Viewpoint

Domain Knowledge Plus AI Fluency: The New Nonnegotiable Competency for the Health Workforce in the AI Era

Seble Frehywot[1], MD, MHSA; Yianna Vovides[2], MA, MS, PhD

[1]Department of Global Health, Milken Institute School of Public Health, The George Washington University, Washington, DC, United States
[2]Centers in New Design in Learning and Scholarship, Georgetown University, Washington, DC, United States

Corresponding Author:
Seble Frehywot, MD, MHSA
Department of Global Health, Milken Institute School of Public Health, The George Washington University
950 New Hampshire Ave NW
Washington, DC, 20037
United States
Phone: 1 7038519645
Email: seblelf@gwu.edu

Abstract

In the ongoing AI boom—now supercharged by large language models and generative AI tools—the rapid integration of AI into health care and public health has created an urgent need for professionals who combine deep domain expertise (clinical, epidemiological, ethical, and contextual expertise in health and public health) with AI fluency (the ability to understand, critically evaluate, use, and oversee AI systems effectively). Without domain guidance, AI risks producing biased, unsafe, or contextually irrelevant outputs that compromise patient safety and exacerbate disparities. Conversely, domain experts lacking AI fluency may underuse or misuse these technologies, missing opportunities to improve care quality, operational efficiency, and population health. A foundational understanding of the modern AI system stack—comprising energy, accelerators, infrastructure, models, and applications—empowers health professionals to make informed decisions on adoption, evaluation, governance, and ethical use without needing to code or build models themselves. This stack framework highlights critical interdependencies and health care–specific challenges, including regulatory compliance, data privacy, sustainability, and resource constraints, particularly in low- and middle-income settings. This paper argues that cultivating domain-AI bilingualism is essential for safe, effective, and equitable AI integration worldwide. It outlines the 5-layer AI system stack applied to health care, key reasons for hybrid expertise, and a practical 4-step pathway: awakening (recognizing AI's transformative role), intrepidity (courageously engaging foundational AI concepts), bilingualism (translating between clinical needs and AI capabilities), and visionary leadership (driving organizational and policy readiness). It outlines major barriers, targeted strategies, and recommendations for educators, institutions, and policymakers. By bridging domain knowledge with AI fluency, health professionals can direct—rather than be directed by—AI technologies, including emerging agentic systems, toward more precise, sustainable, and human-centered outcomes while mitigating risks. This hybrid capability represents a nonnegotiable competency for the health workforce for thriving in the post-2026 AI era.

(JMIR AI 2026;5:e100954) doi: 10.2196/100954

KEYWORDS

artificial intelligence; AI; AI literacy; domain-AI bilingualism; health workforce development; AI system stack; responsible AI integration; digital health

Introduction

In the current AI boom, the rapid integration of AI into health care and public health systems has created an urgent need for professionals who combine deep *domain knowledge* (clinical, epidemiological, ethical, and contextual expertise in health and public health) with *AI fluency* (the ability to understand, critically evaluate, use, and effectively oversee AI tools). This dual competency has become even more critical in the post–large language model (LLM) generative AI era, where LLMs and accessible generative tools have dramatically accelerated AI adoption across clinical, operational, and population health applications. What once required specialized technical teams can now be prompted or fine-tuned by end users in minutes, compressing the timeline for both innovation and potential harm.

XSL•FO
RenderX

Without domain expertise, even sophisticated AI application tools risk generating biased, unsafe, or contextually irrelevant outputs that can compromise patient safety and public health outcomes. Conversely, health domain experts lacking AI fluency may underuse or misuse these technologies, missing critical opportunities to enhance care quality, operational efficiency, and population health. In the ongoing AI boom—now supercharged by generative AI advancements—health and public health practitioners and students require foundational AI knowledge. This is not to transform them into software engineers or coders but to enable them to understand the core layers and components of the modern AI system stack, from data foundations, accelerators (chips), and model architectures to application, deployment, monitoring, and governance.

The imperative holds global relevance. In high-resource settings, domain-AI–fluent professionals are essential for safely scaling precision medicine, hospital operations, and health system optimization. In low- and middle-income countries (LMICs), where health workforce shortages are acute and contextual challenges (such as resource constraints and localized disease burdens) are pronounced, this combination of competencies is vital for equitable and appropriate AI deployment. It helps ensure that globally diffused AI tools do not exacerbate existing disparities but instead strengthen local health systems, support evidence-based decision-making, and address region-specific public health priorities. As generative AI tools continue to proliferate worldwide, building this hybrid expertise across diverse health workforces has become a shared global priority for responsible innovation and sustainable health system strengthening.

This paper examines the key reasons why domain knowledge combined with AI fluency is essential for safe, effective, and equitable AI integration in health and public health; outlines the AI system stack applied to health care and the essential steps required to build domain-AI bilingualism; identifies major barriers to its adoption and practical strategies to overcome them; and offers targeted recommendations for educators, institutions, and policymakers to cultivate this critical workforce capability at scale.

The 5 Components of the AI System Stack Applied to Health Care

The 5 components of the modern AI system stack (Figure 1 [1])—often described as a "five-layer cake" by industry leader NVIDIA's Jensen Huang [1]—apply to the health sector as follows. Each layer builds on the one below it to enable scalable, reliable AI in health care.

Figure 1. The components of the modern AI system stack in health care (image adapted and created by the authors from Huang [1]).

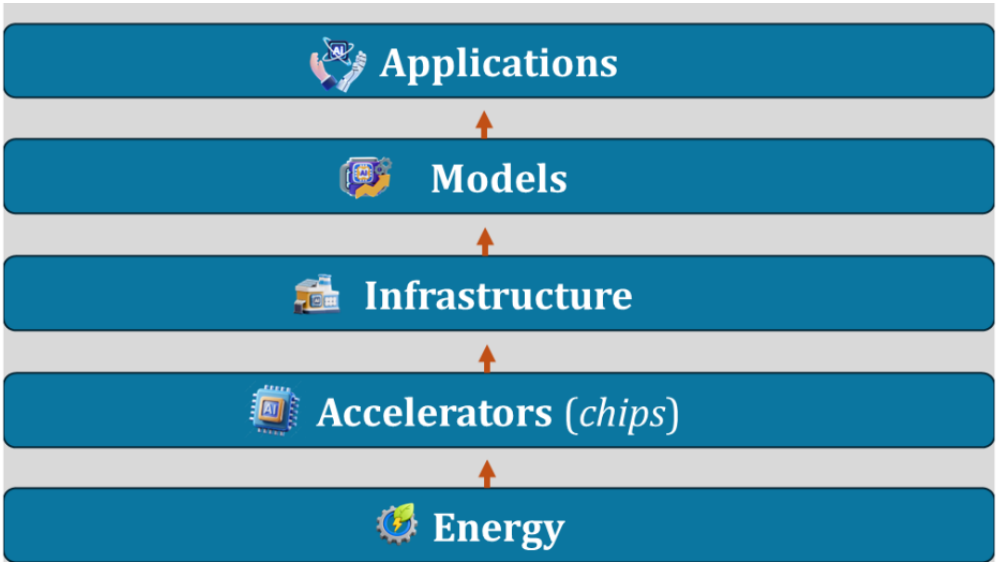

Energy

Energy is the foundational power layer. AI systems in health care (especially large-scale training of models or always-on inference in hospitals) consume massive amounts of electricity. Reliable, sustainable energy sources (including efficient grids, renewable energy, and backup power) are critical to avoid outages in life-critical applications such as real-time diagnostics or intensive care unit (ICU) monitoring while also addressing growing concerns about AI's environmental footprint in medical facilities.

