PUBLIC SUBMISSION

Received: May 25, 2025
Tracking No. mb3-q0kt-42b0
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-0108
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Health AI

General Comment

See attached document

Attachments

AI RFI

Health AI Public Submission to the 2025 National AI Research and Development Strategic Plan

(Docket ID No. NSF-2025-OGC-0001)

Executive Summary

Health AI proposes a two-track federal strategy focused on domains of overlooked urgency: (1) AI infrastructure for industrial stabilization — specifically supply chains, logistics, and smart manufacturing; and (2) biomedical signal decoding at the single-cell level, with applications across neurophysiology, metabolic health, and therapeutic design. Both tracks are actionable within a three-to-five-year window. Both require platforms that are secure, transparent, and non-weaponized. And both offer scalable benefits — not only for innovation, but for stability, equity, and resilience. We submit this proposal not as a conceptual roadmap, but as a call for implementation. Strategic public investment can ensure these tools emerge ethically, openly, and in service of public good — before the moment to do so closes.

About Health AI

Health AI is a U.S.-based research and strategy organization working at the intersection of bio-signal interpretation, public infrastructure, and ethical systems design. Our focus spans biomedical signal decoding, AI-integrated diagnostics, and open models of human-machine feedback. We operate with a commitment to public-good applications of AI: scalable, secure, and non-extractive by design. Our website: healthai.com

Introduction

In 1943, Warren McCulloch and Walter Pitts introduced a groundbreaking mathematical model of neural networks, proposing that the brain's functions could be simulated through logic-based computing elements. This theoretical framework laid the conceptual foundation for modern artificial intelligence. Just three years later, in 1946, the ENIAC computer—originally developed for military artillery and nuclear calculations—came online as the world's first general-purpose digital computer. Though born from warfare, ENIAC was quickly adapted for scientific and medical research, establishing the practical roots of computational medicine. These parallel advances—abstract neural computation and militarized digital hardware—converged into a shared trajectory where logic-based models of cognition could be implemented on increasingly powerful machines. This convergence catalyzed the evolution of AI, automation, and data-driven decision-making that now define Industry 4.0. The origins of today's intelligent systems are not found in consumer devices or tech startups, but in a post-war synthesis of biology, logic, and state-directed computational ambition. It is in homage to this deep historical interplay between military innovation, computational infrastructure, and biomedical advancement that our company, founded in 2019, was named Health AI.

Today, Health AI operates at the nexus of synthetic biology, strategic infrastructure, and American industrial production—three domains that are deeply specialized yet historically siloed. While we cannot disclose details of our current engagements—which span hospitals, military entities, and advanced manufacturing—we can identify two high-impact avenues for future U.S. AI investment: one in biomedical systems, the other in industrial automation. Each area offers transformative potential on its own; together, they represent a quantum leap in national capability, achievable within the three-to-five-year timeframe outlined in this RFI. In the sections that follow, we transcend narrow technical specialization to present clear, integrated recommendations that directly support the government's mission and reignite the spirit of integrated innovation first forged during the Second World War.

Proposal Organization

1) Industry: Challenge, Solution, Outcome
2) Biomedical: Challenge, Solution, Outcome

Industry: Challenge

Over the past decade, U.S. manufacturers in the automotive and heavy-equipment sectors have been squeezed by rising warranty liabilities, opaque demand forecasting, and increasingly stringent regulatory reporting. Warranty-claim adjustments climbed even as quality-assurance recoveries fell, driving parts inventories up by as much as fifty percent and tying up critical working capital. Regulators—from NHTSA to the EPA—now demand granular batch- and failure-rate data, forcing analysts to stitch together siloed ERP records, field reports, and telematics logs in error-prone spreadsheets. In the absence of continuous wear data for advanced composites and electric-vehicle components, premature failures go undetected until costly recalls. Compounding these operational burdens is deep uncertainty over legal liability for AI-driven maintenance and inspection systems: absent a clear framework, manufacturers hesitate to deploy autonomous diagnostics for fear of crippling litigation if algorithms err.

Industry: Solution

To overcome fragmentation and liability concerns, federal support should underwrite development of an on-premises intelligence fabric that harmonizes every data source—from legacy warranty systems and live IoT sensors on the factory floor to telematics feeds and environmental monitors—into a unified data lake. Modular AI copilots, developed through DOE- and DARPA-funded research initiatives and certified under a government-backed "AI Safety Certification" safe-harbor, would learn directly from expert analysts, propose testable hypotheses, uncover hidden correlations, and draft compliance summaries in natural language. A National AI Research Resource would democratize access to high-performance computing for small and mid-sized firms, while streamlined permitting by FERC and DOE would ensure the energy infrastructure needed to power these compute clusters. Immersive AR/VR interfaces would collapse decision cycles, enabling plant managers and field technicians to diagnose and remedy issues in real time.

