

AI in healthcare education

Section 1: *Health education* in the age of AI



ELSEVIER

Advancing human progress together

Table of contents

Chapter 1. AI has landed in the field of health education

I. Data, statistics and evidence	5
II. A new landscape: Changes and windows of opportunity	6
III. The role of AI in personalized learning	8
IV. Doubts, fears and “healthy skepticism”	9

Chapter 2. AI and educational content: Trustworthy design

I. The design challenge: How AI integrates reliable data	10
II. What are hallucinations?	11
III. AI evaluation: Why it is so important	12

Chapter 3. AI and clinical thinking: “Destined to understand each other”

I. How AI influences what students learn	14
II. Critical thinking: The challenge and goal at the heart of the process	16

Contributors



Dr. Ximena Alvira

Clinical & Research Manager (EMEALA and APAC) at Elsevier



Professor Marcus Drake

Professor of Urology at Imperial College London; Editor of the *Medicine in a Day* series, published by Elsevier




Pippa Furey

Senior Lecturer in Paramedic Practice and Emergency Care at Nottingham Trent University; National AI Lead (Education) at the Royal College of Paramedics




David Game

Senior Vice President in Medical Education Partnerships at Elsevier



Raman Kaur

Vice President in Education Technology at Elsevier



Dr. Ruth Paterson

Associate Professor and Head of Nursing at Edinburgh Napier University; Chair of Innovation in the Education Strategic Policy Group at the Council of Deans of Health



Paul Stern

Senior Research and Policy Officer (Professional Standards Team) at the General Osteopathic Council (GOsC)



Gosia Warminska-Marczak

Director of Commercial Healthcare Education (DACH, Eastern Europe and France) at Elsevier



Sundeep Watkins

Education Adviser at the Chartered Society of Physiotherapy; Women's Health Physiotherapy Clinician & Educator



Dr. Philip Xiu

GP; Honorary Senior Lecturer at the University of Leeds; Assessment Editor-in-Chief and AI Board Chair for Elsevier ClinicalKey Student UK

Introduction

Artificial intelligence (AI) in healthcare is **already a reality** and part of the current training landscape for future healthcare professionals. According to various surveys and studies on the subject, more than **92% of students in these fields use AI tools daily**, and institutions and faculty have likewise incorporated them into their teaching plans and strategies. It is worth noting, however, that this rapid adoption should not be confused with a fully consolidated transformation. It still seems too early to speak of substantial changes in health education models, and the coexistence of skepticism among some faculty, alongside very early adoption among others, reflects a landscape that remains inconsistent and under construction.

It could be said that, both in the academic setting and in clinical practice, this technology has had a “warm reception,” although, given the speed with which it has burst onto the healthcare scene, many questions are still being worked out.

All the stakeholders involved in health education agree in highlighting its **advantages**: information synthesis and **immediacy** of access; elements that **facilitate learning** (through the creation of simulations, for example); the availability of **up-to-date**, evidence-based clinical content; and the ability to **personalize** the way students interact with learning materials.

But as the initial “whirlwind” of excitement has passed, challenges of great importance are beginning to emerge in this **new AI-education pairing**. Perhaps the most significant is the need to ensure that AI tools designed for healthcare students **integrate reliable data within a trustworthy design**. Central to all these challenges is the problem of detecting hallucinations — errors or fabrications generated by the AI systems themselves that, despite advances, continue to occur and require rigorous verification mechanisms to guarantee the quality of information students receive.

Aware of these challenges, and drawing on their experience so far, experts in the field have already begun to identify some of the **essential conditions for ensuring effective and safe educational AI**: solid medical and healthcare content, human oversight (indispensable and irreplaceable, both in learning and in subsequent clinical practice) and, above all, constant and consistent evaluation to keep content, quality and safety up to date.

But beyond faster, more intuitive, interactive and up-to-date learning, it is important to remember that the **main goal of implementing AI** in healthcare education is to **promote critical thinking** among students — which is likewise indispensable and irreplaceable. AI’s capabilities can facilitate this as it allows for generating **reflective suggestions**, asking **follow-up questions**, and carrying out **constant self-assessment**, encouraging students to become far more engaged in their own learning. In this sense, AI also represents an opportunity to redefine the role of faculty. By taking on repetitive or synthesis-related tasks, it frees up time and attention that can be redirected toward their true vocation — teaching and closely supporting students’ clinical reasoning. Ultimately, it is about strengthening clinical reasoning — the skill that truly builds experience and is the key to becoming a good healthcare professional.