Accelerators

Accelerators are AI chips. These are specialized hardware such as graphics processing units (GPUs), tensor processing units (TPUs), neural processing units (NPUs), or other AI chips that provide the parallel processing power needed for AI workloads. In health care, these handle compute-intensive tasks such as training models on huge medical datasets or running fast inference for medical image analysis (x-ray, magnetic resonance imaging, computed tomography, and pathology imaging), genomics, real-time edge AI wearables, and portable diagnostic devices.

Infrastructure

This constitutes the physical and virtual backbone, including data centers, cloud or hybrid setups, high-speed networking, cooling systems, storage, and edge servers. In the health sector, this must comply with strict regulations; support secure handling of sensitive patient data; enable hybrid edge-cloud architectures

XSL•FO
RenderX

(eg, on-premises for privacy plus cloud for scale); and power hospital systems, micro–data centers, or AI-ready medical devices.

Models

These are the AI "brains"—foundation models, domain-specific LLMs, computer vision models, or multimodal models fine-tuned or trained on health care data. Examples include models for radiology report generation, clinical decision support, predictive risk scoring, personalized treatment recommendations, or analysis of electronic health records while preserving privacy.

Applications

This is the top layer, where AI delivers real-world value to end users. In health care, this includes clinical tools (eg, diagnostic aids and virtual assistants for physicians), operational solutions (eg, predictive bed or staff management and supply chain optimization), patient-facing apps (eg, wearables for monitoring, chat-based triage, and personal care plans), administrative automation (eg, clinical documentation and billing), and population health tools—all designed to improve outcomes, reduce clinician burnout, and enhance access.

Interdependencies and Health Care–Specific Challenges

The 5 layers of the AI system stack are deeply interdependent in health care, where a weakness in any single layer can cascade into failures across the entire system, directly affecting patient safety, data security, and operational reliability.

Energy and accelerators, for instance, underpin the computational demands of training and inference in large models; unreliable power or insufficient GPU or TPU capacity in hospitals can delay real-time diagnostics or cause outages in ICU monitoring systems. These hardware constraints, in turn, influence infrastructure choices—such as hybrid edge-cloud architectures—where sensitive patient data must be processed locally for speed and privacy while being scaled in the cloud for advanced analytics. Models and applications then rely on this foundation: a multimodal LLM fine-tuned on electronic health records can only deliver trustworthy clinical decision support if the underlying infrastructure ensures secure, compliant data flows and the energy supply remains sustainable.

In resource-limited settings, this interdependence often forces trade-offs between performance (eg, high-accuracy image analysis) and feasibility (eg, edge deployment on portable devices to reduce latency and costs) [2]. Health care introduces unique regulatory and privacy challenges that amplify these interdependencies. Data privacy regulations such as HIPAA (Health Insurance Portability and Accountability Act) in the United States and the General Data Protection Regulation (GDPR) in the European Union (EU) impose strict requirements on patient information across all layers, particularly in infrastructure and models, where vast datasets are ingested for training [3]. Breaches or noncompliant processing—such as using deidentified protected health information in cloud storage or model fine-tuning—can lead to severe penalties and loss of trust [4]. The EU AI Act, which entered into force in 2024, further classifies many health care AI systems (eg, diagnostic tools or predictive risk scorers) as high risk, mandating rigorous conformity assessments, transparency, human oversight, and risk management throughout the stack—from energy-efficient training practices to bias audits in applications [5]. Compliance thus requires domain experts to collaborate closely with technical teams on issues such as secure data governance, explainability of model outputs, and sustainability reporting while navigating extraterritorial implications for global health systems. Without integrated oversight, these challenges risk widening inequities, increasing environmental footprints, or eroding clinician confidence in AI tools. Domain-AI–fluent professionals are therefore essential to navigate this complex landscape responsibly.

Key Reasons Why Domain Knowledge and AI Fluency Are Essential

This foundational understanding of the components of the AI system stack empowers health care professionals to make informed decisions about adoption, evaluation, governance, and ethical use without needing to build models themselves. It assists them in consciously thinking and acting upon the following aspects.

Informed Oversight and Critical Appraisal

AI's impact on the "outcomes of health care" and its "potential to address patient needs" are growing, and this brings ethical and operational challenges [6]. AI models in health care (eg, diagnostic algorithms and predictive tools) can produce errors, hallucinations, or biases if trained on flawed data or misapplied. Domain experts are required to assess accuracy, reliability, validity, and clinical relevance—skills emphasized over pure programming knowledge [7]. Understanding model layers (eg, how transformers process sequences via attention mechanisms) and infrastructure (eg, why GPUs or TPUs are needed for scale) helps identify when outputs are unreliable or when deployment fails due to compute constraints. Without this, there can be overreliance risks, errors, deskilling, or disparities.

Resource and Sustainability Awareness

Training frontier models requires enormous amounts of energy (often equivalent to thousands of households) and specialized hardware [8]. Practitioners need to grasp these layers to advocate for sustainable AI (eg, prioritizing efficient models or edge inference in resource-limited settings) and assess environmental and public health impacts.

Effective Collaboration With Technology Teams

Domain experts with foundational AI knowledge can communicate needs, set realistic expectations, and guide validation—asking precise questions about model architecture, training data, accelerators used, or inference costs—leading to better tools and return on investment.

Ethical, Regulatory, and Risk Management

Knowledge of AI layers helps recognize issues such as hallucinations in LLMs, bias amplification in deep learning, or infrastructure vulnerabilities (eg, data center outages). This supports responsible use, aligning with principles such as human

XSL•FO
RenderX

autonomy and transparency. Emerging concepts such as "critical AI health literacy" [9] empower patients and health care professionals to challenge biased systems.

Building the Bridge: Essential Steps Toward Domain-AI Bilingualism

The journey toward achieving health workforce domain-AI bilingualism requires 4 steps: awakening, intrepidity, bilingualism, and visionary leadership (Figure 2).

Figure 2. Framework for building the bridge: essential steps toward health workforce domain-AI bilingualism (image created by the authors).

