

Public health AI's precision strategy: Globally integrating, institutionalizing, and scaling science, economics, and ethics

Dominique J Monlezun

Monlezun.Dominique@mayo.edu

Universidad Anáhuac México

Nandan Anavekar

Mayo Clinic

Lillian Omutoko

University of Nairobi

Patience Oduor

Africa Bioethics Network

Donald Kokonya

Masinde Muliro University

John James Rayel

Bicol University

Claudia Sotomayor

Georgetown University

Oleg Sinyavskiy

Kazakh National Medical University

Timothy Aksamit

Mayo Clinic

Gary Marshall

Mayo Clinic

David Grindem

Mayo Clinic

Dhairya Jarsania

Mayo Clinic

Keir MacKay

Mayo Clinic

Alberto Garcia

Pontifical Athenaeum Regina Apostolorum

Colleen Gallagher

Honors College, University of Houston

Sagar B Dugani

Mayo Clinic

Ryan Best

Creighton University

Maria Isabel Iñigo Petralanda

Pontificia Universidad Católica Argentina

Gifty Immanuel

Center for AIDS & Antiviral Research

Dinesh Visva Gunasekeran

National University of Singapore

Tien Yin Wong

Tsinghua University

Josiah Halm

UT MD Anderson Cancer Center

Cezar Iliescu

UT MD Anderson Cancer Center

Perspective

Keywords:

Posted Date: August 10th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10531187/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Artificial intelligence is spreading through public health agencies and healthcare systems faster than workforces are trained to govern it. We describe RHAMI (Responsible Health AI Readiness and Maturity Index), the first human-centered automated composite index for health organizations to score AI readiness and maturity as one measurable system of science, economics and ethics, and the physician certification-to-Master's pathway built to teach what it measures—both designed jointly by frontline public health professionals, clinicians, ethicists, engineers, leaders, and policy-makers in low- and middle-income countries and high-income countries.

A trust problem hiding inside a technology problem

Non-communicable diseases (NCDs) such as heart disease and cancer account for close to three-quarters of deaths worldwide, and the toll falls heaviest on low- and middle-income countries (LMICs) which carry over 80% of these deaths—while fielding just a quarter of the health workforce¹. Over 80% of NCDs are driven by modifiable behaviors targeted in public health prevention such as diet and exercise, while healthcare typically treats its acute complications when that fails. But broad critique has been levied that even the world's best-funded national healthcare and public health system illustrates how difficult these problems are to buy one's way out of: the United States has spent more than a trillion dollars on health research through the National Institutes of Health (NIH) over three decades...while life expectancy has only increased by three additional years, without any meaningful narrowing of health disparities²⁻⁴. Meanwhile, most national healthcare and public health systems are underfunded, and most LMICs now spend more on servicing debts than on health and education, even though they will account for most of the world's population growth this century⁵. Further, public trust in health institutions plummeted by as much as half worldwide since the COVID-19 pandemic, modern health's biggest test yet^{6,7}.

Artificial intelligence (AI) lands into this landscape as both a risk (worsening bias, disparities, and irresponsible use) and a promise (democratizing efficiency, productivity, and resilience)⁸. It is already transforming medicine and public health, while recent reports suggest that patients want more of it, not less. But its creators, owners, and infrastructure are remarkably concentrated: the United States and China alone account for more than 80% of global computing power, investment, and leading models⁹. Governments write regulations and health agencies write frameworks to govern AI use. Yet private technology platforms largely decide how AI behaves in clinics and communities at a pace that neither governments nor agencies can match. The Global Initiative on AI for Health by the World Health Organization (WHO) was created to close that gap¹⁰. But a harmonized standard is not the same thing as a workforce able to apply it.

The innovation-governance mismatch is widening quickly. The computing power and task complexity that AI systems can handle is doubling roughly every seven months, resulting in an annual compounding growth rate topping 200%¹¹. In the space of three years, general-purpose AI went from generating simple

text and images on request (generative AI [GenAI]), to autonomously executing complex multistep tasks alongside human collaborators (agentic AI)—and is widely expected as early as the near-term to operate with human-level generality and minimal supervision (artificial general intelligence [AGI]). Despite this technical trajectory, there is still no enforceable global standard for governing even today's simpler generative tools in public health. And roughly 90% of healthcare systems in the United States—the national health ecosystem with the most sophisticated and saturated AI—lack mature, funded AI governance, even though a similar share is already using AI in some form¹². The harder questions, including about AI systems that act with growing autonomy, are only beginning to be debated.

