INNOVAPATH™
The Premier Journal Bridging Discovery & Innovation

Commentary
https://doi.org/10.63501/7dcz6776

Reimagining Health-Professions Education: A Roadmap for the AI University of Health and Medicine

Sohail Rao, MD, MA, DPhil[1]

[1]Texas Institute for Graduate Medical Education, 1917 Ashland Drive, Houston, TX 77008, USA.

[1]Corresponding Author Email: srao@tigme.org

KEYWORDS: *Artificial Intelligence, AI, Health Sciences, Medicine*

A few years ago, many of us in academic medicine still treated artificial intelligence (AI) as an interesting but distant frontier. Today, it is at our desks, in our students' pockets, and increasingly at the bedside. Large language models can draft essays, exam questions, and grant proposals. Multimodal systems can read radiographs, segment pathology slides, summarize guidelines, and generate patient-facing education with a single click.

Higher education is scrambling to respond. A recent essay in The Chronicle of Higher Education asked, "Are You Ready for the AI University? Everything is About to Change," and argued that AI will reconfigure nearly every function of the traditional university, from recruitment to teaching to assessment. Medical and health sciences institutions are not exempt; if anything, we sit at the epicenter of this transformation.

In this article, I want to move beyond the usual debate "ban it" versus "embrace it" and ask a deeper question: *What kind of AI-enabled university do we actually want to build for health and medicine?* Because if we do not answer that question, someone else, usually a vendor, an algorithm, or an accountant, will answer it for us.

Universities have already lived through one revolution: the digital turn. Learning-management systems replaced paper course packs; online libraries displaced card catalogues; "clickers" and Zoom became fixtures of the classroom. Yet the core logic of the university remained recognizable: instructors produced and delivered content, students consumed and reproduced it, and degrees certified that this transfer had occurred.

AI challenges logic at multiple levels. First, content is now abundant and on demand: a motivated student can ask a model to explain the Krebs cycle at a 5th-grade reading level, generate flashcards for USMLE Step 1, or walk through a simulated OSCE station in seconds. Second, assessment can be faked or automated, as essays, discussion posts, and even structured reflections can be generated by machines. At the same time, AI tools can also grade short answers, code OSCE checklists, and flag anomalies in assessment data. Third, professional boundaries are blurring as diagnostic support systems, chatbots, and autonomous agents increasingly handle tasks once considered the exclusive domain of physicians, pharmacists, or nurses.

https://doi.org/10.63501/7dcz6776

If knowledge can be generated on demand and many routine cognitive tasks can be automated, then the traditional "lecture-memorize-examine" model looks increasingly fragile. The AI university, especially in health and medicine, cannot simply be the old university plus a chatbot.

Suggested Design of an AI University of Health & Medicine:

Before we surrender everything to algorithms, we need to be very clear about what must remain human in health-professional education. First, moral and professional formation cannot be automated: no model, no matter how sophisticated, can internalize professional values on a student's behalf. Integrity, empathy, humility, and accountability are formed through relationships, example, and lived experience, not downloaded from a server. Second, judgment under uncertainty must remain a human skill. Clinical life is full of incomplete data, conflicting guidelines, and fragile human beings; AI can calculate probabilities, but it does not bear consequences. Teaching students how to decide, not merely what to choose, is irreducibly human. Third, the therapeutic relationship cannot be ceded to machines. Patients do not simply want answers; they want to be heard, seen, and accompanied. AI may support communication, but it cannot replace the trust that comes from human presence. Finally, curiosity and dissent must be preserved at the core of universities. Their role is to cultivate the capacity to question received wisdom, including the outputs of AI systems themselves. If we begin to accept model outputs as infallible, we abdicate our critical role as educators, clinicians, and scholars.

In designing the AI university of health and medicine, we must therefore treat these domains as non-automatable. AI may assist; it must not displace. If AI will handle much of the "first draft" work, summaries, outlines, and differential diagnoses, then our curricula should pivot from transmission to transformation. The goal should no longer be to deliver information, but to shape how learners think with and beyond AI. These non-automatable domains and their implications for curriculum design are summarized in Table 1.

First, we need to move from content recall to problem framing. Instead of asking, "Can you list the causes of metabolic acidosis?", we might ask, "Given this AI-generated differential, what is missing, what is overemphasized, and what additional history, examination, or investigations would you seek?" The expectation shifts from reproducing lists to interrogating, refining, and contextualizing AI outputs. Students must learn to critique AI, not merely use it.

