



July 7, 2023

# Response to the White House OSTP Request for Information on National Priorities for Artificial Intelligence

Data & Society Research Institute submits this comment in response to the Request for Information (RFI) on National Priorities for Artificial Intelligence published by the White House Office of Science and Technology Policy (OSTP) on May 23, 2023. Data & Society is an independent, nonprofit research institute studying the social implications of data-centric technologies, automation, and artificial intelligence. We produce empirical research that challenges the power asymmetries created and amplified by technology in society.

Data & Society has already submitted other regulatory comments regarding AI, including recent ones on worker surveillance<sup>1</sup> and AI accountability.<sup>2</sup> Rather than reiterate those comments here, we instead zoom out to the broader strategy that we believe the White House should pursue to ensure that AI technology not only is safe, responsible, and trustworthy, but that it is developed in the public interest for the benefit of people, not Big Tech.

We know the Biden-Harris administration understands the real-world harms that AI is already inflicting on people, their families, and their communities. President Biden has spoken forcefully about these concerns, emphasizing his administration's commitment to "safeguarding America's rights and safety, from protecting privacy, to addressing bias and disinformation, to making sure AI systems are safe before they are released."<sup>3</sup>

A national AI strategy centered on people's rights should comprise a key pillar in the administration's "Bidenomics"<sup>4</sup> approach. It should reject a *laissez-faire* market approach and instead recognize that strategic areas of the US economy, such as artificial intelligence, require direct government intervention and guidance to ensure they are developed for the public good.

As we argue below, a national AI strategy should (1) implement OSTP’s Blueprint for an AI Bill of Rights across federal agencies; (2) direct AI development in the public interest towards socially productive outcomes; (3) ban AI where the harms are too great; (4) limit job loss and make “AI factory” jobs good jobs; (5) strengthen the country’s scientific and sociotechnical knowledge; and (6) engage Congress to pass much-needed federal legislation.

## **I. Minimizing the risks of AI systems is necessary, but not sufficient, to address their harms.**

This administration has rightly recognized that addressing the risks of AI and realizing its benefits are not necessarily in conflict.<sup>5</sup> To that end, voluntary frameworks like the National Institute of Standards and Technology’s (NIST) AI Risk Management Framework are important guides for companies to assess their development and deployment of AI.

But many harms of AI extend beyond “risks.” Risks are a probability that an AI system may (for example) unjustly deny someone a benefit, foreclose opportunities, or place them in danger. But some harms are less a probability, more a certainty—and whatever the probability of harm, Americans’ rights must be protected.

Certain AI harms will persist even after the most careful and thorough measures to minimize “risk.” These harms, such as large language models’ titanic energy consumption<sup>6</sup> or AI’s dependence on underpaid, precarious data laborers,<sup>7</sup> are not a probabilistic risk, but a harm *intrinsic to the technology itself*. Thus, risk mitigation alone is insufficient. What is needed is an intentional strategy to ensure that AI serves the public interest.

**AI policy that mitigates risk at the margins while permitting private actors broad discretion to develop and deploy AI is, fundamentally, the kind of trickle-down approach that this administration has rightly rejected.**<sup>8</sup> Instead, the national AI strategy, consistent with this administration’s commitment to grow the economy “from the middle out and the bottom up,” must not only address the risks of AI but drive AI development to enshrine people’s rights and achieve positive outcomes for the many, not the few.

## **II. A national AI strategy must build on the Blueprint for an AI Bill of Rights.**

The administration’s Blueprint for an AI Bill of Rights<sup>9</sup> is a landmark framework for enacting concrete safeguards to protect the public from the harms of AI while realizing its benefits. Its principles and technical companion provide steps that are rooted in empirical research and remain actionable today, applicable even to newer generative AI models. Now, the White House

must build on the Blueprint and move to make its democratic ideals a reality. We propose six actions to operationalize the AI Bill of Rights in the short-term and build an enduring national AI strategy for the long-term.