Chapter 1

AI has landed in the field of health education

I. Data, statistics and evidence

The most recent and relevant research on what AI implementation has and will mean for everything related to the training of future healthcare professionals — across institutions, universities, teachers, students and librarians — makes one thing clear: AI has marked a before and an after in this field.

According to the Elsevier Research 2024* global study on feedback from medical and nursing students on the knowledge and use of AI tools, **76% of medical students and 67% of nursing students were able to name at least one AI product**. Although, at the time of the survey, notable misunderstanding was observed whereby many students confused AI tools with voice assistants, smart devices or robotics.

Both student groups believed that **AI brings value to the entire training program**, and when asked about their requirements — that is, what they demanded from these tools — the responses reflected clear expectations. Both medical and nursing students agreed on **three characteristics** as those they considered most important:

- Reliable sources
- Accuracy
- Ease of use

Meanwhile, data from the second edition of the Higher Education Policy Institute's (HEPI) *Student Generative AI Survey 2025* — produced in collaboration with Kortext (a digital university textbook platform that uses AI to support learning) — revealed that **92% of students used AI in some form during the 2024-25 academic year**, compared to 66% the previous year. In particular, using tools such as **ChatGPT for text generation** was found to be the most popular use of this technology, with student usage rising from 30% to 64%.

Another very significant finding was that **88% of students used generative AI to prepare assessment work**, up from 53% the previous year. Within this, the most popular task was using AI to explain concepts (58%, compared to 36% in the previous survey), though the largest increase was seen in **summarizing articles** (from 24% up to 48%).



18%

of students had included AI-generated and AI-edited text directly in their assessments (*HEPI Student Generative AI Survey 2025*).

*Elsevier Research data corresponds to 2024 and may be slightly different today, given the enormous development in AI implementation in this field in recent times.

Those surveyed in the Elsevier Research 2024 study identified **quick information synthesis** as the dominant use of AI — and the one they perceived as most valuable. This aligns with the pattern seen in the HEPI survey, where students showed strong engagement with both text generation and information processing tasks. Self-assessment and exam preparation followed as secondary uses in the Elsevier Research 2024 study.

The *Student Generative AI Survey 2025* found that **67% of students considered being able to understand and use AI effectively** to be essential in today's world — yet only 36% had received institutional support in this regard. It also indicated that the gap between student demand for institution-provided AI tools (53%) and actual supply (26%) had widened since 2024. However, the **digital literacy** of staff had improved notably, as 42% of students considered faculty to be well-trained and equipped to support them, compared to 18% in the previous survey.

II. A new landscape: Changes and windows of opportunity

The impact of AI has generated a **“ripple effect”** that reaches all areas and stakeholders in the healthcare field. **Sundeep Watkins** explains it this way: “AI is evolving so quickly. It is becoming much more realistic in terms of how you can interact with different AI tools, and the data that’s being utilized to inform the AI tools is becoming more sophisticated. People can interact with AI in multimodal ways, which **replicate clinical scenarios**. So, building skills in communication, clinical reasoning, assessment in terms of questioning, and treatment planning is a fundamental aspect that we can already replicate using AI tools to support learners in that journey from theory to practice.”

For **Dr. Ximena Alvira**, AI will not make us better doctors on its own: “While increasingly sophisticated tools can help us work more efficiently, access information faster, and even improve diagnostic accuracy, they cannot replace the human elements that define good medical practice, including empathy, communication, ethical judgment, and the relationship between doctor and patient. We still need to better understand the impact of integrating AI into so many aspects of healthcare. What matters most is ensuring that healthcare professionals learn how to use these technologies responsibly and effectively, so they can translate their potential into better patient care.”

To illustrate AI’s role in clinical practice, **Dr. Ximena Alvira** offers a practical example: “Today, we already have AI-enabled stethoscopes that can enhance the detection and interpretation of heart sounds. However, the fundamental role of the physician remains unchanged. It is the tools we use that are evolving. That is why I believe we should view AI as an enabler rather than as an entirely new discipline that clinicians must constantly relearn. Approaching it this way can reduce frustration and make its integration into everyday clinical practice much more natural and effective.”

On what AI is bringing — and will continue to bring — to healthcare service users, **Dr. Ruth Paterson** argues that experimentation is necessary to understand what AI platforms can offer patients in terms of clinical and medication information, and how verifiable and transparent that information will be. She adds: “I think there is a **real opportunity for AI to summarize information and make it more understandable** for service users with health literacy difficulties. I’m talking about our learning disability population and our cognitive impairment populations who may be supported in prompts, or in understanding how they use their medicines, or how they approach their health challenge.”