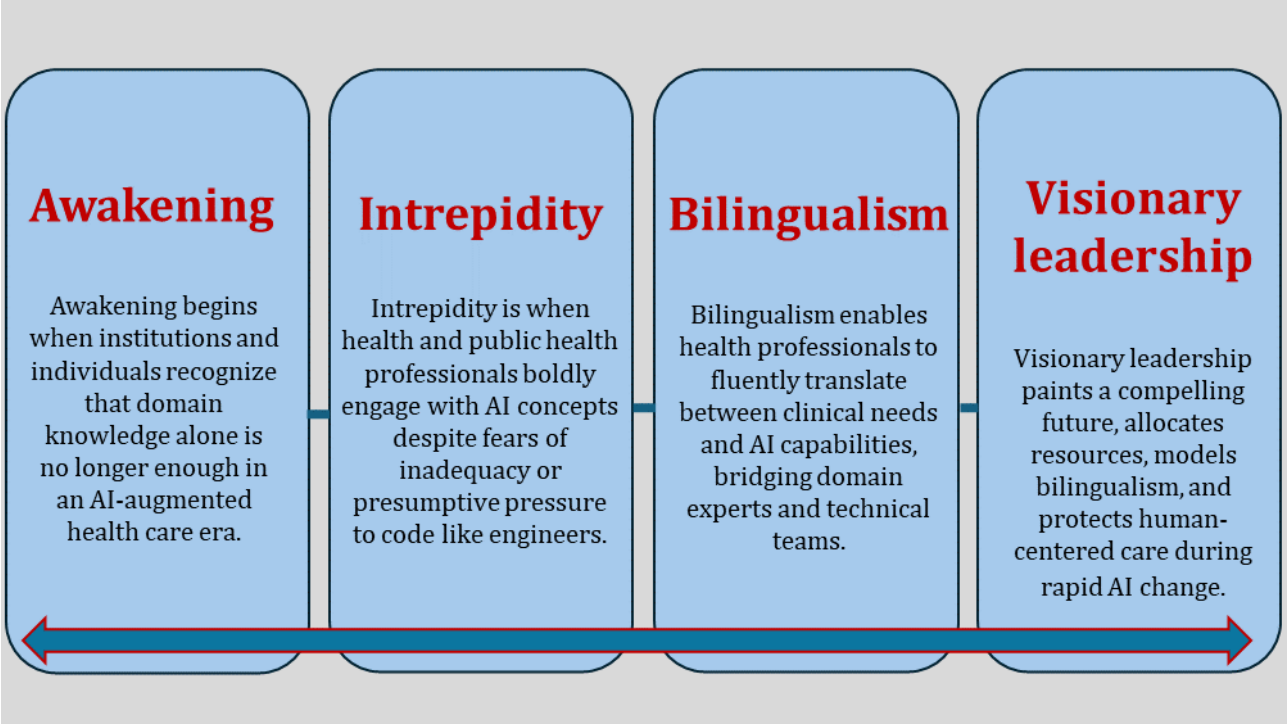

Awakening to This Reality

Awakening begins with institutional and individual recognition that domain knowledge alone is insufficient in an AI-augmented health care era. Many clinicians and public health leaders still treat AI as a peripheral innovation rather than a transformative force reshaping diagnostics, workflows, and population health strategies. This realization creates urgency and overcomes inertia. For example, a 2025 cross-sectional study of European health care professionals revealed persistent low AI literacy, with over 60% citing lack of understanding of AI capabilities and limitations as a major barrier to adoption and only 13.8% feeling that their training actually prepared them for AI integration; institutions that ran targeted awareness workshops saw rapid increases in voluntary engagement with AI tools [10]. Another case is Singapore's national health system, which launched a system-wide AI readiness campaign in 2024-2025, including mandatory orientation modules for all clinical staff, which shifted perceptions from skepticism to proactive exploration, resulting in faster pilot approvals for AI-assisted triage systems and measurable improvements in staff confidence [11,12]. Studies show persistent low AI literacy among health care professionals, with barriers such as lack of understanding of AI's capabilities and limitations, explainability concerns, and fear of professional disruption slowing adoption [13]. Awakening creates urgency, and once there is awareness, motivation emerges for the subsequent steps.

Intrepidity (Courage to Learn Foundational AI Knowledge Without Fear)

Intrepidity involves clinicians boldly stepping into AI concepts despite fears of inadequacy or expectations to code like engineers. In this context, it means learning about foundational AI concepts. This courage turns early adopters into internal champions who reduce organizational resistance. No one expects clinicians or public health officers to build LLMs, but basic fluency lets them critically assess AI outputs, spot errors in clinical decision support, or use tools ethically [14]. A notable example is the Mayo Clinic's faculty development initiative, where physicians participated in short, nontechnical "AI literacy bootcamps" focused on understanding model limitations and bias rather than programming [15]. Similarly, nursing programs incorporated hands-on simulation laboratories where nurses experimented with AI chatbots for patient triage without coding requirements [16]. This fearless experimentation approach can lead to frontline staff proposing practical improvements to vendor tools, demonstrating how intrepidity fosters grassroots innovation.

Bilingualism (Fluency in Both Domain and Foundational AI Language)

Bilingualism equips health professionals to translate seamlessly between clinical needs and AI capabilities, bridging the gap between domain experts and technical teams. This hybrid skill set is the nonnegotiable bridge. Domain experts alone risk

XSL•FO
RenderX

misapplying AI (eg, biased models deployed without clinical context); AI experts alone risk producing clinically irrelevant or unsafe tools [17]. This hybrid skill set prevents misapplication of biased models and ensures clinically relevant tools. At Stanford Medicine, interdisciplinary AI–clinical translation teams pair physicians with data scientists in joint rounds; clinicians learn to articulate requirements in terms of model architecture and data needs while engineers gain contextual understanding of ethical and workflow constraints [18,19]. This has produced more accurate radiology AI tools and reduced deployment failures. In the United Kingdom's National Health Service, a digital health bilingual training pathway for public health officers has enabled professionals to co-design population health dashboards that incorporate epidemiological expertise with multimodal AI models, leading to better-focused interventions during disease outbreaks [20]. Institutions need these translators at every level—clinicians who know the requirements for AI developers and educators who teach foundational AI knowledge–integrated curricula [21].

Visionary Leadership

Visionary leadership articulates a compelling future, allocates resources, models bilingualism, and safeguards human-centered care amid rapid change. Leaders champion organizational readiness for advanced applications such as agentic AI. They pioneer not only AI fluency for individual readiness but also for organizational readiness [22]. They future-proof the institution—preparing the workforce for agentic AI, personalized care and population health, or AI-driven policy simulation—while preserving humanity in care. A compelling case is Cleveland Clinic's AI strategy under visionary executives who integrated AI fluency into leadership competencies and established a dedicated Center for Clinical Artificial Intelligence with cross-functional governance [23,24]. Another example is the World Health Organization's regional initiatives that combine strategic planning with hands-on exposure [25,26]. These efforts help policymakers simulate AI-driven resource allocation scenarios while prioritizing workforce protection, demonstrating scalable visionary impact while preserving humanity in care [27].

These 4 elements—awakening, intrepidity, bilingualism, and visionary leadership—form a progressive pathway that transforms individual mindsets into institutional capability, ensuring that the health workforce thrives alongside AI rather than being displaced by it.

Potential Barriers and How to Overcome Them

Overview

Despite the clear imperative for domain knowledge combined with AI fluency, several barriers continue to slow progress in health care and public health settings. Time constraints in already demanding clinical schedules often leave little room for new learning, while persistent fears of professional disruption or "deskilling" foster resistance among practitioners. Many clinicians worry they lack the technical aptitude to engage with AI concepts, and resource disparities—particularly between well-funded academic centers and underresourced community or LMIC settings—risk widening the digital divide. Infrastructure limitations, such as unreliable energy supplies, legacy systems incompatible with modern accelerators, or stringent data privacy requirements under HIPAA or the GDPR, further complicate deployment of even well-designed models and applications. Additionally, fragmented regulatory environments and the rapid pace of AI evolution can overwhelm institutions, leading to cautious or inconsistent adoption [28].

