Stanford University
Human-Centered
Artificial Intelligence

Recommendations on Updating the National Artificial Intelligence Research and Development Strategic Plan

A WHITE PAPER FOR THE STANFORD INSTITUTE FOR HUMAN-CENTERED ARTIFICIAL INTELLIGENCE

Co-led by Daniel E. Ho, Jennifer King, Russell Wald, and Daniel Zhang

DISCLAIMER

MARCH 2022

Contributors

PRINCIPAL AUTHORS

Daniel E. Ho, J.D., Ph.D. is the William Benjamin Scott and Luna M. Scott Professor of Law, professor of political science, and senior fellow at the Stanford Institute for Economic Policy Research at Stanford University. He directs the Regulation, Evaluation, and Governance Lab (RegLab) at Stanford, and is a faculty fellow at the Center for Advanced Study in the Behavioral Sciences and associate director of the Stanford Institute for Human-Centered Artificial Intelligence (HAI). He received his J.D. from Yale Law School and Ph.D. from Harvard University and clerked for Judge Stephen F. Williams on the U.S. Court of Appeals for the District of Columbia Circuit.

Jennifer King, Ph.D. is the privacy and data policy fellow at the Stanford Institute for Human-Centered Artificial Intelligence (HAI). She completed her doctorate in information management and systems (information science) at the University of California, Berkeley School of Information. Prior to joining HAI, she was the Director of Consumer Privacy at the Center for Internet and Society at Stanford Law School from 2018 to 2020.

Russell Wald is the director of policy at the Stanford Institute for Human-Centered Artificial Intelligence (HAI), where he is responsible for leading the team that advances Stanford HAI's engagement with governments and civil society organizations. He is the coauthor of, "Building a National AI Research Resource: A Blueprint for the National Research Cloud," which is the most comprehensive study on a National Research Cloud to date. Wald has previously held various government relations roles at Stanford University. He is a term member with the Council on Foreign Relations, visiting fellow with the National Security Institute at George Mason University, and a partner with the Truman National Security Project. Wald is a graduate of UCLA.

Daniel Zhang is the policy research manager at the Stanford Institute for Human-Centered Artificial Intelligence (HAI). In this role, he leads the development of the AI Index program and its annual report that aims to measure and evaluate the rapid rate of AI advancement. Daniel also works on the policy team at HAI to develop and manage several policy research programs, building the bridge between technology research and policy communities. Previously, he worked on global AI talent flows and security risks at the Center for Security and Emerging Technology. Daniel holds an M.A. in Security Studies from Georgetown University and a B.A. in Politics & International Affairs from Furman University.

ACKNOWLEDGMENTS

Recommendations on Updating the National Artificial Intelligence Research and Development Strategic Plan

The Stanford Institute for Human-Centered Artificial Intelligence (HAI) offers the following submission for consideration in response to the Request for Information (RFI) by the White House Office of Science and Technology to the Update of the National Artificial Intelligence Research and Development Strategic Plan. Our submission recommends:

- For Strategy 1: Boost non-defense AI R&D budgets, particularly on AI-related infrastructure, to support long-term investments.

- For Strategy 2: Increase support for interdisciplinary and multidisciplinary AI research on human-AI collaboration that expands beyond exclusively technical research.

- For Strategy 3: Please refer to the Stanford HAI letter submitted in January 2022 in response to the White House Office of Science and Technology proposal for an AI Bill of Rights that safeguards the American public against powerful technologies.

- For Strategy 4: Develop appropriate acquisition strategies and update existing procurement regulations to respond to AI procurement and acquisition challenges in the federal government.

- For Strategy 5: Expand government data access to academic researchers to train AI models and develop frameworks for government agencies to evaluate such datasets and their applications in tandem.

- For Strategy 6: Establish a mechanism to evaluate AI models within the exact context of their intended use to ensure safe deployment as well as designate NIST in collaboration with other federal agencies to benchmark AI models in institutional contexts.