But to what extent has AI transformed health education, and in what direction? **David Game** comments: “AI is making significant changes in medical education, but we are only at the very beginning of the beginning. So, it’s far too early to say what substantive changes it makes. What it has shown us is that any given student with access to ChatGPT, Claude or whatever, can answer pretty much any assessment question. The big question for me is, how are we going to respond in terms of assessment? I think that is the question that it has opened up for us.”

Dr. Ruth Paterson highlights the role of AI in the management of **scientific information overload**: “We know that there’s a vast array of health research information that professionals have to sift through to make a judgment on a particular intervention. This can be overwhelming, particularly for our students. Having AI to support some of this work may be of benefit, but we need to encourage people to explore its capacity to do so at their own pace, and this is very different. This is taking a **person-centered and a student-centered approach** to realize that not everyone will develop their knowledge at the same pace and there is no uniform way of doing so.”

“As academics and healthcare professionals, we have a real responsibility to develop rigor in appraisal of AI and rigor in dissemination of its worth. We need to be part of the AI development as healthcare professionals,” continues **Paterson**. “I think there is real strength in working with AI to improve some of the **inefficiencies** in health and education systems — and in considering how we can use AI to perform some of the tasks that take away from the face-to-face care and support which is so important to not only nursing, but allied health and medicine. This needs to be taken forward and explored using the **research skills** that we have all developed and that we teach our students.”

“AI is a great opportunity to get the best information in a format that will be most effective for the learner.”

Professor Marcus Drake

Professor of Urology at Imperial College London;
Editor of the *Medicine in a Day* series,
published by Elsevier

In **Professor Marcus Drake’s** opinion, AI represents a great opportunity, provided that the information is accurate, the best information for the student is made easily available, and that it is presented in the most effective way: “It brings so many opportunities to improve illustrations and clarify text. Furthermore, for the large number of UK medical students and doctors from a range of diverse backgrounds, including linguistic, AI could be a big advantage to ensure the clarity of what is written, which will start out in English, with the opportunity to change to **alternative languages** at the request of the individual, using AI.”

Two first-hand testimonials

These are the real-life experiences of **two experts** regarding the **daily use of AI** in their work:

Professor Marcus Drake, speaking from his experience as a professor and editor of the *Medicine in a Day* medical exam revision textbook series (published by Elsevier): “I found AI particularly helpful for finding appropriate **illustrations**. In fact, I have been very impressed how once the AI has appreciated the sort of image that you’re looking for, it will **quickly develop or offer several options** that it has retrieved from various sources, which is incredibly beneficial. The **sheer speed** with which it can work has been absolutely fantastic.”

“When writing medical textbooks covering various clinical circumstances, you really want to have a coherent and effective group of illustrations to support the information to make it more manageable for the students so they’re not just reading — they’ve also got some **visual elements** to facilitate their learning. I’ve found AI has been spectacularly helpful in recent times for providing me with options of simpler illustrations to present to students. This has had the most notable difference in how I work since incorporating this technology.”

“Furthermore, it is quite useful for **overcoming mental blocks**. If you have a moment where you’re struggling to identify a fact or follow a line of argument, AI can offer suggestions that really help you to unblock and **get back to effective activity**. So, it’s incredibly useful for **productivity, effectiveness and reliability**.”

Dr. Philip Xiu: “I work as a general practitioner in family medicine, and I’ll be honest, AI has really changed my day-to-day work. It is much more than I initially expected and not in ways that I imagined possible.”

“With any clinical practice, I use it to **summarize long discharge summaries** before a specific consultation with a patient. The **precious time** that I get back within a 10-minute consultation really changes the quality and the amount of context I have with the patient in front of me. Within clinical education, I use it to draft case scenarios, to stress test various assessment items and to rephrase possible feedback so it just lands better with a learner who might be struggling.”

“However, there is something key I want to reiterate: I would **never use an AI tool to make any decisions**. Yes, I use it to help me make the decision better, more robust and more defensible, but even so, in a clinical setting, when I see the patient in front of me, the diagnosis remains mine. And the same applies in **any tutorial settings with a student**: the judgements about a specific student are still mine. The important thing is that the relationship I have with both the patients and the students is something I will never outsource. AI simply gives me back some of my time. Now, what I do with that time — whether I listen, teach, notice or care — that’s still the core job and remains firmly within the realms of medicine.”

