

Balancing Risks and Educational Potential: Generative AI Use in Medical Student Learning

Yuxi Huang, Muhammad Tanveer Iqbal*

Henan Medical University (Department of International Education Department)
tanveeriqballegghari@gmail.com (Corresponding author)

Abstract. Artificial intelligence (GenAI) has become a common tool for many medical students learning, even though most curricula have not established clear rules for its application. These tools can clarify complicated biological mechanisms that produce practice focused quiz items stimulate real life clinical dialogues and other instant learner feedback. These practical advantages cannot ignore yet there exist mechanisms that is overlooked drawbacks assured looking outputs can hold factual mistakes biased assumptions or even incorrectly matched citations. This paper reviews material gathered via targeted PubMed searches citation tracking and relevant international guidance published between November 2022 and 13 August 2026. Instead of assigning equal weight to every included study the analysis draws clear separations between learning benefits reported by students themselves technical capabilities of AI models and real-world learning outcomes among learners. Available research favors careful, limited deployment of GenAI for concept clarification, formative feedback, case practice, and custom question development. It provides no solid basis to conduct that regular heavy usage will bring improved learning performance. Hallucinations, automation bias, cognitive offloading, ambiguous authorship status, and unequal access continue to represent key practical worries. How much risk they create depends on the task type, tool quality active participation from teaching staff, and whether learners verify AI produced statements. To untangle these complicated interactions, the paper proposes a Risk-Benefit-Governance framework. The core standpoint of this work is straightforward: GenAI is educationally useful when it supports students independent critical thinking, but harmful when it takes over all critical reasoning work on their behalf. Medical schools need to train learners in verification, disclosure, and reasonable usage boundaries together when firsthand AI competence training. Upcoming research ought to investigate retained knowledge, transfer, clinical reasoning, professional conduct, and patient safety instead of focusing mostly on attitudes or short-lived assessment results.

Keywords: generative artificial intelligence; medical students; medical education; clinical reasoning; AI literacy.

1. Introduction

Generative artificial intelligence spreads into medical education more quickly than many traditional tools. After ChatGPT released for public use in the later part of 2022, students no longer needed to wait for their universities to formally update their policies before using it to directly study, discuss difficult topics or practice clinical decision making.

This easy availability has clearly increased student interest, but it has also raised important questions about how universities, clinicians, and students should balance convenience with responsibility. Medical students must absorb extensive amounts of biomedical information while slowly building skills to work with clinical judgements. Such flexible software can explain complex biological pathways, build realistic clinical practice scenarios or offer feedback can lower obstacles for students learning. Generative system can also deliver learning help beyond formal class and timetables and adjust explanation depth and style according to what the students ask for. These practical functions have encouraged educators to view GenAI as a tutor, stimulation partner, content generator, and resource for independent self-study (Boscardin et al., 2024; Kasneci et al., 2023; Preiksaitis & Rose, 2023; Sallam, 2023).

Fluent prose does not equal reliable information. AI-generated text may sound natural, but it may contain subtle errors or not be applicable to the student's clinical context (Qin, 2026). This is concerning, because students are developing the necessary clinical judgment to identify these subtle errors. Kriplani et al. (2025) found that when students modified their original clinical judgment according to AI-generated recommendations, these changes were accurate. The point is that while genomic AI can enhance student abilities, it can also influence their thinking to the point where they lack the underlying knowledge to verify the information they receive. The debate around genomic AI has become polarized. Some believe that it is the future and that medical education should adopt it quickly. On the other hand, there are concerns about academic dishonesty, misinformation, and the loss of independent critical thinking. None of these approaches benefit students. Banning it outright will only encourage its covert use. In contrast, allowing its unrestricted use increases the speed of rewards accuracy, without indicating whether the student has applied critical thinking. The question of whether we should allow students to use genomic AI is more relevant to how and when it complements rather than replaces it. The purpose of this article is to present a conceptual analysis. It would be irresponsible to conduct a meta-analysis, as the combined data can be misleading. Multiple studies on AI in education address different questions and outcome measures. Survey studies attempt to understand students' use and perceptions of genomic AI. Studies using interventions analyze student performance in a controlled environment. AI studies measure the productivity of algorithms. This, again, has nothing to do with students' knowledge acquisition. Finally, ethical guidelines indicate responsible behaviors, which cannot be considered in a paper about the classroom context (Gordon et al., 2024; Masters, 2023). For that reason, these varied sources contrast to identify consensus, points of conflict, and the root causes behind such divergence.