- For Strategy 7: Update immigration policies to attract talent in AI and other technical fields as well as develop federal programs to hire AI talent and re-skill civil servants with both technical capacity and institutional knowledge.

- For Strategy 8: Strengthen partnerships with academic institutions and build a framework for a public-university-industry AI R&D ecosystem to drive AI development forward.

Strategy 1: Make long-term investments in AI research.

Recommendation: Boost non-defense AI R&D budgets, particularly on AI-related infrastructure, to support long-term investments.

A long-term commitment to sustained federal research and development (R&D) funding in AI is critical to advance the United States' leadership in global innovation. The federal government should increase non-defense investment in AI and basic research to strengthen research in critical fields of AI R&D, including healthcare, education, finance, and more, that underpin economic stability and robust growth. Such investment should reflect a multidisciplinary approach, focused on advancing basic and applied R&D, research on AI governance and norm-setting, and supporting research infrastructure with multi-agency collaboration.

Current federal funding for non-defense AI R&D, however, does not meet the needs of the fast-growing AI field. The public non-defense AI R&D budget requested by 25 federal agencies participating in the Networking and Information Technology Research and Development (NITRD) program and the National Artificial Intelligence Initiative in FY 2022 represents an increase of just 8.8 percent over what was spent in FY 2021.[1] In contrast, the National Security Commission on Artificial Intelligence (NSCAI) recommended in its final report to increase public funding for AI R&D at compounding levels, doubling annually to reach $32 billion per year by FY 2026.[2]

Federal long-term non-defense investments—and high return-on-investment basic research funding[3]—can address the challenges the AI innovation ecosystem is currently facing in the United States. For example, the high cost of compute and the lack of access to critical data are hindering efforts by academic researchers to engage in cutting-edge AI R&D.[4] The federal government is lagging behind the private sector on AI development, and federal standards for technical and ethical AI are sorely needed.[5] Long-term public investment in AI-related infrastructure can strengthen AI R&D by supporting a variety of federal initiatives, including the National Artificial Intelligence Research Resource (NAIRR) that aims to expand access to "critical resources and educational tools that will spur AI innovation and economic prosperity nationwide."[6] Another example of such an initiative is the Multilateral AI Research Institute (MAIRI), recommended by the NSCAI report, that would "facilitate joint efforts to develop technologies that advance responsible, human-centric, and privacy-preserving AI/machine learning (ML) that better societies and allow allies to pool their talents and resources."[7]

---

[1] The Networking & Information Technology R&D Program and the National Artificial Intelligence Initiative Office, "Supplement to the President's FY 2022 Budget," December 2021, https://www.nitrd.gov/pubs/FY2022-NITRD-NAIIO-Supplement.pdf.

[2] Eric Schmidt et al., "Final Report," National Security Commission on Artificial Intelligence, March 2021, https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital.pdf.

[3] Benjamin F. Jones and Lawrence H. Summers, A Calculation of the Social Returns to Innovation (Nat'l Bureau of Econ. Research, Working Paper No. 27863, 2020).

[4] Daniel E. Ho, Jennifer King, Russell C. Wald, and Christopher Wan, "Building a National AI Research Resource: A Blueprint for the National Research Cloud," Stanford Institute for Human-Centered Artificial Intelligence, October 2021, https://hai.stanford.edu/sites/default/files/2022-01/HAI_NRCR_v17.pdf.

[5] David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey, and Mariano Florentino-Cuéllar, "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies," Stanford Law School, February 2020, https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf.

[6] "The Biden Administration Launches the National Artificial Intelligence Research Resource Task Force," The White House (The United States Government, June 10, 2021), https://www.whitehouse.gov/ostp/news-updates/2021/06/10/the-biden-administration-launches-the-national-artificial-intelligence-research-resource-task-force/.

[7] Schmidt et al., "Final Report."

Strategy 2: Develop effective methods for human-AI collaboration.

