July 15, 2025

Dr. Jay Bhattacharya
Director
National Institutes of Health

BY ELECTRONIC SUBMISSION

Re: Request for Information (RFI): Inviting Comments on the NIH Artificial Intelligence (AI) Strategy, notice number NOT-OD-25-117

Dear Dr. Bhattacharya,

This letter is submitted on behalf of Thinqpoint LLC in response to the request for information (the "RFI") on the National Institutes of Health ("NIH") Office of The Director Notice Inviting Comments on the NIH Artificial Intelligence (AI) Strategy. Thinqpoint is a data and analytics consulting organization that provides AI-enabled products and services to safety net organizations. We are the developers and authors of the Public Health Atlas, a public health data dissemination platform powered by generative AI. We deeply appreciate the NIH's commitment to developing a national NIH AI Strategy that prioritizes trust and transparency, and envisions a future where autonomous AI agents can assist in the research, continuous learning, and implementation of credible, sound knowledge.

Our comments were developed using a healthcare care data and analytics lens, but we believe they are broadly applicable to all AI use cases in health and human services.

We applaud the NIH's desire to engage with the community as it develops a short and long-term strategy to transition to AI agents and look forward to continuing to engage with the NIH and other healthcare stakeholders on topics related to artificial intelligence and its uses in a community health setting.

Sincerely,

Bogdan Rau, MPH
Founder & CEO
Thinqpoint

Background

Community Health Centers (CHCs) and Critical Access Hospitals (CAHs) provide crucial health and human services to predominantly underserved and vulnerable populations across the United States. In 2023, over 31 million patients received primary care, behavioral health, oral health, or other specialty or enabling services at, or with the support of a CHC[1]. Despite their substantial (and growing) scope of service, CHCs continue to experience significant barriers in funding, technical infrastructure, and technical skills to leverage population-level data for service delivery, or to develop tailored programs to address evolving community needs[2]. For CAHs, insufficient funding and operating margin have been important contributors to the rate of closures continuously outpacing that of openings since 2017[3], creating substantial and lasting deficiencies in access and outcomes for rural communities[4]. Financial leaders from CAHs have often quoted the need for more credible market data to support Community Health Needs Assessment (CHNA) implementation[5].

While generative AI and large language models (LLMs) show early utility in data analysis and interpretation, only 0.2% of literature focused on healthcare AI describes meaningful community involvement in the development, validation, or implementation of AI applications in healthcare[6]. The exclusion of vulnerable stakeholders in the early stages of machine learning (ML) design and development has been shown to produce significant impacts on fairness of outcomes[7]. In the absence of lived experiences to inform the problem formulation step of solution design, AI/ML capabilities, and by extension, applications leveraging those models, fail to account for peripheral use cases that have been shown to lead to uncontrolled and unintended consequences, including denial of services[8].

We believe agentic AI will enable safety net organizations to unlock unprecedented value from population and patient-level data. Unlike prior technology adoption cycles that required a multitude of technical expertise, the commoditization of generative AI tools has fundamentally lowered the barrier to entry. Although resource-constrained provider organizations will benefit from AI advances at a much faster pace compared to other technologies, effective adoption of AI will necessitate a national strategy that supports and enables:

1. **Participatory Design:** product design and development approaches that incorporate community participation at the onset of problem formulation, and continuous engagement throughout the iterative development lifecycle.
2. **Modern Data and AI Governance:** focused investments to modernize safety net data and AI governance capabilities aimed at unlocking existing and future AI capabilities while mitigating risk.

