

Opportunities and Risks of Artificial Intelligence in the Public Sector

Daniel E. Ho¹
Stanford University

Testimony presented to the U.S. Senate Committee on Homeland Security and Governmental Affairs on May 16, 2023

Chairman Peters, Ranking Member Paul, and Members of the Committee, it is an honor to speak with you today.

I'm a professor at Stanford University, where I serve as Associate Director of Stanford's Institute for Human-Centered AI (or HAI) and Director of the Regulation, Evaluation, and Governance Lab (or RegLab). I also serve as a Member of the National AI Advisory Committee (NAIAC) and the RegLab works with numerous federal agencies, but I speak to you today in my personal capacity.

The U.S. government has an exceptional opportunity. It can seize this moment of AI innovation to modernize federal programs, make agencies more effective, catalyze scientific advancements, and protect civil rights and liberties for the benefit of all Americans. Doing so will strengthen America. But strategic leadership, critical investments in the federal workforce and digital infrastructure, and adapting procurement to AI are preconditions.

I. The Importance of Public Sector AI

¹ William Benjamin Scott and Luna M. Scott Professor of Law; Professor of Political Science; Senior Fellow at Stanford Institute for Economic Policy Research, Stanford University; Associate Director, Stanford Institute for Human-Centered Artificial Intelligence (HAI); Faculty Director, Stanford Regulation, Evaluation, and Governance Lab (RegLab). All views expressed in this testimony are provided in an individual capacity and do not represent the views of any affiliated institution, agency, or the National AI Advisory Commission (NAIAC).

Let me start with what is possible if the government gets this right.

First, government should lead by example and demonstrate how responsible AI can modernize federal programs. In a report to the Administrative Conference of the United States that I co-authored with administrative law scholars, we showed how early AI innovation in nearly half of the largest 142 federal agencies can transform the administration of government benefits like veteran disability compensation, improve monitoring of public health risks and adverse drug effects, and help protect workers, consumers, and the environment.²

Take the Social Security Administration (SSA), which pays benefits to some 18 million Americans annually. The SSA's administrative judges can hear over a half million disability appeals per year.³ With great foresight, SSA began investing in data infrastructure and tools to modernize case adjudication in the 1990s. In one pilot, SSA used AI to reorder cases by similarity, allowing adjudicators to learn complex areas of the law more effectively. The resulting "micro-specialization" increased the speed and accuracy of adjudicators in the pilot.⁴ SSA's early investments culminated in an AI tool that allows judges to check draft decisions for some 30 errors.⁵ Such innovations can expedite and improve agency decision-making to better serve American citizens and some have called the official who pioneered these early investments the "Steve Jobs of the SSA."⁶

² Engstrom, David Freeman, Daniel E. Ho, Catherine Sharkey, and Mariano-Florentino Cuéllar. 2020. "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies." Administrative Conference of the United States.

³ Ames, David, Cassandra Handan-Nader, Daniel E. Ho, and David Marcus. 2020. "Due Process and Mass Adjudication: Crisis and Reform." *Stanford Law Review* 72:1.

⁴ "SSA reported 12% reduction in case processing time and 7.5% reduction in returns from administrative appeal judges to attorneys." Engstrom, David Freeman, Daniel E. Ho, Catherine Sharkey, and Mariano-Florentino Cuéllar. 2020. "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies." Administrative Conference of the United States.

⁵ Glaze, Kurt, Daniel E. Ho, Gerald K. Ray, and Christine Tsang. 2022. "Artificial Intelligence for Adjudication: The Social Security Administration and AI Governance." In *The Oxford Handbook of AI Governance*. Oxford University Press, Handbook on AI Governance.

⁶ See Administrative Conference of the United States, Recommendation 2021-10, "Quality Assurance Systems in Agency Adjudication," 87 Federal Register 1722 (2022).