Second, we must shift from solitary assignments to a transparent process. We should assume AI will be used and design assessments that require students to show their process: prompts, iterations, and reasoning. Viva voce examinations, OSCEs with reflective debriefs, and collaborative case conferences can all be structured to reveal how students think with and against AI, rather than merely presenting polished end products that may or may not be their own work.

Third, curricula should evolve from siloed courses to integrated, data-rich experiences. AI makes it easier to simulate complex, longitudinal cases that cut across disciplines, such as pathology, pharmacology, ethics, and health systems. An AI-enabled curriculum can present evolving virtual patients whose data streams, imaging, and social context unfold over weeks, asking interprofessional teams to manage care together and negotiate trade-offs in real time.

https://doi.org/10.63501/7dcz6776

Table 1: Non-automatable Domains in the AI University of Health and Medicine:

Domain	Why it Cannot Be Automated?	How AI may Assist (without replacing)	Implications for Education and Assessment
Moral and Professional Formation	Professional identity, integrity, empathy, humility, and accountability are shaped through lived experience, role modeling, and relational trust, not pattern recognition.	AI can surface ethical dilemmas, provide scenario prompts, or simulate stakeholder perspectives for reflection.	Preserve mentored clinical exposure, longitudinal faculty–student relationships, and reflective practice as core requirements.
Judgment Under Uncertainty	Clinical life involves incomplete data, conflicting evidence, and trade-offs with real consequences for patients and families; machines do not bear moral responsibility.	AI can provide probabilities, differential diagnoses, and guideline summaries as inputs to human deliberation.	Design assessments (OSCEs, viva voce, case conferences) that explicitly test reasoning under ambiguity and justified overrides of AI.
Therapeutic Relationship	Trust, rapport, and healing presence depend on nonverbal cues, empathy, and human connection that cannot be reduced to algorithms.	AI may support communication (e.g., translation, information tailoring) and documentation to free time for human interaction.	Emphasize direct patient contact, communication skills training, and observed encounters as non-substitutable competency domains.
Curiosity and Dissent	Universities exist to question received wisdom, challenge authority, and generate new knowledge even when it contradicts prevailing models (including AI).	AI can generate alternative perspectives, literature maps, and hypotheses that learners critically appraise.	Assess critical appraisal of AI outputs, encourage student-led inquiry, and protect space for dissenting views and whistleblowing.

Finally, learning spaces should move from passive lecture halls to augmented studios. Classrooms should become environments where human expertise is leveraged to address what machines cannot do: negotiating value conflicts, practicing difficult conversations, rehearsing team dynamics, and analyzing real-world failures. The "content" can be outsourced to AI; the formation cannot.

Perhaps the most significant anxiety triggered by the AI university is among faculty. If AI can generate slide decks, lecture scripts, test banks, and even feedback comments, it is tempting for institutions to view faculty primarily as "cost centers" ripe for automation. This would be a grave mistake. In health and medicine, faculty are not merely content providers. They are role models in navigating ambiguity, ethical tension, resource constraints, and interprofessional collaboration. If we reduce them to content-delivery nodes, then yes, AI will be cheaper and faster. But in doing so, we will have hollowed out the very heart of our profession. Table 2 summarizes these pedagogical shifts and offers concrete examples for health-professional curricula.

Instead, we should redefine academic roles. Faculty can serve as clinical and intellectual coaches, curating AI tools, modeling their responsible use, and helping learners interpret outputs in context. They can act as guardians of standards, leading efforts to validate AI-assisted assessments and ensuring that competencies remain meaningful and not diluted by automation. They can also function as co-investigators with AI, using these tools to generate hypotheses, analyze data, and design translational research that would otherwise be impossible within existing time and resource constraints.

https://doi.org/10.63501/7dcz6776

Table 2: Pedagogical Shifts in an AI-rich Curriculum:

Traditional Focus	AI-enabled Shift	Example in Health/Medical Education	Assessment Implications
Content Recall (e.g., "List causes of X.")	Problem framing and critique of AI outputs	Student receives an AI-generated differential for metabolic acidosis and must identify missing items, overemphasis, and next steps.	Mark quality of critique, justification of additions/omissions, and proposed diagnostic plan.
Solitary, Product-focused Assignments	Transparent process and reasoning with/against AI	Written case analysis submitted alongside AI prompt history, model outputs, and annotations of where the student agreed/disagreed.	Evaluate reasoning trace, decision points, and handling of AI errors, not just final essay.
Siloed Discipline-based Courses	Integrated, data-rich longitudinal experiences	Virtual patient evolving over weeks, linking pathology, pharmacology, ethics, health systems, and interprofessional teamwork.	Use multi-station OSCEs and team-based assessments tied to the same longitudinal case.
Passive Lectures and Note-taking	Augmented studios for practice and reflection	In-class sessions focused on role-plays (e.g., breaking bad news), team huddles, and debriefs using AI-generated complex scenarios.	Direct observation, structured feedback, and reflective writing on interpersonal and team skills.
Faculty as Primary Content Deliverers	Faculty as coaches and curators of AI-enhanced learning	Instructor demonstrates how to select, query, and cross-check clinical AI tools during bedside teaching.	Assess learners on safe, appropriate use of AI tools during supervised clinical encounters.

Rather than asking, "How many lectures can AI replace?" we should ask, "How can AI free faculty to do the uniquely human work that only they can do?" In that reframing lies the possibility of an AI-enabled academy that enhances, rather than erodes, the role of the health professions educator.

The AI university also risks deepening existing inequities between institutions, countries, and students. Well-resourced universities will have access to customized models, integrated data lakes, and robust governance structures. At the same time, underfunded institutions may be forced to rely on generic, opaque systems with limited localization and oversight. Students who own personal devices, have high-speed internet, and can afford paid AI subscriptions will benefit from tailored tutoring and continuous support. In contrast, those without such access are likely to fall even further behind. At a global level, institutions in the Global North may "export" AI-branded curricula to the Global South without adequately engaging local realities, languages, and health-system constraints, thereby reproducing existing asymmetries under the banner of innovation.

An AI-driven revolution in medical and health-professional education that widens the gap between urban and rural, rich and poor, or North and South will be a Pyrrhic victory. Ensuring that AI serves as a force for inclusion rather than exclusion will require deliberate, equity-focused strategies. These include the development of open-access, low-cost AI tools co-designed with partners in low- and middle-income countries; shared governance frameworks that ensure voices from under-resourced settings are represented when drafting institutional AI policies; and systematic research that evaluates who benefits from AI-based interventions in education and who is left out. Key equity risks at the institutional, student, curricular, and global levels, along with potential mitigation strategies, are outlined in Table 3.

https://doi.org/10.63501/7dcz6776

Table 3: Equity Risks and Mitigation Strategies in the AI University:

Level	Equity risk	Illustrative example	Mitigation strategies
Institution	Unequal access to advanced AI infrastructure	Flagship urban university runs custom models on institutional data; peripheral school uses generic tools with poor localization.	Shared regional consortia; open-source / low-cost models; national or philanthropic funding for under-resourced institutions.
Student	Digital divide in devices, connectivity, and subscriptions	Some students have laptops, high-speed internet, and paid AI tutors; others rely on shared devices and limited data.	Device-lending programs; on-campus AI access points; institutional licenses; designing tasks that do not penalize low access.
Curriculum	One-size-fits-all AI content ignoring diversity	Standard AI-generated cases assume Western norms and English language, marginalizing local cultures and health realities.	Co-design cases with diverse stakeholders; include local languages; require context adaptation as a graded learning outcome.
Global	North-South dependency and curricular "export"	Global North school sells "AI-ready" curricula to LMIC institutions with little adaptation to local health-system constraints.	Partnership models with shared governance; local faculty leadership; funding for indigenous AI and curriculum development.
Governance	Exclusion of under-resourced voices from policy decisions	AI policies set by elite universities and vendors with minimal representation from rural or low-income settings.	Multistakeholder governance (students, LMIC partners, community reps); transparency; public comment periods on AI policies.

Continuous Evaluation of AI-rich Learning Environment:

The "AI university" should not remain a speculative headline; it must become an empirically examined, ethically grounded reality. To achieve this, health-professional education needs a focused research agenda that interrogates both the promises and the risks of AI-rich learning environments. These priority domains, key questions, and illustrative methods are summarized in Table 4.

First, research must address learning outcomes and safety. Do AI-enhanced curricula actually produce better clinicians, scientists, and leaders, and how can this be measured without reducing education to a narrow set of test scores or metrics? Studies will need to link AI-enabled learning designs to real-world performance, patient outcomes, and professional behavior.

