July 7, 2023

Office of Science and Technology Policy
Ms. Arati Prabhakar
Eisenhower Executive Office Building
1650 Pennsylvania Ave, NW
Washington, DC 20504

Re: Request for Information National Priorities for Artificial Intelligence

Dear Ms. Prabhakar,

On behalf of the Center for Data Innovation (datainnovation.org), I am pleased to submit this response to the Office of Science and Technology Policy's (OSTP) request for information on national priorities for AI.[1]

The Center for Data Innovation studies the intersection of data, technology, and public policy. The Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, artificial intelligence, and the Internet of Things. The Center is part of the Information Technology and Innovation Foundation (ITIF), a nonprofit, nonpartisan think tank.

We commend OSTP for recognizing the importance of developing a national AI strategy. As we wrote in our 2018 report "Why the United States Needs a National Artificial Intelligence Strategy and What It Should Look Like," a national AI strategy is necessary to bolster U.S. competitiveness, strengthen national security, and maximize the societal benefits that the country could derive from AI.[2] We support OSTP in its ongoing work in this area and offer responses to the following questions in the document below:

---

[1] White House Office of Science and Technology Policy, "Request for Information National Priorities for Artificial Intelligence," news release, May 23, 2023, https://www.whitehouse.gov/wp-content/uploads/2023/05/OSTP-Request-for-Information-National-Priorities-for-Artificial-Intelligence.pdf.
[2] Joshua New, "Why the United States Needs a National Artificial Intelligence Strategy and What It Should Look Like," (Center for Data Innovation, December 2018), https://www2.datainnovation.org/2018-national-ai-strategy.pdf.

CENTER FOR DATA INNOVATION

Sincerely,

Hodan Omaar,
Senior Policy Analyst
ITIF's Center for Data Innovation
homaar@datainnovation.org

Sean Davis Jr,
Senior Policy Analyst
ITIF's Center for Clean Energy Innovation
sdavis@itif.org

Gillian Diebold,
Policy Analyst
ITIF's Center for Data Innovation
gdiebold@datainnovation.org

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?

A single national data privacy framework that preempts state laws, establishes basic consumer data rights, streamlines regulation, and minimizes the impact on innovation will help ensure organizations design, develop, and deploy AI systems responsibly. While the primary responsibility for creating a national privacy law rests with Congress, OSTP can urge Congress to move quickly to stop the growing patchwork of state data protection laws that undermines innovation. In addition, striking a balance between consumer welfare and innovation is important. Consumers should have confidence that those using AI systems will protect their data. However, beyond a baseline of protections, stricter data protection rules do not translate into more digital trust and therefore more digital adoption and use.[3] But stricted data protection rules can deter innovation and harm consumers because overly stringent regulations raise costs and reduce the relative quality of digital technologies, thereby negatively impacting the people who use them.

OSTP's policy guidance for data privacy as it relates to AI systems is overly stringent. Data privacy is one of the five principles in its Blueprint for an AI Bill of Rights, and this principle posits, among other things, that "data collection should be limited in scope, with specific, narrow identified goals." As ITIF argued in its 2019 report "A Grand Bargain on Data Privacy Legislation for America," calls for data minimization significantly limit companies from exploring new data sets that may lead to new or improved products and services.[4] Data minimization negatively impacts start-ups that, at the outset, do not know what data will be most valuable. Data minimization can also hurt existing businesses by limiting their ability to conduct post hoc analyses to develop new types of products and services based on what they learn from the data—even if these organizations use this data in a way that protects individual privacy. And it impacts businesses' future flexibility by limiting those that want to pivot to different business models based on data. Indeed, many of today's AI systems would not have been able to train on existing datasets if organizations had been forced to follow strict data minimization rules. Federal privacy legislation should not include data minimization provisions. OSTP should not discourage organizations from collecting and using data.

---

[3] Alan McQuinn and Daniel Castro, "Why Stronger Privacy Regulations Do Not Spur Increased Internet Use" (ITIF, July 2018), http://www2.itif.org/2018-trust-privacy.pdf.

[4] Alan McQuinn and Daniel Castro, "A Grand Bargain on Data Privacy Legislation for America" (ITIF, January 2019), https://www2.itif.org/2019-grand-bargain-privacy.pdf.

3. Are there forms of voluntary or mandatory oversight of AI systems that would help mitigate risk? Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?

OSTP should be critical of the bountiful analogies being bandied about AI. Some have likened AI to civil aviation and called for international governance of AI similar to that of the International Civil Aviation Organization (ICAO), which has developed common safety standards for civil aviation. Others have said AI is more like nuclear weapons because it is inherently dangerous and a better model of governance is something like the International Atomic Energy Agency (IAEA), meaning the development of the technology and the control its proliferation should be tightly monitored.