Second, government agencies are of course also critical for regulation, be it for FDA’s approval of AI medical devices or the Department of Transportation’s approach to self-driving cars.⁷ The right regulatory approach – striking the appropriate balance between innovation and safeguards – requires expertise in government. Government cannot govern AI if it does not understand AI.⁸ Getting technical talent into the federal workforce is the biggest obstacle to the U.S. government’s internal adoption of AI, effective regulation of its risks, and successful implementation of critical AI policy recommendations from the National Security Commission on AI, NAIAC, and others.⁹

II. Challenges

While much progress has been made, including the Advancing American AI Act, AI Training for the Acquisition Workforce Act, and other important legislation from this committee, we still have a long way to go.

When our research team at Stanford examined the implementation of AI-related legal requirements that EO 13,960, 13,859, and the AI in Government Act placed on federal agencies, we found a critical gap in leadership, strategic planning, and capacity.¹⁰ For example:

⁷ See, e.g., FDA, Artificial Intelligence and Machine Learning in Software as a Medical Device, <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>; DOT, USDOT Automated Vehicles Activities, <https://www.transportation.gov/AV>.

⁸ As an example, FDA approvals of AI medical devices appear to be based on a limited number of hospital sites, and the performance of an algorithm that performs well in hospital A can degrade significantly in hospital B. Technical expertise and domain knowledge are required to ensure that the device approval process is adapted to accounts such complexities with AI systems. See Wu, Eric, Kevin Wu, Roxana Daneshjou, David Ouyang, Daniel E. Ho, and James Zou. 2021. “How Medical AI Devices Are Evaluated: Limitations and Recommendations from an Analysis of FDA Approvals.” *Nature Medicine* 27 (4): 582–84.

⁹ I am not alone in this concern. When Eric Schmidt was asked a few months ago about his biggest concern on Congress’s implementation of the National Security Commission on AI’s recommendations, he singled out one thing: technical talent. Eric Schmidt, Testimony to House Subcommittee on Cybersecurity, Information Technology, and Government Innovation, March. 8, 2023 (“The area that I’m most focused on right now is basically the training problem and I just don’t see the progress in the government to reform the way it hires and promotes technical people.”).

¹⁰ Lawrence, Christie, Isaac Cui, and Daniel E. Ho. 2022. “Implementation Challenges to Three Pillars of America’s AI Strategy.” Stanford HAI-RegLab White Paper. https://dho.stanford.edu/wp-content/uploads/AI_Implementation.pdf

- 88 percent of agencies failed to submit AI Plans to identify regulatory authorities and mechanisms to promote responsible AI and protect Americans’ rights and safety;¹¹
- The implementation of a key transparency measure – agency disclosure of its AI use cases through published inventories – has been inconsistent;¹² and
- The Office of Personnel Management has yet to release a required report – due July of last year – to forecast employment needs and to create an AI hiring line.¹³

This must change. The federal workforce does Herculean work, but faces fundamental challenges developing teams that can design, implement, and regulate AI effectively and responsibly. The most recent AI Index report by Stanford HAI highlights that 65% of AI PhDs land in industry, 28% in academia, and less than 2% in government.¹⁴ Or, in the words of one entrepreneur: “The best minds of my generation are thinking about how to make people click ads.”¹⁵

Strengthening the pipeline of technical talent into the public sector is paramount. So too is ensuring that agencies have the right digital infrastructure, technical directives, and career paths to nurture and grow technical talent. As the National Security Commission on AI noted, it’s not just compensation: “It is the perception, and too often the reality, that it is difficult for digital talent in government to perform meaningful work.”¹⁶ I’ve seen firsthand how the government is failing to recruit and retain technical experts. One Stanford AI Ph.D. student, for instance, became so frustrated by an agency’s decades-old software stack and lack of advanced (GPU) computing that he gave up on government and went back to work in industry.

III. Recommendations

¹¹ See Section 6(c) of EO 13859, OMB M-21-06, and Lawrence, Cai, and Ho, *supra*, note 10.

¹² See Section 5(e) of EO 13960 and Lawrence, Cai, and Ho, *supra*, note 10.

¹³ Section 105 of the AI in Government Act directed the Office of Personnel Management to create an AI occupational series and estimate AI-related workforce needs in each federal agency by July 2022.

