

Response of the International Center for Law & Economics, Request for Information, National Priorities for Artificial Intelligence

Office of Science and Technology Policy

Docket ID: OSTP-TECH-2023-0007

July 7, 2023

Authored by:

Kristian Stout (Director of Innovation Policy, International Center for Law & Economics)

We thank the Office of Science and Technology Policy (OSTP) for this opportunity to provide regulatory commentary on the pivotal subject of artificial-intelligence (AI) regulation. AI technology, already a familiar part of American life, is poised to become among the most consequential technological advancements in the coming years. As the rate of innovation in AI technologies accelerates, there will be greater opportunity for an expanded spectrum of applications that increase social welfare. At the same time, we are cognizant of some potential risks that AI could pose.

The Biden administration has already taken commendable steps toward advancing innovation, safeguarding Americans' rights and safety, and ensuring that the public can benefit from AI. The updated National AI R&D Strategic Plan,¹ the blueprint for an AI Bill of Rights,² and the AI Risk Management Framework,³ among other initiatives, represent thoughtful efforts to grapple with the legal and social implications of AI technologies.

We firmly believe that the prime concern should be to avoid premature regulatory action. Each technology grouped under the broad umbrella of AI is unique and requires careful consideration and understanding on its own terms. It is crucial to take sufficient time to study these important distinctions and appreciate the specific challenges and opportunities inherent in each. Overarching or rushed regulations could stifle innovation, impede economic growth, and inadvertently undermine efforts to realize AI's transformative potential.

Furthermore, when contemplating the adoption of a risk-based regulatory framework, we propose that the OSTP steer clear of overreliance on the precautionary principle. While intended to anticipate potential risks, the precautionary principle can over-index in the direction of caution and, due to its inherently conservative nature, serve as a barrier to innovation and progress. Instead, we recommend an approach that grounds any potential regulation in addressing real harms, with particular focus on preventing or minimizing those harms with a significant likelihood of occurring, that are more comprehensively understood, and that are tangible, rather than based on speculative or nebulous risks.

Developing a comprehensive national AI strategy is, indeed, a commendable undertaking and holds the promise that it could align various stakeholders' interests and offer a holistic approach to address AI's challenges. It is of paramount importance that this strategy remain responsive to the latest AI advances and global changes, considering the dynamic and evolving nature of AI technology. We are confident that the OSTP and the National AI Initiative Office will thoughtfully integrate the

¹ *National Artificial Intelligence Research And Development Strategic Plan 2023 Update*, SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE OF THE NATIONAL SCIENCE AND TECHNOLOGY COUNCIL (May 2023), available at <https://www.whitehouse.gov/wp-content/uploads/2023/05/National-Artificial-Intelligence-Research-and-Development-Strategic-Plan-2023-Update.pdf>.

² *Blueprint for an AI Bill of Rights*, WHITE HOUSE OFFICE OF SCIENCE AND TECHNOLOGY POLICY (2023), available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights>.

³ *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (Jan. 2023), available at <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

inputs provided through this Request for Information (RFI)⁴ to inform the National AI Strategy's development. We look forward to contributing our perspectives and suggestions to this critical dialogue.

Below we answer select questions in the RFI, we wanted to direct attention to a larger set of comments we submitted last month to the National Telecommunications and Information Administration's separate inquiry on this topic.⁵ Those comments are attached in full.

Understanding the Components of AI Must Come Before Regulation

1. What specific measures – such as standards, regulations, investments, and improved trust and safety practices – are needed to ensure that AI systems are designed, developed, and deployed in a manner that protects people's rights and safety? Which specific entities should develop and implement these measures?⁶

Before deciding what standards are necessary to regulate AI, it is necessary to develop some meaningful definition of what "AI" means. The present enthusiasm for AI has led to an oversimplification in the public discourse that can obscure how diverse the underlying technologies and their respective applications actually are. AI, in fact, covers a spectrum of technologies from large language models⁷ to recommender systems⁸ and beyond. These applications differ significantly from some of the more extravagant conceptions of AI, such as artificial general intelligence (AGI). A failure to distinguish among these technologies and their particular use cases can result in what we refer to as "regulatory overaggregation"—that is, a regulatory generalization that clouds the distinct aspects of each technology and may fail to address actual harms due to an inability to adequately address granular subjects.