<u>Recommendation: Increase support for interdisciplinary and multidisciplinary AI research on human-AI collaboration that expands beyond exclusively technical research.</u>

Intentionally building trustworthy AI that is unbiased and supportive of human flourishing is crucial to ensuring the successful development and deployment of human-centered AI. A key part of that effort requires an interdisciplinary and multidisciplinary approach, involving collaboration from a variety of fields to develop the hardware and software, to understand and design for people's behaviors and expectations when interacting with AI in different institutional contexts, and to establish policies and regulations to determine human responsibilities as well as the required domain knowledge for various applications. A human-centric approach to AI calls for wide-ranging collaboration among multiple disciplines. Harnessing the potential of AI with rapidly growing capabilities while addressing its impact on existing structural inequalities and biases cannot rely on the voices of computer scientists and engineers alone.

Yet current policies do not necessarily match this need. The National Science Foundation program on algorithmic fairness, for instance, calls for interdisciplinary perspectives while stating that "this program supports the conduct of fundamental computer science research" and requiring the PI to "bring computer science expertise to the research."[8] But producing fairness-aware algorithms or just understanding the concept of fairness requires knowledge and expertise outside the computer science field to incorporate the social and legal contexts in which AI systems will be deployed.[9] The federal government should expand the support of multidisciplinary AI research for human-AI collaboration to include critical academic fields such as the social sciences, law, ethics, and the humanities—all of which should have a prominent voice in providing the necessary frameworks for understanding AI today and in the future.

Strategy 3: Understand and Address the Ethical, Legal, and Societal Implications of AI.

With respect to recommendations for Strategy 3, please refer to the Stanford HAI letter submitted in January 2022 in response to the White House Office of Science and Technology proposal for an AI Bill of Rights that safeguards the American public against powerful technologies.[10]

Strategy 4: Ensure the Safety and Security of AI Systems.

<u>Recommendation: Develop appropriate acquisition strategies and update existing procurement regulations to respond to AI procurement and acquisition challenges in the federal government.</u>

Public sector AI can rely heavily on contracting and procurement with external vendors to build up technical capacity. Research shows that almost half of identified use cases of federal agencies' use of AI came from external sources, with one-third coming from private commercial

---

[8] "NSF Program on Fairness in Artificial Intelligence in Collaboration with Amazon (FAI)," The National Science Foundation, 2021, https://www.nsf.gov/pubs/2021/nsf21585/nsf21585.htm.

[9] Andrew D. Selbst et al., "Fairness and Abstraction in Sociotechnical Systems," *Proceedings of the Conference on Fairness, Accountability, and Transparency* (January 2019): 59-68, https://doi.org/10.1145/3287560.3287598.

[10] Michele Elam and Rob Reich, "Stanford HAI Artificial Intelligence Bill of Rights," Stanford Institute for Human-Centered Artificial Intelligence, January 2022, https://hai.stanford.edu/white-paper-stanford-hai-artificial-intelligence-bill-rights.

sources via the procurement process.[11] Compared to the internal sourcing of AI systems that may be more policy-compliant and more accountable, the uses of procured AI in government raise several concerns in terms of AI trustworthiness, transparency, and safety. The federal acquisition regulation (FAR), for example, provides strong IP protection for vendors.[12] But such protections can obscure certain information about the inputs their tools use and how the tools operate behind trade secrecy claims, which in turn prevents appropriate analysis, audit, and testing to ensure the fairness of their use.[13] Moreover, those protections may also create uncertainty for acquisition aimed at the black-box nature of AI systems. FAR makes a clear distinction between rights to "software" and rights to "data," but AI systems, particularly machine learning (ML), integrate customer software with the new data generated in the process of training, and current procurement policies do not sufficiently address how rights to that data and the resulting AI systems are to be distributed, and under what constraints and conditions.[14]

The federal government should develop appropriate acquisition strategies and update existing procurement regulations to help address some of the public sector's challenges in evaluating, monitoring, and using AI systems.[15] Specific examples include developing clear standards that call for the disclosure of data and information on the design and operation of contractors' algorithms, requirements that ensure contractors adhere to ethical AI standards, and testing infrastructures that allow for iterative testing and evaluation.[16]

Strategy 5: Develop shared public datasets and environments for AI training and testing.

