



Governing AI Through Acquisition and Procurement

Fei-Fei Li¹

Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University

Testimony presented to the U.S. Senate Committee on Homeland Security and Governmental Affairs on September 14, 2023.

I. Introduction

Chairman Peters, Ranking Member Paul, Members of the Committee, thank you for the privilege of appearing before this prestigious body. As a proud American, few things could make me more honored than offering service to our government leaders.

I have spent my life working in the field of artificial intelligence (AI) with over 25 years studying, developing, and understanding the technology that has just entered the public's consciousness due to recent breakthroughs. I have approached the study and development of AI not only as a pioneer, academic scientist, and teacher, but also as the child of two parents with chronic health issues.

Right now you are hearing wild claims about AI from two ends of the spectrum. Some, propelled by industry hype, have said that AI will fix all of the challenges humanity faces. At the other end of the spectrum, some have claimed AI will lead to the end of the world through a biological or nuclear catastrophe. As a proud, self-proclaimed, "nerd," I, too, love a good science fiction story. However, one goal of my testimony today is to demystify some of these wild claims for you.

II. Benefits of AI

We have arrived at an inflection point in the world of AI, largely propelled by breakthroughs in generative AI, including increasingly sophisticated language models like GPT-4. These models have revolutionized various sectors from customer service to adaptive learning. However, the scope of intelligence is far broader than linguistic capability alone. In my specialized field of computer vision, we have also witnessed remarkable advancements that empower machines to analyze and act upon visual information—essentially teaching computers to 'see.'

I'd like to particularly highlight today how AI advancement is augmenting the capabilities of caretakers and medical professionals. Healthcare is a domain that I have dedicated my entire

¹ Sequoia Capital Professor in Computer Science, Stanford University; Denning Co-Director and Senior Fellow, Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University; Co-Founder and Board Chair, AI4ALL; Member of the National Academy of Engineering, the National Academy of Medicine and American Academy of Arts and Sciences. All views expressed in this testimony are provided in an individual capacity and do not represent the views of any affiliated organization.

career to and I am honored to contribute to the field as a member of the National Academies of Medicine. Earlier, I mentioned that I am the primary caregiver for my aging parents. This personal experience has provided me with invaluable insights into the challenges that caregivers face daily and how AI can transform the healthcare landscape. For example, algorithms can detect anomalies in medical imagery such as X-rays and MRI scans, thereby aiding early diagnosis and treatment.² AI-enabled documentation assistance can reduce the administrative burden on healthcare professionals but also minimize the risk of errors in patient records, thereby improving patient safety and care quality.³

Most importantly, such an innovation presents many benefits to the U.S. government.

First, government entities have long been consumers and stewards of public-use technology to streamline the efficiency of governance. Nearly half of all federal agencies have experimented with AI and related machine learning (ML) tools.⁴ As technology continues to advance at an unprecedented pace, it is imperative that our government remains at the forefront of innovation to better service delivery to its citizens and address complex societal challenges more effectively. When it comes to healthcare, for example, the Department of Health and Human Services has initiated a pilot program aimed at combating Medicare fraud.⁵ This program employs AI-based models to enhance the efficiency of fraud detection within the Centers for Medicare & Medicaid Services (CMS), which handles over one million transactions daily.

Second, the use of AI in healthcare not only improves the quality of life for the elderly population, but also reduces healthcare costs.⁶ Medical AI tools can reduce the frequency of emergency medical interventions and hospital readmissions, which are significant cost drivers in healthcare expenditures. By facilitating proactive, preventive care, the technology can mitigate the financial burden on public healthcare resources, including Medicare and Medicaid, and facilitate more efficient allocation of resources in healthcare settings, enabling medical professionals to focus on other critical areas of public health and welfare that require attention.

III. Harms and Unintended Consequences

However, while AI, like most technologies, promises to solve many problems for the common good, it can also be misused to cause harm and carry unintended consequences. The very same technology that could potentially save countless lives from hospital infections could one day be

² Nikki Goth Itoi. 2022. Could Stable Diffusion Solve a Gap in Medical Imaging Data? Stanford Institute for Human-Centered AI. <https://hai.stanford.edu/news/could-stable-diffusion-solve-gap-medical-imaging-data>.

³ Grace Hong et al. 2020. Clinicians' Experiences with EHR Documentation and Attitudes Toward AI-Assisted Documentation. Stanford University School of Medicine and Google Health. https://med.stanford.edu/content/dam/sm/healthcare-ai/images/Stanford-Google_AI-Scribe_WhitePaper.pdf.