The contemporary urge to overgeneralize the regulation of AI has parallels with the domains of "privacy rights" and "privacy regulation," where sharply divergent potential harms are often

⁴ *Request for Information National Priorities for Artificial Intelligence*, 3270-F1, 88 FR 34194, WHITE HOUSE OFFICE OF SCIENCE AND TECHNOLOGY POLICY (May 26, 2023) ("RFI").

⁵ Kristian Stout et al., *ICLE Response to the AI Accountability Policy Request for Comment*, International Center for Law & Economics (Jun. 2023), <https://laweconcenter.org/resources/icle-response-to-the-ai-accountability-policy-request-for-comment> ("ICLE NTIA Comments").

⁶ RFI at 34195.

⁷ LLMs are a type of artificial-intelligence model designed to parse and generate human language at a highly sophisticated level. The deployment of LLMs has driven progress in fields such as conversational AI, automated content creation, and improved language understanding across a multitude of applications, even suggesting that these models might represent an initial step toward the achievement of artificial general intelligence (AGI). See Alejandro Peña et al., *Leveraging Large Language Models for Topic Classification in the Domain of Public Affairs*, ARXIV (Jun. 5, 2023), <https://arxiv.org/abs/2306.02864v1>.

⁸ Recommender systems are advanced tools currently used across a wide array of applications, including web services, books, e-learning, tourism, movies, music, e-commerce, news, and television programs, where they provide personalized recommendations to users. Despite recent advancements, there is a pressing need for further improvements and research in order to offer more efficient recommendations that can be applied across a broader range of applications. See Deepjyoti Roy & Mala Dutta, *A Systematic Review and Research Perspective on Recommender Systems*, 9 J. BIG DATA 59 (2022), available at <https://journalofbigdata.springeropen.com/counter/pdf/10.1186/s40537-022-00592-5.pdf>.

conflated under the same broad topic. The concept of privacy often invokes an expectation of seclusion or allowing an individual to control their personal information.⁹ This framing, however, is too general and cannot capture all actionable areas of law that implicate privacy, such as "revenge porn" or the unauthorized sale of cellphone location data. Overaggregating these distinct issues under a unified "law of privacy" may lead to regulations that fail to properly address each concern.

On the other hand, the domain of intellectual property (IP) demonstrates a more nuanced approach. Though it covers an array of legal constructs like copyright, patents, and trademarks, each area has specific legislation addressing unique rights, harms, and remedies. This approach fosters legislative richness and avoids the pitfall of overaggregation.

Lessons from both privacy law and intellectual property may be instructive for AI. Overly broad AI regulations risk stifling innovation and technological advancement, while potentially failing to address specific harms. Therefore, rather than a blanket regulatory approach, a detailed understanding of AI's various subdomains is needed to target identifiable harms. This could be aided by OSTP facilitating the development of a comprehensive catalog of AI technologies and their potential risks, which could serve as a reference for regulators and courts.

Emphasize Harm-Based Approaches to AI Regulation and Require Cost-Benefit Analysis

Drawing upon the challenges associated with regulating emergent technologies such as AI, we could begin to explore this domain by considering an analogy to an older technology: photography. If camera technology were nascent, we might project myriad potential harms. But we can reflect from our position of having nearly two centuries of experience with the technology that a universal regulatory framework to manage all aspects of camera technology would be absurd. Instead, existing general laws more adequately address the specific harms that can be facilitated by camera technology, such as infringements on privacy rights arising from covert filming, use in the furtherance of criminal enterprises, or theft of trade secrets. In these instances, it is not the camera technology itself that forms the subject of legal concern, but the illicit actions carried out through its use.

Further, when assessing potential harms facilitated by new technology, a comprehensive analysis must consider the balance between the likelihood of harmful uses and the prospects of beneficial applications. Copyright law, as exemplified in the landmark *Betamax* case,¹⁰ provides an insightful precedent. That case illustrated how law could adapt to new technology, in that instance underscoring the need for copyright law to accommodate "substantial noninfringing uses" of new technologies that could reproduce protected material.¹¹ The decision upheld that, while the

⁹ The prototypical framing of this view is captured by the seminal work by Samuel D. Warren & Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193 (1890).

¹⁰ *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 439 (1984).

