From Illusion of Competence to Cognitive Debt: Confronting AI Hallucination and Dependency in Medical Student

Ilham Kurniawan [1*], Fairuz Haniyah Ramadhani [1], Aulia Umi Rohmatika [1]

[1] Faculty of Medicine, Universitas Pembangunan Nasional Veteran Jawa Timur, Surabaya, 60294, Indonesia

Corresponding Author

Ilham Kurniawan

Faculty of Medicine, Universitas Pembangunan Nasional Veteran Jawa Timur, Surabaya

Email: ilham_kurniawan.fk@upnjatim.ac.id

Preamble

The pages of *Journal of Diverse Medical Research: Medicosphere* have always been a space where medical students learn to think on paper to wrestle a clinical question into a coherent argument, to trace a claim back to its source, to discover that writing is not the transcription of thought but its formation. That space is now being quietly reshaped. Generative artificial intelligence has entered the writing process not as an occasional aid but as a default first responder: the first draft, the first summary, the first reference list. The concern that brings us to this editorial is not that AI exists, nor that students use it. It is that the line between assistance and substitution has blurred and that blurred line is producing a generation of scholarly work marked by fluent falsehood and a deepening dependency that students themselves are beginning to notice.

The most concrete manifestation of this problem is hallucination AI outputs that are plausible, confident, and wrong. In an exploratory qualitative study of London medical students, *every* interviewee recounted experiences with outdated or hallucinated outputs, and 53.3% specifically encountered fabricated research references that could not be traced to real sources. Students described the difficulty of detecting these errors precisely because the fabricated content "often seemed plausible when incorrect" [1]. This is not a fringe phenomenon. A systematic review and meta-analysis of large language model citation generation for biomedical topics found a pooled completely fabricated proportion of 18.6% (95% CI 12.8–26.3), with fully accurate references comprising only 29.4% [2]. Some analyses of medical references generated by widely used AI platforms have placed fabrication rates at more than one-third [3]. When a student submits a manuscript laced with invented citations, the harm is not merely formal. It introduces false building blocks into the literature, erodes trust in the student's work, and most troublingly allows the student to bypass the very exercise that makes scientific writing educational: the disciplined act of finding, reading, and

evaluating real evidence.

But hallucination is only the visible symptom. Beneath it lies a quieter and arguably more consequential process: cognitive offloading and the accumulation of "cognitive debt." Neurophysiological research combining EEG, linguistic analysis, and interviews has shown that early delegation to a large language model is associated with lower neural connectivity in alpha and beta bands during writing, a reduced sense of authorship, and impaired recall few participants could quote their own AI-assisted text minutes later. When the AI was withdrawn, those habituated to it performed *worse* than peers who had not used it. This is what the emerging literature calls cognitive debt: the short-term performance gain that masks a longer-term loss of durable learning [4]. In medical education, where physiology, pharmacology, and clinical reasoning depend on robust encoding and retrieval under uncertainty, this trade-off is not neutral. A medical student who can produce a polished paragraph without being able to reconstruct its argument is not learning to be a doctor; they are learning to be a conduit.

The dependency dimension is increasingly acknowledged even by students. In one survey of medical students, a limited but telling number expressed concern about becoming dependent on ChatGPT, potentially hindering critical thinking skills [5]. The London study likewise identified themes of overreliance and cognitive outsourcing, with some participants describing a self-perceived addiction and a fear that instant answers could diminish perseverance with complex material [1]. A reply in the pathology literature citing a study of 580 university students reported that AI dependence was associated with reduced critical thinking, mediated in part by cognitive fatigue [6]. The pattern is consistent: the tool that promises to accelerate learning can, under conditions of unguided use, accelerate the *erosion* of the capacities that learning is supposed to build.

These concerns acquire a particular urgency in our own institutional context. Many of us have recently participated in AI training designed to expand students' fluency with these tools training that has rightly emphasized capability, efficiency, and the responsible use of AI as a professional skill. But capability without guardrails is not literacy; it is acceleration without direction. **Fairuz Haniyah Ramadhani**, Assistant Editor, put the editorial team's worry sharply: "We have spent months telling students that AI is a tool they must learn to use. We have not spent nearly enough time telling them what AI must never be allowed to replace. The manuscripts arriving at *Medicosphere* suggest that students have heard the first message and missed the second. The result is not fluency. It is a kind of learned helplessness dressed in polished prose."