These barriers are surmountable through targeted, layered strategies that align with the 4 foundational elements. Awakening can be advanced via concise, role-relevant awareness campaigns and success stories shared through grand rounds or internal newsletters, gradually shifting mindsets from skepticism to curiosity. Intrepidity is cultivated by designing low-stakes, modular learning experiences—short online modules, peer-led simulation laboratories, and "AI playground" environments—that emphasize critical evaluation over coding expertise, reducing fear of incompetence. Bilingualism grows when institutions create structured collaboration forums, such as joint AI-clinical rounds or translation workshops, enabling domain experts and technologists to codevelop tools and build mutual understanding—in the process building a critical mass of bilingual peer champions. Finally, visionary leadership is key to overcoming systemic hurdles: leaders can secure dedicated funding; advocate for policy support (eg, aligning with the EU AI Act's high-risk requirements) [29]; and prioritize equitable infrastructure investments, including edge computing solutions for resource-limited environments.

By proactively addressing these challenges with practical, human-centered solutions, health systems can transform potential obstacles into opportunities for sustainable and inclusive AI integration, ensuring that the entire workforce is equipped to harness AI responsibly while preserving the core values of compassionate care.

Navigating Energy Crises, Workforce Shortages, and Hardware Constraints Through Domain-AI Bilingual Leadership

The case for domain-AI bilingualism is strongest in LMICs facing acute workforce shortages [30] and chronic energy crises—blackouts, unstable grids, and limited renewable energy [31]. South Asian studies document specialist shortages and load shedding that force the adoption of lightweight mobile models [32,33]; Latin American and Global South reviews note interoperability failures, technical brain drain, and nonlocal models that underperform on regional disease patterns [32,33]. Bilingualism must augment scarce capacity rather than add administrative layers, using existing human capital instead of paperwork or overtime [34]. AI fluency should be embedded into routine work through passive, just-in-time learning and shared digital infrastructure so that overstretched workers gain oversight without extra load.

GPUs and TPUs need continuous high-wattage power infrastructure, which is rare in LMIC facilities, and cooling that many district hospitals lack [35,36]; hence, large multimodal models become impractical. Clinics often rely on intermittent

XSL•FO
RenderX

grids or diesel, making energy-intensive training impossible and continuous inference unreliable [37]. The result is dependence on foreign models trained on high-resource data, with performance drops and bias against local phenotypes [38]. Edge and TinyML mitigate but do not fully close the indicated gap [39]. Leaders must select lightweight, offline-first architectures on smartphone NPUs or single-board computers drawing under 10 W [40] runnable from small solar panels, car batteries, or weak grids. Options include quantized or distilled models that cut compute by 70% to 85%, on-device inference that avoids continuous connectivity, and hybrids that diagnose locally and sync when power and bandwidth allow [39,41]. The total cost of ownership—diesel vs renewable microgrids—should be assessed, and vendor transparency on energy and hardware should be demanded. Selection then shifts from frontier LLMs to context-appropriate tools that retain utility under instability and stop lower-stack failures from cascading into applications [32]. Cases include offline chest x-ray AI tuberculosis diagnosis in South African mines without cloud or on-site radiologists [42] and India's offline portable primary health systems where bilingual district officers choose low-power models [43]. Leaders who understand energy and accelerator layers scale tools that remain functional when the national grid fails, so bilingualism strengthens rather than strains scarce human resources.

Operationalization of Lower-Stack Hardware and Energy Knowledge by Frontline Clinicians, Public Health Officers, and Nontechnical Professionals

Nontechnical professionals operationalize lower-stack knowledge through simple heuristics in existing workflows, not new software or reports. A nurse manager keeps a laminated power status card beside the AI triage tablet: green (stable grid or solar) indicates the full model, amber (intermittent) indicates the lightweight edge mode, and red (blackout) indicates paper protocol. A community health supervisor who knows that accelerators generate heat schedules AI-assisted screening for cooler morning hours to avoid thermal throttling. These practices are transmitted through brief verbal briefings at shift handover.

Frontline clinicians and public health officers apply the same knowledge in low-frequency, high-stakes decisions instead of continuous monitoring. A rural physician evaluating an AI chest x-ray first asks whether the model runs offline on existing low-power devices or needs cloud access that fails in outages [44]. If a clinician uses a portable ultrasound AI package with an NPU that draws 3 to 5 W when the local solar microgrid can sustain only intermittent 50-W loads, the clinician might reject a higher-accuracy, cloud-dependent model for an offline quantized version usable during the daily brownout [45]. During procurement, a hospital medical director compares total energy costs and backup needs and rejects platforms that would overload limited solar capacity or increase diesel dependency, protecting life-critical systems such as ICU monitoring. A district public health officer evaluating an AI outbreak dashboard recognizes that continuous cloud inference would require diesel backup costing more than the intervention's benefit and, therefore, mandates an offline-first architecture that updates only when connectivity and power coincide. A similar logic is visible in India, where AI-augmented platforms for tuberculosis and infant health have been scaled by embedding tools in existing workflows and keeping frontline officers in command of when inference runs rather than assuming continuous cloud availability [46].

Awakening is typically triggered by one observed failure linked to heat or power—an AI tool unusable during load shedding. After repeated outages, a district nursing supervisor insists that AI triage chatbots include offline and low-power modes and trains staff to fall back to paper protocols when power drops. Intrepidity appears when staff test the heuristic under real conditions or clinicians request vendor demonstrations under simulated brownout. Awareness of accelerator and energy demands then favors quantized edge models over high-accuracy but power-hungry alternatives, preserving availability during night shifts or generator shortages [47]. Bilingualism develops when the nurse manager states the heuristic to procurement in wattage and thermal limits or a primary care physician reviewing vendor proposals asks about inference hardware and energy per encounter and rejects cloud-only solutions that would exclude satellite clinics, yielding language such as "must sustain 4-hour offline inference on a 10 W NPU." Visionary facility leadership formalizes the heuristic as a standard operating procedure and embeds hardware-aware criteria in technology assessment checklists so that every future purchase is energy aware without new administrative burden. These actions—procurement criteria, training content, and contingency planning—translate stack awareness into practical safeguards without coding expertise [48]. Reviews of trustworthy AI adoption in Indian health care argue that frontline criteria—energy reliability, offline fallback, and workflow fit—are as decisive as model accuracy if safety and ethical use are to be sustained [49]. Parallel Nigerian examples show that such operationalization sustains AI utility under chronic infrastructure stress [50].

Agentic AI: Deskilling, Automation Bias, and Governance of Agentic Workflows

Overview

As agentic systems plan and act with less supervision, users trust usually correct outputs—automation bias that turns reviewers into rubber stamps. The handoff deskills oversight: practice migrates to the machine, judgment atrophies, and the human "in the loop" cannot reliably catch errors or override high-stakes steps. Governance cannot treat a token checkpoint as enough; it must keep evaluative agency alive through scoped permissions, irreversible action gates, and training that preserves foundational skills so that accountability is real, not theatrical.

Critical Caveat on Deskilling and Automation Bias With Agentic AI

Visionary claims for agentic AI must be tempered by deskilling and automation bias risks. Multistep autonomous planning [51] leads to progressive deskilling and bias, and repeated AI assistance erodes independent skills [52]. Clinicians who defer to sequential agent recommendations lose the pattern recognition and uncertainty tolerance that constitute clinical intuition [53]. For trainees, the risk is "never-skilling": formative diagnostic

XSL•FO
RenderX

practice is displaced, and cognitive offloading further weakens those skills [52,54]. Psychologically, high-confidence agent outputs produce complacency, reduced vigilance, and eroded self-efficacy—already observed with earlier decision support and amplified by agentic continuity across a care episode [54]. Cognitively, the human-in-the-loop role shifts from active reasoning to passive verification, accelerating decay especially among early-career practitioners [55]. Awareness of this produces mitigation intrepidity, which requires mandatory independent reasoning before viewing agent recommendations, periodic "AI-off" practice, and curricula that treat agents as copilots whose failures bilingual clinicians must detect and override. Without safeguards, agentic advances risk developing a workforce less able to intervene when systems fail or hit edge cases.