⁴ David Freeman Engstrom, 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 Nihal Krishan. 2023. HHS CIO Mathias Says Tree-based AI Models Helping to Combat Medicare fraud. Fedcoop. <https://fedcoop.com/hhs-cio-mathias-says-tree-based-ai-models-helping-to-combat-medicare-fraud/>.

⁶ Milstein A. Haque and L Fei-Fei. 2020. Illuminating the Dark Spaces of Healthcare with Ambient Intelligence. *Nature* 585. <https://www.nature.com/articles/s41586-020-2669-y>, 193–202.

repurposed as an active form of surveillance against people. I know Senator Paul and others on this committee's longstanding concerns related to surveillance and I, too, appreciate these very same concerns as an AI developer.

Let me give you two examples of when the harms of AI could affect how the government approaches AI. First, bias in AI is well-documented.⁷ Take credit scoring as an example. As credit risk scoring tools increasingly leverage AI, research shows that predictive tools used to approve or reject loans are less accurate for low-income, minority groups in the United States due to the lack of data in their credit histories.⁸ To ensure that AI applications deliver reliable results for all Americans, we must ensure the availability of high-quality, representative data sets.

Second, healthcare AI also presents considerable challenges to privacy and data security. In an era of heightened public concern over data collection and misuse, we must build strong privacy and security protocols into these applications from the beginning. Achieving this necessitates a multidisciplinary effort that engages experts across various fields. Developers, policymakers, healthcare providers, and patients should all be proactively involved throughout the entire development and implementation phases to ensure the AI tools are both effective and secure.⁹

This is why we must ensure a diverse ecosystem in the development of AI and why I helped found Stanford's Institute for Human-Centered Artificial Intelligence (HAI), where we study AI and its impact not as a field exclusive to computer science, but instead as a multidisciplinary field that includes the social sciences, engineering, law, medicine, and the humanities. The federal government should adopt a similar approach to properly understand the future of AI.

IV. What Can the U.S. Government Do

The adoption of AI introduces unique complexities and issues, not only for the public bodies acquiring these technologies but also for the communities they impact. It falls upon the U.S. government to spearhead the ethical procurement and deployment of these systems, with the aims of both safeguarding the rights of individuals and communities, as well as encouraging industry innovation through responsible AI guidelines. Indeed, responsible federal acquisitions and procurement have the true potential to set the norms for AI development and ultimately shape the field of responsible AI in a more immediate and direct way than any future regulation that may or may not come from this Congress.

⁷ See David Danks and Alex John London. 2017. "Algorithmic Bias in Autonomous Systems," *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. <https://doi.org/10.24963/ijcai.2017/654>; Joy Buolamwini and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," *Proceedings of Machine Learning Research* 81. <https://doi.org/https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>; Yolande Strengers, Lizhen Qu, Qiongkai Xu, and Jarrod Knibbe. 2020. Adhering, Steering, and Queering: Treatment of Gender in Natural Language Generation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3313831.3376315>.

⁸ Laura Blattner and Scott Nelson. 2021. How Costly is Noise? Data and Disparities in Consumer Credit. arXiv. <https://arxiv.org/abs/2105.07554>.

⁹ Jenna Wiens et al. 2019. "Do No Harm: A Roadmap for Responsible machine learning for health care." *Nature Medicine* 25:9: 1337-1340. <https://pubmed.ncbi.nlm.nih.gov/31427808/>.

I applaud this committee and the work that it has done thus far on AI, including the AI Training Act and the AI Lead Act, which create powerful tools for the federal government to set such norms.¹⁰ The AI Training Act can up-skill procurement officials and equip them with a nuanced understanding of AI capabilities and limitations. I am proud that Stanford HAI has tailored a government education program to fulfill the Training Act requirement specifically for the U.S. General Services Administration, in partnership with the Office of Management and Budget.

As the U.S. government's spending on AI-related contracts has surged, it's more crucial than ever to closely examine these vendors to ensure their goals align with those of the federal government.¹¹ One key component is evaluation, especially in high-risk areas like healthcare, education, and finance. Having created one of the most consequential evaluation datasets for AI models, ImageNet, I firmly believe that evaluation should consider every factor in a holistic way — from accuracy, to fairness and the reliability of models performing under real-world conditions.¹² Second, we must build in transparency measures. Vendors should disclose key information about their systems, including how they collect and annotate datasets; what potential risks their systems pose; and how they mitigate those risks.¹³

V. U.S. Public Sector Investment

For the reasons I have just outlined, the federal government must make the needed critical investments in AI. While the numerous benefits of AI such as efficiency and productivity gains in government that allow it to be more responsive to its citizens make AI attractive, this future is not foretold and it may not be dominated by the United States.

