NATIONAL HEALTH COUNCIL

July 15, 2025

Jayanta Bhattacharya, MD, PhD
Director
National Institutes of Health
9000 Rockville Pike
Bethesda, MD 20892

RE: Request for Information (RFI): Inviting Comments on the NIH Artificial Intelligence (AI) Strategy [NOT-OD-25-117]

Submitted to ai-rfi@nih.gov

Dear Director Bhattacharya,

The National Health Council (NHC) welcomes the opportunity to comment on the National Institutes of Health (NIH) Request for Information (RFI) on the development of an NIH Artificial Intelligence (AI) Strategic Plan.

Created by and for patient organizations over 100 years ago, the NHC brings diverse organizations together to forge consensus and drive patient-centered health policy. We promote increased access to affordable, high-value, equitable, and sustainable health care. Made up of more than 180 national health-related organizations and businesses, the NHC's core membership includes the nation's leading patient organizations. Other members include health-related associations and nonprofit organizations including the provider, research, and family caregiver communities; and businesses and organizations representing biopharmaceuticals, devices, diagnostics, generics, and payers.

As AI technologies become increasingly embedded across the health care continuum, the NIH has a critical opportunity to lead in shaping a future where AI serves as a trustworthy, equitable, and transformative force for public health and biomedical innovation. We urge the NIH to center its strategy on patients' perspectives and real-world outcomes while promoting innovation, equity, transparency, and accountability.

Strategic Architecture

The NIH's AI Strategic Plan should be grounded in a clear architectural framework that articulates the foundational principles, governance structures, and developmental pathways necessary to responsibly advance AI capabilities across the biomedical and public health domains.

Foundational Themes

The NHC has developed foundational principles for the responsible use of AI from the patient perspective and encourages the NIH to use incorporate these themes as a guiding framework throughout its AI strategy.[1]

[1] National Health Council. *Principles on Health AI.* March 25, 2025. https://acrobat.adobe.com/id/urn:aaid:sc:va6c2:1e4c2cb5-5b01-4e65-b386-8eeefea9b6f5.

- **Patient-Centeredness:** AI initiatives must prioritize patient health benefit, incorporate early and sustained patient and family caregiver input, and ensure patients are partners in the research lifecycle.[2]
- **Equity and Inclusion:** AI tools should be developed with an explicit focus on identifying and reducing health disparities, including through representative datasets and equity-focused benchmarks.[3]
- **Trust and Transparency:** Clear disclosure of AI functions, data sources, limitations, and model updates must be standard across NIH-funded AI research.[4]
- **Interoperability and Reproducibility:** Data readiness and FAIR (Findable, Accessible, Interoperable, Reusable) principles should guide AI data infrastructure and reporting.[5]
- **Workforce and Community Readiness:** The NIH should fund education and training initiatives that include patients, family caregivers, community advocates, and non-technical stakeholders, in addition to clinicians and data scientists.[6]

Transition Pathways

The NIH should define staged and ethically grounded transition pathways to guide the development and deployment of AI technologies. This should include the establishment of a multi-stakeholder governance group to define milestones for the validation of semi-autonomous agents in low-risk environments; mandatory implementation of human-in-the-loop safeguards for high-risk or autonomous applications; and the development of objective criteria and audit mechanisms for self-documenting systems to ensure transparency, reproducibility, and accountability.

Research and Innovation Actions

The NIH's AI strategy should articulate a focused research and innovation agenda that leverages artificial intelligence to advance biomedical discovery, improve clinical outcomes, and support public health infrastructure. This agenda must be grounded in high-impact use cases and coupled with methodological standards that ensure the reproducibility, transparency, and equity of AI tools developed with federal support.[7]

High-Impact Use Cases

The NIH should prioritize funding for AI applications that target conditions with significant prevalence, morbidity, mortality, or cost burden, especially in communities experiencing entrenched health disparities. In addition to common chronic diseases, specific emphasis should be placed on AI use cases that improve diagnostic accuracy and care coordination for rare diseases where

---

[2] NHC, *Principles on Health AI*. See Principle 1: Patient Health Benefit; Principle 2: Patient-Centered and Informed; Principle 3: Patient Education and Empowerment; and Principle 13: Ethical Considerations.

[3] NHC, *Principles on Health AI*. See Principle 4: Equity and Accessibility; Principle 7: Bias Mitigation; and Principle 6: Comprehensive Transparency.