Governance, Ethical, and Clinical Validation Protocols Specific to Agentic Workflows

Agentic systems that chain reasoning and action need governance unlike that for passive diagnostic tools [56]. Validation moves from single-output accuracy to end-to-end process evaluation under realistic resource constraints: final accuracy, intermediate decision quality, recovery from partial failures, and energy and accelerator use across the trajectory [51]. The human-in-the-loop role must be substantive (named override authority and predefined escalation thresholds), not nominal [55].

Multiagent interactions create compounded opacity and diffuse accountability; protocols must mandate full-chain explainability, bias and equity audits throughout, transparent disclosure to preserve patient autonomy, and explicit liability [56]. Ethical rules require continuous human override, audit logging of every step, and bias audits of whether multistep plans disadvantage underrepresented populations or low-resource facilities [57]. Governance needs a control plane: agent identity and persona registries, least-privilege tool access, continuous logging of actions and intermediate reasoning, human authorization gates for consequential acts, and life cycle management with decommissioning tied to credential revocation [58]. Frameworks should mandate agentic risk committees that include domain-AI bilingual clinicians and a health workforce empowered to pause or redesign workflows when energy, privacy, or equity thresholds are breached [59]. These controls keep domain experts in meaningful oversight of systems whose autonomy amplifies benefit and harm [58].

Recommended Actions for Educators, Institutions, and Policymakers

Curriculum Integration and Faculty Development Programs

Overview

Educators play a pivotal role in embedding AI fluency into health professions training without displacing core clinical and public health curricula. Rather than teaching programming, programs should integrate modular, foundational AI literacy courses that explain the full AI system stack, model limitations (such as hallucinations and bias), and critical appraisal techniques tailored to health care contexts. Medical, nursing, and public health schools can adopt case-based learning where students evaluate real-world AI tools for diagnostic accuracy, ethical risks, and implications. For instance, simulation exercises using deidentified datasets allow learners to practice "bilingual" translation—articulating clinical needs to hypothetical AI developers—preparing them for collaborative practice. Faculty development is equally essential; institutions should offer accessible workshops similar to those implemented at Mayo Clinic and other centers to help educators themselves achieve AI fluency and redesign curricula progressively [60].

Concrete Competency Matrices for Modular AI Literacy Courses

Modular foundational AI literacy needs concrete competency matrices to fit dense curricula. Medical and public health programs require explicit minimum thresholds by training stage [61]. Drawing on expert consensus and frameworks [61-63], a progressive structure follows.

Training for third-year medical students (preclinical and early clinical stage) should focus on foundational awareness: knowledge (explain the 5-layer AI system stack conceptually, recognize common failures [hallucination, bias, and data shift], and describe basic evaluation [accuracy, fairness, and explainability]), skills (appraise a published AI report or demonstration using a simple checklist and communicate AI use and limits to patients in simulation), and attitudes (maintain human oversight and ethical awareness). Assessment should be via short standardized tests and situational judgment items. No coding or hardware procurement is required.

Training for specialist residents should focus on all prior competencies plus deeper application: integrating AI outputs while documenting independent reasoning, evaluating hardware and energy feasibility for specialty tools, joining local validation or bias audits, and teaching juniors. Technical use includes operating point-of-care devices and recognizing when infrastructure invalidates results. Assessment should be via objective structured clinical examination stations with AI scenarios, critical appraisal portfolios, observed clinical integration, and simulated governance review of an agentic workflow flagging automation bias risks [64].

Awakening embeds the student matrix in existing systems-based courses rather than new modules. Intrepidity uses low-stakes simulation laboratories with offline models. Bilingualism grows when residents colead AI translation rounds with data scientists. Visionary school leadership maps competencies to accreditation standards and gives faculty protected time to reach the resident-level matrix first. Pilots in multiple countries show that graduated thresholds fit dense curricula as modular blocks across existing rotations [65-67].

Organizational Strategies

Institutions must move beyond pilot projects to systematic organizational readiness. Hospitals and public health agencies should establish interdisciplinary AI governance committees that include domain experts at every decision level, ensuring that the AI system stack is evaluated holistically—from

XSL•FO
RenderX

sustainable energy sourcing and compliant infrastructure to bias-mitigated models and user-centered applications. Practical steps include conducting AI maturity assessments, investing in secure hybrid infrastructures that respect region or country AI laws, and creating internal "AI champion" networks drawn from intrepid early adopters. Leadership should also prioritize metrics for success, such as reduced clinician burnout through effective administrative AI and improved population health outcomes. Visionary institutions such as Cleveland Clinic have demonstrated that dedicated AI centers combined with clear ethical frameworks accelerate responsible adoption while protecting core human elements of care [68].

Policy Recommendations

Policymakers and accrediting bodies need to formalize AI fluency as a core competency. This includes updating medical and public health education accreditation standards to require foundational AI knowledge; supporting national AI health literacy frameworks; and incentivizing sustainable, equitable AI deployment through funding and regulatory guidance. The EU AI Act provides a strong model by classifying health care applications as high risk and mandating transparency, oversight, and risk management; similar approaches adapted locally can prevent fragmented adoption [69,70]. Policymakers should also address global disparities by funding open educational resources, cross-border collaboration platforms, and infrastructure grants for low-resource settings. By aligning incentives—such as tying reimbursement or certification to demonstrated AI governance—policymakers can accelerate the cultural shift toward domain knowledge combined with AI fluency as the new nonnegotiable standard, ultimately ensuring that AI amplifies rather than undermines safe, equitable, and human-centered health care.

Conclusions

The AI boom will amplify health care and public health impact only when domain experts actively drive the use of these technologies—ensuring safety, clinical relevance, and human-centered outcomes—rather than allowing technology to lead unchecked.

The modern AI system stack, from energy foundations through accelerators, infrastructure, models, and applications, provides a practical framework for health professionals to understand, evaluate, and oversee AI tools without needing to become coders or engineers. By cultivating awakening, intrepidity, bilingualism, and visionary leadership, students, practitioners, educators, and institutions can bridge the current literacy gaps and build the nonnegotiable hybrid competency required for the AI era. Educators must integrate foundational AI concepts into curricula, institutions should establish robust governance and collaborative structures, and policymakers need to formalize fluency requirements and equitable access frameworks. Addressing barriers such as time constraints, fear, resource disparities, and regulatory complexity through targeted strategies will accelerate this transformation and prevent fragmented or inequitable adoption.

