

INVITED EDITORIAL

Artificial intelligence for medical students: a review of current challenges, opportunities and future directions

Chris Paton

Abstract

As artificial intelligence (AI) technology becomes more widely adopted into clinical practice, it is increasingly important that medical students are aware of how AI technologies work and how to appraise the literature on AI safety and effectiveness. This article reviews the challenges that AI provides for the future health workforce as it re-adjusts to a world where “intelligence” may come from computers rather than humans, and where there are significant risks for patient privacy and safety. However, AI also holds great promise for improving care for patients. As future doctors, medical students should grasp the opportunity now to learn about how to best use AI to provide equitable and effective care for their patients.

Article

Artificial Intelligence (AI) is a term used to describe a set of digital technologies that perform functions that we normally associate with human intelligence such as perceiving, reasoning, and learning.¹ AI systems have been in use now for many decades. There have been several periods of excessive AI hype followed by “AI winters” where the software failed to live up to expectations and the hype died down.

We are currently living a new period of AI hype. This began in the 2010s when computer scientists such as Geoffrey Hinton and Yann Le Cun developed technologies that enabled the application of deep-learning neural networks to solve real-world problems such as image recognition and more accurate prediction models.² The hype reached new levels in 2022 with the launch of ChatGPT, a transformer-based deep-learning neural network model that was able to generate plausible text in response to user inputs.³ The performance of ChatGPT and other similar “Large Language Models” (LLMs) from Google, Meta and others have been so impressive that many (including Hinton) are now speculating that we are entering an age of Artificial General Intelligence (AGI), where AI systems can do everything a human brain can do.⁴ AGI is predicted to proceed to Artificial Super Intelligence (ASI), where AGI systems are able to rewrite their own code, rapidly improving and overtaking human intelligence.⁵

What does this mean for medical students? Will AGI (or ASI) machines replace doctors in the future? How will medical practice change? Or will we enter a new AI winter, where the overblown hype of AGI and ASI will die down as it fails to materialise.

Current AI research seems to indicate that we are still on course for AGI within the next few years. New technologies are being developed that can overcome some of the limitations of LLMs such as “hallucination” (errors) and excessive sycophancy.⁶ However, as Hinton himself has pointed out,⁷ healthcare has unlimited demand.

Until we solve aging, humans will continue to get older and sicker as they progress through life and will want an ever-increasing quantity of healthcare services. Many healthcare systems already lag far behind what most people around the world would consider adequate. This means that even as AI systems take over more and more medical tasks, there will continue to be a role for humans in many aspects of healthcare. The idea of the doctor being the fount of all healthcare knowledge will probably diminish as AI systems become more reliable and informative for patients than talking to a human doctor. However, it is likely that patients will continue to want a human being to care for them in times of vulnerability and illness.⁸

During this transition period, it is important that medical students equip themselves with as good an understanding of how AI works as they have of how the human body works. Without this knowledge, they will be flying blind in a new AI world.^{9,10} Students should learn about how AI systems used in healthcare have been developed: What are the different model architectures and how do they work? How do you conduct a scientific evaluation of an AI system? Above all, they should consider how AI systems will impact patients and they must ensure that AI systems are implemented in a way that meets patient needs.¹¹ Medical schools are beginning to implement courses in AI in healthcare, such as the “Artificial Intelligence in Healthcare” course we now offer as part of the Postgraduate Diploma in Digital Health at the University of Auckland. However, as AI becomes further integrated into all medical specialties, it is likely that clinical educators will need to continuously update their curricula to ensure that the students they teach know how to appraise and use the growing number of AI systems that they will use to care for patients in the future.

References

1. Paton C. Textbook of digital health [Internet]. Chris Paton; 2024 [cited 2025 July 7]. Available from: <https://textbookofdigitalhealth.com/>
2. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–44. <https://doi.org/10.1038/nature14539>
3. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, et al. Language models are few-shot learners. *arXiv*. 2020. <https://doi.org/10.48550/arXiv.2005.14165>
4. Bengio Y, Hinton G, Yao A, Song D, Abbeel P, Darrell T, et al. Managing extreme AI risks amid rapid progress. *Science*. 2024;384(6698):842–5. <https://doi.org/10.1126/science.adn0117>
5. Bostrom N. *Superintelligence: paths, dangers, strategies*. London (GB): Oxford University Press; 2014. p. 432.
6. Wei J, Wang X, Schuurmans D, Bosma M, Chi EH, Xia F, et al. Chain of thought prompting elicits reasoning in large language models. *arXiv*. 2022. <https://doi.org/10.48550/>

arXiv.2201.11903

7. Tecotzky A. The Godfather of AI reveals which jobs are safest — and where “everybody” will get replaced [Internet]. Business Insider; 2025 [cited 2025 July 7]. Available from: <https://www.businessinsider.com/geoffrey-hinton-godfather-of-ai-safe-jobs-2025-6>
8. Morley J, Floridi L. An ethically mindful approach to AI for health care. *Lancet*. 2020;395(10220):254–5. [https://doi.org/10.1016/S0140-6736\(19\)32975-7](https://doi.org/10.1016/S0140-6736(19)32975-7)
9. Paton C. Understanding research on artificial intelligence in healthcare. *BMJ Med*. 2025;4(1):e001614. <https://doi.org/10.1136/bmjmed-2025-001614>
10. Paton C. Welcome to BMJ digital health & AI. *BMJ Digital Health & AI*. 2025;1(1). <https://doi.org/10.1136/bmjdhai-2024-000004>
11. Center for Devices, Radiological Health. Good machine learning practice for medical device development: guiding principles [Internet]. U.S. Food and Drug Administration; 2023 [cited 2025 July 7]. Available from: <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles>

About the author

> Chris Paton is an Associate Professor at the Liggins Institute and Department of Epidemiology and Biostatistics at the University of Auckland. He trained in medicine in the United Kingdom and worked in both the NHS and New Zealand healthcare system prior to entering academia. He is a current member of the Health Systems Collaborative at the University of Oxford, the Editor in Chief of *BMJ Digital Health & AI*, and the author of the *Textbook of Digital Health*.

Correspondence

Chris Paton: chris.paton@auckland.ac.nz