Industry: Outcome

Phased over three to five years, this approach yields a monetizable, permissioned "data ocean" licensed to suppliers, regulators, and OEM partners. In Year 1, legacy reports convert to continuous data streams and safe-harbor liability guidelines are established. In Year 2, centralized data lakes support advanced analytics and pilot subscription models. By Years 3–5, a self-sustaining data marketplace—funded by subscription and pay-per-query revenues—scales nationally, funding its own compliance, R&D, and infrastructure. Regulatory provenance and audit trails become automated outputs of the intelligence fabric, transforming compliance from a cost center into a competitive advantage and securing U.S. industrial leadership in the AI-driven economy of the future.

Biomedical: Challenge

Until recently, probing the molecular choreography inside living cells—or across tissues—meant disrupting biology, losing spatial context, or relying on bulk measurements that flatten out critical cell-to-cell variation. Biomarkers in glycans, lipids, nucleic acids, and proteins have remained hidden behind the limits of conventional stains and genetic reporters. Simultaneously, FDA, NIH, HIPAA, and CLIA frameworks impose stringent provenance and audit-trail requirements on any in-vivo diagnostic or sampling method, fragmenting data workflows and stalling innovation in single-cell analysis. The net result is a bottleneck in drug discovery, precision diagnostics, and early-warning biosurveillance—despite the existence of powerful AI methods capable of decoding cellular signals.

Biomedical: Solution

We propose federal investment—via ARPA-H and NIH challenge prizes—to establish prototype centers that fuse bioorthogonal chemistry with AI-driven protein engineering. These centers would develop reagents that introduce orthogonal chemical handles (azides, cyclooctynes) into metabolites or biomolecules and ligate imaging or enrichment probes in situ, paired with bespoke molecular sensors—designer pores, receptors, and allosteric switches—that transduce binding or enzymatic events into electrical, optical, or conformational readouts at the single-cell level. On-chip AI pipelines, funded through SBIR/STTR and other transaction authority contracts, would process this high-fidelity data in real time. All systems would operate under a government-underwritten liability safe-harbor, assuring adopters that they can deploy advanced biosensing without exposing themselves to unmanageable legal risk.

Biomedical: Outcome

Over a three-to-five-year horizon, these integrated platforms yield a permissioned "single-cell data ocean" in which millions of compliant, timestamped, and context-tagged annotations—glycome, lipidome, proteome—are licensed to pharmaceutical R&D, academic consortia, and diagnostic labs. Phase 1 delivers robust biosensing pipelines and audit-trail frameworks; Phase 2 scales the data marketplace, funding continuous reagent improvement, AI model refinement, and on-premises deployments. By transforming siloed workflows into a self-funding ecosystem, the U.S. regains its leadership in AI-embedded biosensing, accelerating precision medicine and fortifying health security.

Integrated Barriers & Enablers

Across both domains, the greatest barriers are regulatory overhead, risk-averse funding, liability uncertainty, and looming energy constraints. We urge the Strategic Plan to harmonize agency requirements into a centralized audit-trail framework that accepts AI-generated compliance artifacts; to channel federal R&D dollars directly to agile startups through expanded SBIR/STTR, ARPA-style programs, and milestone-based Other Transaction Authority contracts; to underwrite AI liability for publicly funded projects via a national safe-harbor or insurance mechanism; and to modernize energy regulations—streamlining data-center permitting, incentivizing renewables and storage, and deploying AI-driven grid optimization—so that compute capacity scales sustainably with demand.

Conclusion

The infrastructure for applied AI in the United States is still under construction — but the window for shaping its direction is narrowing. The proposals outlined here focus on two domains that remain underrepresented in current federal AI planning: next-generation industrial platforms and biomedical signal interpretation. Both are technically feasible within a three-to-five-year timeframe, and both offer high public value with relatively low regulatory risk.

As with GPS, mRNA, and the internet, meaningful public benefit will depend not just on innovation, but on intentional architecture: systems that are auditable, secure, and non-weaponized from the outset. AI that serves supply chain resilience, neurophysiological health, and decentralized diagnostics must be protected from short-term optimization and speculative capture.

Health AI is a U.S.-based research and strategy group focused on building public-benefit infrastructure at the intersection of signal science, applied AI, and ethical systems design. We integrate biological signal modeling, human-machine feedback, and open AI protocols across both biomedical and industrial domains. Our work is grounded, scalable, and already in development. With federal alignment, the systems proposed here can be implemented quickly — and in ways that reflect national values, public trust, and long-term resilience.

Dr. Lawrence Meche
Chief Technology Officer, Health AI

244 Fifth Avenue, Suite 2588
New York, NY 10001

“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.”