[1] Health Resources and Services Administration. 2023 National Report: Health Center Program UDS Data. U.S. Department of Health and Human Services; 2023. Accessed May 10, 2025.
[2] Martin et al. (2025). How Community-Based Organizations View and Use Data: Practices, Challenges, and Opportunities. *Journal of Health Care for the Poor and Underserved*, 36(2), 458–479.
[3] KFF. *Key facts about hospitals.* Rural Hospitals Openings and Closings. Retrieved June 3, 2025, from https://www.kff.org/key-facts-about-hospitals/?entry=rural-hospitals-rural-openings-and-closings.
[4] Khushalani et al. (2022). Impact of Rural Hospital Closures on Hospitalizations and Associated Outcomes for Ambulatory and Emergency Care Sensitive Conditions. *The Journal of Rural Health*, 39(1), 79–87.
[5] Boothe, A. L. (2020). *The Cost of the Community Health Needs Assessment: A Qualitative Analysis* (dissertation).
[6] Loftus et al. (2024). Community-Engaged Artificial Intelligence Research: A Scoping Review. *PLOS Digital Health*, 3(8).
[7] Obermeyer et al. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, 366(6464), 447–453.
[8] Eubanks, V. (2019). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* Picador, St. Martin's Press.

3. **Trust and Transparency:** policies and/or guidance that expand beyond the concept of observability, and include local adaptation, community representation, and procedural transparency.

Participatory Design

Community-engaged approaches to the design and implementation of interventions and health services have shown effectiveness in a range of settings, from HIV prevention in youth experiencing homelessness[9], to the development of a community pharmacy model for enhancing screening and management of atrial fibrillation[10]. Despite growing interest in community engagement among AI researchers and practitioners, "community participation" within the context of AI/ML product development is not well defined, and consequently, has had limited application.

We recommend that NIH develop guidelines and recommendations for meaningful community participation in AI product design and development. Rather than prescriptive mandates that may prove counterproductive, these guidelines would provide AI researchers, practitioners, and community stakeholders, with a foundational framework to articulate the goals of participation, transparently define the scope and desired effects of engagement, and identify the right forms of community involvement. In association with these recommendations, NIH could develop criteria to evaluate AI products based on their level of community engagement (consult → include → collaborate → own)[11], contributing to enhanced trust and transparency, while establishing a common language describing "community participation" in an industry-agnostic way.

Modern Data and AI Governance

While safety net organizations stand to gain significantly from the benefits of AI, chronic and systemic underfunding and resource limitations leave CHCs/CAHs disproportionately exposed to AI-related risks, such as bias, discrimination, or adverse effects due to ineffective/inappropriate use of AI. These risks materialize in organizations that: a) lack the capacity or workforce necessary to properly evaluate AI systems; and b) have not implemented data and AI governance structures that establish rights and accountabilities for effective ownership of data and AI platforms. We anticipate that the implementation of H.R. 1 – One Big Beautiful Bill Act (OBBB) may exacerbate gaps in technical infrastructure and workforce capacity, as safety net providers will be forced to further reduce investments in infrastructure in favor of sustaining frontline activities.

NIH should establish a set of shared AI governance strategic objectives aimed at ensuring equitable AI access across communities and healthcare settings. In addition, NIH should consider establishing a national, safety net-focused AI investment strategy aimed at strengthening data and AI governance capabilities within CHCs, CAHs, as well as other community-based safety net providers. A laser-focused, national approach to investing in data and AI governance in the safety net could include incentivizing:

[9] Wilder et al. (2021). Clinical Trial of an AI-Augmented Intervention for HIV Prevention in Youth Experiencing Homelessness. *Proceedings of the AAAI Conference on Artificial Intelligence,* 35(17), 14948-14956.
[10] Sabater-Hernandez et al. (2018). A Stakeholder Co-Design Approach for Developing a Community Pharmacy Service to Enhance Screening and Management of Atrial Fibrillation. *BMC Health Services Research*, 18(1).
[11] Delgado et al. (2023). The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 37, 1-23.