[4] NHC, *Principles on Health AI*. See Principle 5: Trustworthiness; Principle 6: Comprehensive Transparency; Principle 8: Privacy and Security; and Principle 12: Accountability.

[5] NHC, *Principles on Health AI*. See Principle 7: Bias Mitigation; Principle 10: Continuous Improvement; and Principle 11: Safety and Effectiveness.

[6] NHC, *Principles on Health AI*. See Principle 3: Patient Education and Empowerment; Principle 9: Human Collaboration; and Principle 13: Ethical Considerations.

[7] NHC, *Principles on Health AI*. See Principle 4: Equity and Accessibility; Principle 6: Comprehensive Transparency; and Principle 7: Bias Mitigation.

traditional research models face recruitment and data limitations.[8] Investment should also support the development of predictive and preventive care models—including those that integrate social determinants of health—and AI systems that optimize clinical workflows, including screening, referral, and follow-up processes. Furthermore, the NIH should promote AI-enabled innovation in clinical trial infrastructure, including tools for eligibility screening, site selection, participant engagement, real-time data monitoring, and adaptive trial design. These applications have the potential to reduce costs, enhance generalizability, and improve enrollment equity in biomedical research.

Reproducibility and Benchmarking

To preserve scientific integrity and facilitate public trust, the NIH must establish minimum standards for the reproducibility and independent evaluation of AI systems across its funded portfolio. All AI models supported by the NIH should undergo external validation using diverse, demographically representative datasets that reflect real-world populations. Results should be stratified by key sociodemographic variables such as race, ethnicity, age, gender, language, disability status, family caregiver status, and geography to assess differential performance and prevent exacerbation of health inequities.[9] The NIH should also support the development of standardized performance benchmarks and require the publication of AI model metadata, source data descriptions, and validation results in accessible, centralized registries. An example of this is the Coalition on Health AI (CHAI) Model Card Framework.[10] Finally, the NIH should encourage grantees to adopt open model-sharing practices and contribute to public repositories that enable re-analysis, reproducibility checks, and cumulative knowledge-building across research programs.

Intramural–Extramural Synergy

The success of the NIH AI Strategic Plan will depend in part on the agency's ability to facilitate collaboration and knowledge exchange between its intramural research programs and the broader extramural research community. Establishing a cohesive, interoperable research ecosystem will promote consistency, reduce redundancy, and maximize the impact of federally funded AI initiatives across the biomedical enterprise. The NIH should prioritize mechanisms that enable mutual learning, joint stewardship, and sustainable access to data and tools developed under its auspices.

Bidirectional Innovation

The NIH should take concrete steps to ensure that AI-related resources developed through intramural programs—such as software tools, annotated datasets, and pretrained models—are made publicly available whenever feasible, ideally in accordance with FAIR (Findable, Accessible, Interoperable, Reusable) data principles. These resources should be published through open-access platforms, accompanied by clear documentation and appropriate licensing frameworks to support broad usability. Simultaneously, extramural grantees should be incentivized to contribute derivative tools, methodological enhancements, and real-world performance insights back into NIH-managed repositories. Such bidirectional innovation will strengthen the national AI infrastructure, support cumulative progress, and reduce barriers to entry for under-resourced institutions and early-stage investigators.

---

[8] NHC, *Principles on Health AI*. See Principle 1: Patient Health Benefit and Principle 2: Patient-Centered and Informed.

[9] NHC, *Principles on Health AI*. See Principle 4: Equity and Accessibility; Principle 6: Comprehensive Transparency; and Principle 7: Bias Mitigation.

[10] "Applied Model Workgroup," *Center for Health AI (CHAI)*, accessed July 15, 2025, https://www.chai.org/workgroup/applied-model.

Governance Models

The NIH should establish robust, transparent governance frameworks to oversee the development, maintenance, and dissemination of AI tools and assets produced through both intramural and extramural programs. These governance models must explicitly prioritize open science, protect against undue privatization of publicly funded innovations, and uphold the rights and interests of patients whose data underpin AI development.[11] Multi-stakeholder governance bodies—comprising researchers, ethicists, technical experts, and patient and caregiver representatives—should be convened to guide policy on data sharing, intellectual property, model retraining, decommissioning, and equity risk mitigation. Moreover, the NIH should build in formal feedback loops to ensure that communities impacted by AI tools have opportunities to inform ongoing oversight and improvements.