2. Analytical Approach

PubMed queried to retrieve publications released between November 2022 and 13 August 2026. The following search strings applied: ("generative artificial intelligence" OR "large language model" OR ChatGPT) AND ("medical student" OR "medical education") AND (learning, teaching, assessment, OR feedback OR "clinical reasoning" OR ethics OR "academic integrity"). Titles and abstracts screened for direct applicability to undergraduate medical learning. Higher priority went to empirical studies involving medical students, research synthesis papers, and papers that delivered meaningful educational or governance-related insights. Studies focusing on health-profession training in wider contexts kept if their results could apply to medical education. Citation chasing of key publications helped locate extra relevant sources. WHO guidance on large multimodal models. UNESCO guidance on GenAI in

education along with research included because educational practice cannot fully be disconnected from privacy, transparency, accountability, and human oversight.

Literature centered purely on clinical diagnostic accuracy with no educational component mostly ruled out, as were marketing-style accounts lacking analysis of learning outcomes or hazards. Five core research questions shaped like this analysis: What are students using GenAI to do? Is the stated advantage merely subjective or empirically proven? Which part of the student's cognitive labor supported or replaced? What hazards arise concerning knowledge acquisition, reasoning ability, integrity, privacy, and equity? What protective measures match the demands of each given task? Available evidence contrasted across the learner, learning activity, institution, and future clinical environment. Recurring tradeoffs appear speed against productive intellectual effort, personalization against consistent reliability.

Recurring core tension speed versus productive struggle, personalization versus reliability, assistance versus substitution, and access versus equity—formed the core organizing ideas behind this framework.

The approach stays interpretive in nature. It did not conduct dual independent screening, so relevant work published beyond PubMed may be missing. Furthermore, AI model capabilities evolve more rapidly than academic publishing timelines. The paper therefore refrains from framing short-term observations as fixed educational truths and regards its framework as a logical proposal open to real-world empirical validation.

3. Applications for Generative AI in Medical Student Learning

Medical students use generative AI in a wide range of low- and high-risk activities, from seeking clarification on isolated concepts to performing tasks that mimic diagnostic or therapeutic reasoning. Common uses include studying complex topics, retrieving annotated notes, developing memorization techniques, generating self-assessment questions, verifying explanations, and organizing study schedules. In clinical reasoning activities, students can explore differential diagnostic options, compare treatment algorithms, practice communicating with patients, or design realistic cases. According to a Canadian multicenter survey, students used generative AI to review medical information, summarize clinical guidelines, and make differential diagnoses, thus blurring the lines between academic and clinical activities (Tran et al., 2025).

Table1. Educational uses of GenAI and their conditions for value

Learning use	Possible educational contribution	Main vulnerability	More defensible practice
Concept explanation	Adapts wording, depth, examples, and language	Confident simplification or factual error	Compare with a textbook or guideline and ask the student to restate the concept
Practice-question generation	Increases retrieval practice and variation	Ambiguous stems, incorrect keys, poor difficulty calibration	Faculty review or source-linked verification before use

Case simulation	Supports diagnostic reasoning and communication rehearsal	Unrealistic cases, premature closure, unsafe recommendations	Require justification, alternatives, and debrief with a teacher or trusted source
Formative feedback	Provides rapid comments when faculty time limited	Generic feedback and reinforcement of misconceptions	Use rubrics, retain human review for high-stakes judgments
Summarization	Reduces initial burden of long material	Omissions, loss of nuance, invented citations	Read the source and audit the summary against it
Writing support	Help organization and language development	Authorship ambiguity and loss of individual voice	Disclose material use and retain responsibility for every claim

Depending on the task design, AI-based interactions can facilitate or hinder the learning process. An application to describe the renin-angiotensin-aldosterone system in varying levels of detail can foster a deeper understanding of concepts. Conversely, customizing the writing of a comprehensive article submitted without rigorous review will diminish the educational value of that work. Similarly, using GenAI to draft a clinical case and then critically evaluate the results is not the same as accepting the answer without question. Therefore, the same tool has the potential to enhance or hinder learning, depending on how the task is designed.