III. The role of AI in personalized learning

When considering how personalized learning has evolved with the integration of AI, **Dr. Philip Xiu** starts by reflecting on the greatest fear when AI first arrived in healthcare education — that it would replace the educator: “In my opinion, the opposite is happening. **AI is actually restoring the role of the educator**. For a long time now, faculty have been buried under marking, assessments, content updates and responding to student queries. But now, AI is actually shouldering a lot of the load which means that faculty can return to the part of teaching that truly inspired them to do it in the first place. That includes **mentorship, modelling, professional**

identity, and, for me personally, sitting with a struggling learner and overseeing their own development and their own pathway. That’s what **personalized learning** means to me. It means much more than the computer system picking your next question. Now, hopefully, it means there is a **continuous conversation** between the learner, the content that they’re trying to tackle and the educator themselves, and AI is really the connective tissue between them all.”

“As a faculty member, I can see where my students get stuck even before they tell me, and AI now gives me the opportunity to potentially intervene earlier and with much better data to back it up. Also, **AI saves time for the faculty** because students have often self-resolved any queries by their own AI searches anyway.”

“AI presents an opportunity for students to engage with and individualize their interaction with learning content and materials. It can moderate how they interact with learning content. But we want to keep in mind that AI should be utilized in support of teaching delivery rather than as a teaching replacement.”

Sundeep Watkins

Education Adviser at the Chartered Society of Physiotherapy; Women’s Health Physiotherapy Clinician & Educator

Sundeep Watkins highlights simulations — an increasingly relevant tool in the healthcare curriculum — and describes how AI can create safe and personalized learning experiences through this resource: “Simulation presents the opportunity to provide a reproducible, safer environment to test out knowledge, skills and theory, and put it into practice. It can create really tailored and realistic interactions, which can be multi-layered in terms of the presentation — providing that intersection between the different aspects of an individual’s care needs. You could also vary the experiences in simulation to match the ability of the learner at different stages of their development...but the expertise of the facilitators in

those simulation environments is key and that debrief is essential to ensuring that learners get the best out of simulation experiences.”

However, **Watkins** emphasizes that, despite the advances AI has brought in this area and the promising outlook for the future of simulation, **clinician expertise remains essential** — and always will be: “We want to use those simulation experiences to **enable students to hit the ground running** when they get into practice settings, and to take out some of the uncertainty they may have about the situations they will face and build their confidence. That is why we absolutely need clinicians with expertise informing these simulation experiences, providing feedback and feeding into the development of those algorithms to make sure that they are as representative of real-life scenarios as possible.”

IV. Doubts, fears and “healthy skepticism”

Alongside the great advantages, new opportunities and advances it has brought, the arrival of the “AI whirlwind” has also given rise to doubts, suspicion and even fears among all players in the healthcare world — a completely logical reaction, given the speed of this technology and the fact that its implementation **requires a rethink of teaching systems that have been in place for a long time.**

For **Dr. Ruth Paterson**, this is **healthy skepticism**, and she explains why: “I think **skepticism around AI is good** and it’s okay to feel that way, and it is really important for us all to hear those concerns. We must understand and respect the professionals who are hesitant about the use of AI. Everybody likes a shiny new thing, and there’ll always be those who are early adopters of it, but there are also people that are very, very hesitant, and there are lots of reasons for this hesitancy. For example, if we look at the **sustainability argument against AI**, there are strong reasons why we shouldn’t be using it.”

Paterson insists that, despite these reservations, the **efficiency and support** AI offers healthcare professionals cannot be overlooked, and we should approach AI through curious, evidence-based exploration: “As health professionals, we’re there to protect the public in terms of privacy, data sharing, confidentiality, and we need to understand what AI can do to enable that, but also the associated risks. I think we need to take a research

approach to this. As academics, we’re all trained in the research process and making sure that we’ve got rigor, transparency and applicability when using AI is important. So, we need to think about **exploring, testing, and building evidence** for its reliable capability.”

In the context of AI-powered simulation, **Sundeep Watkins** explains how we can **build trust in this technology** by acknowledging that it should complement, but **never replace, real-world clinical practice**: “We need to use simulation to build skills for our future practitioners but we don’t want simulation to fully replace practice-based learning in that context. We want it to reproduce real-world scenarios, ensure there is a space to safely practice in advance of treating real patients and gain exposure to complicated scenarios to build confidence and resilience in our learners. It’s there to **prepare, enhance, and augment that practice-based learning experience.** We absolutely want learners to feel confident, but we also need them to understand that that confidence may not fully translate into the practice setting until they’ve had those real-life experiences with people.”