Operational Excellence

As the world's foremost biomedical research agency, the NIH should set a high standard for the ethical and effective use of AI within its own operations. Our NHC principles call for AI to be developed and used in accordance with ethical principles, including respect for human dignity, autonomy, and privacy.[12] Embedding AI into internal workflows offers significant opportunities to improve efficiency, transparency, and user experience for applicants, reviewers, grantees, and other stakeholders. The NIH's internal deployment of AI should be guided by the same principles it promotes externally—namely transparency, accountability, fairness, and human oversight. Lessons learned from these operational pilots can inform best practices for the broader research and health care communities.

Customer-Facing AI Improvements

The NIH should prioritize the implementation of AI tools that streamline high-burden administrative processes and improve stakeholder-facing services. In particular, AI should be deployed to assist in preparing and processing grant applications, identifying and flagging potential conflicts of interest in peer review, and analyzing applicant feedback to inform continuous process improvement. The agency should also invest in AI-driven platforms that enhance engagement with stakeholders—especially patient and caregiver communities—by enabling more responsive, accessible, and personalized communication channels. All such tools should be developed with built-in human oversight, transparency safeguards, and usability testing that includes input from diverse end users.

Evaluation Metrics

To ensure responsible and effective implementation, all operational AI initiatives should be subject to rigorous and multidimensional evaluation. The NIH should define and track performance metrics that reflect technical performance (e.g., accuracy, explainability), operational gains (e.g., reduction in processing time or administrative burden), and user experience (e.g., stakeholder satisfaction, accessibility). These metrics should be disaggregated where possible to identify differential impacts across stakeholder groups. CHAI has developed a responsible AI checklist that would be useful in creating appropriate metrics. In addition, the NIH should establish formal mechanisms for periodic review, error auditing, and public reporting to maintain accountability and foster trust in its use of AI.[13]

---

[11] NHC, *Principles on Health AI*. See Principle 2: Patient-Centered and Informed; Principle 8: Privacy and Security; and Principle 13: Ethical Considerations.

[12] NHC, *Principles on Health AI*. See Principle 8: Privacy and Security; Principle 12: Accountability; and Principle 13: Ethical Considerations.

[13] NHC, *Principles on Health AI*. See Principle 6: Comprehensive Transparency; Principle 10: Continuous Improvement; and Principle 12: Accountability.

Facilitating and Validating AI in Healthcare Delivery

Given the growing integration of AI into clinical workflows, the NIH plays a vital role in ensuring that such technologies are developed, validated, and implemented in a manner that protects patients and promotes equitable health outcomes. The NIH's strategy should support robust evaluation frameworks for clinical AI tools, advance regulatory science in collaboration with federal partners, and facilitate transparency and accountability through publicly accessible validation infrastructure. This is another area where the CHAI responsible AI checklist can be a guide.[14] The goal should be to ensure that AI tools deployed in healthcare settings are not only technically sound but also aligned with patient needs, clinical realities, and ethical standards.[15]

To this end, the NIH should invest in the development of shared testbeds—jointly operated with agencies such as the Food and Drug Administration (FDA), Department of Veterans Affairs (VA), and Centers for Medicare & Medicaid Services (CMS)—to enable controlled evaluation of clinical AI tools across a variety of settings and use cases. The NIH should also fund the creation of validated simulation environments capable of stress-testing models under atypical or edge-case scenarios, including those that may exacerbate disparities or result in harm to underserved populations. Furthermore, the NIH should support the development of open-access toolkits that facilitate third-party auditing, including bias detection tools, safety validation protocols, and model explainability assessments.

Critically, the NIH should adopt a policy that prohibits funding for clinical deployment of AI tools unless there is independent, peer-reviewed evidence of safety, effectiveness, and equity across diverse, real-world patient populations. This requirement should extend to any NIH-supported implementation studies, demonstration projects, or translational initiatives involving clinical AI. In doing so, the NIH will help to establish a robust standard for evidence generation that reinforces public trust and safeguards patient welfare.

Reproducibility & Trust

Reproducibility and public trust are foundational to the responsible advancement of AI in health care and biomedical research. To ensure that AI technologies supported by the NIH uphold scientific integrity and foster confidence among stakeholders, the NIH must establish and enforce standards that promote transparency, accountability, and inclusive engagement.[16] These standards should apply across the AI research lifecycle, from data acquisition and model development to deployment and ongoing oversight.

