

HOW DO WE PREPARE THE GOVERNMENT FOR AI?

MATTHEW MITTELSTEADT

Research Fellow, Mercatus Center at George Mason University

Draft memorandum titled *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (AI)*

Agency: Office of Management and Budget

Comment Period Opens: November 3, 2023

Comment Period Closes: December 5, 2023

Comment Submitted: December 5, 2023

Docket No. 2023-24269

INTRODUCTION

As President Biden said in the AI executive order he announced on October 30, “Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure.”¹ The promise of this technology, for both private and public application, cannot be understated. As the OMB revises its guidance, it is essential that the office both embrace and disseminate an understanding that, for public-sector application, AI is a collection not of “what ifs” but of real, transformative applications that can and must be employed today. Already, we are seeing the tip of the iceberg for truly impactful and beneficial relevant public-sector application of this technology:

- **In some cases, AI can supercharge key government services.** In November, DeepMind introduced GraphCast, a weather-forecasting algorithm that can produce higher accuracy predictions in one minute than the National Weather Service’s gold standard models produce in hours.² Implementation of these AI models would have the added bonus of dramatically lowering the National Weather Service’s staggering carbon toll.
- **In other cases, AI can help further government priorities.** Using AlphaFold, DeepMind solved the notoriously complex *protein folding problem*, providing crucial biological understanding that may supercharge rapid, ultra-targeted vaccine and drug development during future pandemics.³

¹ Exec. Order No. 14110, 88 Fed. Reg. 75191 (October 30, 2023).

² Google DeepMind, “GraphCast: AI Model for Faster and More Accurate Global Weather Forecasting,” (blog), November 14, 2023.

³ Google DeepMind, “Technologies, AlphaFold” (website), accessed December 4, 2023, <https://deepmind.google/technologies/alphafold/>.

- **Critically, AI is also helping officials manage the now unwieldy bloat of government rules.** To distill and manage its policies—a cumbersome body of text longer than 100 copies of Leo Tolstoy’s 500,000-word novel *War and Peace*—the Department of Defense developed GAMECHANGER, a program that parses and navigates these policies, a task no longer humanly possible.⁴

Beyond these already mature, relevant public-sector applications, there are countless other potential use cases: Powerful agency chatbots could cost-effectively serve constituents on a 24/7 basis; conduct consular interviews; qualitatively deepen census data; monitor budgetary waste, fraud, and abuse; and provide instant translation for non-English speakers. To harness this potential, and the many mature AI applications, I urge the administration and the OMB to pivot toward a “diffusion-centric AI strategy,” focused not on rules-for-rules’ sake but on guidance aimed at building trust, fostering engagement, and pushing the broad, rapid diffusion of this technology across federal agencies and society where applicable.⁵

The OMB should also recognize that the significance of getting this right for the federal government goes beyond just seizing these AI opportunities. While 71 percent of government executives believe the benefits of emerging AI applications outweigh the risks,⁶ a sense of AI skepticism and fear has taken root in the broader workforce. Today’s government workers will be tomorrow’s decision-makers, and their view of AI will have deep impacts on future policies and implementation efforts. To counter AI skepticism and fear of the AI unknown, AI diffusion is essential. Only through direct hands-on experience with AI can the government workforce understand its benefits. If public-sector AI diffusion lags, we will miss this hands-on experience: interest in AI’s promise will likely fade and long-term efforts will likely fail.

With this important goal of paid AI diffusion in mind, below I highlight certain concerns and opportunities the OMB should consider to improve this guidance and serve a diffusion-centric implementation strategy:

1. Provide education on AI potential, use cases, and mature applications.
2. Consider opportunity costs for minimum practices for safety-impacting or rights-impacting AI.
3. Recognize excessive regulation as a barrier to responsible use.
4. Measure AI diffusion.

1. PROVIDE EDUCATION ON AI POTENTIAL, USE CASES, AND MATURE APPLICATIONS

Broad AI success requires energy, engagement, and understanding of the possible. Today, many in the government workforce may not understand how AI can be used and why this technology may matter in their work; for most, the concept of “AI” is often falsely limited to chatbots like ChatGPT.

For agencies and their staff to successfully identify and imagine use cases, they need to understand the wide range of AI capabilities beyond chatbots. I recommend that the OMB require each agency to hold educational sessions for staff at all levels on mature AI technologies—both their range and ability to

⁴ Ken Klippenstein, “Pentagon’s Budget Is So Bloated That It Needs an AI Program To Navigate It,” *The Intercept*, September 20, 2023.

⁵ Jordan Schneider and Matthew Mittelsteadt, “The Key to Winning the Global AI Race,” *Noema*, September 19, 2023.

⁶ Fedscoop, “Government Gears Up to Embrace Generative AI,” October 17, 2023.

improve work. If implemented well, such sessions should help generate a sense of engagement and excitement, emboldening agencies and their staff to *want* AI rather than to feel compelled to use it.