Chapter 2

AI and educational content: Trustworthy design

I. The design challenge: How AI integrates reliable data

The fact that approximately **75% of healthcare students** already use AI tools today raises a critical question: what needs to be done to ensure this **technology is used with reliable, trustworthy content in a safe learning environment?**

On this, **Professor Marcus Drake** points to the need for **strong validation of AI-generated information**: “It is quite unimpressive how often AI ends up hallucinating a fact and portraying it in a way that’s suggested as being authoritative when, in reality, it is a complete fabrication.”

Drake recognizes that, although these hallucinations are becoming less common as AI improves and begins to **incorporate internal checks and validations**, the risk of these tools conveying inaccurate information is still at large: “Strong emphasis on **validation, cross-checking and elimination of false information** has to be a priority. Regrettably, AI can, of course, call on any information that’s published on the internet, whether it’s true or not.

If it was possible to flag information as false, that would be very valuable, but also very unrealistic.”

These concerns further highlight the importance of AI tools designed for healthcare students integrating **reliable data**. The question is: **how?**

From her experience in education technology, **Raman Kaur** shares some key insights on the subject: “AI in medical education works best when it’s grounded in trusted medical content and not just open internet data. This has already been studied and proven — that connecting our education tools with trusted content creates **epistemic trust** with our educators and our students, and that’s a key competitive advantage for any tool. In practice, the system combines generative AI with curated educational resources such as peer-reviewed textbooks, clinical references, anatomy content and assessment materials. Along with that, we incorporate educator-reviewed learning experiences, as the **real-life experiences of an educator** play a key role in making AI in medical education more grounded and trustworthy.”

For **Kaur**, the value of these AI tools lies in combining AI’s capabilities with trusted educational content and student engagement data to create a very powerful, tailored and curated experience.

Kaur emphasizes that **trustworthy AI systems for healthcare** rely heavily on:

- Grounded medical content
- Citations and traceability
- Human oversight
- Consistent evaluations
- Safety monitoring

Two very illustrative examples

Raman Kaur points to two educational tools developed by Elsevier that illustrate, “how we’re approaching AI integration with medical education, with epistemic trust as the fundamental pillar:”

Osmosis AI — “The AI-driven user experience helps students reinforce concepts through videos, flashcards, question banks, adaptive learning support and a conversational AI, which sits on top of all these capabilities.”

StudyFinder AI (in ClinicalKey Student) — “Through this tool, students can access evidence-based medical references, as well as textbooks, videos and clinical answers in a single ecosystem, enjoying an all-in-one experience rather than having to navigate through different tabs, links and content delivery formats. In this way, AI can help surface the most relevant information faster, and support more personalized learning pathways.”

II. What are hallucinations?

The *Student Generative AI Survey 2025* report defined **hallucinations** as the generation of false outputs such as made-up facts, statistics or citations. When students were asked about how often the tools they used produced this type of error, 39% said rarely; 30% said quite often; 12% said never; and only 8% said it occurred very often.

For the study’s authors, a striking comparison between the 2025 report and the previous year is the dramatic drop in students who say they don’t know how often the AI tools they use produce hallucinations (from 35% in 2024 to 12% in 2025). This suggests students have **rapidly acquired a much deeper understanding of**

how AI tools work and how reliable their outputs are. “Digital maturity has increased notably in just one year,” the authors conclude.

But despite these advances, hallucinations continue to “loom over” AI in the field of medical and healthcare education and can **hinder its effective use.**

As **Raman Kaur** warns, hallucinations are “fabricated, put together pieces of information that sound about right, but are not entirely accurate and, if read quickly and not reviewed carefully, you might miss that it is a ‘**hallucinated artifact**’ and not accurate information.”

Kaur gives a series of **examples of AI hallucinations** that can occur in medical education: “It could look like an **invented clinical citation** for research, suggesting outdated treatment guidance or providing an **inaccurate explanation** of a disease process because two similar sounding symptoms were connected into one. Or even **mislabeling anatomical structures**, such as labeling an epicondyle next to a bicep, because the human brain processes more context than AI can, and that gap is where hallucinations can occur.”

Detecting hallucinations is, therefore, one of the main challenges facing AI in health education: “The evaluation process is highly critical because individuals who are consuming our healthcare AI education materials are not experts yet, so we wouldn’t want them to ground their learning in inaccurate systems.” notes **Kaur**.

“One of the biggest risks in healthcare AI is not that the system sounds uncertain — it’s actually when the system sounds extremely confident while being wrong. This is why having a special evaluation framework for hallucinations is so important.”