The deeper problem is not a knowledge gap that one training course can fix. It is an integration gap. Public health agencies and healthcare systems are offered an ever-expanding menu of AI tools, frameworks, and guidelines, built by different disciplines. Yet these are rarely validated against each other and almost never unified into something a working health professional or leader can actually scale, especially in LMICs^{13–16}. Health leaders describe an “interconnected readiness gap” spanning workforce skill, data infrastructure, and their governance, but with no proven road map to close that gap¹⁷. Scholars have called for embedding ethics “throughout the AI lifecycle,” and a recent review of 35 governance frameworks proposed a maturity pathway for healthcare systems with different resource levels^{18,19}. What is still missing is a unified, practical, ethical instrument that ties these efforts together empirically—and a curriculum that empowers health professionals to use it.

RHAMI: one empirical and ethical instrument for AI readiness, maturity, and governance

Responsible AI in health is usually treated as either an ethics problem to be solved by a principles document or a technical problem to be solved by a benchmark. Neither framing is sufficient alone. Any AI system that touches population health simultaneously raises three connected questions across diverse domains: what the technology can do and become (science), who can create and afford it (economics), and whether its use enhances or undermines the common good (ethics). Treated separately, these three questions correspond to fragmented and failed initiatives commonly seen across public health: tools that measure technical readiness without governance, or governance maturity without accounting for affordability or ethics^{20–22}. Treated together, they can equip the full spectrum of health ecosystem partners with a shared language for evaluating the same integrated AI solutions and capabilities (not just isolated pilots and frameworks).

Those observations led our international team—spanning frontline physicians, public health professionals, bioethicists, engineers, leaders, and policymakers across eleven countries in both LMICs and high-income countries (HICs)—to propose the Responsible Health AI Readiness and Maturity Index (RHAMI)²³. To our knowledge, it is the first tool to measure and advance this convergence empirically rather than describe it aspirationally. RHAMI scores and guides health AI systems along two axes (readiness to deploy and scale as well as the maturity of an organization's integration of it) and across four domains: responsible AI (transparency, explainability, accountability, and education), compliance

(ethics, policy, law, and regulation), organizational integration (workflow fit, executive engagement, workforce training, and digital infrastructure), and sustainability (health outcomes, financial return, operational risk, and environmental impact)—ultimately enabling end-to-end health AI governance strategically and operationally. This matters most where governance has the least precedent: increasingly autonomous, multi-agent AI systems that plan and adapt across health organizations' workflows. Existing regulations were largely written for traditional software that predictably performs defined functions programmed for it, rather than emergent software that adaptively reasons and executes increasingly on its own. Whatever pace of development anyone expects for such systems, the governance vocabulary and organizational muscle needed to supervise them cannot be improvised after the fact. It must be built into how the workforce is currently trained and workflow is structured.

RHAMI is also structurally different from prior scorecards (minted then filed away). We built it as an automated, API-embeddable tool that health organizations can run continuously against their own health and operational data in real time, paired with a customizable education layer. Therefore, the same instrument that scores and accelerates an organization's AI maturity can also upskill its staff, even where technical infrastructure is limited.

But ethics is where RHAMI departs most from existing tools. "Responsible AI" is one of the most repeated phrases in this field, but one of the least consistently defined. Regulators, engineers, ethicists and executives converge on the same words while applying them so differently that agreement in principle collapses into disagreement in practice. Rather, RHAMI grounds its ethical criteria explicitly in the 1948 *Universal Declaration of Human Rights (UDHR)*, the moral foundation from which the WHO's own constitution, international law, and most major AI governance frameworks already derive. It then translates that shared foundation into criteria specific enough to be scored and prescriptive enough to then guide. The index goes one step further, by detailing and operationalizing the *multicultural moral convergence*, not just the *pragmatic political consensus* that produced the *UDHR*. RHAMI does so by incorporating Personalist Social Contract ethics, which extends to AI the convergence across the world's culturally diverse belief systems (manifested by their varying political economic structures): protecting each person's dignity that is fulfilled through the common good, constituted by the good of each individual⁸. It is not a cure all for all disagreements. But it does give health professionals and policymakers a shared vocabulary robust enough to debate *and* decide about specifics of AI implementation.