**i. Immediately implement the protections of the AI Bill of Rights across the federal government.**

In the short-term, the protections of the AI Bill of Rights should be implemented across the federal government through the upcoming Office of Management and Budget guidance on AI, as also called for by a coalition of civil rights organizations.<sup>10</sup> This guidance should include all agencies, without exceptions for law enforcement, and should also cover eligibility for government grants. Government money must not be spent on technology that violates people’s civil rights—or that simply doesn’t work. Agencies with enforcement authority over existing civil rights laws, and other agencies with authority that permits oversight of AI in specific sectors, must use that authority to govern the safe use of AI and to remedy AI-driven discrimination. Some of these actions are already underway, including through Executive Order 14091 (Sections 4b, 8b, and 8f), which prohibits algorithmic discrimination in the government’s use of automated systems and calls on agencies to enforce against such harms.<sup>11</sup>

However, more must be done to further equity in the use of automated systems, including by adding bias testing and impact assessments to the annual AI use case inventory required by Executive Order 13960<sup>12</sup> and by ensuring agencies’ full compliance with the inventory. Currently, the inventory requires only basic information about AI systems such as a short summary and includes only a small subset of the automated systems of concern. For example, the Department of Justice failed to report<sup>13</sup> on its use of an AI-based risk assessment tool,<sup>14</sup> despite its own National Institute of Justice report identifying racial bias in the tool.<sup>15</sup> Similarly, the Department of Homeland Security failed to report<sup>16</sup> on its use of facial recognition within the Customs and Border Patrol One App, which has proven ineffective for Black asylum seekers.<sup>17</sup> Without a full accounting of the use of AI by the federal government, and full enforcement of Executive Order 14091’s mandate that such systems do not discriminate, the public cannot have confidence that the federal government is using AI safely, effectively, and in a non-discriminatory manner.

**ii. Drive AI development towards socially beneficial outcomes.**

Beyond implementation of the AI Bill of Rights across federal agencies in the short-term, the US government must enact a national AI strategy that proactively shapes AI’s development in

the public interest. **A national AI strategy should advance specific societal objectives—ones driven by democratic values, not the prerogatives of powerful tech companies—and present enforceable actions to protect Americans’ rights.**

Given the potentially transformative nature of AI and its well-documented, real-world harms—in policing,<sup>18</sup> housing,<sup>19</sup> and health care,<sup>20</sup> among other areas—the US government should play a direct, hands-on role in its development, steering the technology in ways that will benefit all Americans, not only private industry. The US government cannot simply aspire for the tech industry to voluntarily adopt frameworks to address risks. It should affirmatively establish strategic priorities for the development, deployment, and use of AI for the public good.

**A strategy that focuses on social outcomes puts people, not technology, first.** OSTP and the Biden-Harris administration should articulate the specific outcomes that AI should seek to achieve. Without presenting an exhaustive list, we offer two examples of purposeful, people-first AI objectives. First, AI might be used to support the President’s “Cancer Moonshot” initiative, which itself leads with clear societal goals such as preventing cancer deaths. If AI can help achieve this goal, based on the best expertise of computer scientists, medical professionals, and researchers, then the federal government should ensure that the nation’s development and use of the technology (in both the public and private sectors) align to support it. Second, AI might be used to support crop management and resilience to global warming. If AI can safely and responsibly help to achieve this goal, with attention to broader socio-ecological factors that may impact food security,<sup>21</sup> the US government should strategically orient AI deployment (through mechanisms like public funding, regulatory guidance, and incentives for private industry) to support it.

In each case, the US should begin with the desired social outcome, not the blind expansion of a new technology. In exploring the particular steps to achieve any desired outcome, the strategy should contemplate how AI may be used effectively *alongside* people, not in their displacement. **Critically, if AI cannot plausibly achieve the desired societal goals, the government should not support its use at all.**

### iii. **Prohibit AI where the harms are too great.**

AI should never be used in a way that violates or limits people’s rights, opportunities, or access. Accordingly, the White House’s national AI strategy should enforce bright-line prohibitions where AI fundamentally conflicts with core values like people’s freedom of association, freedom of thought, freedom of movement, their civil rights, and more.

For example, the national AI strategy should prohibit law enforcement from using affective computing, or the purported detection of a person’s emotional state especially through voice, image, or video-based AI.<sup>22</sup> Research suggests that such technology is unlikely to ever

work<sup>23</sup>—and if it *did* someday work, such a capability would be extremely damaging to freedom of thought.