Raman Kaur

Vice President in Education Technology at Elsevier

When it comes to **AI safety monitoring**, **Raman Kaur** identifies this as something that is probably uniquely important to healthcare education tools compared to, for example, an AI tool for sales, “because you want to make sure that the information provided by our tool is relevant,

up-to-date and represents the correct method for potentially building the knowledge of a future treatment or surgery intervention.”

While **Pippa Furey** feels that “hallucination” is an imperfect term — humanizing a process that is fundamentally statistical — she argues that detecting them is not as new a challenge as it may seem: “I remember when I started studying at university, we were taught how to know whether the academic paper you’re reading is unbiased or if it’s relevant, and that we must look at recent publications.”

Furey points to the “**BREAD framework**” — **Bias, Relevance, Evidence, Author and Date** — as a tool for evaluating information sources, whether that’s traditional academic publications or AI outputs: “We’ve always needed to think about the motivator of publication, who’s benefiting from putting these things out there and who is not. It’s not new. It’s something we’ve always done with sources, but it translates really well when applied to AI outputs and can serve as a great starting point for that conversation with students.”

In **Gosia Warminska-Marczak’s** view, reliable AI in healthcare education is not down to a single factor; it is a combination of **trusted content, skilled users, clear rules, appropriate checks** before implementation and **awareness of the risks**.

On this first point, she highlights the importance of the content behind the tool: “AI has to be grounded in trusted, peer-reviewed content. This is where publishers and academic institutions have a real role because the quality of the output depends on the quality of what goes in.”

Warminska-Marczak goes on to warn about **AI hallucinations**, stressing that they are more than a mere annoyance in healthcare settings: “Hallucinations are **one of the biggest risks**. A system can generate something that sounds very confident but is simply **wrong**. In everyday life, this is annoying, but in a medical context it can be **truly dangerous**.”

To combat issues like this, **Warminska-Marczak** identifies AI literacy as another key factor for the implementation of **trustworthy AI** in healthcare

education: “I really believe AI literacy needs to become a core clinical skill, not just an add-on. Students need to understand how AI works, where it fails, how to question what it produces and to not accept the answer just because it looks professional.”

III. AI evaluation: Why it is so important

“Once evaluation takes place for an AI tool in medical education, we don’t stop there. This evaluation happens on an ad hoc and consistent basis because medical education is a rapidly evolving construct. AI systems and healthcare education cannot be just set and forget. Continuous and consistent evaluation is crucial.”

Raman Kaur

Vice President in Education Technology at Elsevier

Given the risk of **hallucinations** and the critical importance of **reliable content and trustworthy design**, evaluation has become one of the cornerstones of responsible AI integration in healthcare education.

But what exactly does evaluation mean in this context, and what does it involve?


“Evaluation is essentially our **human-in-the-loop process**. As we introduce more AI capabilities to our medical education platforms, there is a constant human-in-the-loop AI evaluation to make sure that we **identify and catch any type of inaccuracy** or potential bias from happening before it lands in the hands of the students,” explains **Raman Kaur**.

Kaur sees that evaluation is, in fact, one of the most important parts in ensuring that an AI tool is **well-grounded** and provides responses that look at the information constructively and are free from any type of hallucination. The system must provide responses aligned with trusted references and **medical standards**.

To carry out this evaluation process, **Kaur** notes that organizations must assess aspects such as **factual accuracy, clinical safety, citation quality and bias across patient populations**. “We look for all these types of challenges that could arise when AI provides a response to ensure our tools remain grounded and trusted and that they deliver ethically sound and education-aligned answers.”

An important point to bear in mind is that this evaluation is **continuous**: “Once the evaluation takes place for a version of an AI tool in medical education, we don’t stop there. This evaluation happens on an ad hoc and consistent basis because medical education is a rapidly evolving construct,” notes **Kaur**.

Dr. Ruth Paterson focuses on the importance of **applying research rigor to AI evaluation**: “We need to test hypotheses in relation to AI. We need to use reproducible methodologies to investigate AI interventions and use precision to inform the direction of future practice in this area. I believe that if we do this collectively, as a collaborative effort across the professions, with our students, our healthcare users and ourselves, this will put us in a really strong position when it comes to the use of AI.”