But an index, however sound, cannot by itself give a health organization the skilled people needed to act on what it measures. That gap is what led our team to build the curriculum described below, translating RHAMI into something health professionals and leaders can embrace and deploy.

From index to curriculum: teaching what RHAMI measures

An index typically tells health leaders where they stand. It does not teach them or the health workers under them where the next needed step is. Our team thus created what is to our knowledge the first

standardized, sequential, streamlined AI executive certifications-to-Master's pathway designed specifically for public health professionals and physicians (issued by Universidad Anáhuac México according to graduate level-accreditation standards). It translates RHAMI's convergence into the first structured, adaptive, and personalized curriculum to be directly integrated into graduate medical education (GME), backed by global partnerships for inclusive hands-on-learning, and guided by GME's first dedicated AI Associate Program Director (D.J.M.), based in the Department of Cardiovascular Medicine in Mayo Clinic (which pioneered the first GME program a century prior).

This skills-based program is built on the respective public health and healthcare credentialing requirements by the Council on Education for Public Health (CEPH) and the Accreditation Council for Graduate Medical Education (ACGME). While it is available in a condensed continuous medical education (CME) pathway for working health professionals, it is principally designed within GME curriculums as a longitudinal AI sub-curriculum. Both pathways are productized as three progressive eight-week certifications (Practitioner, Innovator and Expert). Each features assigned readings, faculty-produced video lectures, live case discussions with AI experts, and structured faculty feedback, all culminating in a capstone project. The Practitioner stage builds foundational AI literacy for health workers with no data science background, teaching them to understand how a given AI platform operates and critically interrogate its safety, equity, and governance gaps before it ever reaches a patient. The Innovator stage extends that fluency across healthcare, public health, and global health governance simultaneously, positioning the health professional not merely as a user of AI but as someone accountable for governing it within the organization. The Expert stage addresses senior leadership directly, empowering learners to build a board-ready strategic roadmap for their own institutions that integrates a governance charter, a political economic risk analysis, and measurable implementation milestones, mirroring the same readiness-maturity-governance-sustainability logic intrinsic to RHAMI.

Several elements of this index-based curriculum are also practically without known precedent in health AI education. The three certifications can culminate in a Master's degree in Global AI Bioethics and Governance. It is the first such pathway that is simultaneously integrated into clinicians' GME and backed by internationally accredited universities spanning both LMICs and HICs. It is deliberately priced with online options to be widely accessible for both (rather than centralized as nearly all comparable health AI trainings currently are in HICs). So unlike AI's compute and infrastructure already concentrated in HICs, democratizing its training and governance can empower the health practitioners and leaders in local communities to care for their communities with world-leading capabilities wherever they are in the world^{1,4}. This end-to-end pathway is meant to connect AI development to real-world deployment within a unified, inclusive, continuously updated governance framework (rather than treating research, products, and policy as separate silos, balkanized into different economic strata). It also incorporates from its inception a co-designed curriculum created jointly by an international team of frontline health professionals, ethicists, data scientists, lawyers, leaders, and policymakers. And structurally expanding an AGI technique, this approach also is meant to enable 'open and responsible recursive self-improvement'. The underlying course materials, an open-source no-code deployment of a RHAMI-based

governance tool, and their prospective international cohort study (with embedded Bayesian adaptive clinical trials) are made widely accessible, so that the same instrument built to evaluate AI governance can also be used to teach it globally.

Why health professionals, and why now?

It is fair to ask why this convergence of science, economics and ethics should be taught to health professionals specifically, rather than left separately to data scientists, health economists, or policy specialists. The reason is structural not territorial. Physicians bear the clinical, ethical, and legal accountability for what happens to a patient in the room—regardless of whether a recommendation came from a colleague, a guideline, or an algorithm. An AI system that fails a patient becomes, operationally and morally, the treating physician and hospital's responsibility long before it becomes the vendor's, the engineer's, or the regulator's. Physicians who know AI only as a tool to be used—without understanding the economics, ethics and governance determining it—are being asked to accept broad accountability for systems they have no systematic way to evaluate. And without translating the centuries-long insights from public health into physician training, our health workers can treat (expensively) but still struggle to prevent diseases and disparities (much more affordably) with it.