That is why with its democratic values, commitment to the rule of law, and spirit of innovation and entrepreneurship, America must lead in AI. And while the U.S. industry currently enjoys unique advantages in AI, there is a deep imbalance in the U.S. AI innovation ecosystem that hinders leadership in the field.

Because of the vast amounts of compute and data required to train these systems, only a select few industry players can currently work at the frontiers to shape the future of the technology,

¹⁰ See S.2551 - 117th Congress (2021-2022): AI Training Act, S.2551, 117th Cong. 2022. <https://www.congress.gov/bill/117th-congress/senate-bill/2551/summary/55>; S.2293 - 118th Congress (2023-2024): AI LEAD Act, S.2293, 118th Cong. 2023. <https://www.congress.gov/bill/118th-congress/senate-bill/2293>.

¹¹ See Nestor Maslej et al. 2023. "The AI Index 2023 Annual Report." AI Index Steering Committee, Stanford Institute for Human-Centered AI.

¹² Jia Deng et al. 2009. ImageNet: A Large-scale Hierarchical Image Database. IEEE Conference on Computer Vision and Pattern Recognition. <https://ieeexplore.ieee.org/document/5206848>, 248-255; Rishi Bommasani, Daniel Zhang, Tony Lee, Percy Liang. 2023. Improving Transparency in AI Language Models: A holistic evaluation. Foundation Model Issue Brief Series. Stanford Institute for Human-Centered AI. <https://hai.stanford.edu/foundation-model-issue-brief-series>;

¹³ See Emanuel Moss et al. 2021. Assembling Accountability: Algorithmic Impact Assessment for the Public Interest. Data & Society Report. <https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/>; Margaret Mitchell et al. 2018. Model Cards for Model Reporting. arXiv. <https://arxiv.org/abs/1810.03993>; Rishi Bommasani, et al. 2023. Ecosystem Graphs: The Social Footprint of Foundation Models. arXiv. <https://arxiv.org/abs/2303.15772>.

leaving an imbalance in the AI innovation ecosystem that lacks the diverse voices of academia and government labs. In fact, of the 32 significant industry breakthroughs in AI last year, only three originated from academia, and none from government labs.¹⁴

The lack of public sector investment in AI means that not only thoughtful regulation but also proper federal procurement and acquisition processes are at risk. Without the ability to train AI talent, the federal government will not have the necessary human capital to create meaningful regulation, ensure ethical AI procurement, and be the true AI leader it has the potential to be.

In June I personally shared with President Biden how I believe the United States is not prepared for this imminent AI moment and how the federal government needs to change its thinking about AI and adopt a moonshot mentality. If the United States is to truly lead in AI, we must not only adopt a robust regulatory and procurement framework, but must also invest in scientific AI research.

This is why I am unequivocally a strong supporter of the CREATE AI Act, a strong bipartisan legislation introduced this summer in both chambers.¹⁵ The CREATE AI Act will establish a National AI Research Resource which will ultimately provide the needed computational infrastructure and data resources to allow academic researchers to innovate and train the next generation of AI leaders.

What we need right now is a coordinated moonshot effort for the nation to ensure America's leadership in AI for the good of humanity. And that moonshot will have to include infrastructure investments such as the National AI Research Resource, as well as national labs focusing on solving the hardest problems. This task will be no small feat, but with meticulous coordination, significant investment, and robust collaboration across government, academia, and industry, we can rise to meet this challenge and ensure America's leadership in AI is both impactful and enduring.

Thank you to the Chairman, Ranking Member, and all the Members of the Committee for allowing me to testify today.

¹⁴ Nestor Maslej et al. 2023.

¹⁵ See S.2714 - 118th Congress (2023-2024): CREATE AI Act of 2023, S.2714, 118th Cong. 2023. <https://www.congress.gov/bill/118th-congress/senate-bill/2714>; H.R.5077 - 118th Congress (2023-2024): Creating Resources for Every American To Experiment with Artificial Intelligence Act of 2023, H.R.5077, 118th Cong. 2023. <https://www.congress.gov/bill/118th-congress/house-bill/5077>.