But framing the regulation of AI like the regulation of any one application of a technology is illogical. AI technology is not a tangible thing like a power source (such as nuclear energy) or a physical object (such as a jet engine). Instead, AI refers to the ability of computer systems to solve problems, make decisions, and perform tasks similar to what a human can do, which means that rather than being like nuclear energy, AI is more like the ability to do nuclear physics.[5] And rather than being like a jet engine, AI is more like the ability to fly a plane. When viewed through this lens, the incoherence of popular comparisons becomes clearer. Governments don't create regulatory structures to govern the field of nuclear physics, even though malicious actors might apply that knowledge to make nuclear weapons. Instead, they regulate the risky use case of this knowledge: nuclear weapons. Governments also don't regulate knowledge about how to fly. Governments do regulate pilots' licenses to determine who can legally operate an aircraft, but they don't impose transparency requirements to disclose the neural activity of pilots to ensure they are operating the aircraft correctly.[6] When it comes to AI, the same should apply: the emphasis for regulations should be on applications, not the technology itself.

8: How does AI affect the United States' commitment to cut greenhouse gases by 50-52% by 2030, and the Administration's objective of net-zero greenhouse gas emissions no later than 2050? How does it affect other aspects of environmental quality?

AI holds great promise as a tool for reducing greenhouse gas (GHG) emissions. Indeed, Microsoft and PwC found in a 2020 report that AI-enabled decarbonization technologies could reduce the carbon intensity of the global economy between 1.5 and 4 percent by 2030. The application areas are broad; AI could help create more accurate climate models, optimize wind farms, and power the

---

[5] Joshua New, "Bad Analogies for AI Make for Bad Policy," Linkedin post, June 20, 2023, https://www.linkedin.com/pulse/bad-analogies-ai-make-policy-joshua-new/.
[6] Ibid.

development of digital twin models that enable researchers to explore the impact of clean energy innovations.

There are many potential applications of AI for smart cities that would increase energy efficiency, including transportation, the electric grid, buildings, and city operations. By implementing AI solutions, cities can significantly reduce their environmental footprint. But as ITIF explains in its 2021 report "How Congress and the Biden Administration Could Jumpstart Smart Cities With AI," even the most capable cities struggle to evolve into smart cities because cities are ill-equipped to overcome the key challenges limiting smart city development.[7] The first challenge is that research in the underlying technologies for smart cities is a public good. Few want to bear the costs of going first, when the benefits mainly accrue to others. Second, few cities have the tools to share data with one another, which hampers the development of accurate AI models. Third, cities have little incentive to bear all the risk of failure involved in adopting technology fueled by emerging technologies. Finally, without a federal data privacy law, cities struggle to address unchecked privacy fears.

The federal government should play a role in helping U.S. cities overcome these challenges. It is able to provide funding and coordination on a larger scale than cities working individually. While the federal government has undertaken an array of activities to support the development of smart cities, these efforts have mostly been uncoordinated, and the government has had no strategic vision for AI research, development, and deployment (RD&D) of smart city technologies. Smart cities offer an important opportunity to address both infrastructure needs and strained state and local budgets at the same time. Much like Singapore has focused on smart cities as one of the focus areas of its national AI strategy, so too should OSTP include a focus on developing AI-enabled smart cities in its national AI strategy.[8]

9: What are the opportunities for AI to enhance equity and how can these be fostered? For example, what are the potential benefits for AI in enabling broadened prosperity, expanding economic and educational opportunity, increasing access to services, and advancing civil rights?

While much of the reporting on AI systems has focused on these systems' potential for bias, there are many ways these systems can help reduce inequality and improve individuals' access to opportunity, especially for individuals for whom opportunities have been historically limited. For

---

[7] Colin Cunliff, Ashley Johnson, and Hodan Omaar, "How Congress and the Biden Administration Could Jumpstart Smart Cities With AI," (ITIF, March 2021), https://itif.org/publications/2021/03/01/how-congress-and-biden-administration-could-jumpstart-smart-cities-ai/.

[8] Smart Nation Singapore, "National AI Strategy," accessed July 6, 2023, https://www.smartnation.gov.sg/files/publications/national-ai-strategy.pdf.

instance, using algorithms can improve equity in decisions about allocating scarce resources in healthcare settings and expanding access to educational opportunities. Algorithms can also replace biased human decision-making for tasks like home appraisals.[9] However, even when algorithms can do good by making existing processes more efficient and equitable for consumers, public backlash and opaque implementations can erode the trust needed for them to achieve impact. OSTP should actively foster and support the adoption of AI in the areas it can make outcomes more fair.

