

In AI we trust? A promised panacea in the consult culture of modern medicine

Christopher J Peterson, MD, MS

The hype around AI comes with lots of promises. Quick answers, rapid data processing, new solutions, new possibilities. This is certainly appealing to a medical community that is inundated with massive amounts of data, heavy patient loads, a veritable mountain of medical literature, and no shortage of administrative demands. To physicians, the promise is this—AI will process this mass of information and reduce it to clear, concise recommendations based on the best evidence to guide clinical practice. While this seems revolutionary, there is something quite familiar in these promises to streamline, summarize, and simplify. And that's because this already exists—as society guidelines, clinical algorithms, and medical reference texts. Millions of peer-reviewed articles, including highly focused expert reviews, are readily available in medical databases.¹ Clinical reference guides like UpToDate and DynaMed provide highly focused clinical recommendations to practicing clinicians.² Institutional protocols have further adapted workflows and processes to meet local or in-house needs and fill in gaps in societal guidelines. The Physicians' Desk Reference, once a protean symbol of medical knowledge and authority on a physician's bookshelf, now effectively rests in one's pocket.

And yet, despite all of this readily available information, physicians continue to rely heavily on expert consultation.^{3–5}

There are multiple reasons for a medical culture that relies heavily on specialist consultation. Young physicians and physicians-in-training may consult out of uncertainty; they may have the knowledge (or know where to find it), but having a consultant provide a stamp of approval provides reassurance to a clinician who

otherwise has limited clinical experience. Physicians overwhelmed with a large census list may consult out of necessity; a hospitalist with a list of 30 patients may have the cognitive ability to address most medical concerns, but not the time or energy. Hospital systems may require auto-consultation for a particular problem (e.g., *S. aureus* bacteremia), regardless of the primary physician's skill or expertise. Physicians may consult out of fear of "missing something". This is likely driven, in part, by a highly litigious and metrics-oriented medical system. The argument is that an expert opinion is needed to determine whether an expert-driven intervention is appropriate (e.g., "They probably won't be a candidate for surgery, but we'll give surgery the chance to decline"). Nevertheless, specialists have undoubtedly been consulted for their "blessing" or "to weigh in" on a particular course of action, though the primary team likely already knows what to do and is doing just that. Instead, they want a consultant to assume a level of responsibility for which the consultant's expertise provides additional protection for the primary team. Practicing medicine is an increasingly complex endeavor, with more and more subspecialties being created.^{6,7} Finally, there are some consultations made simply (and unfortunately) out of pure convenience.

Regardless of the reason, reliance on experts has and continues to be a critical practice of medicine. More than just a byproduct of the medical system, consulting experts have an emotional and psychological element. For example, physicians will not necessarily take advantage of readily available medical knowledge outside of their expertise, as doing so takes time and requires them to act and assume responsibility outside of their experience. While guidelines are abundant, they do just that—*guide* physicians. They don't dictate practice (though some may feel otherwise), may not apply to the specific patient at hand, and are often insufficiently nuanced for highly complex clinical scenarios. What's more, these guidelines, for better or

Corresponding author: Christopher J Peterson
Contact Information: Christopher.Peterson@unchealth.unc.edu
DOI: 10.12746/swjm.v14i59.1631

worse, may be missed, ignored, or rejected by physicians to some degree.⁸ Physicians, therefore, don't just want information—they want experience, reassurance, and evaluation of the patient in front of them.

AI will change none of this.

First, the information provided by AI isn't necessarily better, and is arguably worse, than what is already available to physicians.⁹ While AI may be able to provide summative information, it does so without any human oversight. Guidelines committees, while at times criticized for heavy reliance on expert opinions, are nonetheless developed after careful deliberations among experts, critical analysis of the evidence, and these results are available for public scrutiny and discussion (at least in theory).¹⁰ There is also a self-awareness of the quality of evidence ("strong recommendation, moderate-quality evidence"). Healthcare-focused AI algorithms like Open Evidence have none of this—no real-world clinical experience, no expertise, no self-awareness of the quality or limitations of the evidence (other than that which is already stated in the literature). While AI may cite sources, its sources are often few (there are some exceptions to this—Open Evidence now offers a "DeepConsult™" which provides, as the name suggests, a more in-depth analysis of a subject).¹¹ These "outputs" (a more fitting descriptor than "recommendations") are not vetted by experts—they are simply a compilation of what is already available. The output may also change depending on how the question is asked or based on the latest iteration of the AI model. If physicians struggle to apply current guidelines or to search readily available clinical information, it seems unlikely that cursory summaries by AI will be more persuasive.

Second, AI can't take any responsibility for its recommendations. This is both from a legal standpoint and a philosophical one. For example, AI companies include disclaimers and terms of use that certainly distance themselves from any responsibility (concerningly, these disclaimers have begun to decline).¹² Companies will not want to assume responsibility for an AI giving inaccurate advice, especially when there is a human operator (a physician) making the ultimate call (and able to take the ultimate blame). From a philosophical standpoint, the AI itself cannot take responsibility for a

recommendation because it is not a conscious entity. Large language models produce predictive results based on the data they are trained on; they do not possess an understanding of the data they process or an awareness of what is being recommended. In other words, AI can't be morally responsible for its outcomes any more than Microsoft Word is responsible for a bad grade on your term paper. While it may be just as easy (or easier) to "blame" an inanimate object than a person, it would be impossible to hold AI responsible in the same way a physician is held responsible for negligence. How would AI defend "itself" in court? Or articulate the inner reasoning behind its output? To this point, AI remains a "black box" of reasoning, with the predictions used to create the output typically unavailable to the user (and even the programmers).¹³ While the server may be able to recall the output for a particular recommendation, it will not be able to reveal its inner "reasoning" for this recommendation. It can't explain why it gave the result it did. Proving that an AI model acted irresponsibly (such as giving advice outside of the standard of care) will be challenging¹⁴. Even with the evolving legal and ethical landscape of AI, developers will almost certainly minimize their responsibility and maximize the physician's.¹⁵