¹¹ *Id.* In this case, the Supreme Court imported the doctrine of "substantial noninfringing uses" into copyright law from patent law.

technology may facilitate some infringement, it would be inappropriate to apply a broad presumption against its use.¹² Moreover, the case stressed the importance of examining each circumstance on a case-by-case basis.¹³

Regulation and accountability in the realm of AI should echo this approach, emerging organically through bottom-up, case-by-case processes that examine the relevant facts of any given situation and how they alter (or do not alter) our legal system's baseline assumptions. New legislation, if required, should be incremental, guided by well-defined principles, and focused on identifiable harms, thus allowing law to fit specific circumstances without conflicting with established legal and regulatory principles.

AI, like any tool, can be misused, and any such misuse should incur legal consequences. Yet, the legal analysis should focus primarily not on the AI itself, but on the malefactors' actions and the resulting harms. Attempting to construct a foolproof regulatory framework that precludes the misuse of AI may prove futile and could potentially stifle the development of socially beneficial tools.

Moreover, the fact that AI technology remains largely in the research and development phase complicates regulatory decisions. Proactive regulation based on the precautionary principle might thwart unforeseen benefits that could emerge as these technologies mature and find unique applications.¹⁴ Even in high-risk industries like nuclear power, precautionary regulation often results in net social harms.¹⁵

When imagining the harms that could occur, it is crucial to distinguish two broad categories of AI-related concerns. First is the largely theoretical fear associated with AGI—the understandable apprehension many feel about inadvertently creating a superintelligence that could potentially extinguish human life.¹⁶ If it is even possible to create AGI, about which there remains significant doubt, it is crucial to emphasize that current AI technologies are far from AGI. AI technologies today are essentially sophisticated prediction engines for dealing with text or pixels.¹⁷ It is highly unlikely

¹² *Id.*

¹³ *Id.*

¹⁴ See ADAM THIERER, PERMISSIONLESS INNOVATION: THE CONTINUING CASE FOR COMPREHENSIVE TECHNOLOGICAL FREEDOM (2016).

¹⁵ See, e.g., Matthew J. Neidell, Shinsuke Uchida, & Marcella Veronesi, *The Unintended Effects from Halting Nuclear Power Production: Evidence from Fukushima Daiichi Accident*, NBER WORKING PAPER 26395 (2022), <https://www.nber.org/papers/w26395> (Japan abandoning nuclear energy in the wake of the Fukushima disaster led to decreased energy consumption, which in turn led to increased mortality).

¹⁶ See, e.g., Eliezer Yudkowsky, *Pausing AI Developments Isn't Enough. We Need to Shut it All Down*, TIME (Mar. 29, 2023), <https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough>.

¹⁷ See, e.g., Will Knight, *Some Glimpse AGI in ChatGPT. Others Call It a Mirage*, WIRED (Apr. 10, 2023), <https://www.wired.com/story/chatgpt-agi-intelligence> ("GPT-4, like its predecessors, had been fed massive amounts of text and code and trained to use the statistical patterns in that corpus to predict the words that should be generated in reply to a piece of text input.")

that we will accidentally stumble onto AGI by merely chaining thousands of these prediction engines together.

The second, more realistic set of concerns pertain to the misuse of AI technologies to perpetuate illicit activities. Specifically, these very impressive technologies might be misused to further discrimination and crime, or could have such a disruptive impact on areas like employment that they will quickly generate tremendous harms. When contemplating harms that *could* occur, however, it is also necessary to acknowledge the many significant benefits *also* could be generated. Moreover, as with earlier technologies, economic disruptions will provide both challenges and opportunities. It is easy to see the immediate effect on the jobs of content writers posed by ChatGPT, for example, but less easy to measure the benefits that would be realized by firms that can deploy this technology to “in-source” tasks. Thus, static analyses of AI’s substitution power are likely to miss the bigger picture of social welfare that could be realized as organizations improve their efficiency through the adoption of AI tools.