Second, there is a pressing need to investigate bias, transparency, and accountability in AI systems used for admissions, progression decisions, or exam proctoring. Work is needed to understand how these tools affect equity and fairness, what biases they introduce or amplify, and which mechanisms of appeal or review are appropriate when an algorithm's decision is contested.

Third, human-AI teaming in clinical training deserves scrutiny. Researchers should explore what an OSCE station looks like in which a student must decide when to trust an AI-generated triage recommendation and when to override it, and how best to train learners for such a relationship. This includes understanding how AI alters clinical reasoning, situational awareness, and shared decision-making.

https://doi.org/10.63501/7dcz6776

Table 4: Priority Research Domains for the AI Health University:

Domain	Key Questions	Example Methods/Study Designs	Potential Outcomes/Metrics
Learning Outcomes and Safety	Do AI-enhanced curricula produce better clinicians, scientists, and leaders? How do we evaluate benefit without over-reliance on test scores?	Randomized or quasi-experimental curriculum trials; longitudinal cohort studies; competency-based OSCE and workplace-based assessment data.	Clinical performance, patient outcomes, error rates, OSCE scores, progression and remediation patterns.
Bias, Transparency, and Accountability	How do AI tools in admissions, progression, and proctoring affect equity and fairness? How should appeals of algorithmic decisions work?	Audit studies; disparate impact analysis; policy and legal reviews; mixed-methods studies of candidate and faculty experience.	Differences in admission/progression by demographic group; documented appeals; perceived fairness/trust.
Human–AI Teaming in Clinical Training	How do learners decide when to trust or override AI recommendations? What competencies define effective human–AI teams?	Simulation and OSCE stations with embedded AI tools; think-aloud protocols; observational studies in real clinical settings.	Appropriateness of overrides; time to decision; teamwork ratings; safety and quality indicators.
Well-being and Professional Identity	How does learning with pervasive AI affect burnout, imposter syndrome, cynicism, or moral injury among students and faculty?	Cross-sectional and longitudinal surveys; qualitative interviews; diary studies; intervention trials (e.g., AI literacy programs).	Burnout indices; measures of professional identity formation; qualitative themes on meaning, autonomy, and trust.
Institutional Models and Governance	Which governance structures and incentive schemes support responsible AI adoption? Which experiments fail, and why?	Comparative case studies of institutions; policy analysis; realist evaluation of AI implementation projects.	Adoption patterns; sustainability of initiatives; incidents/near-misses; stakeholder satisfaction and trust.

Fourth, studies should examine well-being and professional identity in AI-saturated learning environments. Questions include how constant exposure to AI tools affects burnout, imposter syndrome, or cynicism among students and faculty; whether continual comparison to machine performance erodes confidence; and under what conditions AI can instead be used to support reflection, growth, and resilience.

Finally, there is a need to analyze institutional models for AI adoption in universities and teaching hospitals. This includes identifying governance structures, incentive schemes, and funding mechanisms that promote responsible, mission-aligned use of AI, as well as learning from failed experiments and understanding why they failed. Together, these lines of inquiry can help ensure that the emerging AI health university is not only technologically sophisticated but also just, humane, and worthy of the professions it serves.

A Call to Courageous Imagination:

The worst response we can have to the AI university is paralysis: a patchwork of bans, loopholes, and quietly tolerated workarounds. The second-worst is uncritical enthusiasm: outsourcing our thinking to the latest platform and calling it innovation. What we need instead is courageous imagination anchored in professional responsibility.

https://doi.org/10.63501/7dcz6776

That means the courage to admit that our existing models of education were already under strain before AI arrived. Courage to resist purely financial logics that treat faculty and students as interchangeable units of input and output. And courage to experiment, fail, learn, and iterate in public, sharing both successes and missteps so that others can benefit, rather than hiding behind polished narratives. AI will not wait for us to get comfortable. In health and medicine, everything is already changing at the bedside, in the lab, and in our students' phones. The real question is not whether an AI university will emerge, but whose values will shape it.

The hope is that health-professional schools, teaching hospitals, and scholarly communities will help ensure that the AI university we build deepens rather than diminishes our commitment to science, compassion, justice, and human dignity. That future is not inevitable; it is a choice. And the time to make it is now.

Acknowledgment of AI Assistance: AI-assisted tools were used only for grammar and language editing. All concepts, interpretations, and conclusions are entirely those of the author(s), who take full responsibility for the content.