Looking forward, the rise of agentic AI—autonomous systems capable of planning, reasoning, and executing multistep tasks—promises profound advancements in personalized medicine. Agentic AI could integrate multimodal patient data, real-time sensor inputs from wearables, and population-level insights to generate dynamic, adaptive care plans tailored to an individual's genetics, lifestyle, environment, and evolving health status. In public health, these systems could simulate policy scenarios at scale, optimizing resource allocation during outbreaks or chronic disease prevention campaigns. However, realizing this vision demands vigilant attention to disparity: without deliberate design, agentic AI risks exacerbating disparities if trained predominantly on data from high-resource settings, deployed unevenly across regions, or governed without diverse domain expertise. Energy sustainability, robust privacy protections across the full AI stack, and inclusive governance will determine whether these technologies narrow or widen global health gaps. Visionary leaders and bilingual professionals must therefore champion ethical guardrails, continuous oversight, and human-in-the-loop frameworks that preserve compassion, autonomy, and trust at the heart of care. Ultimately, the future of health and public health lies not in replacing human expertise with AI but in empowering domain-fluent professionals to direct AI tools toward more precise and humane outcomes for all.

Funding

The authors declared no financial support was received for this work.

Generative AI Use

The authors declare that generative AI was not used in the preparation of this manuscript.

Authors' Contributions

Conceptualization: SF
Writing—original draft: SF
Writing—review and editing: SF, YV

Conflicts of Interest

None declared.

References

XSL•FO
RenderX

1. Huang J. AI is a 5-layer cake. NVIDIA. Mar 10, 2026. URL: https://blogs.nvidia.com/blog/ai-5-layer-cake/ [accessed 2026-03-13]
2. Frimpong V. Artificial intelligence investment in resource-constrained African economies: financial, strategic, and ethical trade-offs with broader implications. World. May 20, 2025;6(2):70. [doi: 10.3390/world6020070]
3. Klonoff DC, Scheideman AF, Shao MM, Ho CN, Ayers AT, Jendle J, et al. The need for clear medical data ownership laws. J Transl Med. Jan 03, 2026;24(1):5. [FREE Full text] [doi: 10.1186/s12967-025-07486-z] [Medline: 41484623]
4. 2026 healthcare cybersecurity and AI benchmarking study executive summary. Censinet. URL: https://www.censinet.com/2026-healthcare-cybersecurity-ai-benchmarking-study-executive-summary [accessed 2026-05-02]
5. Artificial intelligence in healthcare. European Commission. URL: https://health.ec.europa.eu/ehealth-digital-health-and-care/artificial-intelligence-healthcare_en [accessed 2026-05-02]
6. Hamsal M, Binsar F. Toward a new era of healthcare services: the role of artificial intelligence in shaping tomorrow's landscape. Cogent Bus Manag. Oct 01, 2025;12(1):2566441. [doi: 10.1080/23311975.2025.2566441]
7. Ahmed MI, Spooner B, Isherwood J, Lane M, Orrock E, Dennison A. A systematic review of the barriers to the implementation of artificial intelligence in healthcare. Cureus. Oct 04, 2023;15(10):e46454. [FREE Full text] [doi: 10.7759/cureus.46454] [Medline: 37927664]
8. Why AI uses so much energy — and what we can do about it. Penn State Institute of Energy and the Environment. URL: https://iee.psu.edu/news/blog/why-ai-uses-so-much-energy-and-what-we-can-do-about-it [accessed 2026-05-04]
9. Campos H, Salmi L. Critical AI health literacy as liberation technology: a new skill for patient empowerment. NAM Perspect. Dec 8, 2025;2025:10.31478/202512a. [doi: 10.31478/202512a] [Medline: 41970622]
10. Chatzichristos C, Chatzichristos G, Borremans I, Gruyaert S, De Vos I, De Vos M, et al. Bridging the AI-literacy gap in health care: qualitative analysis of the Flanders case study. J Med Internet Res. Dec 08, 2025;27:e76709. [FREE Full text] [doi: 10.2196/76709] [Medline: 41359863]
11. AI accelerate 2025: Synapxe boosts AI capabilities to transform HealthTech in Singapore. Synapxe. Jun 16, 2025. URL: https://www.synapxe.sg/media-releases/artificial-intelligence/ai-accelerate-2025 [accessed 2026-05-08]
12. Amadi C, Ojo A. Building trustworthy AI in healthcare. IEEE Access. Dec 25, 2025;14:1182-1212. [doi: 10.1109/access.2025.3648410]
13. Dolezel D, Lambert E, Watzlaf V, Morton M, Lalani K, Sand J, et al. Assessing artificial intelligence literacy: what do health information professionals know about AI? Adv Health Inf Sci Pract. Nov 3, 2025;1(2):KLVV3078. [FREE Full text] [doi: 10.63116/KLVV3078] [Medline: 41333329]
14. Frehywot S, Vovides Y. Artificial intelligence algorithmic literacy: gaining and deepening the artificial intelligence knowledge of global health workforce education in the Fifth Industrial Revolution. AIH. Dec 05, 2025;3(2):025420090. [doi: 10.36922/AIH025420090]
15. Mayo Clinic Harper Family Foundation artificial intelligence education in medicine program. Mayo Clinic. URL: https://education.mayo.edu/mayo-clinic-harper-family-foundation-artificial-intelligence-education-in-medicine-program/ [accessed 2026-05-05]
16. Wei Q, Pan S, Liu X, Hong M, Nong C, Zhang W. The integration of AI in nursing: addressing current applications, challenges, and future directions. Front Med (Lausanne). Feb 11, 2025;12:1545420. [FREE Full text] [doi: 10.3389/fmed.2025.1545420] [Medline: 40007584]
17. Chustecki M. Benefits and risks of AI in health care: narrative review. Interact J Med Res. Nov 18, 2024;13:e53616. [FREE Full text] [doi: 10.2196/53616] [Medline: 39556817]
18. Unleashing the power of AI in primary care. Stanford Medicine. URL: https://domannualreports.stanford.edu/unleashing-the-power-of-ai-in-primary-care/ [accessed 2026-05-08]
19. New frontiers in team science: empowering patients with AI-driven e-consults. Stanford Medicine. URL: https://domannualreports.stanford.edu/new-frontiers-in-team-science-empowering-patients-with-ai-driven-e-consults/ [accessed 2026-05-08]
20. Topol digital fellowships. NHS. URL: https://topol.digitalacademy.nhs.uk/digital-fellowships [accessed 2026-05-05]
21. Gazquez-Garcia J, Sánchez-Bocanegra CL, Sevillano JL. AI in the health sector: systematic review of key skills for future health professionals. JMIR Med Educ. Feb 05, 2025;11:e58161. [FREE Full text] [doi: 10.2196/58161] [Medline: 39912237]
22. Frehywot S, Vovides Y. Contextualizing algorithmic literacy framework for global health workforce education. AIH. Nov 28, 2024;2(2):41-46. [doi: 10.36922/aih.4903]
23. Bryant M. Cleveland Clinic creating center for AI in healthcare. Healthcare Dive. Mar 14, 2019. URL: https://www.healthcaredive.com/news/cleveland-clinic-creating-center-for-ai-in-healthcare/550370/ [accessed 2026-05-05]
24. Bean R. How Cleveland Clinic is innovating in healthcare with data, analytics, and AI. Forbes. Oct 06, 2024. URL: https://www.forbes.com/sites/randybean/2024/10/06/how-cleveland-clinic-is-innovating-in-healthcare-with-data-analytics-and-ai/ [accessed 2026-05-05]
25. Muralidharan V, Ng MY, AlSalamah S, Pujari S, Kalra K, Singh R, et al. Global Initiative on AI for Health (GI-AI4H): strategic priorities advancing governance across the United Nations. NPJ Digit Med. Apr 23, 2025;8(1):219. [FREE Full text] [doi: 10.1038/s41746-025-01618-x] [Medline: 40269244]