Standards and Auditability

The NIH should require all AI models developed with federal funding to include detailed, transparent documentation of model architecture, data sources, assumptions, intended uses, and known limitations. These requirements should be embedded in the NIH's funding solicitations and progress reporting mechanisms. In addition, grantees should be required to provide reproducibility protocols, including access to training and validation datasets, as well as version-controlled codebases when feasible. Audit trails documenting model evolution, performance shifts, and retraining triggers should be maintained in publicly accessible registries. To promote accessibility, a summary of these disclosures should be provided in language understandable to non-expert stakeholders, including patients and caregivers.

---

[14] "Responsible AI Checklists (RAIC)," *Center for Health AI (CHAI)*, accessed July 15, 2025, https://www.chai.org/workgroup/responsible-ai/responsible-ai-checklists-raic.

[15] NHC, *Principles on Health AI*. See Principle 1: Patient Health Benefit; Principle 11: Safety and Effectiveness; and Principle 13: Ethical Considerations.

[16] NHC, *Principles on Health AI*. See Principle 4: Equity and Accessibility; Principle 6: Comprehensive Transparency; and Principle 12: Accountability.

Trust Building

The NIH should promote trust-building by mandating early and ongoing involvement of patients, providers, and affected communities in the development, testing, and refinement of AI systems. Stakeholder input should inform key design choices, validation criteria, and deployment conditions.[17] In addition, the NIH should require plain-language disclosures about the presence and function of AI in any research or clinical tool supported by its funding.[18] These disclosures should clearly state how the AI system influences decisions, what limitations exist, and what recourse is available in cases of model error. By embedding these practices into its strategy, the NIH can help ensure that the deployment of AI in health care reflects ethical principles, centers patient agency, and maintains the public's trust in science.

Partnerships & Ecosystem Building

A robust and ethically grounded AI ecosystem in biomedical research and health care requires sustained collaboration across a diverse array of stakeholders. The NIH should take an active role in convening and supporting strategic partnerships that amplify the voices of patients and foster an inclusive, transparent, and accountable research environment. In particular, the NIH must prioritize the integration of patient organizations as co-equal partners in AI development, governance, and implementation.[19] These partnerships are essential to ensure that AI tools address real-world patient needs and to legitimize and democratize innovation in this emerging field.

Collaboration Modalities

The NIH should expand and formalize its collaborations with patient organizations, particularly those that represent underserved and high-need populations. These partnerships should span the research continuum—from early-stage design and use case selection to deployment, post-market surveillance, and long-term governance. The NIH should also strengthen cross-agency collaborations, including with the FDA, CMS, VA, and the HHS Office for Civil Rights, to align federal activities related to AI ethics, regulation, data stewardship, and civil rights protections. Public-private partnerships involving academic and industry actors must be subject to enforceable guardrails that protect data integrity, ensure transparency, and safeguard patient rights. The NIH should clearly articulate expectations regarding data access, intellectual property sharing, conflict of interest management, and dissemination of research findings.

Governance Approaches

The NIH should fund and support the creation of multi-stakeholder governance bodies that include meaningful representation from patient organizations, caregivers, and community-based groups. These bodies should have clearly defined responsibilities in the oversight of AI research, including ethical review, evaluation of health equity impacts, and review of consent processes. The NIH's governance frameworks must be aligned with open science principles, support national security and privacy protections, and incorporate community-driven standards for responsible data use. By ensuring that patient organizations have a formalized, ongoing role in these processes, the NIH will enhance accountability and help ensure that the benefits of AI innovation are broadly and equitably distributed.

---

[17] Ensuring Health AI Supports Human Flourishing," *Center for Health AI (CHAI)*, January 24, 2025, https://www.chai.org/blog/ensuring-health-ai-supports-human-flourishing.

[18] NHC, *Principles on Health AI*. See Principle 3: Patient Education and Empowerment; Principle 5: Trustworthiness; and Principle 6: Comprehensive Transparency.

[19] NHC, *Principles on Health AI*. See Principle 2: Patient Health Benefit and Principle 9: Human Collaboration.

Conclusion

Thank you again for the opportunity to respond to this RFI. Please do not hesitate to contact Kimberly Beer, Senior Vice President, Policy & External Affairs at kbeer@nhcouncil.org or Shion Chang, Senior Director, Policy & Regulatory Affairs at schang@nhcouncil.org, if you or your staff would like to discuss the NHC's input in greater detail.

Sincerely,

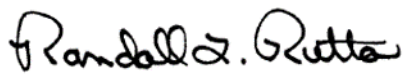

Randall L. Rutta
Chief Executive Officer