Table 2. How different forms of evidence should be interpreted.

Evidence type	What it can show	What it cannot establish by itself	Analytical implication
Student perception survey	Patterns of use, confidence, concerns, and willingness to receive training	Improvement in knowledge, reasoning, or clinical competence	Enthusiasm is evidence of acceptability, not effectiveness
Model examination study	Whether a named model can answer a particular set of questions	Whether students learn by consulting the model	Machine performance should not be reported as student learning
Short educational intervention	Immediate change under a defined teaching condition	Retention, transfer, or safe independent performance	Promising results require delayed and authentic assessment
Trust or error-detection experiment	Susceptibility to misinformation and automation bias	Frequency of harm in ordinary study	Verification skills deserve direct teaching and assessment
Evidence synthesis	Recurring applications, benefits, and concerns across studies	A stable effect when studies and models are heterogeneous	Conclusions should remain conditional and context-sensitive

Ethics or governance guidance	Normative expectations for privacy, autonomy, accountability, and equity	Educational effectiveness of a particular classroom activity	Pedagogical benefit does not cancel professional obligations
-------------------------------	--	--	--

4. Educational Potential

4.1 Personalized explanation and access

One of the most promising aspects of GenAI in education is its responsiveness. I can request a more comprehensive answer, a step-by-step explanation, a comparison chart, or a simple and accessible summary for the patient without having to wait until the next meeting with the professor. This could be especially valuable for students who need to communicate in a foreign language or who, for other reasons, are hesitant to ask simple questions in front of their classmates. Another way to look at it is not as a guarantee of perfect answers, but as greater access to explanations. Educational value only comes into play when personalization is appropriate and presents appropriate challenges.

4.2 AI-supported formative feedback and deliberate practice

GenAI can provide rapid diagnostic feedback on medical reasoning narratives, short answers, and communication scripts. Çiçek et al. (2025) conducted a randomized trial with first-year medical students comparing the effectiveness of formative feedback generated by ChatGPT-3.5 to written feedback from experts. The two groups showed no significant difference in overall performance on the immediate and delayed tests. That said, the expert group outperformed the GenAI group on an unusually complex diagnostic case. This is a promising but important finding. This suggests that AI-assisted feedback can be integrated into specific training contexts, but its value is reduced in very specific or important domains. Rapid feedback allows for faster iterative cycles through more opportunities to refine scenarios and responses. However, feedback alone is not instructional. An effective way to benefit from feedback (human or otherwise) is to have students first draft their responses, explain their reasoning, review the feedback, and revise their work. Presenting AI-generated solutions before a human solves them undermines the intellectual discipline that is needed to foster the kind of reflective practice required to engage in the work.




Fig.1: Immediate and delayed clinical-reasoning scores following expert-written or AI-supported feedback. The AI condition used ChatGPT-3.5. Data from Çiçek et al. (2025).

4.3 Case-based learning and clinical reasoning

Generative AI tools can be used to modify case details, participate in dialogues as a patient, and formulate clarifying questions. This flexibility facilitates better learning opportunities by addressing the lack of exposure to real-world clinical practice. Problem-based learning with ChatGPT has been associated with greater progress in a subset of clinical competencies, and students appear to benefit

from access to AI-generated resources. However, good performance on AI-supported learning tasks cannot be automatically generalized to clinical practice, where competencies often require considering multiple factors, contexts, constraints, and downstream effects that are not always integrated within the context of a well-structured vignette.

