Mathematica
Progress Together

Response to the Request for Information on the

Development of an Artificial Intelligence Action Plan

March 15, 2025

Submitted to:

NITRD National Coordination Office
U.S. National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22315
Attention: Faisal D'Souza

Submitted by:

Mathematica
1100 First Street, NE, 12th Floor
Washington, DC 20002-4221
Phone: (202) 484-9220
Fax: (609) 228-4958

Introduction

Mathematica is a nonpartisan research and data analytics organization that delivers evidence-based solutions to optimize programs and policies for efficiency, cost savings, and measurable impact. The rapid development and deployment of artificial intelligence (AI) presents enormous opportunities and risks. Mathematica's deep interdisciplinary expertise—including our experience with AI tools—positions us well to anticipate many of the opportunities and risks and to consider the ways policy might support those opportunities and mitigate those risks while encouraging innovation, enhancing economic competitiveness, and protecting U.S. national security.

We applaud the interest of the Office of Science and Technology Policy and the Networking and Information Technology Research and Development (NITRD) National Coordination Office, under the direction of Executive Order 14179, in having this conversation now: the pace of AI development and deployment is so rapid that favorable and unfavorable world-changing consequences are possible over a brief period. In the following pages, we draw on the broad and deep expertise of Mathematica staff across disciplines and policy domains to address many of the relevant AI policy topics identified by Office of Science and Technology Policy and NITRD National Coordination Office. Our own use of AI across contexts—from public health to education to employment and labor—to inform public decision making helped us build our response. For example, we have demonstrated in our previous work several relevant skills:

- We explored the value of wastewater data to anticipate surges in public health emergencies.
- We used machine-learning techniques to identify academically at-risk students.
- We designed and trained AI solutions to predict unplanned hospital admissions for a challenge sponsored by the Centers for Medicare & Medicaid Services (CMS).
- We predicted fatal opioid overdoses using geospatial analytics.

In these and other projects, we have shown how AI can be a powerful tool to support decisions and augment services that improve the lives of Americans. Our experience with AI tools has also informed us about many of the associated risks, including the ways that a lack of or poor data collection can skew results, such as spotty internet access or no centralized wastewater processing in rural areas. **Understanding where the gaps in the data are and how AI can strategically project data is key to success when using AI.**

Our response is informed by more than half a century of rigorous analysis of organizational systems for implementing policy and delivering services in domains such as health, education, and employment. The effects of AI on these systems will likely vary substantially, depending not only on the particular AI tools in question but also on the differing institutional structures and features. Mathematica knows these institutional features well, enabling us to anticipate some of the ways the impact of AI might differ across sectors.

Response

Protecting rights, safety, and national security

Measures taken to protect privacy rights and safety must account for the fact that many types of AI systems exist and that AI systems evolve over time. Over the past decade, we have seen an increase in the use of AI to guide decision making and resource allocation in sectors in which protecting people's rights and safety are of the utmost importance, such as healthcare, education, transport, employment, and welfare. To date, most of these AI systems have used classification and ranking algorithms for which research and frameworks exist to support responsible design, development, and deployment. In contrast, far less guidance is available on how to responsibly build and deploy AI systems that use large language models such as GPT-4, which are being deployed rapidly to execute customer support, content creation, translation, and more.

Looking to the future, people's rights to privacy and safety must be considered in even broader applications of AI that do not yet exist: artificial general intelligence (AGI) (Stuart Russell 2019). Although there is no consensus about when AGI will arrive, the possibility of a misalignment of interests between humanity and a super-intelligent AGI is real and demands the attention of policymakers well before the creation and deployment of AGI. The possibilities of AGIs motivate a need for investing in AI alignment research with the aim of ensuring that safe solutions are built into future AGIs when they launch.

Appropriate regulations require improved understanding on the part of policymakers and the public of AI's risks and safety. Philanthropies and government agencies should therefore invest in effective information dissemination to inform decision makers and the public about AI's risks and safety. Increased collaboration between academic researchers, policymakers, and the technology sector could promote translational work, which could play a key role in making clear what AI systems look like in practice and provide policymakers with enough understanding so they can collaborate with researchers and technologists to develop feasible standards and relevant non-burdensome regulations.