Chapter 3

AI and clinical thinking: “Destined to understand each other”

I. How AI influences what students learn

For AI to fully harness its potential and transform learning in healthcare, certain key conditions must be met to ensure its responsible and effective integration in healthcare education. From his perspective as a regulator in osteopathy, **Paul Stern** highlights what these key conditions are:

Appropriate innovation — “First, I think there needs to be space to allow for innovation. However, that space must sit within the existing education and training standards that we set as a regulator, as well as within our graduate outcomes so that graduates have the knowledge and skills needed to meet our Standards (the Osteopathic Practice Standards) and can practice autonomously at the point of registration.”

Critical thinking — “Students and staff need to make sure they understand that AI outputs can be inaccurate and biased. The focus needs to be on **AI tools enhancing what they do**, rather than relying on the AI to provide the answer. Students need the

skills to be able to evaluate AI outputs rather than accepting them uncritically.”

Accountability — “Students and staff need to be accountable for how they use AI, and that is really important in education,” says **Stern**. “But then also, for students, understanding that accountability and being able to continue to take that into their practice post-graduation is crucial to ensuring that they keep patients safe.”

Safeguards — “The fourth condition is around having the right safeguards in place. So, making sure that there’s **data protection and security**, and healthcare education providers will obviously set their own policies around what’s acceptable in relation to AI use and what’s not acceptable. Having these types of policies in place is a way of ensuring academic integrity.”

Transparency — “The fifth condition is transparency. Making sure there’s an **open, collaborative relationship** with the institutions that we work with — one where we, and they, can share knowledge on what’s working well, what isn’t working well, and learning from each other.”

Patient voice — As his final condition, **Stern** highlights the importance of “making sure that patients and the public have a voice in how AI is integrated into the educational process.”

“Health education has shifted from simply presenting information to helping students interrogate it. I think now we’re really asking: ‘Where does this come from? Can I trust it? What does it mean clinically?’ — and that shift feels really fundamental.”

Pippa Furey

Senior Lecturer in Paramedic Practice and Emergency Care at Nottingham Trent University; National AI Lead (Education) at the Royal College of Paramedics

On what has changed the most in teaching since generative AI, Pippa Furey notes “it is not what we teach, but how openly we talk about knowledge itself and organic learning. Since generative AI became more widely available, we’ve had to **explicitly address some concepts** that have previously sat beneath the surface in education that we’ve known about and alluded to, but we maybe haven’t had as explicit topic types that we’ve needed to engage with so willingly. Things like cognitive offloading, bias, what counts for organic learning and understanding, and how do we guarantee that students actually *know* what they know?”

Judgment beyond technology

Drawing on her experience as a paramedic, **Pippa Furey** offers the following analogy to illustrate how the innovation that AI represents should be approached in the training of future healthcare professionals: “Let’s take **blood pressure measurement** as an example: we were taught how to do it manually and we were also taught how to do it with an automated blood pressure machine. But we didn’t stop teaching manual blood pressures when the technology came out because you still need to **understand the underlying physiology**, what’s happening and how to interpret the results.”

“Also,” continues the expert, “if the equipment isn’t working or you’ve got a patient who shouldn’t have an automatic blood pressure taken, that knowledge is still necessary. With the automated blood pressure, we were redistributing the cognitive effort. AI is similar — it can generate information quickly, but if you haven’t got that foundational knowledge, then the output is just text on the screen and what you make of it is going to be very different.”

Sundeep Watkins highlights the **potential role of AI as a personal tutor**: “AI will enable a more **flexible and asynchronous approach** to accessing learning pathways, providing **individualized feedback** for those learners and interacting with them on an individualized level, even just in terms of the timing that they would like to access content they’re engaging with. AI can adapt to a student’s pace, to their learning preferences, and it could also transform how we assess those learners in the process of their development. We’re seeing that AI could potentially act almost as a personal tutor or as a critical friend in the process of a learning program.”

When defining the **pedagogical implications of AI in healthcare education** and, specifically, what it brings (and what it does not) to the way students learn and teachers teach, **Pippa Furey** notes that AI works “as a thinking partner, not a shortcut,” and makes the following reflection: “We need to foster an environment where students feel comfortable to talk about how they’re using AI, and also educating on use versus misuse, which is really difficult if it seems like something that students shouldn’t be using.”

Furey continues: “If faculty aren’t open about it, students won’t be open back because a lot of the narrative is around **academic integrity** and the risk of cheating and creating work that’s not their own.”