4.4 Efficiency and self-regulated learning

Students often report that the time saved represents a significant benefit. This time can be educational by removing unnecessary but necessary distractions, such as providing additional illustrative examples that go beyond the basics once a certain level of proficiency is reached. However, it poses significant risks by eliminating valuable intellectual effort. Intellectual engagement during information retrieval, synthesis, and presentation is an essential element of knowledge construction. The key question is whether the reduction in task completion time, facilitated by general AI, will be devoted to deeper reflection, practice, and skill development, or to more comprehensive information acquisition.

4.5 Preparation for AI-enabled clinical practice

Newly graduating medical students will be exposed to an ever-increasing amount of information and medical decision support tools processed by AI. The ability to critically evaluate these tools is a professional skill that will serve them well beyond their training. Teaching AI literacy is an investment in critical thinking that graduates will use long after graduation. Understanding these tools requires an awareness of data limitations, predictor variables, algorithmic biases, patient confidentiality, and professional responsibility. Finally, it is important to instill in medical students a sense of professional responsibility for the use of these technologies and the consequences they produce.

5. Risks and Challenges

5.1 Error, fraud, and unstable output

Complex language models produce passages of text that seem plausible rather than based on proven medical evidence, as is customary. They allow for the invention of false references, the confusion of different medical recommendations, the omission of essential contradictions, or the generation of contradictory answers to well-researched questions. Improvements in later versions of these models do not address this fundamental problem. In an educational setting, the danger posed by these answers is not limited to easily identifiable false claims, as incorrect conclusions or suggestions can go unnoticed and be retained by students.

5.2 Automation bias and misplaced confidence

Students can be misled by AI responses, as they seem highly persuasive. The article by Kriplani et al. (2025) presents experimental results that support the idea that students' beliefs changed after being exposed to an AI response that turned out to be incorrect. Furthermore, the authors found that students were unable to change these beliefs during subsequent group discussions. Therefore, the educational risk is explained by the students' ability to interact with the technology and the inherent technical flaws of that interaction. The student's ability to accept or reject the response should be an important factor in future research. Consequently, it should be measured on a large scale and compared with the accuracy rate of AI.

5.3 Cognitive offloading and weakened independence

All learning tools shift certain cognitive burdens away from the learner. The critical concern concerns what kinds of mental labor get outsourced. A calculator can take over arithmetic tasks once a learner grasps the underlying principles. GenAI can offload the more central acts of picking supporting evidence, linking pathophysiological mechanisms building logical arguments, along with weighing competing options. When students routinely request completed explanations before attempting a problem, they may lose opportunities for retrieval, error recognition, and metacognitive monitoring. By

contrast, using AI after an independent attempt can make the same system a source of comparison and feedback. Accordingly, timing is therefore an important safeguard.

5.4 Academic integrity, authorship, and assessment validity

GenAI blurs well-recognized boundaries between supportive aid and original authorship. Blank restrictive policies are hard to enforce in practice and may create unfair disparities if learners utilize AI secretly while others follow institutional rules. At the same time, undisclosed AI-generated content prevents accurate evaluation of a learner's actual knowledge base. Assessment frameworks ought to prioritize traceable reasoning, oral viva examinations, process-oriented evidence, supervised practical demonstrations, along with assignments rooted in real-local clinical scenarios. Disclosure policies should distinguish acceptable editing or brainstorming from generation that replaces assessed work.

5.5 Confidentiality and Privacy

Students may input patient data or images into publicly accessible AI applications and, without knowing what happens to that information, provide it to other users and allow it to be stored, processed, and used for other purposes. This includes the potential for sets of identifying information to be used to re-identify patients. Students should not input patient information into unauthorized genomic AI applications. The university should ensure that students receive procedures, approved AI applications where appropriate, and examples of anonymization. Students should receive confidentiality training before they begin their clinical practice, not after a data breach has occurred.

5.6 Prejudice, language, and unequal access

AI models are built using training datasets and specific construction decisions made by their developers. Lack of representation of certain demographic groups, local languages, healthcare systems, or disease states can lead to disparities in the quality of education. Furthermore, students with financial resources will have an advantage in using more advanced systems over those who rely on free versions of the same software. The gap between the two will widen even further due to differences in the ability to respond to cues. This can lead to disparities in both access to education and the ability to use such technology effectively. Therefore, an institution that promotes the use of AI-assisted learning should offer a basic level of access and education, common to all students, as well as the option to use a different, non-commercial platform.