Take the example of the Boston public school system, which in 2018 proposed using an algorithmic system to improve school busing in ways that would cut costs by millions of dollars a year, help the environment, and better serve students, teachers, and parents.[10] The district had two aims, the first of which was to cut transportation costs. More than 10 percent of the public school system's budget goes toward busing children to and from school—the district's annual cost per student is the second highest in the United States.[11] The district's second goal was to to reconfigure school start times so that high school students could get more sleep, as early school starts for teenagers has been linked to serious health issues such as decreased cognitive ability, increased obesity, depression, and increased traffic accidents. Indeed, the American Academy of Pediatrics recommends that teenagers not start their school day before 8:30 AM, but only 17 percent of U.S. high schools comply according to a 2015 report from the Center of Disease Control and Prevention (CDC).[12]

Boston public school officials engaged researchers from the Massachusetts Institute of Technology (MIT) to build an algorithm to achieve its twin goals, which they did. The *Boston Globe* called its solution a "marvel."[13] The algorithm helped the district optimize bus routes, cutting 50 of the 650 school buses used, $5 million off the budget, and 20,000 pounds of carbon emissions each day while also optimizing bell times. Importantly, the algorithm's solution for bell times redressed equity. In the past, the district manually staggered start and end times, but its approach predominantly provided wealthier and whiter schools with later start times while schools with poorer and minority students disproportionately shouldered earlier times. In contrast, the algorithm's solution distributed

[9] Debra Kamin, "Home Appraised With a Black Owner: $472,000. With a White Owner: $750,000," The New York Times, August 25, 2022, https://www.nytimes.com/2022/08/18/realestate/housing-discriminationmaryland.html.
[10] David Scharfenberg, "Computers Can Solve Your Problem. You May Not Like The Answer," The Boston Globe, September 21, 2018, https://apps.bostonglobe.com/ideas/graphics/2018/09/equity-machine/.
[11] Ellen P. Goodman, "The Challenge of Equitable Algorithmic Change," The Regulatory Review (2019), https://www.theregreview.org/wp-content/uploads/2019/02/Goodman-The-Challenge-of-EquitableAlgorithmic-Change.pdf.
[12] Centers for Disease Control and Prevention, "Most US middle and high schools start the school day too early," news release, August 6, 2015, https://www.cdc.gov/media/releases/2015/p0806-school-sleep.html.
[13] David Scharfenberg, "Computers Can Solve Your Problem. You May Not Like The Answer."

advantageous start times equally across major racial groups, while significantly improving them for students in all of those groups.[14]

Despite everything the algorithm offered, the district had to scrap the algorithm due to the swift and strong public pushback. NTIA Senior Advisor for Algorithmic Justice and Rutgers law professor Ellen Goodman describes how disgruntled parents carried signs at a school committee meeting that read "families over algorithms," and "students are not widgets" in her 2019 paper The Challenge of Equitable Algorithmic Change.[15] But it is not clear the algorithm was really the problem. Indeed, Goodman describes the pushback as a case of "algorithmic scapegoating," which Cornell researchers explain is where the algorithm "stood in for substantive issues around equity and disruptive change that were really at stake (though potentially more contentious to discuss) and might well have been at stake even without an algorithm in the picture. The tragedy of the case is that the algorithm could have provided the flexibility to involve the public in choosing among multiple trade-offs. If implemented, it might have created a more equitable system than what existed originally."[16]

The takeaway for OSTP from this episode is twofold: One, algorithmic systems can reduce inequality from human decision-making when they are designed well. Two, these beneficial solutions "gain legitimacy not through their mathematical exactitude but through community engagement throughout the process," as Goodman puts it.[17] OSTP should make clear that AI is not "inherently biased" and encourage organizations to find ways to use the technology to increase equity and promote opportunity.

16. What steps can the United States take to ensure that all individuals are equipped to interact with AI systems in their professional, personal, and civic lives?

Bolstering AI literacy is a critical first step to ensuring that the American public is equipped to interact with and understand AI systems. Many Americans lack a concrete understanding of AI and instead have become fearful of these systems thanks to a perpetuation of AI "myths" circulating online and in the media. Improving AI literacy would enable more Americans to understand and use AI-enabled technologies in their professional, personal, and civic lives.

---

[14] Ellen P. Goodman, "The Challenge of Equitable Algorithmic Change."
[15] Ibid.
[16] Karen Levy, Kyla E. Chasalow, & Sarah Riley, "Algorithms and Decision-Making in the Public Sector," Annual Review of Law and Social Science (2021): Volume 17, https://www.researchgate.net/publication/352208766_Algorithms_and_DecisionMaking_in_the_Public_Sector.
[17] Ellen P. Goodman, "The Challenge of Equitable Algorithmic Change."

