MERCATUS CENTER
George Mason University

POLICY BRIEF

Lost in (Machine) Translation: A Case Study in Enabling Flexible Government AI Risk Management

Matthew Mittelsteadt

September 2024

In April 2024, the Office of Management and Budget (OMB) released its AI governance, innovation, and risk management guidelines for federal agencies.[1] While advertised as a light touch, what these guidelines offer is a top-down, rules-bound approach to federal AI risk management. To comply with these rules, most federal agencies with AI systems will have to check off 15 new process requirements and standards on top of the already onerous layers of pre-existing regulations that bind federal IT.[2] While well meaning, the rules reflect a common misjudgment in AI and tech policy: the hope that with the right processes in place, policymakers can ensure technology is used judiciously and its risk managed.

What this blunt, top-down approach misses, however, is context. This is a problem. As noted in the National Institutes of Standards and Technology (NIST) AI risk management framework,[3] "AI systems are not inherently bad or risky" but "it is often the *contextual environment* that determines whether or not negative impact will occur [emphasis added]." True risk management is not just sorting tech between "good AI" and "bad AI," but instead demands consideration of the context in which the AI application functions—the "who, what, when, why, and where"—that most often must be done on the ground. Unfortunately, top-down rules like the OMB's preclude such in-context risk analysis in favor of bureaucratic treatments that assume all systems are "guilty until proven innocent" and incentivize procedural compliance over safety, reliability, and quality.

The result of this bureaucratic approach? AI diffusion is stifled, actual risk management effectively barred, and services and safety degraded.

Unfortunately, this scenario is already playing out in real-world cases. To understand how such blunt top-down rules constrain risk management, consider the case study of machine translation, a capable technology now lost in the crosshairs of new rules.

A *Pretty Good* Technology Meets a *Pretty Bad* Crisis

Though uniquely useful, AI machine translation isn't new, and years of familiarity with it may have obscured how much recent progress has been made in its development, showing dramatic improvements from the nonsensical translations that persist in popular memory. According to the Association of Language Companies, in 2023, almost two-thirds of translation companies used machine translation in their workflow, while 70 percent of them said they plan to accelerate the adoption of automated processes by 2025. The quality translators have recognized is backed up by empirical results. A 2020 study published in *Nature Communications* found that English–Czech machine translation surpassed human translators at preserving the meaning of translated news articles.[4] This remarkable feat not only predates the ChatGPT revolution but did so hampered by the relative scarcity of quality Czech–English training data.

Despite machine translation's utility, the new OMB AI risk management policy targets it with a uniquely hard-edged regulatory treatment. Per the policy, any system ". . . providing live language interpretation or translation, without a competent interpreter or translator present, for an interaction that directly informs an agency decision or action" is subject to the OMB's extensive new rules-based compliance process. Even after machine translation systems are approved for use, the rules further require continuous monitoring by a competent interpreter, thereby capping potential.

Though well-intentioned, these rules will only hinder government as it attempts to address its current translation capacity crisis. Although the US government serves a pluralistic, multilingual country where 6.7 million adult citizens have limited English proficiency, federal investments in linguistic accommodation are minimal. The Office of Workers' Compensation Programs (OWCP) is a good example. Despite operating four disability compensation programs in which speakers of nine languages are regularly encountered, OWCP has yet to translate its website or most of its key documents (beyond a handful of brochures) into languages other than English. This crisis extends beyond retail-level services. A 2017 GAO report on the Department of State's linguistic capacity identified major gaps, finding that 23 percent of foreign service officers did not possess the language skills and capabilities required for their roles. These agencies are hardly unique. Across most agencies, such translation gaps are exceedingly common, limiting the federal government's ability to serve its role and citizens.

These gaps are big, and AI has evolved to the point where it can fill at least some of them. Unfortunately, this potential is lost in the OMB's flat, top-down rules, inhibiting agencies' ability to weigh the technology's potential and risks against the risk of providing no translation at all. For offices like the OWCP, the cost of inaction can be denial of cash benefits. For others, the cost may result in circumstances that are truly life or death.

Safety Costs

In 2017, Puerto Rican citizens felt the pain of translation capacity gaps firsthand. That year, Hurricane María inundated Spanish-speaking areas of the territory with floodwaters while resulting power grid failures left millions without electricity for months. In this moment of crisis, language barriers stood in the way of help. Analyzing the response of the Federal Emergency Management Agency (FEMA), a 2022 US Commission on Civil Rights investigation[5] noted that "FEMA did not have enough Spanish-speaking employees to accommodate the Island" and "sign linguists sent to Puerto Rico often only knew ASL [American Sign Language]–English, rather than the Spanish variant used generally throughout the island." As a result, "Spanish-speaking Puerto Ricans received disproportionately lower amounts of assistance for María recovery than English-speaking mainland Americans received." The situation was so egregious that the Department of Justice ruled FEMA was in "clear violation of federal court precedent and EEOC Guidance," and that their inability to provide for Puerto Ricans represented a general failure to comply "with the principles underlying Title VI" of the Civil Rights Act. Not only did capacity failures cause harm, but they also resulted in violations of the law.

