. "Where AI entails risk, agencies should consider the potential benefits and costs of employing AI, when compared to the systems AI has been designed to complement or replace."

This represents a wise governance default for fast-moving AI markets. Increasingly, policymakers are rethinking their approach to the regulation of many technologies because of the so-called "pacing problem," which refers to the growing gap between the rate of technological innovation and the ability of public policy to keep up with it.[10] AI markets are unpredictable and constantly evolving. It is impossible to forecast the future and plan for all potential outcomes, good or bad. Public policies for emerging technology should not be premised on worst-case thinking that views each new development as a potential crisis. Instead, technological change should be viewed as an opportunity to improve existing social and economic systems and realities.[11] If problems do develop, agencies already possess considerable authority to deal with concerns in specific contexts. Until then, "agencies should consider forgoing regulatory action," as the *AI Guidance* advises. This approach exemplifies regulatory humility in action.

THE NEED FOR A FLEXIBLE, PRO-INNOVATION FRAMEWORK

When agencies do consider policy interventions for AI applications, regulators should keep the other essential governance virtues in mind: flexibility and forbearance. "Fostering innovation and growth through forbearing from new regulations may be appropriate," the *AI Guidance* correctly notes. However, the *AI Guidance* does not foreclose the possibility of agency action. The memorandum continues on to specify a set of 10 "principles for the stewardship of AI applications," as well as several potential "non-regulatory approaches."

The *AI Guidance* strikes a sensible balance by noting that agencies can "calibrate approaches concerning these principles and consider case-specific factors to optimize net benefits." This represents policy flexibility. The *AI Guidance* wisely does not impose a straitjacket on agencies that would completely foreclose their ability to respond to potential policy concerns. At the same time, the *AI Guidance* discourages rash action and requires that agencies consider a fuller range of potential responses instead of just immediate regulatory constraints.

This approach is very much consistent with executive orders implemented in previous administrations as well as past guidance from OMB. For example, OMB's *Circular A-4*, issued in September 2003, provides guidance to regulatory agencies to help guide their rulemaking activities.[12] The new *AI Guidance* is essentially an extension of these well-established procedures. Among the 10 principles itemized in the *AI Guidance*, the document stresses the importance of risk analysis and evaluation of costs and benefits. The *AI Guidance* rightly specifies, "the need for agencies, consistent with their authorities, to evaluate the benefits, costs, and distributional effects associated with any identified or expected method for accountability." This principle has animated

10. Adam Thierer, "The Pacing Problem and the Future of Technology Regulation," *The Bridge*, August 8, 2018.
11. James Broughel and Adam Thierer, "Technological Innovation and Economic Growth: A Brief Report on the Evidence" (Mercatus Research, Mercatus Center at George Mason University, Arlington, VA, February 2019).
12. Office of Management and Budget, *Circular A-4*, September 17, 2003.

cost-benefit analysis efforts for many decades now and should apply to any AI applications that federal agencies consider regulating.

Equally important is the call in the *AI Guidance* for "a risk-based approach . . . to determine which risks are acceptable and which risks present the possibility of unacceptable harm, or harm that has expected costs greater than expected benefits." The memorandum wisely observes that "[i]t is not necessary to mitigate every foreseeable risk" because "a foundational principle of regulatory policy is that all activities involve tradeoffs." Once again, this is consistent with the requirements of past executive orders and OMB guidance on regulatory policymaking.

In sum, the *AI Guidance* leaves the door open for policy interventions while recommending regulatory humility to ensure that innovation continues.

SOFT LAW ALTERNATIVES MAKE SENSE

The *AI Guidance* identifies some of the specific concerns raised by advocates of more precautionary regulation of AI applications. Those issues include fairness and nondiscrimination, transparency, and various safety matters. The *AI Guidance* advises agencies to consider balanced responses to such concerns. It also identifies "non-regulatory approaches that may be appropriate to address the risk posed by certain AI applications."

Such nonregulatory approaches can help agencies seek out less restrictive remedies to complex social and economic problems before resorting to top-down proposals that might deter important AI innovations.[13] Specifically, the *AI Guidance* encourages agencies to consider three particular nonregulatory models: (1) sector-specific policy guidance or frameworks, (2) pilot programs and experiments, and (3) voluntary consensus standards.