Finally, AI is neither an expert nor a colleague, and can't provide a satisfactory reassurance to those practices in the real world. While working overnight as a resident, I would on occasion ask my colleagues to weigh in on a clinical situation. I could reference guidelines, I'd studied papers and textbooks, but there was something reassuring about being able to discuss the case with another physician—someone who understands the mental, emotional, and logistical challenges of medicine. Someone who could contemplate the nuances of the case, draw on their real-world experience, and tell me what they would do in this specific situation. Someone who had felt the weight of a similar scenario on their back. This was not only reassuring, but proved to be a crucial exercise for treating complex patients, whose care often involves no easy answers or algorithms to guide physicians. There was comfort in a shared understanding made possible by a shared humanity. In contrast to my colleagues, AI has no sense of the world we live in, no understanding of what it's like to be a physician,

and can offer no real-world experience (other than the experiences of others that it has consumed in model development). It can offer quick answers to particular questions, but it's doubtful these answers will be much more reassuring than the clinical references that we already have, already ignore, or already don't trust (or at least, don't feel are sufficient without the help of colleagues). While we can ask AI anything, this doesn't mean it will tell us what we need or want to hear, nor that we'll trust it.

In the end, AI does not change or solve the current medical climate where physicians seek out experts for their knowledge and reassurance. AI does not yet provide a reliable, trustworthy alternative to society guidelines or expert consultation. It cannot take responsibility for these recommendations. It cannot be considered an expert. Even more important, it cannot provide the insight, experience, and trust that colleagues and specialists provide. While it may provide improved searches and summaries of the medical literature, it will likely change our efficiency, not our fundamental practice. In today's highly subspecialized, metric-driven, and litigious medical climate, AI will not provide physicians the answers or reassurance they seek. It may help provide information more quickly, but when there is a true question or concern, there is unlikely to be an alternative to consulting and collaborating with a fellow physician.

Article citation: Peterson CJ. In AI we trust? A promised panacea in the consult culture of modern medicine. *The Southwest Journal of Medicine*. 2026;14(59):41–43

From: Department of Internal Medicine, University of North Carolina School of Medicine, Chapel Hill, NC (CJP)

Conflicts of interest: none

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

REFERENCES

- Ghasemi A, Mirmiran P, Kashfi K, et al. Scientific publishing in biomedicine: a brief history of scientific journals. *Int J Endocrinol Metab*. 2023;21(1):e131812. DOI: 10.5812/ijem-131812.
- Ventola CL. Mobile devices and apps for health care professionals: uses and benefits. *P t*. 2014;39(5):356–64.
- Hailu R, Wilcock AD, Zachrisson KS, et al. National trends in the use of specialty consultations in emergency department visits, 2009 to 2019. *Annals of Emergency Medicine*. 2023; 82(5):634–5. DOI: 10.1016/j.annemergmed.2023.06.021.
- Barnett ML, Song Z, Landon BE. Trends in physician referrals in the United States, 1999–2009. *Arch Intern Med*. 2012;172(2):163–70. (In eng). DOI: 10.1001/archinternmed.2011.722.
- Jordan MR, Conley J, Ghali WA. Consultation patterns and clinical correlates of consultation in a tertiary care setting. *BMC Res Notes*. 2008;1:96. (In eng). DOI: 10.1186/1756-0500-1-96.
- Cassel CK, Reuben DB. Specialization, subspecialization, and subspecialization in internal medicine. *Mass Medical Soc*. 2011:1169–73.
- Karadacic R, Chan DC, Landon BE, et al. Subspecialization of surgical specialties in the US. *JAMA Health Forum*. 2025;6(9):e253192. DOI: 10.1001/jamahealthforum.2025.3192.
- Cabana MD, Rand CS, Powe NR, et al. Why don't physicians follow clinical practice guidelines?: A framework for improvement. *Jama*. 1999;282(15):1458–65.
- Takita H, Kabata D, Walston SL, et al. A systematic review and meta-analysis of diagnostic performance comparison between generative AI and physicians. *NPJ Digit Med*. 2025;8(1):175. (In eng). DOI: 10.1038/s41746-025-01543-z.
- Shekelle PG, Woolf SH, Eccles M, et al. Clinical guidelines: developing guidelines. *Bmj*. 1999;318(7183):593–6. (In eng). DOI: 10.1136/bmj.318.7183.593.
- OpenEvidence, the Fastest-Growing Application for Physicians in History, Announces \$210 Million Round at \$3.5 Billion Valuation. *Open Evidence* 2025.
- O'Donnell J. AI companies have stopped warning you that their chatbots aren't doctors. *MIT Technology Review* 2025.
- Xu H, Shuttleworth KMJ. Medical artificial intelligence and the black box problem: a view based on the ethical principle of "do no harm". *Intelligent Medicine*. 2024;4(1):52–7. DOI: 10.1016/j.imed.2023.08.001.
- Elamin S, Duffourc M, Berzin TM, et al. Artificial intelligence and medical liability in gastrointestinal endoscopy. *Clinical Gastroenterology and Hepatology*. 2024;22(6):1165–9.e1. DOI: 10.1016/j.cgh.2024.03.011.
- Mello MM, Guha N. Understanding liability risk from health-care AI. Policy Brief, Stanford University Human-Centered Artificial Intelligence <https://hai-production.s3.amazonaws.com/files/2024-02/Liability-Risk-Healthcare-AI.pdf> 2024.