Recommendation: Expand government data access to academic researchers to train AI models and develop frameworks for government agencies to evaluate such datasets and their applications in tandem.

While there are publicly available resources for AI development, there is still more work to be done to promote the open and collaborative non-commercial use of training and testing data and environments. For example, access to data resources sufficient for training AI systems is increasingly limited to large private companies, which in turn direct resources toward developing applications with a focus on private profit instead of public interest.[17] Because large platforms have unequaled access to data for AI development, smaller actors, some of which may legitimately lack the financial resources to invest in building training data from scratch, are incentivized to mine the public sphere for data, violating individual privacy expectations and creating privacy risks both for individuals and society at large. At the same time, there are

[11] Engstrom et al., "Government by Algorithm."
[12] "Federal Acquisition Regulation Part 27 - Patents, Data, and Copyrights," U.S. General Services Administration, accessed March 2022, https://www.acquisition.gov/far/part-27.
[13] Deirdre K. Mulligan and Kenneth A. Bamberger, "Procurement As Policy: Administrative Process for Machine Learning," *Berkeley Technology Law Journal* 34 (2019), https://dx.doi.org/10.2139/ssrn.3464203.
[14] Ken Farber, Kristine Lam, and Ellery Taylor, "From Ethics to Operations: Current Federal AI Policy," Advanced Technology Academic Research Center, October 4, 2021, https://atarc.org/wp-content/uploads/2021/10/Current-Federal-AI-Policy-Assessment-20211004-for-public-comment.pdf.
[15] Laura Gerhardt and Mark Headd, "Digital Service Delivery: Why We Love Modular Contracting," 18F, April 9, 2019, https://18f.gsa.gov/2019/04/09/why-we-love-modular-contracting/.
[16] Engstrom et al., "Government by Algorithm."; Lavi M. Dor and Cary Coglianese, "Procurement as AI Governance," *IEEE Transactions on Technology and Society* 2, no. 4 (2021): pp. 192-199, https://doi.org/10.1109/tts.2021.3111764.
[17] Jathan Sadowski, "When Data Is Capital: Datafication, Accumulation, and Extraction," *Big Data & Society* (January 2019), https://doi.org/10.1177/2053951718820549; Ho et al., "Building a National AI Research Resource."

significant barriers for interagency and external researchers to access a rich portfolio of public sector data (e.g., employment, healthcare, education).[18]

As a starting point, the executive branch can use the NAIRR as an opportunity to make more and better quality government data available to the research community at no cost.[19] In doing so, the federal government should weigh considerations such as privacy, security, and fairness, and it should begin to develop its own frameworks for evaluating such datasets and their applications in tandem, informed by important developments under the Foundations for Evidence-Based Policymaking Act of 2018 and the National Secure Data Service. Considering the lack of a standardized framework to test and evaluate AI models for safety and security, the government should pursue creating a testing environment for this purpose, especially for government-procured AI systems.[20]

Strategy 6: Measure and evaluate AI technologies through standards and benchmarks.

Recommendations:
- Establish a mechanism to evaluate AI models within the exact context of their intended use to ensure safe deployment.
- Designate NIST in collaboration with other federal agencies to benchmark AI models in institutional contexts.