2. CONSIDER OPPORTUNITY COSTS FOR MINIMUM PRACTICES FOR EITHER SAFETY-IMPACTING OR RIGHTS-IMPACTING AI

Responsibility is essential when implementing AI technologies, but the OMB's guidance must recognize that attaining true AI safety and responsibility is not equivalent to providing a list of rules. Opportunity costs, implementation burdens, and the potential negatives of low-tech alternatives also matter. While it's unclear whether the eight new minimum practices required for AI applications listed in this guidance will be burdensome or not, I urge the OMB to err on the side of a light touch. If process burdens are flexible, agencies will have the freedom to experiment and choose the most responsibility-enhancing IT paths forward.

To illustrate, consider the opportunity cost of forgoing machine-translation services. While human translators may be more familiar and therefore more trusted, by over-relying on this limited human labor the government will necessarily forgo the ability to make translation services accessible in real time to a diverse audience. In this case, delaying or blocking AI application may indeed be the irresponsible option.

Recognition of such opportunity cost means embracing that AI may indeed be the safest, most efficient, low-cost, and responsible option. Recognition of opportunity costs is not tantamount to blind AI boosterism; embracing this principle, likewise, means recognizing that in many cases standard digital options might indeed be the most responsible best path forward. These balanced considerations of AI responsibility and opportunity costs, however, may not be possible if excessive processes create a burden of longstanding IT implementations, bureaucratic hoops, and millions of taxpayer dollars. Faced with high effort and high costs, the ability to weigh AI benefits against traditional capacities may be removed from decision-makers, forcing them to rely on older technologies and capacities that might limit their ability to successfully carry out government functions and reap the rewards of innovation.

3. RECOGNIZE EXCESSIVE REGULATION AS A BARRIER TO RESPONSIBLE USE

I commend the guidance's emphasis on removing needless barriers to responsible AI. Building AI capacity indeed requires recognizing that business as usual, and existing capacities, are likely insufficient. What's missing, however, from the OMB's list of barriers is perhaps the biggest barrier of all: bureaucratic process burdens. As I argued in *Noema Magazine*,

Today's bureaucratic culture is often structured like a waterfall: Policymakers at the top of the falls make one-way decisions that flow down on top of developers and project managers who, operating under rigid strictures, struggle to build successful systems. The result is a culture of IT inflexibility, where developer ingenuity is bound by rules written by often inaccessible decision-makers. Rather than succeed at policy outcomes, systems are instead designed to follow rules.

Former Obama-Era Deputy Chief Technology Officer and Code for America founder Jennifer Phalka has concurred, arguing that the key to successful policy and IT implementation success lies in "peeling back the layers of process that have accumulated on top of policy."⁷

⁷ Jennifer Phalka, "Culture Eats Policy," Niskanen Center, June 21, 2023, <https://www.niskanencenter.org/culture-eats-policy/>.

To achieve AI success, therefore, the OMB must adjust this guidance to recognize regulatory process as a barrier. The OMB must emphasize to agencies that flexible rules and the removal of needless processes will be the key to experimentation, appropriate design, and rapid deployment.

To build on this emphasis, one potential opportunity the OMB should consider is requiring agencies to create “regulatory burden notification channels” between the low-level staff who pen the code and contracts and the high-level officials who have power to adjust the processes. With such channels in place, process barriers can be revealed and systems adjusted, allowing the system to accommodate implantation needs. In support of the administration’s emphasis on AI responsibility, providing such flexibility will allow implementors to experiment, build, and iterate on systems to maximize safety and results, rather than build to maximize adherence to rules.

4. MEASURE AI DIFFUSION

Measurement of AI diffusion will be essential for an effective rollout. In the OMB’s guidance, as well as in the executive order, agencies are required to annually take stock of their AI systems inventory. What these inventories miss, however, are key metrics relating to the “depth of adoption.” Are these systems being used, or are they collecting dust? Are staff failing to adapt to new AI tools? Are certain departments or offices lagging in AI uptake? Is there perhaps an overreliance on AI automation? All of these questions are important to identify capacity gaps, opportunities for change, and, for the public, transparency about the depth of AI automation in government services.

To capture essential metrics and provide the nuanced detail needed to support whole-of-government efforts, the OMB should require agencies to develop depth of adoption metrics, to illustrate not only what systems they have, but when and how deeply they are being used. Naturally, pushing for the development of such metrics may be beyond the expertise of certain agencies. To aid this process, I recommend the OMB, and perhaps expert agencies like the National Institute of Standards and Technology (NIST), develop default depth of adoption metric recommendations individual agencies can use as a template for their own metrics.

SEIZING OPPORTUNITY

The federal government must embrace the unique nature of the moment. AI is here, yet the AI future remains uncertain. While driving in the dark of AI policy, the OMB must embrace the fact that present uncertainty demands future flexibility. We cannot say what future technologies will emerge and problems will arise in the coming years. Just as the European Union’s AI Act failed to predict the rise of ChatGPT, any rule written today will no doubt fail to account for future developments. Prediction remains a challenge, but what we can say is that AI diffusion and harnessing the positive benefits of proven AI technology must be the end goals. Charting this course therefore means building an AI-ready government, with the the talent, process flexibility, and measured data needed to deploy this tech, experiment, and meet unexpected challenges as they emerge.

Working towards this goal and managing unexpected issues require building an AI-ready government, armed with an AI-educated workforce, low process burdens, and measurements needed to understand, confront, and nimbly solve future AI-related challenges.