That asymmetry is widening quickly. The AI systems capable of autonomous, multistep reasoning are already moving from single clinical decision support toward coordinated, multi-agent orchestration including in cardiology and oncology, with some developers explicitly detailing a trajectory toward artificial super intelligence (ASI) superior to human and AGI performance. Whatever one's view of how near that capability actually is, the governance questions it raises remains immediate: who decides how such systems are validated, who is trained to supervise them, and whose health needs shape their design. Waiting for more capable AI to arrive before building the workforce to govern it is a wager the health ecosystem cannot responsibly make, least of all on behalf of LMIC populations—historically the last to benefit from innovation and the first to bear its missteps.

There is a further urgent and practical reason physicians need this training now: most of them and their patients are already using AI tools, which are entering practice faster than most licensing boards, GME curricula, and organizational governance have adapted to address. A physician who has never been taught to ask what data trained a model, who owns its outputs, what a vendor's liability actually covers, or how an AI recommendation should be documented and overridden is not protected from those questions simply by declining to engage with the technology that is already in the room. The structured training that our RHAMI-based curriculum attempts to provide turns this from an ad hoc risk individually absorbed by each clinician into an institutionally governed capability and core competency.

A global partnership, not a commercial export

A defining feature of both RHAMI and the curriculum is its partnership-by-design: healthcare and public health, HICs and LMICs, theorists and practitioners, workers and leaders. Our diverse team spanning

them built these solutions based on our earlier analyses and guidance on responsible digital health transformation for LMICs that the WHO solicited from us²⁴. This is a strategic design requirement, not a stylistic preference. Health AI frameworks built exclusively within high-income institutional contexts regularly assume levels of infrastructure, financing, and regulatory capacity that most of the world's public health agencies and healthcare systems do not have. An index or curriculum calibrated only to that context will misjudge readiness and advise governance that does not generalize.

A genuinely global instrument and curriculum, by contrast, can generate scalable insights a narrower one may miss (i.e., that a low-cost, smartphone-based AI malnutrition assessment tool validated in underserved areas may matter more for scaling high-impact public health nutrition across more of the world than a far more sophisticated precision nutrition-based platform built for a well-resourced neonatal intensive care unit)^{1,23}. Responsible AI for public health is not the same as responsible AI for tertiary academic medicine, and a workforce trained only in the latter will consistently underserve the former. RHAMI and its curriculum strive to unite both ends of the health spectrum needed for global health.

A human-centered precision strategy: responsible AI for truly public health

We do not present RHAMI or the curriculum as finished solutions, but hopefully concrete steps in the right direction. Both are crafted to adapt in real time to global critique and technological developments^{8,23}. And both are meant to further evidence how the health ecosystem globally will not scale responsible AI simply by publishing more principles, benchmarking more models, or running more pilots. It will achieve that scale when the professionals accountable for population health outcomes are trained and collaborate with the same rigor in human-centered AI economics and governance that we already expect of them in health and management sciences.

We therefore ask three things of the field. First, that health AI readiness and maturity indices build financial sustainability including for LMICs into their core design; an index that only wealthy systems can act on is not a global index (as innovation accelerates proportional to the networked size of the ecosystem). Second, that health AI training moves out of isolated modules and into structured, credentialed pathways integrated into GME, so that governance competence becomes an embedded professional expectation rather than an optional enthusiasm. Third, that both indices and curricula for health AI be built from day one as genuine partnerships between the Global North and Global South, not as instruments manufactured in one and subsequently exported to the other. Anything less will keep producing AI that is impressive in its capability, but uneven—sometimes dangerously so—in whom it actually touches.

We entered the 21st century equipped with a genome map enabling precision medicine (tailored to the person's uniqueness through biology down to the genetic level). By the next decade, precision public health emerged (expanding the former through behaviors to the population level). By our decade, we

propose a precision strategy uniting both to integrate, institutionalize, and scale effective, sustainable, and responsible AI transformation of our shared health ecosystem globally.

The smarter our methods and metrics for health AI become, the more they should return us to the oldest fundamental question in the field: not just what science can do, but first, for whom is it for? Building a workforce able to keep asking and answering that question, in the room where the decision is made, is the task this Perspective aims to begin addressing—personally, precisely, and strategically. Not just for the health of a few, but for all of us.

Declarations

Competing interests

The author declares no competing financial or non-financial interests relevant to this Perspective's subject matter.

Funding source:

No funding was used in this study.