Understanding the true in-domain accuracy, or the accuracy of AI systems' deployment in specific contexts (e.g., in different industries, with different subpopulation groups), is crucial for the federal government to capture the capabilities of the technology and ensure safe deployment. Many current performance evaluations do not comprehensively assess how AI systems would perform in a real-world context.[21] Object recognition systems, for example, are often evaluated against large-scale benchmark datasets to validate their performance, but such datasets are limited in their coverage of non-Western contexts and temporally bounded, capturing only parts of the real world.[22] The same gap is observed in language models where these state-of-the-art systems amplify human bias and discriminate against minority users, and their performance degrades when given out-of-domain text.[23] All told, the accuracy of AI systems in one domain does not automatically translate to its uses in other domains, and changing context can significantly impact performance.

[18] Amy O'Hara and Carla Medalia, "Data Sharing in the Federal Statistical System: Impediments and Possibilities," *The Annals of the American Academy of Political and Social Science* 675 (2018): 138–50, https://www.jstor.org/stable/26582286.

[19] Ho et al., "Building a National AI Research Resource," 35.

[20] Note that we recommend government agencies withhold certain data to allow sufficient testing. For example, NIST's Face Recognition Vendor Test (FRVT) challenge does not provide training images because "the tests seek to mimic operational reality and, there, algorithms are almost always shipped and used 'as is' without any training or adaptation to customer data." Read "Face Recognition Vendor Test Ongoing–Frequently Asked Questions (FAQs)," National Institute of Standards and Technology, November 6, 2019, https://www.nist.gov/system/files/documents/2019/04/22/frvt_frequently_asked_questions.pdf; Avrim Blum and Moritz Hardt, "The Ladder: A Reliable Leaderboard for Machine Learning Competitions," arXiv.org, February 16, 2015, https://arxiv.org/abs/1502.04585. For creating a testbed for the federal government, see Tina Huang, "Creating an AI Testbed for Government," Institute for Progress, January 2022, https://progress.institute/wp-content/uploads/2022/01/Creating-an-AI-Testbed-for-Government-_final.pdf.

[21] Inioluwa Deborah Raji et al., "AI and the Everything in the Whole Wide World Benchmark," arXiv.org, November 26, 2021, https://arxiv.org/abs/2111.15366.

[22] Raji et al., "AI and the Everything in the Whole Wide World Benchmark."

[23] Kawin Ethayarajh and Dan Jurafsky, "Utility Is in the Eye of the User: A Critique of NLP Leaderboards," *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (EMNLP), 2020, https://doi.org/10.18653/v1/2020.emnlp-main.393.; Thomas Manzini et al., "Black Is to Criminal as Caucasian Is to Police: Detecting and Removing Multiclass Bias in Word Embeddings," *Proceedings of the 2019 Conference of the North*, 2019, https://doi.org/10.18653/v1/n19-1062.

The White House should consider a proposal to charge the National Institute of Standards and Technology (NIST), in collaboration with federal agencies that have regulatory oversight for AI-powered products, such as the U.S. Food and Drug Administration (FDA), Consumer Financial Protection Bureau (CFPB), and the National Highway Traffic Safety Administration (NHTSA), to develop improved AI benchmarking protocols.[24] Such benchmarks should explicitly address and incorporate the institutional contexts in which the AI systems are developed (e.g., commercial settings) and deployed (e.g., border control). NIST should also consider how to measure the effectiveness of deploying AI in real-world settings where enabling technologies (e.g., cameras, microphones, computing hardware) and other factors may vary. For instance, one field experiment of an earlier generation of predictive policing algorithms found that models that worked well in the lab did not perform well in the field, failing to reduce crime when used in context.[25] However, it should do so from a human-centered perspective grounded in ethical frameworks, such as privacy by design, and not focus exclusively on identifying technical benchmarks. This requirement suggests a potentially different set of roles and expertise than NIST has had in the past.

Strategy 7: Better understand the national AI R&D workforce needs.

Recommendations:
- Update immigration policies to attract talent in AI and other technical fields.
- Develop federal programs to hire AI talent and re-skill civil servants with both technical capacity and institutional knowledge.