These nonregulatory models represent "soft law" approaches to technological governance. Soft law mechanisms are informal, collaborative, and constantly evolving governance efforts.[14] While not formally binding like "hard law" rules and regulations, soft law efforts nonetheless create a set of expectations about sensible development and use of technologies.[15] Soft law can include multistakeholder initiatives, best practices and standards, agency workshops and guidance documents, educational efforts, and many other governance strategies.

Soft law has become the dominant governance approach for emerging technologies because it is often better able to address the pacing problem. Not only do traditional legislative and regulatory hard-law systems struggle to keep up with fast-paced technological change, but oftentimes those older mechanisms are just too rigid and unsuited for new sectors and developments. That is definitely the case for AI, which is multidimensional in nature and defies easy definition.[16] Soft law offers a more flexible, adaptive approach to learning on the fly and cobbling together principles and policies that can address new policy concerns as they develop in specific contexts, without derailing potentially important innovations.[17]

13. Adam Thierer and Trace Mitchell, "Technological 'Governance' Requires a Balanced Approach," *The Bridge*, March 12, 2019.
14. Ryan Hagemann, Jennifer Huddleston, and Adam Thierer, "Soft Law for Hard Problems: The Governance of Emerging Technologies in an Uncertain Future," *Colorado Technology Law Journal* 17, no. 1 (2018): 37–129.
15. Jennifer Huddleston, Adam Thierer, and Ryan Hagemann, "'Soft Law' Is Eating the World," *The Bridge*, October 11, 2018.
16. Pei Wang, "On Defining Artificial Intelligence," *Journal of Artificial General Intelligence* 10, no. 2 (2019): 1–37.
17. Adam Thierer, "Trump's AI Framework & the Future of Emerging Tech Governance," *Medium*, January 8, 2020.

ADDRESSING THE DOWNSIDES OF SOFT LAW EFFORTS

The informality of soft law creates some dangers, however. When regulatory agencies take steps to influence private activities, those policy actions need to be transparent and accountable. This is why a strict set of procedural requirements guides federal rulemaking activity. Those requirements are set forth in the Administrative Procedures Act, the Federal Register Act, the Freedom of Information Act, the Federal Advisory Committee Act, and the OMB guidelines already mentioned. Such traditional rulemaking activities require the publication of any proposed rules in the US *Code of Federal Regulations*, typically followed by hearings, the opportunity for affected parties to present evidence, and a notice-and-comment process that welcomes public participation.

Of course, these formal procedural requirements also limit the flexibility of traditional hard-law mechanisms and make policy interventions harder to implement and less likely able to be responsive and timely. This is why agencies are tapping soft law mechanisms to address various emerging technology policy concerns. However, the informality of soft law leaves it open to abuse, unaccountability, or nontransparency.

As the governance of AI applications shifts toward soft law mechanisms, it will be important the agencies be transparent about their informal governance efforts. OMB should remind agencies of their obligations under Executive Order 13891, on "Promoting the Rule of Law Through Improved Agency Guidance Documents," which President Trump signed last October.[18] That executive order requires that agencies "treat guidance documents as non-binding both in law and in practice," and demands that each agency "establish or maintain on its website a single, searchable, indexed database that contains or links to all guidance documents" that the agency issues.

This is a wise policy because, as the executive order correctly notes, "Americans deserve an open and fair regulatory process that imposes new obligations on the public only when consistent with applicable law and after an agency follows appropriate procedures." Nonetheless, guidance documents and other soft law instruments can play a crucial role in filling governance gaps left by the absence of more formal regulatory enactments. Soft law efforts need to be more trackable and accountable.

SOFT LAW COORDINATION NEEDED

James Broughel of the Mercatus Center has recommended that, beyond "simply posting guidance on their own websites, agencies could house guidance on a central government-wide database."[19] He recommends that the federal government use a centralized tracking system modeled after Virginia's guidance document site, which is simple and easy to navigate.[20] The *AI Guidance* should be expanded to include such a system.

