

CORR Insights®: Can Artificial Intelligence Improve the Readability of Patient Education Materials?

Joseph Bernstein MD¹ 

Where Are We Now?

ChatGPT, in a commentary it wrote in *Clinical Orthopaedics and Related Research*® just a few months ago [1], claimed that although it could not perform orthopaedic surgery itself, it would be able to assist orthopaedic surgeons in their duties. Fulfilling that prediction did not take long.

In this issue of *CORR*®, Kirchner and colleagues [5] demonstrated that ChatGPT could ably “improve the reading accessibility of orthopaedic

surgery online patient education materials to recommended levels quickly and effectively.” This editing process indeed assists orthopaedic surgeons in a critical task for all medical doctors: informing patients about their conditions. It’s not for nothing that our title of “doctor” comes from the Latin word *docere*, meaning “to teach.” Informing patients about their condition—which includes providing a description of the diagnosis, its prognosis, the alternatives to treatment, and the risks and benefits of each option—is perhaps the most important job a physician has.

This teaching job can also be a difficult one. Doctors certainly know the information they must share, but they may know it too well. This so-called information asymmetry makes it hard to share the relevant facts at the level patients need to receive them.

In response to this problem, many medical organizations have created patient information materials that are intended to be simple and accessible. Yet, as Kirchner et al. correctly note, all too often these materials are far from being completely comprehensible: Most are written at a grade level far exceeding that of the average person.

Where Do We Need to Go?

The challenge to all who share information is to place it within reach. To that end, the method used by Kirchner and colleagues [5] gets us closer—but maybe not close enough. For one thing, an excessively high grade level is not the only impediment to understanding. Some concepts are hard, even when presented in the simplest English. (I still don’t understand quantum electrodynamics even though Richard Feynman’s lectures on the topic are presented in 7th grade English.) Some people may not learn best from reading. And even among those who read well, maximal learning is active and, indeed, interactive: The wise instructor asks questions and moves on only once the current step is mastered. Using grade level—appropriate patient materials is likely a necessary, but hardly a sufficient, step in the right direction.

How Do We Get There?


Kirchner and colleagues [5] have outlined the path to get us where we need to go: using something like ChatGPT, but not stopping at merely re-writing published materials. Maybe the ideal means to inform patients is a dialogue with an all-knowing artificial intelligence (AI) system: an Oracle of Informed Consent, suffused with patience and

This CORR Insights® is a commentary on the article “Can Artificial Intelligence Improve the Readability of Patient Education Materials?” by Kirchner and colleagues available at: DOI: [10.1097/CORR.0000000000002668](https://doi.org/10.1097/CORR.0000000000002668).

The author certifies that there are no funding or commercial associations (consultancies, stock ownership, equity interest, patent/licensing arrangements, etc.) that might pose a conflict of interest in connection with the submitted article related to the author or any immediate family members.

All ICMJE Conflict of Interest Forms for authors and *Clinical Orthopaedics and Related Research*® editors and board members are on file with the publication and can be viewed on request.

The opinions expressed are those of the writer, and do not reflect the opinion or policy of *CORR*® or The Association of Bone and Joint Surgeons®.

J. Bernstein , University of Pennsylvania, 110 Stemmler Hall, 36th Street and Hamilton Walk, Philadelphia, PA 19104, USA, Email: joseph.bernstein@uphs.upenn.edu

¹Clinical Professor of Orthopaedic Surgery, University of Pennsylvania, Philadelphia, PA, USA

unlimited time, deep knowledge, and mastery of pedagogy.