XSL•FO
RenderX

26. Harnessing artificial intelligence for health. World Health Organization. URL: https://www.who.int/teams/digital-health-and-innovation/harnessing-artificial-intelligence-for-health [accessed 2026-05-06]
27. Sriharan A, Kuhlmann E, Correia T, Tahzib F, Czabanowska K, Ungureanu MI, et al. Artificial intelligence in healthcare: balancing technological innovation with health and care workforce priorities. Int J Health Plann Manage. Jul 2025;40(4):987-992. [doi: 10.1002/hpm.3927] [Medline: 40075189]
28. Faveri B, Shank C, Whitt R, Dawson P. The need for and pathways to AI regulatory and technical interoperability. Tech Policy Press. Apr 16, 2025. URL: https://www.techpolicy.press/the-need-for-and-pathways-to-ai-regulatory-and-technical-interoperability/ [accessed 2026-05-06]
29. AI Act. European Commission. URL: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai [accessed 2026-05-06]
30. Anwar S. The global health workforce crisis: are task-shifting strategies sustainable? BMJ Glob Health. Jun 16, 2025;10(6):e019264. [FREE Full text] [doi: 10.1136/bmjgh-2025-019264] [Medline: 40527528]
31. Avordeh TK, Salifu A, Quaidoo C, Opare-Boateng R. Impact of power outages: unveiling their influence on micro, small, and medium-sized enterprises and poverty in Sub-Saharan Africa - an in-depth literature review. Heliyon. Jun 27, 2024;10(13):e33782. [FREE Full text] [doi: 10.1016/j.heliyon.2024.e33782] [Medline: 39040250]
32. Al-Ganad A, Al-Shahdhi A, Al-Dhaifi O, Hajeb E, Hajeb H, Al-Motarreb A. Deploying medical AI in low-resource settings: a scoping review of challenges and strategies. Front Digit Health. Apr 1, 2026;8:1743634. [FREE Full text] [doi: 10.3389/fdgth.2026.1743634] [Medline: 42022499]
33. Lima HA, Trocoli-Couto PH, Zaman M, Engelmann DC, Parkes-Ratanshi R, Junck L, et al. Spotlighting healthcare frontline workers´ perceptions on artificial intelligence across the globe. Npj Health Syst. 2025;2(1):28. [doi: 10.1038/s44401-025-00034-3] [Medline: 40757093]
34. Chen H, Zeng D, Qin Y, Fan Z, Ng Yu Ci F, Klonoff DC, et al. Large language models and global health equity: a roadmap for equitable adoption in LMICs. Lancet Reg Health West Pac. Oct 18, 2025;63:101707. [FREE Full text] [doi: 10.1016/j.lanwpc.2025.101707] [Medline: 41158957]
35. Dennstädt F, Hastings J, Putora PM, Schmerder M, Cihoric N. Implementing large language models in healthcare while balancing control, collaboration, costs and security. NPJ Digit Med. Mar 06, 2025;8(1):143. [FREE Full text] [doi: 10.1038/s41746-025-01476-7] [Medline: 40050366]
36. Chen X, Wang X, Colacelli A, Lee M, Xie L. Electricity demand and grid impacts of AI data centers: challenges and prospects. arXiv. Preprint posted online on September 8, 2025. [doi: 10.48550/arXiv.2509.07218]
37. Vinci S, Kwong LH, Miles SB, McCord R, Cady G, Tang K, et al. Reliable electricity to advance quality health care in low-resource settings. Lancet Glob Health. Apr 2026;14(4):e626-e639. [FREE Full text] [doi: 10.1016/S2214-109X(25)00537-6] [Medline: 41856145]
38. Chen RJ, Wang JJ, Williamson DF, Chen TY, Lipkova J, Lu MY, et al. Algorithmic fairness in artificial intelligence for medicine and healthcare. Nat Biomed Eng. Jun 2023;7(6):719-742. [FREE Full text] [doi: 10.1038/s41551-023-01056-8] [Medline: 37380750]
39. Suwannaphong T, Jovan F, Craddock I, McConville R. Optimising TinyML with quantization and distillation of transformer and Mamba models for indoor localisation on edge devices. Sci Rep. Mar 24, 2025;15(1):10081. [FREE Full text] [doi: 10.1038/s41598-025-94205-9] [Medline: 40128553]
40. Samson HH. Lightweight transformer architectures for edge devices in real-time applications. arXiv. Preprint posted online on January 5, 2026. [doi: 10.48550/arXiv.2601.03290]
41. Bhushan CM, Koppuravuri P, Prasanthi N, Gazi F, Hussain MM, Abdussami M, et al. Deploying TinyML for energy-efficient object detection and communication in low-power edge AI systems. Sci Rep. Dec 05, 2025;15(1):44299. [FREE Full text] [doi: 10.1038/s41598-025-27818-9] [Medline: 41350337]
42. Terespolsky S, Yassi A, Ehrlich R, Bruton J, Lockhart K, Wang H, et al. The case for local AI development: lessons from computer-aided detection of tuberculosis and silicosis in Southern Africa's ex-miners. Ann Glob Health. Feb 12, 2026;92(1):16. [FREE Full text] [doi: 10.5334/aogh.5064] [Medline: 41694799]
43. Sharma RS, Rohatgi A, Jain S, Singh D. The Ayushman Bharat Digital Mission (ABDM): making of India’s Digital Health Story. CSIT. Mar 31, 2023;11:3-9. [doi: 10.1007/s40012-023-00375-0]
44. Giri L, Regmi PR, Gurung G, Gurung G, Aryal S, Mandal S, et al. Diagnostic accuracy of an offline CNN framework utilizing multi-view chest x-rays for screening 14 co-occurring communicable and non-communicable diseases. Diagnostics (Basel). Dec 24, 2025;16(1):66. [FREE Full text] [doi: 10.3390/diagnostics16010066] [Medline: 41515561]
45. Kumar A. Integrating artificial intelligence into resource-constrained clinics in South Asia. Asian J Med Sci. Jul 31, 2026;17(8):2-3. [doi: 10.71152/ajms.v17i8.5492]
46. Sivasubramanian S, Raval A. Artificial intelligence-augmented public health interventions in India. Health Aff Sch. May 07, 2025;3(5):qxaf097. [FREE Full text] [doi: 10.1093/haschl/qxaf097] [Medline: 40443397]
47. Xie Y, Fang Q. An energy-aware generative AI edge inference framework for low-power IoT devices. Electronics. Oct 17, 2025;14(20):4086. [doi: 10.3390/electronics14204086]
48. Steve TS, Yeluripati GR. LiteGuard: efficient offline language model for resource-constrained healthcare. In: 2025 Conference on Information Communications Technology and Society (ICTAS). New York, NY. IEEE; 2025.