Similarly, people should not be subject to constant monitoring and surveillance as a precondition for access to basic needs. As rightly identified in the AI Bill of Rights, “Continuous surveillance and monitoring should not be used in education, work, housing, or in other contexts where the use of such surveillance technologies is likely to limit rights, opportunities, or access.”<sup>24</sup> To that end, a national AI strategy should ban such uses.

These are only two examples. A national AI strategy should consider the current and possible use cases of AI and enforce prohibitions (including through executive action, regulation, and legislation) where AI presents such an intolerable degree of harm that risk mediation and technical improvements cannot suffice.

#### **iv. Limit job loss and make AI factory jobs good jobs.**

AI stands to displace many workers (or, at least, to dramatically change the tasks that workers are hired to do).<sup>25</sup> Yet in many cases, AI will likely do a worse job than people, and the American public will distrust AI to make appropriate, considered decisions. A national AI strategy should mandate human oversight and intervention in cases where such systems impact people’s rights, opportunities, or access. This might include regulatory or legislative action to require a doctor’s expert opinion for any medical decision, call center staffing for any benefits decisions, or a hiring manager’s final say on any candidate ranking. Such actions protect the public from harm while simultaneously stemming job loss.

Within the “AI factory” itself, people are involved in all aspects of AI. Workers enter data, label photos, answer prompts, or fill in spreadsheets. Later in the process, workers rate the results of a system, providing feedback to improve a system’s accuracy. Many systems have a (sometimes hidden) “human-in-the-loop,” where humans intervene during AI use to perform tasks the system is unable to or otherwise aid the technology.<sup>26</sup> And some AI companies present a human alternative to a technical system (as the AI Bill of Rights recommends), including by employing workers in call centers to assist users.

The problem with these jobs is that they generally have abysmal working conditions and are contracted out to low-wage workers overseas.<sup>27</sup> A national AI strategy should ensure that any such AI factory jobs are good, well-paying American jobs and that workers have full employment and labor rights. Protecting people from AI harm via requiring human alternatives, consideration, and fallback can go hand-in-hand with job creation and workers’ rights.

**v. Invest in the country’s scientific and sociotechnical knowledge.**

Just as the 1960s space race led to a massive investment in education, including a billion dollars for science curriculum redesign in addition to other resources,<sup>28</sup> the US government must meet the moment of AI by investing in the country’s base of scientific knowledge. From elementary schools to universities to professional training programs, a computer science education based on foundational science—and not company training programs—should be available to all students. The government should fund computer science experts to develop curricula at all levels, fund teacher training programs, and create incentives and pathways for computer science teachers to be hired in schools across the country. Such a generational investment, including job retraining programs, can stem the tide of AI-driven job loss while improving students’ education across the nation.

It’s not just the hard sciences. The success of a national AI strategy that strives for particular societal outcomes depends on investments in *sociotechnical* expertise. Sociotechnical research is a field of knowledge that assesses the use of technologies not in a vacuum but atop human, material, and cultural infrastructures. Because AI is not simply a technical domain—it affects human beings, communities, work and labor relations, private and state systems of surveillance, the environment, and more—broad sociotechnical expertise, drawing from the humanities and other fields, must be core to a national AI strategy. Accordingly, Data & Society echoes calls by OSTP,<sup>29</sup> NIST,<sup>30</sup> and the National AI Advisory Committee<sup>31</sup> for a sociotechnical approach to AI. The federal government must strengthen public and private bases of sociotechnical expertise through federal funding and hiring, and include sociotechnical understanding as part of its educational investment.

**vi. Engage Congress to pass AI-specific legislation.**

Finally, the administration should be working directly with Congress to shape the fall legislative agenda and the many forthcoming bills focused on artificial intelligence. The administration has supported data privacy legislation, which is key to protecting the public from the harms of AI, but needs to do more to guide Congress on AI-focused legislation. With its subject matter expertise and forward vision, the Biden-Harris administration is well-positioned to work with Congressional leadership to ensure that any future legislation is guided by the principles of the AI Bill of Rights and aligns with other goals of the pending national AI strategy.