Finally, it is important to remember that dynamic competition—where technology is continually evolving and firms are competing to provide consumers with innovative products and services—drives far more economic growth than static competition. As the economist Joseph Schumpeter noted, competition thrives not merely on price but on the advent of disruptive new commodities, technologies, and supply sources.¹⁸

Regulation of AI must be seen in the same light. To this end, we advocate a regulatory regime for AI that encourages sector-specific rules to emerge when regulators discover that their existing rules are inadequate for new AI-augmented technologies. This approach should be harm-based, rather than risk-based. In other words, regulations should focus on mitigating the known and likely harms caused by the misuse of AI rather than trying to predict and prevent every possible risk associated with it. A clear-eyed cost-benefit analysis should guide this process.

Rather than preemptively stifling innovation with burdensome regulations based on hypothetical risks, a more nuanced approach would be to respond to actual harms as they arise, carefully weighing the potential harms against the prospective benefits of AI technologies. Such a balanced approach would not only protect society from misuse of AI but would also allow for the continued development and beneficial application of these transformative technologies.

Adopting this approach will require an ongoing dialogue among all stakeholders and an openness to adjust our regulatory frameworks as our understanding of AI and its societal impact deepens. A harm-based, case-by-case approach to AI regulation is consistent with our common-law tradition and promises to be the most effective and flexible approach to guide the development and application of AI technologies.

¹⁸ JOSEPH A. SCHUMPETER, CAPITALISM, SOCIALISM AND DEMOCRACY 74 (1976).

The Implications of a Centralized Regulator for AI: Risks to Competition and Innovation

1. ... Which specific entities should develop and implement these measures?¹⁹

The prospect of creating a centralized regulator for emergent technologies like AI raises important concerns, particularly those relating to market competition. A central regulator may inadvertently favor established industry players like OpenAI, as new entrants might be hindered by regulations and compliance costs, which incumbents could manipulate to increase rivals' costs.²⁰ The strategic promotion of a strong central regulator can thus serve to maintain or increase incumbents' market dominance.

In recent U.S. Senate hearings, some witnesses and senators proposed a central regulator to create and administer a licensing regime for AI.²¹ While licensing might be necessary for certain AI applications, such as military weaponry, it is broadly inadvisable due to the diverse nature of AI technologies. Developers of AI tools face numerous challenges, including assuring data collection and management, anticipating downstream usage of tools, and managing the complex chain of AI-system development and deployment. A centralized AI regulator would struggle to understand the nuances of each distinct industry, leading to ineffective or inappropriate licensing requirements.

Unlike such sectors as railroads and nuclear power, which have dedicated regulators, AI is more akin to a general-purpose tool, like chemicals or combustion engines. Different agencies regulate the use of these tools as appropriate for their context, without a central regulator overseeing every aspect of development and use. A licensing requirement could introduce undesirable delays into the process of commercializing AI technologies, significantly impeding technological progress and innovation, and potentially leaving the United States behind in the global AI race.

A better advisable approach would be to create product-centric and harm-centric frameworks that other sectoral regulators or competition authorities could incorporate into their rules for goods and services. For example, safety standards for medical devices should be maintained, whether or not AI is involved. But a thoughtful framework might raise questions that the Food and Drug Administration (FDA) finds are necessary to consider when implementing new regulations. This product-centric regulatory approach would ensure safety, quality, and effectiveness without stifling

¹⁹ RFI at 34195.

²⁰ This competition concern is one that is widely shared across the political spectrum. See, e.g., Cristiano Lima, *Biden's Former Tech Adviser on What Washington Is Missing about AI*, THE WASHINGTON POST (May 30, 2023), <https://www.washingtonpost.com/politics/2023/05/30/biden-former-tech-adviser-what-washington-is-missing-about-ai> (Tim Wu noting that he's "not in favor of an approach that would create heavy compliance costs for market entry and that would sort of regulate more abstract harms").

²¹ *Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. on Privacy, Technology, and the Law of the S. Comm. on the Judiciary*, 118th Cong. (2023) (statement of Sam Altman, at 11), <https://www.judiciary.senate.gov/download/2023-05-16-testimony-altman>.

innovation. With their deep industry knowledge, sectoral regulators are better positioned to address the unique challenges posed by AI technology within their spheres of influence.

By contrast, there is a risk that a centralized regulator, operating with an overaggregated concept of AI, might design rules that slow or prevent AI-infused technologies from coming to market if they cannot navigate the complex tradeoffs among interested parties across all such technologies.²² This could make society worse off and strengthen the position of global competitors. Therefore, it is crucial to approach the regulation of AI with careful consideration of its impacts on competition and innovation, advocating for a framework that encourages diversity and flexibility.