Table 3. Risk domains and proportionate safeguards.

Risk domain	Educational consequence	Student safeguard	Institutional safeguard
Hallucination	Incorrect knowledge becomes encoded	Verify claims against named authoritative sources	Provide approved resources and teach evidence checking
Automation bias	Students adopt an incorrect recommendation	Commit to an answer before viewing AI output	Use discrepancy cases along with assess justification
Cognitive offloading	Reduced Retrieval and independent reasoning	Use AI after an initial attempt	Design assessments that expose the reasoning process
Integrity ambiguity	Authorship and competence become unclear	Disclose how the tool used	Publish task-specific rules and redesign vulnerable assessments

Confidentiality	Patient information may expose	Never enter identifiable clinical data	Restrict clinical use to approved systems and train students early
Bias and inequity	Uneven or discriminatory learning support	Compare perspectives and recognize missing populations	Audit tools, ensure baseline access, and retain non-AI routes

6. Balancing Educational Potential and Risk

This risk-benefit balance cannot be described internally. It depends on the characteristics of the user, the task for which it is used, the capabilities of the model, and the educational context. The risk profile is different between a novice student requesting comparative analogies to understand a basic physiological process and an advanced student requesting treatment options for a particular patient. A student may be better than his or her peers at recognizing an incomprehensible recommendation, but a higher level of competence does not eliminate the bias of automation. With this in mind, governance should be adapted to the identified risks.




Fig.2: Benefit-Risk Matrix For Genai-Supported Learning

Figure 2 Benefit–risk matrix: educational value rises when cognitive engagement and oversight are both strong. Four criteria can be used to distinguish between appropriate and inappropriate use of generative artificial intelligence (GAI) in medical education. First, instruction should be distinguished from complete solutions, because in the former, the tool merely supplements the work of students and teachers, while in the latter, it replaces them. Second, formative methods should be separated from summative assessments, i.e. the GAI should not be used to create final exams. Third, it is important to distinguish between general educational content and content containing patient data. Finally, fluent but unauthentic responses should be recognized as such. These standards are fundamental to effective governance, as they allow for addressing specific situations rather than issuing prohibitions or permits.

6.1 Perceived usefulness is not proven learning.

While positive attitudes have emerged as a result of education, they are not the same as its source. As research findings suggest, structured AI education and ethical concerns were in demand (Jackson et al., 2024). Similarly, Tran et al. (2025) highlight widespread use and interest in learning about GenAI

despite awareness of potential errors and biases. These findings present a paradox, as evidenced by the fact that knowledge of deception does not necessarily translate into the ability to detect it in practice, especially in complex cases. While the survey results are informative for designing a competent curriculum, they cannot, on their own, demonstrate that GenAI education is intrinsically linked to higher levels of competence.

6.2 Model competence is not the result of learning

Results from certification exam-style questions suggest that complex language models have the potential to generate clinically relevant responses (Gilson et al., 2023; Kung et al., 2023). This is important, but it is important to note that this does not necessarily mean that these responses are beneficial for learning. A correct answer should be evaluated in the context of the learning process itself. A student who studies and critically analyzes the explanation provided by the AI will benefit more than one who simply copies the correct answer and becomes accustomed to relying on external help.

6.3 The value of the aid changes the moment it is given

The timing of students' interaction with AI assistance plays a crucial role in determining whether such technology promotes or hinders knowledge acquisition. When working on a given problem, receiving AI assistance after attempting to solve the task independently allows students to develop initial solutions to compare with the AI's proposed solution. This makes it easier to identify inconsistencies and potential errors. In contrast, interacting with an AI assistant that provides a solution to a specific problem at the beginning of the task leads to a cognitive bias known as anchoring. In fact, Kriplani et al. (2025) found that students may incorporate incorrect information from AI assistants into their solutions. Therefore, the order of the task is far from a secondary consideration in learning. A practical implication of this study is the following sequence of actions: attempt to reason independently, receive AI assistance, and finally, compare your solution with the solution suggested by the AI assistant.

