

PUBLIC SUBMISSION

Received: May 29, 2025 Tracking No. mba-6j8k-7pd0 Comments Due: May 28, 2025 Submission Type: API
--

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0326
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Association for Health Learning and Inference

General Comment

On behalf of the Association for Health Learning and Inference (AHLI), we are submitting a formal response to the NSF's Request for Information on the 2025 National Artificial Intelligence Research and Development Strategic Plan. Our comment highlights the importance of targeted investment in AI for health research, including research funding, data and computing infrastructure, education, and community-building. Please see the attached PDF for our full response.

Attachments

AHLI Response to OSTP AI R and D RFI May 2025

AHLI Response to the Request for Information

on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

Submitted to:

Office of Science and Technology Policy (OSTP)
National Artificial Intelligence Initiative Office

Submitted by:

Association for Health Learning and Inference (AHLI)
www.ahli.cc

Date:

May 29, 2025

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

Executive Summary

As the United States sets priorities for artificial intelligence (AI) research and development, health should be a leading focus. Health impacts the daily lives of Americans profoundly, and AI in this field has the potential to accelerate diagnoses,¹ improve population health,² and identify medical errors.³ These advances are especially important in light of rising healthcare costs⁴ and projected workforce shortages.⁵

Health AI R&D differs fundamentally from commercial AI development, as it is deeply interdisciplinary, and must meet rigorous clinical and regulatory standards⁶— factors that may extend development timelines and delay return on investment. Without targeted public investment, robust health AI will remain inaccessible to Americans. More importantly, high-impact innovations in health AI risk harming patients⁷ by compromising patient safety;⁸ documented harms include hallucinations,⁹ reinforcing race-based medicine,¹⁰ and inconsistent demographic performance.¹¹

Strategic investments in health AI will not only advance national goals of economic competitiveness, public well-being, and global technological leadership, but also position the United States at the forefront of responsible AI innovation. Below we list four specific recommendations to improve AI in health: 1) Investing in AI and Health research broadly, 2) Targeting data and computing resources needed, 3) Improving training and education programs, and 4) Supporting Health/AI thought leadership.

¹ [Yala et al. \(2019\)](#) developed a deep learning model for breast cancer screening that outperformed the current clinical standard in predicting future cancer risk.

² [Obermeyer et al. \(2019\)](#) found evidence of racial bias in a widely used clinical algorithm, leading to eligible Black patients being overlooked for additional care.

³ [Seyyed-Kalantari et al. \(2021\)](#) demonstrated that AI systems may contribute to underdiagnosis in underserved populations.

⁴ A [2022 Bloomberg report](#) states that U.S. healthcare spending surpassed \$4.3 trillion, despite underwhelming health outcomes compared to peer nations.

⁵ [Kirch and Petelle \(2017\)](#) cite demographic shifts and rising demand as key contributors to physician shortages.

⁶ [Higgins and Madai \(2020\)](#) outline a three-phase framework that highlights the clinical, regulatory, and technical complexity contributing to longer timelines in health AI R&D.

⁷ A [2023 case](#) involving the Tessa chatbot, used by the National Eating Disorders Association, showed how poorly designed health AI can offer harmful advice to vulnerable users, underscoring the risks of deploying unvalidated systems in clinical or wellness settings.

⁸ [Chen et al. \(2021\)](#) highlight key ethical challenges in health AI, including bias, accountability, transparency, and the risk of patient harm due to misaligned algorithmic objectives.

⁹ *Hallucination* in AI refers to when a system makes up information that sounds believable but is actually false or inaccurate ([Koenecke et al., 2024](#)).

¹⁰ [Omiye et al. \(2023\)](#) argue that large language models can perpetuate race-based medicine, leading to biased clinical recommendations and contributing to harmful, inequitable care.

¹¹ [Zack et al. \(2023\)](#) found that GPT-4 consistently generated clinical content and reinforced demographic stereotypes, including biased diagnoses and treatment plans based on race, ethnicity, and gender.

Specific Recommendations

AHLI offer four key recommendations to guide the National AI R&D Strategic Plan:

1. Invest in World-Class AI Research in Health:

AI for health is a high-impact research domain with the potential to advance scientific discovery. Private-sector investment in AI is often directed toward high-return commercial applications such as consumer personalization, advertising, and automation.¹² In contrast, AI for public health, medical research, and clinical equity often lacks immediate commercial incentives yet holds profound societal value. To solidify U.S. leadership and accelerate innovation in research, we recommend:

- New and sustained R&D funding for interdisciplinary, health-centered AI research, including those focused on clinical goals, on general improvement of technical AI methods to be used in health, and deployment-focused benchmarks and evaluations. Funding opportunities should prioritize public-interest R&D areas that are less likely to be addressed by industry, including underserved populations and rare diseases.
- Targeted funding for early career researchers and practitioners to enter the interdisciplinary AI and health space, funding new work in this rapidly evolving ecosystem.