In the case of large language models, the pressing need is for research that establishes best practices for responsible design, development, and deployment. As shown by OpenAI's grant program for such research, there are knowledge gaps in the way these systems can threaten people's rights and safety as well as in the best practices that can prevent or mitigate those threats.

Although the 118th Congress proposed several bills that would regulate generative AI, filling the knowledge gaps is key to ensuring that these bills reflect the varied ways in which generative AI systems can threaten people's rights and safety.

Voluntary oversight of AI systems is not sufficient, given the collective action problem: individual firms know there is substantial private value in being first to develop new AI technologies, and society bears much of the risk. This is a prototypical case for regulation by government, though exactly how government should regulate them is a much harder question.

Some agencies, such as the Food and Drug Administration, have policies and procedures to regulate certain AI-enabled products, but most auditing to date is done reactively by academics and activists.

Although these audits are high impact, government agencies must rely on mechanisms other than private actors to gather the evidence they need to open an investigation into an algorithm, especially as AI systems continue to evolve.

One potential form of oversight might be to give these audits a more formal role with a seat at the table in a government agency. The government could establish a bureau whose mission is to use audits as a means of enforcing legislation that regulates the use of AI. For this form of oversight to be feasible, the United States could establish concrete rules regarding the levels of risk posed by the different types of AI systems and the allowability of such systems based on their level of risk. In the early stages, the government could contract with a third-party organization to conduct these audits and then build in-house capabilities over time.

In addition to downstream audits of existing AI systems, another approach is to focus on upstream risk management. For example, regulators could push for more public assessments of data quality, such as the work Mathematica conducts on behalf of CMS to assess Transformed Medicaid Statistical Information System data.

Driving effective solutions and innovation

AI-driven solutions must effectively address the problems they seek to solve. As of today, few AI solutions are evaluated in a rigorous way to determine effectiveness. One example of a well-evaluated AI platform is Bayesian Health's clinical platform, which is backed by research showing reductions in mortality. Researchers and regulators should promote frameworks on how to evaluate the effectiveness of AI solutions, such as rapid-cycle evaluations, because large experimental designs are not always possible—especially for large-scale consumer-facing AI tools.

After determining their effectiveness, fostering the benefits of AI-driven solutions depends on affordability, acceptability, and accessibility, with sustained maintenance and improvement over time. For example, intensive human tutoring is known to be academically effective, but it is expensive and difficult to access because of the limited availability of tutors. AI tutoring, if effective, could address educational gaps if it is useful, accessible, and affordable to students, but because curricula and knowledge change over time, the AI solution must continue to update and improve over time.

Using healthcare as an example, Mathematica used AI to predict unplanned hospital admissions and mortality for the CMS AI Health Outcomes Challenge. We worked with clinical and patient advocate partners to develop the model and address concerns over its utility and accessibility because doctors were concerned about yet another point-and-click solution. We drew on open-source data tools from federal agencies (namely the Agency for Healthcare Research and Quality and CMS) to reduce the cost of ongoing maintenance and align with definitions used in industry.

A major consideration in reducing errors in AI output is to clearly define who is accountable and responsible for safeguards throughout the AI life cycle. In Mathematica's experience, no single team or government agency can achieve this alone, and therefore the work must be interdisciplinary and involve public–private partnerships. Similar to what some federal initiatives that Mathematica has participated in now use, federal agencies should consider a standing socio-technical working group and a field working group. The socio-technical working group represents a broad set of perspectives to develop frameworks that support regulation, and a field working group would evaluate existing applications to identify risks. Depending on the scope of the

AI, these interdisciplinary teams would include sociologists, technical experts, policy analysts, legal experts, and social workers (the socio-technical working group) as well as organizations and individuals who would be affected by the technology, such as community-based organizations, civic and religious leaders, educators, and community members interested in AI (the field working group). We see opportunities for adapting policy analysis techniques like rapid-cycle evaluation and Learn, Innovate, Improve (LI^2) that can mitigate errors and support the American people.