6.4 Performance can mask cognitive decline.

While time savings are positive, there are some drawbacks to medical training that should not be eliminated. Remembering a procedure, setting up a differential diagnosis, and explaining why one manifestation is less likely than another is not a waste of time, as it allows for the synthesis of information to yield useful results. General AI is useful when it eliminates unnecessary pain. It ceases to be useful when it eliminates the need to perform the work required to learn. Therefore, efficiency should not be measured by the time saved, but by the ability to retain learning.

Table 4. A practical traffic-light model for student use

Category	Examples	Expected control
Green: supportable	Generating extra practice prompts; changing explanation level; language clarification; simulated low-stakes dialogue	Student verification and disclosure when required
Amber: conditionally supportable	Feedback on reasoning; case generation; literature summarization; drafting study materials; differential-diagnosis practice	Independent attempt first, source checking, faculty guidance, no patient identifiers

Red: inappropriate or unsafe	Uploading identifiable patient data; relying on unverified treatment advice; undisclosed completion of assessed work; using AI where explicitly prohibited	Do not use; apply institutional academic or clinical governance procedures
---------------------------------	--	--

7. Proposed Risk-Benefit-Governance Framework

This proposed framework focuses education on educational goals rather than means. Each student comes with different levels of subject matter expertise, artificial intelligence (AI) knowledge, personal involvement, and sensitivity to misconceptions, while each task presents different levels of risk and cognitive demand. Generative AI offers teachers and students personalized support, increased productivity, formative feedback, and opportunities for imitation, but it also presents risks such as promoting misinformation, reinforcing biases, creating dependency, invading privacy, and devaluing intellectual performance.

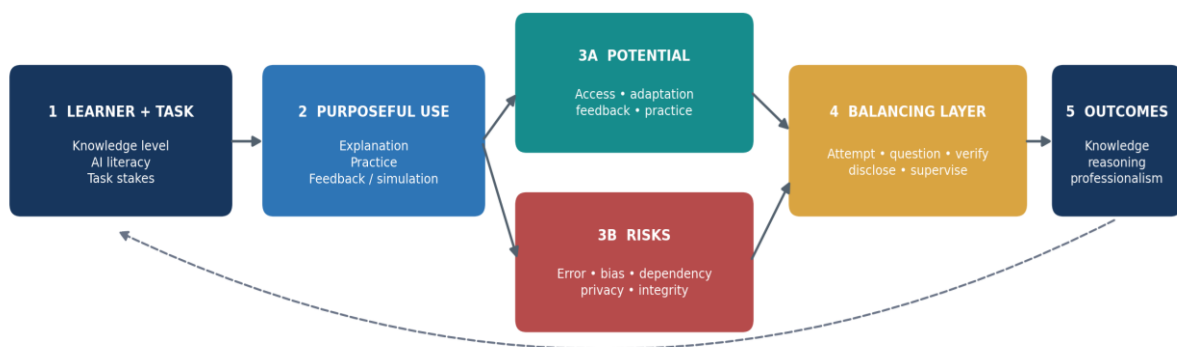


Fig.3: Risk–Benefit–Governance Framework for GenAI-supported medical learning

Critical engagement involves solving the task independently before asking for AI’s help, explaining the reasoning rather than simply providing the answer, comparing AI’s answers with credible sources, acknowledging knowledge gaps or uncertainties, and critically considering inconsistencies in information. Governance includes integrating AI into curricula, teacher training programs, formally approving specific AI tools, ensuring data privacy, establishing transparency requirements, and adopting evaluation and outcome monitoring. With critical engagement and good governance, creative AI tools can help expand knowledge, develop clinical reasoning skills, teach professional competencies, and promote self-directed learning. Without them, time-saving automation can mask superficial learning or unnecessary reliance on AI responses.