In the aftermath of the 2012 attack on the US Embassy in Benghazi, the State Department's Accountability Review Board found a similar glaring failure, noting that "the lack of Arabic skills among most American personnel assigned to Benghazi and the lack of a dedicated, locally employed staff interpreter and sufficient local staff served as a barrier to effective communication and situational awareness at the mission." Limited in its abilities to communicate in the local language, the embassy was helpless to anticipate and perhaps prevent the tragic death of many, including US ambassador J. Christopher Stephens.

These two cases illustrate the steep opportunity costs of inaction. While machine translation tech may not have been mature in 2012 at the time of the Benghazi Embassy attack, in 2024 neither the State Department nor FEMA have yet fully addressed and resolved the translation capacity issues that were, and may still be, matters of life or death.

Through excessive concern about theoretical safety risks, the federal government's top-down, one-size-fits-all rules about how AI is to be used are failing to prevent, or even mitigate, very real, very present safety risks. In the absence of any other good alternatives for contingencies of translation, AI could offer a workable, if sometimes imperfect, solution.

Enabling In-Context Risk Management

To suggest that AI tools like machine translation are perfect and should go ungoverned without concern for risk would be unwarranted. But when regulation is formed from the top down, as the OMB's ruleset has been, risk management is constrained, not strengthened. As an alternative approach, government should make fewer one-size-fits-all rules, avoiding them when possible, instead enabling and fostering context-aware risk management at lower levels. Policymakers at

higher levels of government should take on the role of facilitator, not rule maker, for lower-level agencies adopting and implementing AI systems into their work and programs. In the OMB's rules, we see glimmers of what this could look like. Rules and policies that encourage interagency coordination and require each agency to have an AI officer responsible for developing and maintaining an agency AI plan enable agencies to take action on AI risk management while remaining flexible on the details of those actions. Through such rules and policies, the OMB is not dictating how agencies should manage risk, but lightly pushing them to facilitate the process throughout their operation.

While lower-level agencies have better use-case contextual knowledge, what can be lacking is effective information about the tech at hand, its systemic risks, and best practices for using it. To empower lower-level actors, governments should focus on producing voluntary risk management and implementation processes. Public sector–tailored versions of the NIST Risk Management Framework, for instance, could be a worthy start. Coordination is also key. At the top of government, the executive branch should focus on facilitating coordination and communication between agencies and offices, disseminating threat analysis and compilations of best practice, augmenting specific on-the-ground understanding with broader situational awareness, and enabling these disparate actors to learn from each other.

Government should also consider enabling responsible adoption at lower levels. When a technology has clear utility, officials should assume it will be used despite efforts to limit application. Many employees will be familiar with AI applications, so it should be assumed that they *will* be using work-related AI tools in the office, though perhaps not in the safest way. This risk can be best managed by providing staff with the best, most secure AI tools available and appropriate training to ensure they understand both the benefits and the risks that come from using this technology.

Conclusion

As governments consider new processes, facilitating risk management, not altogether eliminating it, should be the first-order goal. Rules and regulations are always needed and will indeed be part of any effective approach. But if we truly want the risk-based approach the OMB and other rule makers advocate, policymakers must consider empowering, not binding, implementers to understand and assess these technologies by digging in deep. While it will come with inevitable missteps, a looser leash will delegate and encourage the kind of ownership of risk, thoughtful application, and creative experimentation that will unlock not only AI safety but the untold benefits of rapid, responsible AI diffusion.

About the Author

Matthew Mittelsteadt is a research fellow and technologist for the AI and Progress Project at the Mercatus Center at George Mason University. His work highlights the importance of AI diffusion and of ensuring that emerging AI technologies yield a net benefit. His research focuses on AI

regulatory design and measurement, critical infrastructure and cybersecurity, and the unleashing of rapid AI diffusion. As a scholar, he places special emphasis on grounding policymaking through education and helping policymakers understand AI through seminars and Mercatus's *AI Policy Guide*.

Mittelsteadt's work has been published by *The Hill*, *Noema Magazine*, and the Federal Judicial Center. It has been cited by media including *The New York Times*, *Bloomberg*, *Foreign Policy*, and *Politico*. Before joining Mercatus, he was a fellow at the Institute for Security Policy and Law, where he focused on the intersection of AI and national security. He holds a BA in economics from St. Olaf College, an MPA from Syracuse University, and an MS in cybersecurity from New York University.

Notes

Owen Yingling provided research assistance for this brief.

1. Shalanda D. Young, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence" (memorandum), M-24-10 (March 28, 2024), Office of Management and Budget, Washington, DC.
2. Matthew Mittelsteadt, "Binding Public Sector AI Diffusion: Will New OMB AI Safety Policies Do More Harm Than Good?" *Digital Spirits*, April 09, 2024.
3. National Institute of Standards and Technology, "AI Risk Management Framework: Second Draft," August 18, 2022, Washington, DC.
4. M. Popel et al., "Transforming Machine Translation: A Deep Learning System Reaches News Translation Quality Comparable to Human Professionals," *Nature Communications* 11, no. 4381 (2020), https://doi.org/10.1038/s41467-020-18073-9.
5. U.S. Commission on Civil Rights, "Civil Rights and Protections During the Federal Response to Hurricanes Harvey and María" (report), September 21, 2022, Washington, DC.