For larger applications of AI, regulators could consider this kind of interdisciplinary review of AI applications to be a required part of the product development process, just like prototyping or initial design for a product. To encourage innovation, the federal government could design structures that facilitate access to this kind of mitigation review for organizations without the bandwidth to implement this on their own.

Promoting economic growth and good jobs

According to AI expert Suchi Saria, the most effective AI is developed for domain-specific use cases. **So, although there is a need for clear and transparent cross-cutting regulation, sector-specific regulations are important to meet the domain-specific AI use cases.**

On a specific use-case-by-use-case basis, a major measure to consider is monitoring compliance; thus, enforceability is a major concern. One idea is to create a certification program similar to how electronic medical record software is regulated. Examples of how to monitor for compliance for specific industries and populations could look at Mathematica's work with the U.S. Department of Labor's (DOL) Office of Federal Contract Compliance to examine the implementation and effects of new training on the office's mission of compliance. **In addition, monitoring could examine how agencies support employers to comply with AI standards** (see Mathematica's evaluation of DOL's effectiveness in bringing employers into compliance with labor standards). New standards surrounding AI will call for the same type of rigorous evaluation and continuous improvement. As new policies and standards around AI emerge, effective training on new compliance standards and evaluations of the strategies used to enforce those rules and regulations will be key ingredients in an effective government response.

The conversation around AI and jobs often focuses on job opportunities that could be lost because of technology, but there is also opportunity to create or improve jobs through AI. We suspect this opportunity brings risks that the American worker is not prepared for, and we see policy as having a role in facilitating the job-to-job transitions that will enable AI to achieve the goal of generating opportunity.

To ensure that all Americans can gain the skills needed to thrive in an economy influenced by AI, policymakers should ensure that resources are available to workers to make a variety of career paths viable. This includes preparing for the new economy at the early stages of education and for adult workers who might want to transition to new work. Curriculums are likely to continue to include more topics on technical and complimentary skills; skills such as critical thinking and strong communication have become increasingly important for hybrid jobs that combine technical skills with complimentary soft skills. For workers directly affected by AI or automation, policies such as retraining assistance or temporary income support are worth exploring.

New curriculums and training programs will need iteration as AI continues to evolve. Assessing the efficacy of our education and training will enable the focusing of resources into those programs shown to be effective and efficient.

Ensuring cost-effective services

The federal government can leverage AI across programs and services to effectively carry out their missions. AI could reduce the administrative budget for citizens and government employees, add new features to programs, or assist in decision making. **The overarching principal should be to use AI that is designed and evaluated for the specific use case to address major barriers to services.** Building on their experience with the opioid epidemic, governments at all levels should consider how AI tools can assist at each stage of a disaster or emergency. We see the possibility of such tools as especially strong in the context of expanding incidents, in which the nature of a disaster is unclear or changing and the level of impact appears to increase. We previously produced a report on human services and disasters showed that a lack of good data—namely, where those displaced by a disaster went and the human services supports that the receiving locales needed—hampered disaster response and early recovery efforts. By the time these gaps rise to the level of government actors, the resources of local community-based organizations are stretched. Using predictive analytics and related tools, we believe emergency managers can more proactively position needed resources. Sun and colleagues (2020) gave specific suggestions on the types of AI tools that might be appropriate at each stage of the disaster life cycle.

As we saw in our work on public health threats and wastewater monitoring, **AI tools are most effective when the appropriate data sets are appropriately integrated into the analytics.** To combat public health threats, it seems obvious that detecting the pathogen was crucial. What we learned through our work is that linking the emergency to its cause is not as simple as linking a public health threat to wastewater: we must situate that causal component within a larger context of local public health policies and other community factors. That is when truly useful insights emerged, and it is when the most vigilance is needed to ensure AI is deployed legally and effectively.

About Mathematica

Mathematica delivers evidence-based solutions to optimize programs and policies for efficiency, cost savings, and measurable impact. We're committed to outcomes that enhance daily lives while responsibly stewarding public, private, and philanthropic investments.

Mathematica Inc.

Our employee-owners work nationwide and around the world.

Find us at **mathematica.org** and **edi-global.com**.