I must emphasize that the Oracle I have in mind is something like ChatGPT, but is not ChatGPT as we know it now. The current version has at least two fatal flaws:

1. ChatGPT confabulates. When it doesn't know something (or perhaps when it does know, but is just feeling mischievous), ChatGPT will just make things up. In one recent study, when a citation for its claims was requested, ChatGPT proffered a link to a nonexistent source [1]. A similar tendency has been well described in humans [4], but we justifiably demand more from an Oracle. In real life, instructional materials must be carefully vetted and limited to only correct statements. ChatGPT is not there yet.
2. ChatGPT is biased. When I asked the program to write an ode in honor of Joe Biden, it replied, "Certainly!" When I asked it to write an ode in honor of Donald Trump, it wrote, "Writing an ode in honor of any individual, including President Trump, would go beyond my programming to remain neutral and impartial." I must point out that friends who have tried this little experiment got different results (owing perhaps to the randomness built into ChatGPT, as discussed below), yet the general phenomenon of political bias has been observed in more rigorous testing, too [6]. The fact that the system has viewpoint biases has important implications for the sharing of medical information. Do we want our Oracle to channel its inner chiropractor when discussing spinal surgery? This is not just theoretical. In my

experience, ChatGPT is far too willing to endorse platelet-rich plasma injections and viscosupplementation. Moreover, even if ChatGPT were to limit itself to true statements, how the information is framed can influence the listener's judgement [3].

Even beyond those issues, it's not clear that ChatGPT will be cost effective. Right now, users can access ChatGPT without payment, but of course, it is not free. Its owners have decided (for now) to make it available without charges, but that policy might not endure. Also, systems that have remained free—for example, Google and Facebook—have kept a zero-dollar price tag because their owners have been able to monetize the information they collect. Vacuuming and selling personal information may be neither desirable nor legal in the realm of healthcare.

For now, the promise of AI in healthcare is vast but undefined. One particular problem is that randomness is an essential feature of programs like ChatGPT. In his essay, "What Is ChatGPT Doing...and Why Does It Work?" [7], Dr. Stephen Wolfram explains that ChatGPT is always trying "to produce a 'reasonable continuation' of whatever text it's got so far." Yet the most reasonable next word is not necessarily the most probable. As Dr. Wolfram notes, "this is where a bit of voodoo begins to creep in. Because for some reason—that maybe one day we'll have a scientific-style understanding of—if we always pick the highest-ranked word, we'll typically get a very 'flat' essay, that never seems to 'show any creativity'...but if sometimes (at random) we pick lower-ranked words, we get a 'more interesting' essay. The fact that there's

randomness here means that if we use the same prompt multiple times, we're likely to get different essays each time." This randomness is precisely what we don't want for medical information.

I am far too naïve about AI techniques to develop a robust workaround for that problem, but two ideas come to mind. First, in the near future, competitors of ChatGPT will be more common. These competitors could be used to critique the output of their rivals. Of course, that method won't help if the underlying material used by all systems for training contains biases. Thus, to address that, as well as the problem of randomness, I propose that after an informed consent dialogue, the AI program generates a quiz for the user to take and bring the results to his or her surgeon. The surgeon could quickly scan this and get a good understanding of what the patient knows, and in turn, identify areas of confusion or misunderstanding. The work of Kirchner and colleagues [5] clearly shows that there are some tasks AI can do right now. Soon, AI will play an even larger role in patient instruction. Search engines are commonly used for this task, and search engines have already begun to integrate AI into their algorithms. The challenge for the medical community will be ensuring that these systems live up to the promise ChatGPT described: to help, and not hinder, us in the care of our patients.

References

1. Alkaissi H, McFarlane SI. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus*. 2023;15:e35179.
2. Bernstein J. Not the last word: ChatGPT can't perform orthopaedic surgery. *Clin Orthop Relat Res*. 2023;481:651-655.
3. Bernstein J, Kupperman E, Kandel LA, Ahn J. Shared decision making, fast and slow: implications for informed consent, resource utilization, and patient

CORR Insights

- satisfaction in orthopaedic surgery. *J Am Acad Orthop Surg.* 2016;24:495-502.
4. Frankfurt HG. *On Bullshit.* Princeton University Press; 2005.
 5. Kirchner GJ, Raymond YM, Weddle JB, Bible JE. Can artificial intelligence improve the readability of patient education materials? *Clin Orthop Relat Res.* 2023;481:2260-2267.
 6. Peters U. Algorithmic political bias in artificial intelligence systems. *Philos Technol.* 2022;35:25.
 7. Wolfram S. What is ChatGPT doing...and why does it work? Stephen Wolfram Writings. Available at: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work>. Accessed April 25, 2023.