Author Contribution

D.J.M. conceived the perspective and drafted the manuscript. D.J.M. and N.A. conceived the framework and performed the literature synthesis. D.J.M., N.A., L.O., P.O., D.K., J.J.R., C.S., and O.S. analyzed current policy trends. D.J.M., N.A., T.A., G.M., D.G., D.J., K.M., A.G., C.G., S.B.D., R.B., M.I.I.P., G.I., D.V.G., T.Y.W., J.H., and C.I. provided critical intellectual feedback, structured specific thematic sections, and substantively revised the manuscript. All authors read and approved the final version.

Data availability

Data sharing is not apply to this Perspective, as no new datasets were generated or analyzed.

References

1. World Health Organization. Fact sheets. <https://www.who.int/news-room/fact-sheets>
2. United States Congress. National Institutes of Health (NIH) funding: FY1996–2026. <https://www.congress.gov/crs-product/R43341>
3. Centers for Disease Control and Prevention. Life expectancy. <https://www.cdc.gov/nchs/fastats/life-expectancy.htm>

4. The Commonwealth Fund. Health disparities report.
<https://www.commonwealthfund.org/publications/fund-reports/2026/apr/commonwealth-fund-2026-state-health-disparities-report>
5. United Nations. Financing development: external flows of financial capital to developing countries and their cost. <https://unctad.org/publication/financing-development-external-flows-financial-capital-developing-countries-and-their>
6. Guo, C. et al. The effect of COVID-19 on public confidence in the World Health Organization: a natural experiment among 40 countries. *Globalization and Health* 18, 77 (2022).
7. Perlis, R. H. et al. Trust in physicians and hospitals during the COVID-19 pandemic in a 50-state survey of US adults. *JAMA Network Open* 7, e2424984 (2024).
8. Monlezun, D. J. *AI Platforms as Global Governance for the Health Ecosystem: The Future's Global Hospital* (Morgan Kaufmann/Elsevier, 2026).
9. Stanford University Institute for Human-Centered Artificial Intelligence. The 2026 AI Index Report. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
10. Muralidharan, V. et al. Global Initiative on AI for Health (GI-AI4H): strategic priorities advancing governance across the United Nations. *npj Digit. Med.* 8, 219 (2025).
11. RBC. The exponential leap: understanding AI's real impact.
<https://www.rbccm.com/en/insights/2026/01/the-exponential-leap-understanding-ais-real-impact>
12. HFMA. Healthcare system readiness for artificial intelligence. <https://elicitinginsights.com/hfma-system-readiness-for-ai-study-2025/>
13. Durlach, P. et al. The AI Maturity Roadmap: a framework for effective and sustainable AI in health care. *NEJM AI* (2024).
14. Gupte, N. Responsible AI Readiness Index: a tool to track the progress and impact of responsible AI. International Monetary Fund (2024).
15. Schwalbe, N. & Wahl, B. Artificial intelligence and the future of global health. *Lancet* 395, 1579–1586 (2020).
16. Ciecierski-Holmes, T. et al. Artificial intelligence for strengthening healthcare systems in low- and middle-income countries: a systematic scoping review. *npj Digit. Med.* 5, 162 (2022).
17. Durlach, P. et al. At the frontier: gauging health care's readiness for agentic AI innovation. <https://ai.nejm.org/doi/full/10.1056/AI-S2501336>
18. Antao, E. M. et al. Reason and responsibility as a path toward ethical AI for (global) public health. *npj Digit. Med.* 8, 329 (2025).
19. Hussein, R., Zink, A., Ramadan, B. et al. Advancing healthcare AI governance through a comprehensive maturity model based on systematic review. *npj Digit. Med.* 9, 236 (2026).
20. Gupte, N. Responsible AI Readiness Index: A Tool to Track the Progress and Impact of Responsible AI (2024).

21. Schwalbe, N. & Wahl, B. Artificial intelligence and the future of global health. *Lancet* 395, 1579–1586 (2020).
22. Ciecierski-Holmes, T. et al. Artificial intelligence for strengthening healthcare systems in low- and middle-income countries: a systematic scoping review. *npj Digit. Med.* 5, 162 (2022).
23. Monlezun, D. J. et al. The Responsible Health AI Readiness and Maturity Index (RHAMI): applications for a global narrative review of leading AI use cases in public health nutrition. *Nutrients* 18, 38 (2026).
24. Monlezun, D.J. et al. Digitalization of health care in low- and middle-income countries. *Bulletin of the World Health Organization* 103, 2, 148–154.