Future U.S. leadership in AI hinges on the country having the necessary talent-generation process and hiring pipeline—as well as the ability to attract and retain talent that already exists. The country needs individuals who are not only equipped with the skills to build AI systems, but who also know when, where, and how to ask the right questions when such systems pose risks to individuals and society and/or break the law. While not every AI expert has or should have a technical background, understanding the technology is essential to designing and implementing accountable AI initiatives. Such talent could also be a useful tool of accountability, helping design and maintain transparent, auditable, and responsible systems as well as engaging with stakeholders to ensure trustworthiness in AI systems.[26] This could include expediting the hiring process for certain kinds of AI talent,[27] updating immigration policies to attract and retain

[24] For FDA, see "Artificial Intelligence and Machine Learning (AI/ML) Medical Devices," U.S. Food and Drug Administration (FDA, September 22, 2021), https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices; for CFPB, see Patrice Alexander Ficklin, Tom Pahl, and Paul Watkins, "Innovation Spotlight: Providing Adverse Action Notices When Using AI/ML Models," Consumer Financial Protection Bureau, July 7, 2020, https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/; for NHTSA, see "Automated Vehicles for Safety," National Highway Traffic Safety Administration, accessed March 2022, https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety.
[25] Daniel E. Ho, Emily Black, Maneesh Agrawala, and Li Fei-Fei, "Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains," Stanford Institute for Human-Centered Artificial Intelligence, November 2020, https://hai.stanford.edu/sites/default/files/2020-11/HAI_FacialRecognitionWhitePaper_Nov20.pdf; Priscillia Hunt, Jessica Saunders, and John S. Hollywood, "Evaluation of the Shreveport Predictive Policing Experiment," National Institute of Justice, 2014, https://nij.ojp.gov/topics/articles/evaluation-shreveport-predictive-policing-experiment.
[26] Engstrom et al., "Government by Algorithm."
[27] See, e.g., problems hiring young tech talent: Jack Corrigan, "The Government's Struggle to Hire Young Tech Talent Is Worse Than You Thought," *Nextgov*, December 1, 2017, https://www.nextgov.com/cio-briefing/2017/12/governments-struggle-hire-young-tech-talent-worse-you-thought/144225/.

technical talent,[28] and expanding opportunities for permanent residency to those with AI and other technical degrees, as well as opportunities for entrepreneurs.[29]

Improving public sector capacity can be essential to accountability and oversight to ensure AI systems are (or are not) built and deployed in ways that promote, rather than degrade, the public interest. Currently, the federal government faces numerous challenges in hiring AI talent, including competing with private-sector salaries and benefits, competing with shorter private-sector hiring timelines, and enticing applicants who face an onerous and often confusing federal hiring process versus an often faster and easier one in industry or civil society.[30] The White House should work on its own, as well as with Congress and with industry partners, to identify the biggest challenges in attracting AI talent to the government and ways to potentially resolve them. Several agencies have initiated such efforts. In 2019, the Office of Personnel Management (OPM) established a classification of information technology (IT) positions to ease the hiring burden of the federal government's competitive service and the "direct hire" appointing authority for several STEM and IT positions for agencies with critical hiring needs.[31] In 2021, the U.S. General Services Administration (GSA) launched a two-year fellowship aimed at placing early-career software engineers, data scientists, and others with technical skills in federal agencies.[32]

Finally, it bears mentioning that the government should take a long-term view to developing future AI talent and expertise by investing in technical skills development in primary and secondary education through STEM-focused educational initiatives.

Strategy 8: Expand public-private partnerships to accelerate advances in AI.

Recommendation: Strengthen partnerships with academic institutions and build a framework for a public-university-industry AI R&D ecosystem to drive AI development forward.

The federal workforce to date can still lack appropriate training in AI to use it effectively in public operations and to discharge regulatory responsibilities. A Government Business Council and Accenture survey found that more than 60 percent of federal employee respondents worry about the "lack of technical support and user training" for public AI deployment.[33] An evaluation of the U.S. Customs and Border Protection (CBP)'s facial recognition program used at air exit by the Government Accountability Office (GAO) found that agents on the ground received little

[28] Tina Huang and Zachary Arnold, "Immigration Policy and the Global Competition for AI Talent," Center for Security and Emerging Technology, June 2020, https://cset.georgetown.edu/publication/immigration-policy-and-the-global-competition-for-ai-talent/.