1. **AI-ready metadata**: prioritized metadata management for AI-use cases, data annotation and data dictionaries for AI agents, development of data context and lineage for AI model ingestion.
2. **Stewardship:** development of subject matter expertise, and SME-driven curation of data and AI capabilities, equipped with objective, measurable guidelines to assess data that is "fit for AI," as well as enterprise expectations for transparency, observability, and oversight.
3. **Practical data and AI governance:** guidelines, guardrails, and standards to manage bias, fairness, privacy, and regulatory adherence, helping define organizational posture and expectations of vendor-driven AI solutions.

Trust and Transparency

While no single definition of "trustworthy AI" exists to date, most frameworks converge on core technical principles of explainability (understanding AI decision-making), reproducibility (consistent replication of results), and fairness (equitable treatment and bias mitigation). The emerging field of AI evaluation and observability continues to mature, providing methodologies that address inference-level performance metrics – accuracy, precision, and reliability – as well as system-level anomaly detection and drift monitoring. However, technical excellence alone is insufficient for healthcare contexts, and trust in AI extends beyond algorithmic performance to encompass issues of procedural fairness, community engagement, and cultural appropriateness. Given the direct and indirect impact AI can have on patients and the care delivery process, a comprehensive trust and transparency framework must address both technical and human-centered governance elements that build confidence among healthcare providers, and seed trust in the patients and communities they serve.

We believe NIH should establish guidance on the essential components necessary to characterize trustworthy AI systems across the landscape of healthcare organizations and use-cases. This framework should establish requirements for clear evidence and help define minimum expectations from key stakeholder groups. For example, this framework could:

- Define or inform requirements for standardized assessment criteria and measures of performance that **health systems** can use to systematically evaluate AI technologies and their application.
- Require transparent documentation and communication of limitations, appropriate use-cases, and observed unintended or unexpected "side effects" to inform **health and human services providers** within the context of service delivery.
- Incentivize and/or require transparent communication about the role AI plays in the service delivery process, enabling **participants, patients, and community groups** to understand how AI informs decisions "about them," and how to seek recourse.
- Clarify requirements for documentation, validation, bias testing, and ongoing monitoring that must be met by **AI vendors** for healthcare deployment.

In addition to technical assessment criteria (performance metrics, measures of reliability and bias reduction), we recommend that a comprehensive trust and transparency framework also addresses:

1. **Risk disclosure and communication:** similar to pharmaceutical warning labels, proactively identify and communicate failure modes, as well as associated risks to providers, patients, or communities.

2. **Procedural transparency:** promote visibility into process, methods, and decision-making workflows, but look beyond explainability of individual AI decisions to also address oversight and accountability.
3. **Community engagement**: ensure meaningful participation of community groups and the inclusion of lived experience in the design and development of AI capabilities.
4. **Local adaptation and cultural responsiveness:** recognize and communicate the level of customization applied to AI capabilities to account for differences in context, culture, and population characteristics.

Several frameworks have emerged that independently address portions of a broader, more comprehensive approach we are recommending. The CLeAR (**C**omparable, **Le**gible, **A**ctionable, **R**obust) documentation framework[12] addresses high level concerns related to well-defined criteria, formats, and outputs that ensure comparability, accessibility, and utility of AI documentation for a wide array of audiences. More explicitly, the Coalition for Health AI (CHAI) Applied Model Card[13], designed specifically for healthcare AI applications, addresses measures of performance, bias, and various ethical factors.

Conclusion

We appreciate the opportunity to contribute our perspective through this RFI, and to provide context from a viewpoint beyond that of direct patient care. As healthcare AI applications continue to expand and evolve, we understand the significant challenge posed to federal and state agencies, and believe that a transparent, predictable, and inclusive engagement process will result in frameworks that accelerate innovation while safeguarding the well-being of healthcare workers, patients, and communities across the nation.

---

[12] Chmielinski et al. (2024). *The CLeAR Documentation Framework for AI Transparency*. https://shorensteincenter.org/wp-content/uploads/2024/05/CleAR_KChmielinski_FINAL.pdf.
[13] Giannikopoulos et al. (2024). *The CHAI Applied Model Card*. https://chai.org/workgroup/applied-model.