The model follows the following logic: socialization of student performance → intentional use of created AI → emerging benefits and risks of learning → meaningful student participation as well as community leadership → outcomes related to academic and professional practice → continuous iterative evaluation. The results of the evaluation should inform the further development of the supporting infrastructure as well as economic and social policies. This is due to the increasing influence of large language models and the need to update learner behavior accordingly.

8. Implications for Medical Education

Medical schools must move beyond a superficial, general understanding of AI to achieve practical, applied literacy in the field. Students take advantage of the opportunity to analyze flawed AI outputs, differentiate between model-generated text and reference work, and explain unsafe AI recommendations. While mentoring in a master of law program can yield satisfying results, the ability of a mentor is no substitute for certification or professional accountability.

Developing teachers' skills is equally important. Teachers cannot advise students to use AI to support their learning if they themselves do not understand how to use it. Furthermore, varying instructor expectations create confusion among students. Schools recommended providing examples of permitted, prohibited, and acceptable uses of generative AI for specific tasks. Teachers should consider alternative assessments that prevent AI from generating responses that are indistinguishable from a real student. Oral exams, iterative written tasks, observation of skill acquisition, individual case studies, and reflective written justification of reasoning are assessment options that maintain validity while acknowledging the reality of these AI tools.

Final Enforcement Policy areas to consider include mandatory disclosure, patient confidentiality, certified artificial intelligence (AI), authorship standards, liability for AI errors, and penalties for unauthorized use. Furthermore, appropriate educational experiences require institutional protection. Otherwise, students will be reluctant to report academic misconduct or seek professional help. Most importantly, effective institutional policies should be reasonably comprehensive to ensure the feasibility of implementation, while also being precise enough to guide practical behavior, thus striking a balance between constraints and oversight.

9. Future Research Directions

The current research is limited by the fact that it primarily examines the effectiveness of interventions through cross-sectional behavioral surveys and model performance tests on isolated outcome measures. Important domains for research include retention of learned information, transfer of acquired knowledge to new domains, calibration of diagnostic certainty, the nature of uncertainty-based dialogue, and unsupported student performance following instruction. To enable accurate comparisons between studies, it is important to report in detail all of these variables, including the name and version number of the AI model, the method of inclusion, the domain, the level of students, the control intervention, the follow-up period, and the teacher support.

Without this information, it is difficult to interpret and compare the results of individual studies. Therefore, the researchers concluded that, without this information, it is difficult to interpret and compare the results of individual studies.

Further studies are needed to determine which student groups benefit from generative AI tools and which are negatively impacted. Differences in benefits may be linked to variations in students' prior language skills, academic readiness, access to resources, disabilities, and digital self-efficacy.

Further qualitative studies are needed to examine how students integrate the use of AI-generated information and analyze changes in their study habits over time.

Finally, additional longitudinal studies are needed to understand whether reliance on AI in the early stages of learning affects students' attitudes, responsibilities, and behaviors related to patient safety during practical training.

10. Limitations

This conceptual review is based on an evolving literature and therefore cannot offer definitive conclusions. Most of the published evidence comes from individual perspectives, single-center studies, short-term trials, and results from the launch of specific AI models. Positive findings may reflect the novelty of the model and may not generalize to other curricula and healthcare sectors. The conceptual framework presented in this article is preliminary and requires further research. It will likely be refined based on longer-term and comparative studies in the future.

11. Conclusion

Generative AI can offer real educational benefits. However, speed, fluency, and prevalence are not reliable indicators of its true value. Combined with appropriate educational goals, it can provide

students with greater access to explanations, additional practice, and more effective feedback, so the potential for misuse is real. Unconditionally accepting or completely rejecting generative AI is not the solution.

In medical training, generative AI systems should be considered as aids to supervised learning rather than as final authorities. Students should be able to identify which types of questions are appropriate for generative AI tools, challenge them when they provide incorrect information, verify their answers are correct, report when they use them, and refrain from using them when they recommend inappropriate actions. Faculty should be able to conduct intellectually rigorous activities that clarify clinical reasoning, while institutions should take responsibility for protecting confidentiality and ensuring that procedures are fair and pedagogically effective. It is important to note that the goal is not to train future doctors who are able to produce well-crafted automated responses. Medical training should teach students to evaluate the reliability of responses, defend their clinical decisions, and understand that they are personally responsible for their actions.