Aulia Umi Rohmatika, Reviewer, has seen this pattern from the closest vantage point the evaluation of submitted manuscripts: "In a growing number of submissions from medical students, the AI use is not incidental; it is structural. The prose is fluent, the structure is conventional, and the citations are plentiful but when you follow the citations, they dissolve. What is missing is not intelligence but *engagement*: the student has not read the sources, has not grappled with the contradictions, has not made the argument their own. Reviewing such work is not reviewing scholarship; it is reviewing an output. The authors are often unaware that this is what they have submitted, because the tool has given them the illusion of competence without the experience of it."

This is the heart of our concern. The problem is not that medical students are using AI. The problem is that unguided, unfiltered, front-loaded AI use is producing a counterfeit of scholarly formation a performance of writing without the cognitive work that writing is supposed to do. And the counterfeit is hardest to detect precisely where it matters most: in the early stages of training, when students lack the domain knowledge to recognize a hallucinated citation or a subtly incorrect claim.

What, then, should *Medicosphere* do? We do not believe the answer is prohibition. AI tools are already woven into students' study routines, often through peer demonstration and the sheer accessibility of device-agnostic platforms. A journal that pretends otherwise would be irrelevant. Nor do we believe the answer is uncritical celebration. The answer must be structured, sequenced, and accountable use what the cognitive evidence supports as a "first think and draft on your own, only afterwards use AI to check, expand, or compare" model [4]. AI literacy belongs in the curriculum, but it must be taught as *verification literacy* and *authorship literacy*, not merely as prompting proficiency.

For *Medicosphere*, this means four commitments. **First**, we will require explicit disclosure of AI use at submission, specifying the tool, the purpose, and the extent of use, in line with ICMJE guidance that AI cannot be listed as an author and that authors retain full accountability for all content, including references [7]. **Second**, we will treat reference verification as a non-negotiable author responsibility: every citation must be traceable to a real, accessible source, and reviewers will be instructed to spot-check references for fabrication and metadata accuracy. **Third**, we will not accept manuscripts in which the core intellectual work problem formulation, argument construction, and critical synthesis has been delegated to an AI system; AI may assist with language, formatting, and exploratory search, but the scholarly judgment must remain human. **Fourth**, we will use this editorial as an opening, not a verdict: we invite students, mentors, and reviewers to engage with us in

building a culture where the question is not “Did you use AI?” but “Can you defend every claim, every citation, and every sentence as your own?”

The task of medical education is not to produce efficient text generators. It is to produce physicians who can reason under uncertainty, evaluate evidence, and communicate with precision and integrity. Writing is one of the primary sites where that capacity is built. If we allow AI to hollow out that site in the name of convenience, we will not have saved students time; we will have stolen from them the very practice that makes them clinicians and scholars. *Medicosphere* chooses to hold that line not against AI, but for the human work that AI cannot do for them.

Declaration of Generative AIin Scientific Writing

The authors affirm that the usage of generative artificial intelligence tools, including DeepL and DeepSeek, is intended solely for language refinement and grammatical correction throughout the manuscript drafting process. These tools were not used for generating content, data analysis, or interpretation. The authors are fully responsiblefor the accuracy and integrity of the content as well as the conclusions presented in this study

References

1. Alazzawi M, Lam K. Uptake of Large Language Models by London Medical Students: Exploratory Qualitative Interview Study. JMIR Form Res 2026;10:e82828; doi: 10.2196/82828.
2. Malikian R. Fabricated and inaccurate citations generated by large language models for biomedical topics: a systematic review and meta-analysis. Zenodo 2026; doi: https://doi.org/10.5281/zenodo.23584772.
3. The Royal College of Surgeons of England. More than a third of medical references generated by some AI platforms may be fabricated, new study finds [Internet]. London: The Royal College of Surgeons of England; 2026 Apr 14 [cited 2026 Nov 15]. Available from: https://www.rcseng.ac.uk/news-and-events/media-centre/press-releases/ai-chatbots-paper-2026/
4. Tenajas R, Miraut D. Rapid response: When and how to use AI in studying medicine: insights from cognitive research. BMJ 2025;390:r1418/rr-4. doi: 10.1136/bmj.r1418.
5. Abdelhafiz AS, Farghly MI, Sultan EA, et al. Medical students and ChatGPT: analyzing attitudes, practices, and academic perceptions. BMC Med Educ 2025;25(1):187; doi: 10.1186/s12909-025-06731-9.
6. Tazelaar HD, Aubry M-C, Roden AC, et al. In Reply. Arch Pathol Lab Med

2026;150(7):499–500; doi: 10.5858/arpa.2026-0139-LE.

7. International Committee of Medical Journal Editors (ICMJE). Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals. 2025.