Coordinating and centralizing related guidance efforts is particularly important for AI-related matters because of the cross-cutting nature of AI technologies and policy concerns. Consider why coordination is important as it relates to just one specific AI application, autonomous vehicles, and the set of concerns surrounding it (security). In October 2016, the National Highway Traffic Safety Administration (NHTSA) released "nonbinding guidance to the

18. Exec. Order No. 13891, 84 Fed. Reg. 55235 (October 9, 2019).
19. James Broughel, "Shining a Light on Agency Guidance Practices," *Morning Consult*, October 18, 2019.
20. "Guidance Documents in Effect," Virginia Regulatory Town Hall, accessed February 17, 2020, https://townhall.virginia.gov/L/GDocs.cfm.

automotive industry for improving motor vehicle cybersecurity," which included various best practices for the sector.[21] A year later, in June 2017, the Federal Trade Commission (FTC) and NHTSA hosted a joint workshop on the "Benefits and Privacy and Security Issues Associated with Current and Future Motor Vehicles," or what was billed as a "Connected Cars Workshop."[22] The agencies were jointly concerned about the "hackability" of these vehicles, as well as the safety and privacy concerns that could result. At the beginning of 2018, the FTC issued a "staff perspective" that summarized the findings from that workshop and then again recommended a set of best practices to address security concerns.[23]

Other security-oriented soft law governance efforts have been ongoing that affect driverless car security. Over the past five years, the Department of Transportation (DOT) has been cobbling together informal "rules of the road" for driverless cars through informal guidance documents. These documents have been "versioned" as if they were computer software (e.g., version 1.0, 2.0, 3.0). Version 4.0 of the DOT guidance for automated vehicles was released earlier this year.[24] The security of autonomous systems has been a major focus throughout all four iterations of the guidance.

Meanwhile, in 2018, the National Institute of Standards and Technology (NIST) developed a "Framework for Improving Critical Infrastructure Cybersecurity," which "consists of standards, guidelines, and best practices to manage cybersecurity-related risk."[25] The first iteration of that framework was issued in early 2014, and then a revised version was released in April 2018. Like the FTC, NHTSA, and DOT, the NIST cybersecurity framework relied on a combination of workshops and agency reports to establish and constantly reinforce a set of cybersecurity best practices. Many of the principles found in the NIST framework would be applicable to autonomous vehicles.

In other words, these soft law activities, which span four agencies, all have an effect on the future of AI applications in autonomous vehicles (and the security of those systems, in particular). Yet the public and affected parties do not have access to a single "one-stop" portal for these various guidance documents related to driverless cars. This highlights the need for better coordination and traceability of soft law activities and documents.

CONCLUSION

The *AI Guidance* represents a continuation and extension of a growing trend toward more flexible, adaptive governance approaches for emerging technologies. The framework offers a pragmatic vision that builds on the policies and paradigms of the past while also encouraging fresh thinking about how best to balance the concerns associated with AI innovation alongside its ability to do profound good for society.[26]

---

21. National Highway Traffic Safety Administration, "Cybersecurity Best Practices for Modern Vehicles" (Report No. DOT HS 812 333, National Highway Traffic Safety Administration, Washington, DC, October 2016).
22. Federal Trade Commission, "FTC and NHTSA Seek Input on Benefits and Privacy and Security Issues Associated with Current and Future Motor Vehicles," press release, n.d., https://www.ftc.gov/system/files/attachments/press-releases/ftc-nhtsa-conduct-workshop-june-28-privacy-security-issues-related-connected-automated-vehicles/notice_connected_cars_workshop_with_nhtsa_1.pdf.
23. Federal Trade Commission, *Connected Cars Workshop: Staff Perspective*, January 2018.
24. National Science & Technology Council and US Department of Transportation, *Ensuring American Leadership in Automated Vehicle Technologies: Automated Vehicles 4.0*, January 2020.
25. National Institute of Standards and Technology, *Framework for Improving Critical Infrastructure Cybersecurity: Version 1.1*, April 16, 2018.
26. Adam Thierer and Connor Haaland, "The Clinton-Bush-Obama-Trump Innovation Vision," *The Bridge*, November 21, 2019.

A policy approach rooted in humility, flexibility, and forbearance will help ensure that America's regulatory policies continue to promote both innovation and the public good. We applaud the administration for adopting this sensible approach to technological governance.