Respectfully submitted,

**Brian J. Chen**, Policy Director

**Sorelle Friedler**, Senior Policy Fellow

**Serena Oduro**, Senior Policy Analyst

## Endnotes

- 1 Brian J. Chen, Livia Garofalo, Alexandra Mateescu, Aiha Nguyen, and Eve Zelickson. “Response to the White House OSTP’s Request for Information on Automated Worker Surveillance and Management.” Data & Society Research Institute. June 21, 2023. <https://datasociety.net/announcements/2023/06/21/response-to-the-white-house-ostps-request-for-information-on-automated-worker-surveillance-and-management/>.
- 2 Serena Oduro, Ranjit Singh, and Jake Metcalf. “Response to the NTIA’s Request for Comment on AI Accountability Policy.” Data & Society Research Institute. June 12, 2023. <https://datasociety.net/announcements/2023/06/12/response-to-the-ntias-request-for-comment-on-ai-accountability-policy/>.
- 3 The White House. “Remarks by President Biden on Seizing the Opportunities and Managing the Risks of Artificial Intelligence.” June 20, 2023. <https://www.whitehouse.gov/briefing-room/speeches-remarks/2023/06/20/remarks-by-president-biden-on-seizing-the-opportunities-and-managing-the-risks-of-artificial-intelligence/>.
- 4 The White House. “Bidenomics Is Working: The President’s Plan Grows the Economy from the Middle Out and Bottom Up—Not the Top Down” (hereinafter Bidenomics). June 28, 2023. <https://www.whitehouse.gov/briefing-room/statements-releases/2023/06/28/bidenomics-is-working-the-presidents-plan-grows-the-economy-from-the-middle-out-and-bottom-up-not-the-top-down/>.
- 5 A senior administration official stated, “As President Biden has underscored, in order to seize the benefits of AI, we need to start by mitigating its risks. And this core principle has guided our work on AI from the beginning.” The White House. “Background Press Call on New Artificial Intelligence Announcements.” May 4, 2023. <https://www.whitehouse.gov/briefing-room/press-briefings/2023/05/04/background-press-call-on-new-artificial-intelligence-announcements/>.
- 6 Mack DeGuerin. “Move Aside, Crypto. AI Could Be The Next Climate Disaster.” *Gizmodo*. April 3, 2023. <https://gizmodo.com/chatgpt-ai-openai-carbon-emissions-stanford-report-1850288635>.
- 7 Josh Dzieza. “AI Is a Lot of Work.” *The Verge*. June 20, 2023. <https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>.
- 8 “Bidenomics” (noting the “failed trickle-down theory that supported . . . shrinking public investment in critical priorities”).
- 9 White House Office of Science and Technology Policy. “Blueprint for an AI Bill of Rights”