AI literacy is the ability to understand how AI-enabled technologies like machine learning work, how they can be used for problem-solving, and the consequences of the technology. AI-literate members of society can make more informed decisions about how and when to use AI and better understand its development. AI literacy can be considered an evolving component of digital literacy, an area that also includes data/statistical literacy and media literacy.

All Americans need AI literacy, and OSTP plays a critical role in facilitating the advancement of this type of skill. OSTP should oversee and guide the National AI Initiative Office to potentially publish a five-year AI education strategic plan, distinct from any plans housed in the Department of Education. Under AI education, which should include things like upskilling and reskilling, OSTP should ensure that AI literacy is one of the priorities. The Committee on STEM Education (CoSTEM) has such a plan for STEM education, with one of its agencies focusing on digital literacy initiatives. Likewise, an interagency working group for AI education could focus on AI literacy.

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?

The Center for Data has written extensively about how AI can provide enormous societal and economic benefits in areas such as education in our report "How AI Can Improve K-12 Education in the United States;" in healthcare in our report "The Promise of Data-Driven Drug Development;" and transport in our ITIF report "How Congress and the Biden Administration Could Jumpstart Smart Cities With AI."[18] However, realizing these benefits requires more than just public and private investments in the development of AI; the United States also needs a multipronged national AI adoption strategy to ensure the potential opportunities are translated into all the areas where they can make a positive difference in people's lives. That is exactly what OSTP's effort with this RFI can help achieve.

---

[18] Gillian Diebold and Chelsea Han, "How AI Can Improve K-12 Education in the United States," (Center for Data Innovation, April 2022), https://datainnovation.org/2022/04/how-ai-can-improve-k-12-education-in-the-united-states/; Joshua New, "The Promise of Data-Driven Drug Development," (Center for Data Innovation, September 2019), https://www2.datainnovation.org/2019-data-driven-drug-development.pdf; Colin Cunliff, Ashley Johnson, and Hodan Omaar, "How Congress and the Biden Administration Could Jumpstart Smart Cities With AI," (ITIF, March 2021).

One area of focus to this end in OSTP's strategy should be transforming government itself with AI (further details given in the next question). Another is for each federal agency to develop its own AI strategy and appoint a chief AI officer. One reason certain agencies devote so little attention to AI is that it is generally not formally recognized as part of agency agendas or strategic plans. Each agency should explicitly identify specific steps for how it will connect its data, users, and mission priorities to support AI transformation, much like the Department of Defense, Department of Veteran Affairs, and Food and Drug Administration already have. To really coordinate and drive implementation of AI, each federal agency should consider appointing a chief AI officer, like the Department of Health and Human Services has done.

24. How can the Federal Government effectively and responsibly leverage AI to improve Federal services and missions? What are the highest priority and most cost-effective ways to do so?

Policy for government adoption of AI has so far has predominantly focused on getting federal workers to better understand and care about the process of AI innovation, and while this is important, structural factors play a much more important role in limiting and enabling innovation across the federal enterprise as the Center for Data Innovation explains in its 2022 report "U.S. AI Policy Report Card."[19] Most federal workers know innovation is useful, but they don't innovate with AI because there are few rewards and many barriers. This strategy should therefore focus on overcoming structural factors related to culture; financing; metrics and incentives; procurement; and oversight and review. In order to be effective, the Biden administration should concurrently work on reviving the federal data strategy, an effort that has been stalled for several years. Without an effective strategy to harness data—a fundamental enabler of AI innovation—the administration will struggle to implement a successful federal strategy for AI.[20] Finally, the Biden administration should leverage AI to improve federal services and missions, as the opportunities and use cases for AI in government are numerous. Generative AI use cases in government range from the standard—creating a really good chatbot—to the more ambitious—designing an end-to-end digital service.[21] But generally, sensible applications of this technology can simplify access to and organization of information, automate bureaucratic processes and data entry, and contribute to design and content production. These possibilities will expand as generative AI progresses.

---

[19] Hodan Omaar, "U.S. AI Policy Report Card," (Center for Data Innovation, July 2022), https://www2.datainnovation.org/2022-ai-report-card.pdf.

[20] Eric Egan, "Reviving and Reimagining the Federal Data Strategy for Mission Success," (ITIF, June 2023), https://itif.org/publications/2023/06/05/reviving-and-reimagining-the-federal-data-strategy-for-mission-success/.

[21] Eric Egan, "Generative AI Offers Federal Agencies Common-Sense Opportunities to Simplify and Improve How They Work," (ITIF, June 2023), https://itif.org/publications/2023/06/28/generative-ai-offers-federal-agencies-common-sense-opportunities-to-simplify-and-improve/.