

# Stop Worrying, and Let A.I. Help Save Your Life

Jan. 19, 2026

Erik Winkowski

By Robert Wachter

Dr. Wachter is the chair of the department of medicine at the University of California, San Francisco.

We physicians have a long tradition of the curbside consult — when we bump into specialists or more seasoned colleagues in the hospital cafeteria and ask for their advice on a vexing clinical case. Over my 35 years of practice, I used to track down other doctors for a couple of curbsides during morning rounds each day.

These days, I'm getting far more curbsides, but they are not with colleagues. They're with A.I. Sometimes I consult with ChatGPT; other times I turn to OpenEvidence, a specialized tool for physicians. I find A.I.'s input is virtually always useful. These tools provide immediate and comprehensive answers to complex questions far more effectively than a traditional textbook or a Google search. And they are available 24/7.

To be clear, A.I. isn't perfect. For my curbside consults, the

answers are not as nuanced as the ones I'd hear from my favorite hematologist or nephrologist. On rare occasions, they're just plain wrong, which is why I review them carefully before acting.

Some people argue that A.I.'s imperfections mean that we shouldn't use the technology in high-stakes fields like medicine or that it should be tightly regulated before we do. But the biggest mistake now would be to overly restrict A.I. tools that could improve care by setting an impossibly high bar, one far higher than the one we set for ourselves as doctors. A.I. doesn't have to be perfect to be better. It just has to be better.

Many people — patients, clinicians and policymakers — are dissatisfied with the current state of health care. American medicine delivers miracles every day, but the system itself is a mess, chaos wrapped in mind-boggling paperwork and absurdly high prices. It's in desperate need of transformation.

Sign up for the Opinion Today newsletter Get expert analysis of the news and a guide to the big ideas shaping the world every weekday morning.

A.I. can support this transformation, but only if we stop disproportionately focusing on rare bad outcomes, as we often do with new technologies. While research now

demonstrates that driverless cars are safer than those with human drivers, a serious accident involving a robotaxi is deemed highly newsworthy and often cited as a reason to take driverless cars off the road, whereas one accident involving a human driver may hardly leave a media ripple.

A.I.-based mental health assistants are being subjected to similar scrutiny. A handful of tragic cases involving harmful mental health chatbot responses have made national headlines, spurring several states to enact restrictions on these tools.

These cases are troubling and demand scrutiny and guardrails. But it's worth remembering that millions of patients are now able to receive counseling via bots when a human therapist is impossible to find or impossibly costly.

At U.C.S.F. Medical Center, where I work, many of our physicians now use A.I. scribes that, with the patient's permission, "listen" to doctor-patient conversations and automatically create summaries of the appointment. A.I. can also quickly review and summarize patients' medical records, a huge boon when [one in five](#) patients has a record longer than "Moby-Dick." In both cases, the A.I. isn't flawless, but it can outperform our previous system, which had physicians working as glorified data entry clerks.

As A.I. becomes more commonplace in health care, we need

to develop strategies to determine how much to trust it. As we measure error rates and harms from A.I., we need frameworks to make apples-to-apples comparisons between what human doctors do on their own today and what A.I.-enabled health care does tomorrow. In these early days, we should favor a “walk before you run” strategy, starting with using A.I. to handle administrative paperwork tasks before focusing all our energy on higher-stakes tasks like diagnosis and treatment.

But as we consider the full range of areas in which A.I. can make a positive impact and design strategies to mitigate its flaws, delaying the full adoption of medical A.I. until some mythical state of perfection is achieved will be unreasonable and counterproductive.

Imagine a world in which a young woman with vision problems and numbness visits her doctor. An A.I. scribe captures, synthesizes and documents the patient-physician conversation, a diagnostic A.I. suggests a diagnosis of multiple sclerosis, and a treatment A.I. recommends a therapy based on her symptoms, test results and the latest research findings. The doctor would be able to spend more time focusing on confirming the diagnosis and treatment plan, comforting the patient, answering her questions and coordinating her care. Based on my experience with these tools, I can tell you that this world is within reach.

I am not arguing that we shouldn't aspire to perfection or that A.I. in health care should receive a free pass from regulators. A.I. designed to act autonomously, without clinician supervision, should be closely vetted for accuracy. The same goes for A.I. that may be integrated into machines like CT scanners, insulin pumps and surgical robots — areas in which a mistake can be catastrophic and a physician's ability to validate the results is limited. We need to ensure patients are fully informed and can consent to A.I. developers' intended use of their personal information. For patient-facing A.I. tools in high-stakes settings such as diagnosis and psychotherapy, we also need sensible regulations to ensure accuracy and effectiveness.

But as the saying goes, "Don't compare me to the Almighty; compare me to the alternative." In health care, the alternative is a system that fails too many patients, costs too much and frustrates everyone it touches. A.I. won't fix all of that, but it's already fixing some of it — and that's worth celebrating.