2. Target Data and Compute Resources for Health/AI:

Progress in AI for health is constrained by limited access to high-quality, representative datasets and affordable computing infrastructure. Projects like the NIH-supported All Of Us¹³ have begun to address some of these gaps, but to date are limited and require expansion. Resources invested in by other countries have spurred massive innovation; the UKBiobank has over 38,000 registered users (84% outside of the UK) and has enabled around 10,000 published studies with more than 2.5 million citations.¹⁴ To democratize innovation, the federal government should:

- Elevate health AI as a priority area in national data availability, with ethically-sourced, de-identified data as a focus. Such data can be used to create AI benchmarks to evaluate progress in key AI and health target areas, evaluating where state-of-the-art performance in clinical subspecialties lies.
- Create a subsidized national high performance computing infrastructure accessible to researchers focused on high-impact health use cases. This cluster could be linked to the national de-identified health datasets for broader evaluation.

¹² [Stanford HAI \(2025\)](#) provides an overview of global AI trends and the need for responsible oversight.

¹³ The NIH's [All of Us Research Program](#) is a national initiative aiming to build one of the most diverse health databases in history to support precision medicine and equitable health research across populations.

¹⁴ [UK Biobank](#) has become a global data resource, demonstrating the impact of open, large-scale health datasets on scientific output and international collaboration.

- Fund the development of open, interoperable health AI infrastructure in all health-adjacent sectors, including electronic health records systems, patient-facing chatbots, and incident reporting systems.

3. Improve Training and Education Programs in Health/AI:

A future-proof AI workforce must be interdisciplinary, inclusive, and grounded in ethical and equitable practice. While there exist many training programs and summer schools for other AI specialties¹⁵, to date there has been a notable lack of such opportunities for AI and Health.¹⁶ This lack of opportunity reduces the overall pipeline of qualified individuals to support innovative health/AI research. The U.S. government should:

- Fund health AI summer schools and year-round training initiatives that integrate technical, clinical, and policy perspectives to prepare the next generation of leaders.
- Support career development programs in health AI, especially through partnerships with academic institutions, health systems, and community-based organizations.
- Invest in K-12 education programs that build early awareness and foundational skills, creating an AI-ready workforce.

4. Support Health AI Thought Leadership:

Community building is essential for field formation in the interdisciplinary areas that AI and health span, especially with the rapid growth of our field in the past five years. Computer science conferences such as NeurIPS and ICML have increased the number of health-related submissions, leading to the creation of AI and health-focused events.¹⁷ In the past four years, the AHLI community grew to 900+ members, and submission volumes increased by 78%.¹⁸ Several AI-focused health journals have also been launched, including the [New England Journal of Medicine \(NEJM\) AI](#), [Lancet Digital Health](#), [Nature Digital Medicine](#) and [Journal of the American Medical Association \(JAMA\) AI](#).

We recommend:

- Federal funding for conferences to establish and elevate U.S. leadership in health AI research on the global stage, and shape the health AI research agenda.
- Support for informal, smaller meetings focused on themes like entrepreneurship in health AI, to accelerate the translation of research into real-world solutions, stimulate

¹⁵ [Climate Change AI Summer School](#)

¹⁶ The National Academy of Medicine noted a shortage of targeted training initiatives to equip the health workforce with AI competencies in its [2019 report](#).

¹⁷ Originally launched as a NeurIPS workshop, the Machine Learning for Health Symposium became a standalone event in 2021 in response to the growing community. Source: <https://ahli.cc/ml4h/symposium-overview/>

¹⁸ Internal metrics shared by the [Association for Health Learning and Inference](#), as of May 2025. Detailed data available upon request.

economic activity, and contribute to long-term growth in the U.S. health technology sector.

- Directly funding organizations that provide necessary services and resources to mobilize communities and drive transformative innovation in health AI. This includes fostering interdisciplinary dialogue and public engagement in health AI, including those leading to develop clear, actionable standards for safe and ethical deployment of AI in clinical and public health.

About AHLI

The Association for Health Learning and Inference (AHLI) is a non-profit organization committed to building an inclusive, responsible, and forward-looking health AI ecosystem. We convene researchers, practitioners, and policymakers to shape the future where AI meaningfully improves health outcomes and is leveraged for health's greatest good. AHLI's work is centered around four key priority areas: research, data, education and thought leadership.