XSL•FO
RenderX

49. Chettri SK, Deka RK, Saikia MJ. Bridging the gap in the adoption of trustworthy AI in Indian healthcare: challenges and opportunities. AI. Jan 13, 2025;6(1):10. [doi: 10.3390/ai6010010]
50. Olawade DB, Clement David-Olawade A, Gore MN, Wada OZ. Perceptions and challenges of artificial intelligence adoption in Nigerian public healthcare: insights from consultant doctors across five tertiary hospitals. Clin Epidemiol Glob Health. 2025;36:102200. [doi: 10.1016/j.cegh.2025.102200]
51. Nisa U, Shirazi M, Saip MA, Pozi MS. Agentic AI: the age of reasoning—a review. J Autom Intell. Mar 2026;5(1):69-89. [doi: 10.1016/j.jai.2025.08.003]
52. Heudel PE, Crochet H, Filori Q, Bachelot T, Blay J. Artificial intelligence in medicine: a scoping review of the risk of deskilling and loss of expertise among physicians. ESMO Real World Data Digit Oncol. Mar 19, 2026;12:100693. [FREE Full text] [doi: 10.1016/j.esmorw.2026.100693] [Medline: 41890350]
53. Natali C, Marconi L, Dias Duran LD, Cabitza F. AI-induced deskilling in medicine: a mixed-method review and research agenda for healthcare and beyond. Artif Intell Rev. Aug 27, 2025;58:356. [doi: 10.1007/s10462-025-11352-1]
54. Ke Y, Jin L, Ong JC, Thirunavukarasu AJ, Car J, Cheung CY, et al. AI-induced never-skilling in medical education. Nat Med. Jun 2026;32(6):1997-2006. [doi: 10.1038/s41591-026-04438-y] [Medline: 42174254]
55. Lazaros K, Vrahatis AG, Kotsiantis S. Human-in-the-loop artificial intelligence: a systematic review of concepts, methods, and applications. Entropy (Basel). Mar 26, 2026;28(4):377. [FREE Full text] [doi: 10.3390/e28040377] [Medline: 42072503]
56. Xie Z, Wang H, Dai L, Wang Z, Song H, Qian J. Ethical issues in multi-agent AI systems for healthcare: a narrative review. Front Public Health. Apr 28, 2026;14:1792627. [FREE Full text] [doi: 10.3389/fpubh.2026.1792627] [Medline: 42130876]
57. Nguyen TT, Fournier H, Jackson P, Itoh M, Freeman S, Richard R, et al. Position: multi-agent algorithmic care systems demand contestability for trustworthy AI. arXiv. Preprint posted online on March 21, 2026. [doi: 10.48550/arXiv.2603.20595]
58. Prakash C, Lind M, Sisodia A. Agentic AI governance and lifecycle management in healthcare. arXiv. Preprint posted online on January 22, 2026. [doi: 10.48550/arXiv.2601.15630]
59. Zhu L, Lu Q, Ding M, Lee SU, Wang C. Designing meaningful human oversight in AI. AI Ethics. May 04, 2026;6:286. [doi: 10.1007/s43681-026-01147-7]
60. Zand J, Suvakov M, Overgaard SM, Kozikowski CG, Jimison RC, Nelson KB, et al. Building artificial intelligence and data literacy across the health care research workforce in the Mayo Clinic Center for Clinical and Translational Science: why it matters and what we are doing about it. Mayo Clin Proc. Jan 2026;101(1):17-25. [doi: 10.1016/j.mayocp.2025.11.008] [Medline: 41485858]
61. Liaw W, Kueper JK, Lin S, Bazemore A, Kakadiaris I. Competencies for the use of artificial intelligence in primary care. Ann Fam Med. 2022;20(6):559-563. [FREE Full text] [doi: 10.1370/afm.2887] [Medline: 36443071]
62. Gong M, Li J, Ma Y, Jin B, Chen W, Hou Y, et al. A Chinese expert consensus on the artificial intelligence proficiency of medical students: competencies and the multi-modal assessment. Health Care Sci. Feb 17, 2026;5(1):49-57. [doi: 10.1002/hcs2.70049] [Medline: 41767172]
63. Gong MC, Pan H, Liu H, Tong K, Li J, Ma YH, et al. Expert consensus on the artificial intelligence proficiency competency list and assessment framework for medical students (2025 edition). Zhongguo Yi Xue Ke Xue Yuan Xue Bao. Feb 2026;48(1):13-23. [FREE Full text] [doi: 10.3881/j.issn.1000-503X.17210] [Medline: 41645431]
64. Zafar I, Aboueisha H, Elhassan I, Shersad F, Caliskan SA, Syeda AF, et al. Ten tips for utilizing AI to generate high quality OSCE stations in medical education. Front Med (Lausanne). Jan 16, 2025;12:1744657. [FREE Full text] [doi: 10.3389/fmed.2025.1744657] [Medline: 41625731]
65. Park TI, Seo J, Yoon HJ, Lee KE. Institutionalizing convergence education for medical artificial intelligence. Biomed Eng Lett. Nov 01, 2025;15(6):1073-1083. [doi: 10.1007/s13534-025-00523-2] [Medline: 41280148]
66. Levingston H, Anderson MC, Roni MA. From theory to practice: artificial intelligence (AI) literacy course for first-year medical students. Cureus. Oct 02, 2024;16(10):e70706. [doi: 10.7759/cureus.70706] [Medline: 39493023]
67. Zheng L, Xiao Y. Refining AI perspectives: assessing the impact of ai curricular on medical students' attitudes towards artificial intelligence. BMC Med Educ. Jul 25, 2025;25(1):1115. [FREE Full text] [doi: 10.1186/s12909-025-07669-8] [Medline: 40713579]
68. Siwicki B. Cleveland Clinic's advice for AI success: democratizing innovation, upskilling talent and more. Healthcare IT News. Mar 22, 2024. URL: https://www.healthcareitnews.com/news/cleveland-clinics-advice-ai-success-democratizing-innovation-upskilling-talent-and-more [accessed 2026-05-06]
69. van Kolfschooten H, van Oirschot J. The EU Artificial Intelligence Act (2024): implications for healthcare. Health Policy. Nov 2024;149:105152. [FREE Full text] [doi: 10.1016/j.healthpol.2024.105152] [Medline: 39244818]
70. Aboy M, Minssen T, Vayena E. Navigating the EU AI Act: implications for regulated digital medical products. NPJ Digit Med. Sep 06, 2024;7(1):237. [FREE Full text] [doi: 10.1038/s41746-024-01232-3] [Medline: 39242831]

Abbreviations

EU: European Union
GDPR: General Data Protection Regulation
GPU: graphics processing unit

XSL•FO
RenderX

HIPAA: Health Insurance Portability and Accountability Act
ICU: intensive care unit
LLM: large language model
LMIC: low- or middle-income country
NPU: neural processing unit
TPU: tensor processing unit

Edited by F Prasser; submitted 10.May.2026; peer-reviewed by F Binsar, FA Causio; comments to author 14.Aug.2026; revised version received 01.Sep.2026; accepted 15.Oct.2026; published 30.Oct.2026

Please cite as:
Frehywot S, Vovides Y
Domain Knowledge Plus AI Fluency: The New Nonnegotiable Competency for the Health Workforce in the AI Era
JMIR AI 2026;5:e100954
URL: https://ai.jmir.org/2026/1/e100954
doi: 10.2196/100954
PMID:

©Seble Frehywot, Yianna Vovides. Originally published in JMIR AI (https://ai.jmir.org), 30.Oct.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR AI, is properly cited. The complete bibliographic information, a link to the original publication on https://www.ai.jmir.org/, as well as this copyright and license information must be included.

XSL•FO
RenderX