- (hereinafter “AI Bill of Rights”). Oct. 2022. <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.
- 10 The Leadership Conference on Civil and Human Rights, et al. “Next Steps to Advance Equity and Civil Rights in Artificial Intelligence and Technology Policy.” June 13, 2023. <https://civilrights.org/resource/next-steps-to-advance-equity-and-civil-rights-in-artificial-intelligence-and-technology-policy/>.
  - 11 “Further Advancing Racial Equity and Support for Underserved Communities Through The Federal Government.” 88 Fed. Reg. 10825. Feb. 22, 2023. <https://www.federalregister.gov/documents/2023/02/22/2023-03779/further-advancing-racial-equity-and-support-for-underserved-communities-through-the-federal>.
  - 12 “Promoting the Use of Trustworthy Artificial Intelligence.” 85 Fed. Reg. 78939. Dec. 8, 2020. <https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government>.
  - 13 Department of Justice. “Response to AI Use Case Inventory.” Available at: <https://www.justice.gov/open/page/file/1517316/download> via <https://www.ai.gov/ai-use-case-inventories/>.
  - 14 Carrie Johnson. “Flaws Plague a Tool Meant to Help Low-Risk Federal Prisoners Win Early Release.” *NPR*. Jan. 26, 2022. <https://www.npr.org/2022/01/26/1075509175/flaws-plague-a-tool-meant-to-help-low-risk-federal-prisoners-win-early-release>
  - 15 National Institute of Justice. “2021 Review and Revalidation of the First Step Act Risk Assessment Tool.” National Institute of Justice NCJ 303859. Dec. 2021. <https://www.ojp.gov/pdffiles1/nij/303859.pdf>.
  - 16 Department of Homeland Security. “Response to AI Use Case Inventory.” Available at: [https://www.dhs.gov/data/AI\\_inventory](https://www.dhs.gov/data/AI_inventory) via <https://www.ai.gov/ai-use-case-inventories/>.
  - 17 Melissa del Bosque. “Facial Recognition Bias Frustrates Black Asylum Applicants to US, Advocates Say.” *The Guardian*. Feb. 8, 2023. <https://www.theguardian.com/us-news/2023/feb/08/us-immigration-cbp-one-app-facial-recognition-bias>.
  - 18 Kashmir Hill and Ryan Mac. “Thousands of Dollars for Something I Didn’t Do.” *The New York Times*. March 31, 2023. <https://www.nytimes.com/2023/03/31/technology/facial-recognition-false-arrests.html>.
  - 19 Heather Vogell. “Pressure Grows on Real Estate Tech Company Accused of Colluding With Landlords to Jack Up Apartment Rents.” *ProPublica*. Nov. 14, 2022. <https://www.propublica.org/article/yieldstar-rent-increase-realpage-lawmakers-collusion>.

- 20 Casey Ross and Bob Herman. “Denied by AI: How Medicare Advantage Plans Use Algorithms to Cut Off Care for Seniors in Need.” *STAT News*. Mar. 13, 2023. <https://www.statnews.com/2023/03/13/medicare-advantage-plans-denial-artificial-intelligence/>.
- 21 Tzachor A, Devare M, King B, Avin S, Ó hÉigeartaigh S. “Responsible Artificial Intelligence in Agriculture Requires Systemic Understanding of Risks and Externalities.” *Nature Machine Intelligence* 4 (Feb 2022): 104-109. <https://www.nature.com/articles/s42256-022-00440-4.epdf>.
- 22 Alex Engler. “Why President Biden Should Ban Affective Computing in Federal Law Enforcement.” Brookings Institute. Aug. 4, 2021. <https://www.brookings.edu/articles/why-president-biden-should-ban-affective-computing-in-federal-law-enforcement/>.
- 23 Barrett LF, Adolphs R, Marsella S, Martinez AM, Pollak SD. “Emotional Expressions Reconsidered: Challenges to Inferring Emotion from Human Facial Movements.” *Psychological Science in the Public Interest* 20, no. 1 (July 2019): 1-68. <https://journals.sagepub.com/stoken/default+domain/10.1177%2F1529100619832930-FREE/full>.
- 24 “AI Bill of Rights” at 6.
- 25 Elizabeth Napolitano. “AI Eliminated Nearly 4,000 jobs in May, Report Says.” *CBS News*. June 2, 2023. <https://www.cbsnews.com/news/ai-job-losses-artificial-intelligence-challenger-report/>.
- 26 Mary L. Gray and Siddarth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (Harper Collins, 2019).
- 27 Billy Perrigo. “OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic.” *Time*. Jan. 18, 2023. <https://time.com/6247678/openai-chatgpt-kenya-workers/>.
- 28 Larry Abramson. “Sputnik Left Legacy for US Science Education.” *NPR*. Sept. 30, 2007. <https://www.npr.org/2007/09/30/14829195/sputnik-left-legacy-for-u-s-science-education>.
- 29 “AI Bill of Rights” at 14 (describing “practical technical and *sociotechnical* approaches to protecting rights, opportunities, and access”) (emphasis added).
- 30 National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework*. Jan. 26, 2023. 1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> (“AI systems are inherently socio-technical in nature[.]”).
- 31 National Artificial Intelligence Advisory Committee. *Year 1 Report*. May 2023. 37. <https://www.ai.gov/wp-content/uploads/2023/05/NAIAC-Report-Year1.pdf> (recommendation to “[d]evelop a research base and community of experts focused on sociotechnical research in the AI R&D ecosystem”).