Since its establishment in 2021, AHLI has:

- Built a diverse and distinguished Board of Directors composed of leading faculty and researchers from top institutions including Stanford University, MIT, Harvard Medical School, University of Chicago, Duke University, Emory University, and UC Berkeley & UCSF.
- Hosted seven scientific conferences, drawing hundreds of researchers globally and supporting 284 peer-reviewed publications.
- Launched a thought leadership retreat designed to catalyze collaboration, curate research tools and resources, develop educational content, and shape the direction of the field.

References

1. Yala, A., Lehman, C., Schuster, T., Portnoi, T., & Barzilay, R. (2019). A deep learning mammography-based model for improved breast cancer risk prediction. *Radiology*, 292(1), 60–66. [doi:10.1148/radiol.2019182716](https://doi.org/10.1148/radiol.2019182716)
2. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. [doi:10.1126/science.aax2342](https://doi.org/10.1126/science.aax2342)
3. Seyyed-Kalantari, L., Zhang, H., McDermott, M. B. A., Chen, I. Y., & Ghassemi, M. (2021). Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine*, 27, 2176–2182. [doi:10.1038/s41591-021-01595-0](https://doi.org/10.1038/s41591-021-01595-0)
4. Tozzi, J. (2022, December 14). US health spending edges past \$4.3 trillion as Covid drags on. *Bloomberg News*. <https://www.bloomberg.com/news/articles/2022-12-14/us-health-spending-edges-past-4-3-trillion-as-covid-drags-on>
5. Kirch, D. G., & Petelle, K. (2017). Addressing the physician shortage: The peril of ignoring demography. *JAMA*, 317(19), 1947–1948. [doi:10.1001/jama.2017.2714](https://doi.org/10.1001/jama.2017.2714)
6. Higgins, D., & Madai, V. I. (2020). *From bit to bedside: A practical framework for artificial intelligence product development in healthcare*. *Advanced Intelligent Systems*, 2(10), 2000052. [doi:10.1002/aisy.202000052](https://doi.org/10.1002/aisy.202000052)
7. Heilweil, R. (2023, June 5). *A mental health chatbot was taken offline after it gave harmful advice*. *Wired*. <https://www.wired.com/story/tesa-chatbot-suspended/>
8. Chen, I. Y., Joshi, S., Ghassemi, M., & Wallach, H. (2021). Ethical machine learning in health care. *Annual Review of Biomedical Data Science*, 4, 123–144. [doi:10.1146/annurev-biodatasci-092820-114757](https://doi.org/10.1146/annurev-biodatasci-092820-114757)
9. Koenecke, A., Choi, A. S. G., Mei, K. X., Schellmann, H., & Sloane, M. (2024). Careless whisper: Speech-to-text hallucination harms. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. [insert page numbers if available]). ACM. [doi:10.1145/3630106.3658996](https://doi.org/10.1145/3630106.3658996)
10. Omiye, J. A., Lester, J. C., Spichak, S., Rotemberg, V., & Daneshjou, R. (2023). Large language models propagate race-based medicine. *NPI Digital Medicine*, 6, 195. [doi:10.1038/s41746-023-00970-z](https://doi.org/10.1038/s41746-023-00970-z)
11. Zack, T., He, B., Mohan, A., & Rajkomar, A. (2023). Coding inequity: Assessing GPT-4’s potential for perpetuating racial and gender biases in healthcare. *bioRxiv*. [doi:10.1101/2023.07.13.23292577](https://doi.org/10.1101/2023.07.13.23292577)
12. Stanford Institute for Human-Centered Artificial Intelligence. (2025). *Artificial Intelligence Index Report 2025*. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
13. National Institutes of Health. (n.d.). *All of Us Research Program*. <https://allofus.nih.gov/>

14. Allen, N. E., Lacey, B., Lawlor, D. A., Pell, J. P., Gallacher, J., Smeeth, L., ... & Collins, R. (2024). Prospective study design and data analysis in UK Biobank. *Science Translational Medicine*, 16(729), eadf4428. [doi:10.1126/scitranslmed.adf4428](https://doi.org/10.1126/scitranslmed.adf4428)
15. Climate Change AI. (2024). *Climate Change AI Summer School 2024*.
https://www.climatechange.ai/events/summer_school2024
16. Matheny, M., Thadaneysrani, S., Ahmed, M., & Whicher, D. (Eds.). (2022). *Artificial intelligence in health care: The hope, the hype, the promise, the peril*. National Academy of Medicine. [doi:10.17226/27111](https://doi.org/10.17226/27111)
17. Association for Health Learning and Inference. (n.d.). *ML4H Symposium overview*.
<https://ahli.cc/ml4h/symposium-overview/>