Furey lists some of the ways in which **AI can contribute its full potential** in the academic training of future professionals:

- AI is particularly effective for prompting **idea generation to get the ball rolling**. As **Furey** explains: “It shouldn’t create content for them, but it can absolutely give them ideas on how to start something when they don’t know where to begin.”
- It is also very useful for **feedback and informative learning**. “It’s one thing to ask the system to create a multiple-choice questionnaire for them to complete, and quite another thing to then verify what the AI thinks the correct answer is to create a learning opportunity there,” notes the expert.
- Its use is also important for helping students **explore multiple perspectives and find different voices**, getting out of their own echo chamber. **Furey** notes: “We all, in life, surround ourselves with people that are like us and reflect our values and actually, AI gives us the opportunity to find some different voices in that.”

II. Critical thinking: The challenge and goal at the heart of the process

When experts are asked which human skill future healthcare professionals must strengthen most in a landscape of unstoppable AI advancement, the answer is unanimous: **critical thinking**.

“I believe that, as humans, we have a unique ability to take a critical eye — seeing different points of view and trying to eliminate bias. And really, rather than believing everything that we see in AI or that is in AI, we need to appraise it. We need to develop **structured models to appraise the information provided by AI** that protects the public from bias, inappropriate information disclosure and unreliable information. That is absolutely vital,” comments **Dr. Ruth Paterson**.

“Next, we need to think about **how we can integrate AI into our systems** and ask ourselves: “What is the practical application of AI to the problem we are seeking to address?” Whether it is to develop more robust and up-to-date teaching plans, to summarize large amounts of clinical information or to verify information that patients have curated,” continues the expert.

And throughout this entire process, **critical thinking is a key and irreplaceable element**. As **Paterson** reiterates, the essential human skills in the face of advancing AI

are: “Asking the right questions around AI and how it can support our work, appraising the answers and outputs it provides, and reflecting on whether that information is reliable enough to integrate into our healthcare systems.” Once we have appraised the AI, **Paterson** sees **continuous evaluation as the next step**, “to ensure that we have a truly integrated system that sits alongside what we do, which is essential to health professionals as they manage the caring, treatment and human aspect of healthcare delivery.”

Along the same lines, **Pippa Furey** highlights the role of AI in **reflective practice** as a tool to deepen learning and reinforce that critical thinking, while also warning of the potential risk that it may oversimplify the reflective process if not used properly: “AI works best when it creates a space for clinicians to focus on patients, not when it sits between them. I don’t think AI should be present in the patient interaction — which should be **human**. For instance, if AI writes a reflection instead of the doctor, it’s fair to say that a reflection has not actually occurred. Let’s say a clinician puts into the AI: ‘Man has fallen over and broken his leg. Please reflect for me,’ the AI output will be inhuman in tone and won’t speak to the clinician’s learning.”



“When I’ve seen reflection using AI or tried it myself, it really simplifies ideas and doesn’t create learning. AI can be very effective as a prompt generator, helping students move beyond descriptive accounts into deeper analysis, but the thinking must come from the student.”

Pippa Furey

Senior Lecturer in Paramedic Practice and Emergency Care at Nottingham Trent University; National AI Lead (Education) at the Royal College of Paramedics

“Therefore,” continues **Furey**, “if we want to make the most of this technology in our students and clinicians, I think we should be able to use AI as an unscheduled **reflective prompt generator**. Unlike more traditional analog, process-driven reflection models, **AI encourages ongoing engagement** because it asks follow-up questions. For example, when you finish inputting your information for the reflection, the AI will ask: ‘*What do you think about this? Have you explored that? Do you want more suggestions?*’ and it is then up to the user to respond. AI wants you to use it, to engage with it and to keep asking it more questions. Applying this to the field of teaching, AI provides an opportunity to create a safe space for students to reflect and explore their feelings and knowledge without having to share it with another person.”

A thinking aid (not a substitute for thinking)

Pippa Furey emphasizes that the problem with AI in healthcare education lies in replacing the human thinking element in students: “If students are **using AI to generate answers without engaging**, then learning becomes really superficial. Therefore, if we’re going to create a task that’s AI-enabled, we need to make it really clear to the student that it’s not just about generating an output — it’s about what that output actually means. So, the pedagogical implication of AI is that it amplifies the existing teaching design, and it can make teaching more accessible to students that struggle to engage. However, strong pedagogy will become stronger with AI, while weak pedagogy becomes more visible and this gap can be widened.”



Scan code to explore additional resources on AI in health education



References

Elsevier. (2024). *AI use cases: Medicine & nursing students. Global Research Study*. [Internal research report].

Freeman, J. (2025). *Student Generative AI Survey 2025*. Higher Education Policy Institute & Kortext. www.hepi.ac.uk/reports/student-generative-ai-survey-2025/