[29] The House version of the America COMPETES Act of 2022 is one possible starting point. See "H.R. 4521 - Bioeconomy Research and Development Act of 2021 [America COMPETES Act of 2022]," House of Representatives Committee on Rules, February 2022, https://rules.house.gov/bill/117/hr-4521.

[30] Joan Timoney, "Building a Federal Civil Service for the 21st Century: The Challenge of Attracting Great Talent to Government Service," Pan-Organizational Summit on the US Science and Engineering Workforce: Meeting Summary (U.S. National Library of Medicine, 2003), https://www.ncbi.nlm.nih.gov/books/NBK36382/.

[31] Margaret M. Weichert, "Delegation of Direct-Hire Appointing Authority for IT Positions," Office of Personnel Management, April 5, 2019, https://chcoc.gov/content/delegation-direct-hire-appointing-authority-it-positions.

[32] Chris Kuang, "Introducing the U.S. Digital Corps: A New Path to Public Service for Early-Career Technologists," U.S. General Services Administration, August 30, 2021, https://www.gsa.gov/blog/2021/08/30/introducing-the-us-digital-corps-a-new-path-to-public-service-for-early-career-technologists.

[33] Kristen Vaughan, Britaini Carroll, and Michael R. Gavin, "Federal Workers Ready to Thrive in the Age of AI," Accenture, February 15, 2019, https://www.accenture.com/us-en/insights/us-federal-government/ready-thrive-ai.

training to use a feature of the system.[34] Not only does this undermine agencies trying to carry out their missions, but individuals using AI systems without appropriate training can also create or exacerbate threats to privacy, civil liberties, and even safety. Supporting AI education and research in the university environment can help address some of the government's talent problems and help fill the talent pipeline for the public sector.

This kind of government-academic collaboration in scientific and technological areas can also fuel innovation. For example, after World War II, the U.S. Department of Veteran Affairs (VA) collaborated with academic institutions (specifically medical centers) to meet the increasing medical needs of returning veterans.[35] The collaboration between academic medicine and the VA helped to revolutionize VA healthcare and spurred innovation in healthcare at many levels, including, for instance, the invention of pacemakers and CAT scan prototypes.[36]

<p align="center">* * * *</p>

As lead authors, we proudly submit this response on behalf of our colleagues and the Stanford Institute for Human-Centered Artificial Intelligence (HAI).

Daniel E. Ho, J.D., Ph.D.
William Benjamin Scott and Luna M. Scott Professor of Law, Stanford University; Faculty Associate Director, Stanford Institute for Human-Centered Artificial Intelligence (HAI)

Jennifer King, Ph.D.
Privacy and Data Policy Fellow, Stanford Institute for Human-Centered Artificial Intelligence (HAI)

Russell C. Wald
Director of Policy, Stanford Institute for Human-Centered Artificial Intelligence (HAI)

Daniel Zhang
Policy Research Manager, Stanford Institute for Human-Centered Artificial Intelligence (HAI)

---

[34] Adam Hoffman et al., "CBP and TSA Are Taking Steps to Implement Programs, But CBP Should Address Privacy and System Performance Issues," U.S. Government Accountability Office, September 2020, https://www.gao.gov/assets/gao-20-568.pdf.
[35] "75th Anniversary of VA's Academic Mission," U.S. Department of Veterans Affairs, January 12, 2021, https://www.va.gov/OAA/75th_anniversary.asp.
[36] Rob Marek et al., "Actions Needed to Help Better Identify Agency Inventions," U.S. Government Accountability Office, April 2018, https://www.gao.gov/pdf/product/691501.