References

- Boscardin, C. K., Gin, B., Golde, P. B., & Hauer, K. E. (2024). ChatGPT and generative artificial intelligence for medical education: Potential impact and opportunity. *Academic Medicine*, 99(1), 22–27. <https://doi.org/10.1097/ACM.0000000000005439>
- Çiçek, F. E., Ülker, M., Özer, M., & Kiyak, Y. S. (2025). ChatGPT versus expert feedback on clinical reasoning questions and their effect on learning: A randomized controlled trial. *Postgraduate Medical Journal*, 101(1195), 458–463. <https://doi.org/10.1093/postmj/qgae170>
- Gilson, A., Safranek, C. W., Huang, T., Socrates, V., Chi, L., Taylor, R. A., & Chartash, D. (2023). How does ChatGPT perform on the United States Medical Licensing Examination? The implications of large language models for medical education and knowledge assessment. *JMIR Medical Education*, 9, e45312. <https://doi.org/10.2196/45312>
- Gordon, M., Daniel, M., Ajiboye, A., Uraiby, H., Xu, N. Y., Bartlett, R., Hanson, J., Haas, M., Spadafore, M., Grafton-Clarke, C., Gasiea, R. Y., Michie, C., Corral, J., Kwan, B., Dolmans, D., & Thammasitboon, S. (2024). A scoping review of artificial intelligence in medical education: BEME Guide No. 84. *Medical Teacher*, 46(4), 446–470. <https://doi.org/10.1080/0142159X.2024.2314198>
- Jackson, P., Sukumaran, G. P., Babu, C., Tony, M. C., Jack, D. S., Reshma, V. R., Davis, D., Kurian, N., & John, A. (2024). Artificial intelligence in medical education: Perception among medical students. *BMC Medical Education*, 24, 804. <https://doi.org/10.1186/s12909-024-05760-0>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kirpalani, A., Grimmer, J., & Wang, P. Z. T. (2025). Exploring medical students trust in generative artificial intelligence (ChatGPT) versus peers in team-based learning. *Medical Science Educator*, 35(2), 671–674. <https://doi.org/10.1007/s40670-025-02289-9>
- Kung, T. H., Cheatham, M., Medenilla, A., Sillos, C., De Leon, L., Elepaño, C., Madriaga, M., Aggabao, R., Diaz-Candido, G., Maningo, J., & Tseng, V. (2023). Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health*, 2(2), e0000198. <https://doi.org/10.1371/journal.pdig.0000198>
- Masters, K. (2023). Ethical use of artificial intelligence in health professions education: AMEE Guide No. 158. *Medical Teacher*, 45(6), 574–584. <https://doi.org/10.1080/0142159X.2023.2186203>

Preiksaitis, C., & Rose, C. (2023). Opportunities, challenges, and future directions of generative artificial intelligence in medical education: Scoping review. *JMIR Medical Education*, 9, e48785. <https://doi.org/10.2196/48785>

Qin, Y. (2026). Reconceptualizing Standards for Identifying Copyright Misuse in the Age of Artificial Intelligence: Innovations and Controversies. *Journal of System and Management Sciences*, 16(1), 78-89. <https://doi.org/10.33168/JSMS.2026.0105>

Sallam, M. (2023). ChatGPT utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. *Healthcare*, 11(6), 887. <https://doi.org/10.3390/healthcare11060887>

Tran, C., Hryciw, B. N., Moore, S. W., Chaput, A., & Seely, A. J. E. (2025). Perceptions and use of generative artificial intelligence in medical students: A multicenter survey. *Journal of Medical Education and Curricular Development*, 12, 23821205251391969. <https://doi.org/10.1177/23821205251391969>

UNESCO. (2023). Guidance for generative AI in education and research. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>

World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. <https://www.who.int/publications/i/item/9789240029200>

World Health Organization. (2025). Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. <https://www.who.int/publications/i/item/9789240084759>