¹⁴ Nestor Maslej, Loredana Fattorini, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Vanessa Parli, Yoav Shoham, Russell Wald, Jack Clark, and Raymond Perrault, “The AI Index 2023 Annual Report,” AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2023.

¹⁵ Ashlee Vance, “This Tech Bubble Is Different,” *Bloomberg* (Apr. 14, 2011).

¹⁶ Schmidt, Eric, Bob Work, Safra Catz, Steve Chien, Chris Darby, Kenneth Ford, Jose-Marie Griffiths et al. “National Security Commission on Artificial Intelligence (AI), Final Report.” 2021.

Let me conclude with four recommendations necessary for the U.S. government to maintain American leadership in AI innovation and responsible AI.

First, strategic leadership from the federal government is required to coordinate and drive forward AI innovation and trustworthy adoption. Congress should borrow a page from the bipartisan Evidence Act, which required the appointment of officials responsible for data and evaluation, and empower Chief AI Officers to ensure that senior leadership within agencies is driving forward responsible AI innovation.¹⁷ The White House, too, must also be organized and staffed to rise to this challenge.¹⁸

Second, Congress should establish new pathways and trajectories for technical talent in government. We need better models – building on the US Digital Service, public-private partnerships, and academic-agency partnerships – to attract AI talent to public service and build cross-functional teams. Retaining AI talent requires giving them meaningful positions related to their expertise.

Third, an effective procurement system should capitalize on American innovation and spur developments of rights-preserving, privacy-enhancing technologies. The AI Training Act is a fantastic step in the right direction, but we cannot rely on procurement officials alone. We need upskilling of business units and to enable more modular forms of contracting – which DOD has illustrated – that enables more effective development, acquisition, assessment, and auditing of AI systems.¹⁹

Last, we have to invest in digital infrastructure. The federal task force proposal for a National AI Research Resource would give AI researchers, students, and small businesses secure access to

¹⁷ This recommendation is made by the National AI Advisory Commission's Year 1 Report. <https://www.ai.gov/wp-content/uploads/2023/05/NAIAC-Report-Year1.pdf>

¹⁸ See National AI Advisory Commission's Year 1 Report; Schmidt, Eric, Bob Work, Safra Catz, Steve Chien, Chris Darby, Kenneth Ford, Jose-Marie Griffiths et al. "National Security Commission on Artificial Intelligence (AI), Final Report." 2021.

¹⁹ Raji, Inioluwa Deborah, Peggy Xu, Colleen Honigsberg, and Daniel Ho. 2022. "Outsider Oversight: Designing a Third Party Audit Ecosystem for Ai Governance." In Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, 557–71.

high-quality administrative data and computing resources to level the playing field between industry, academia, and government.²⁰ AI requires high-fidelity data and many of the negative impacts of AI we've observed stem from training large models on anything on the web, including unverified and harmful information. Government data, which is higher-quality, more representative, and more reliable is an important part of the solution. When the U.S. Geological Service made Landsat satellite imagery free to researchers in 2008, it generated 3 to 4 billion dollars in benefits annually, catalyzing discoveries in habitat modification, climate change, and poverty.²¹ That is the promise of getting the public sector innovation infrastructure right.

The U.S. government should act expeditiously to foster responsible AI adoption. AI that embodies American values and helps agencies better serve Americans equitably can also build public trust and confidence. Thank you for the opportunity to speak. I'm looking forward to your questions.

²⁰ National Artificial Intelligence Research Resource Task Force, "Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource," Jan. 2023, <https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf>. Ho, Daniel, Jennifer King, Russell Wald, and Christopher Wan. 2021. "Building a National AI Research Resource: A Blueprint for the National Research Cloud." Stanford HAI White Paper. https://hai.stanford.edu/sites/default/files/2022-01/HAI_NRCR_v17.pdf.

²¹ Straub, Crista L., Stephen R. Koontz, and John B. Loomis. "Economic Valuation of Landsat Imagery." 2019.