17. What will the principal benefits of AI be for the people of the United States? How can the United States best capture the benefits of AI across the economy, in domains such as education, health, and transportation? How can AI be harnessed to improve consumer access to and reduce costs associated with products and services? How can AI be used to increase competition and lower barriers to entry across the economy?²³

The advent of AI promises transformative potential across various domains, heralding numerous benefits for the people of the United States and beyond. Foremost, AI can drastically improve worker efficiency. Advanced AI algorithms could handle repetitive tasks swiftly and accurately, allowing employees to focus on more complex and strategic aspects of their jobs. In sectors ranging from manufacturing to health care to customer service, AI-driven automation can accelerate processes, minimize errors, and enhance productivity, ultimately leading to improved business performance and growth.

For instance, in health care, AI can help practitioners analyze complex medical data rapidly, improving diagnostic accuracy and speed. In manufacturing, AI-powered machines can manage labor-intensive tasks, reducing the possibility of human error and occupational injuries. These efficiencies can reduce costs, with the potential for savings to be passed on to consumers.

Furthermore, AI technology, like many disruptive technologies before it, may be capable not only of augmenting existing workforces but also of fostering new types of industries and opportunities. As AI becomes more sophisticated, we anticipate the emergence of entirely new job categories, similar to how the advent of the internet spurred professions in web design, digital marketing, and e-commerce.

AI can also improve consumer access to, and reduce costs associated with, various products and services. For instance, we have already seen AI-powered recommendation systems personalize the shopping experience, allowing consumers to find relevant products with ease. And in education,

²² This is a well-known problem that occurs in numerous regulatory contexts. See, e.g., Raymond J. March, *The FDA and the COVID-19: A Political Economy Perspective*, 87(4) S. ECON. J. 1210, 1213-16 (2021), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8012986> (discussing the political economy that drives bureaucratic agencies' incentives in the context of the FDA's drug-approval process).

²³ RFI at 34196.

we've seen AI personalize learning for individual students, tailoring educational content to match each learner's needs and pace and, in turn, improving educational outcomes and accessibility.

The promise of AI extends to increasing competition and lowering barriers to entry across the economy. By providing businesses with more information and greater efficiency, AI can give rise to more effective business strategies and models. It could level the playing field for small and medium-size enterprises, allowing them to compete with larger corporations by offering cost-effective solutions that previously required significant capital or resources.

19. What specific measures – such as sector-specific policies, standards, and regulations – are needed to promote innovation, economic growth, competition, job creation, and a beneficial integration of advanced AI systems into everyday life for all Americans? Which specific entities should develop and implement these measures?²⁴

As noted above, we believe that specific measures to promote innovation and the safety of advanced AI systems are best approached with a sector-specific focus. Due to the diverse nature of AI applications and the varying impacts on diverse industries, sector-specific policies and standards will be more effective and beneficial than broad, sweeping regulations.

For instance, in the health-care sector, safety and privacy standards must be upheld when deploying AI tools for diagnosing diseases or managing patient data. In such cases, regulators like the FDA or the Department of Health and Human Services could leverage their expertise to develop and implement targeted regulations that ensure safety without stifling innovation.

Similarly, in the automotive sector, where AI is used for autonomous vehicles, transportation authorities could create guidelines and standards to ensure road safety, while also promoting innovation. In finance, where AI algorithms are used for trading, credit scoring, and risk management, the Securities and Exchange Commission (SEC) and other relevant financial regulators can establish rules to prevent unfair practices and ensure market stability.

Conclusion

We again thank the OSTP for initiating this important and timely inquiry into AI regulation. It is through dialogues like these that we can collectively explore AI's impacts on society. It is crucial to reiterate that regulation, while necessary, should be formulated with a nuanced understanding of the technology. Being eager to impose regulations prematurely could stifle the very innovation that we seek to cultivate and the potential benefits that we aim to harvest. AI has the potential to be a transformative force for the United States and the world, providing a multitude of benefits, and empowering us with the tools to address some of the most pressing challenges of our time. A measured and informed approach to AI regulation would further reinforce our nation's position as a global leader in technological innovation.

²⁴ RFI at 34196.