

Medicine Needs to Get Serious About AI

Ezekiel J. Emanuel, Vinod Khosla

For as long as ChatGPT has existed, Americans have been asking it, “Could that funny little spot on my toe be cancer?” And as AI has advanced over the past few years, doctors have begun to use specialized models to ask much the same questions in clinics. In the face of these developments, professional medical societies have presented a unified front: Doctors must always have the final say in clinical care.

We disagree. As the medical societies allow, physicians’ enhancing their work with AI tools may be a perfectly good way to deliver medicine. (One of us—Khosla—has in fact invested in a primary-care start-up, co-founded by his son, in which doctors supervise AI assistants.) But for Americans to get the best health care possible, we need to be willing to experiment with models that put AI in control.

This notion has caused a stir in the medical community. When we published a commentary in the *Journal of the American Medical Association* last month arguing that AI-only care could soon become superior to care provided by human physicians, the CEO of the American Medical Association denounced us as undermining the patient-doctor relationship. The organization’s position was unchanged: AI should be relegated to a supporting role in clinical care. It’s an argument that ignores the available evidence and could hold untold numbers of Americans back from better health.

Much of modern medicine—including managing hypertension, determining the right treatment for many cancers, and addressing chemotherapy side effects—is largely algorithmic. A physician needs to ask a patient targeted questions, determine possible diagnoses, order tests to identify the correct one, offer a treatment plan, and adjust it depending on how the patient responds. That series of actions is steered by detailed guidelines from professional societies.

Yet doctors—despite the best of intentions—make plenty of mistakes when following those steps. Studies indicate that up to 15 percent of American patients are misdiagnosed each year, leading to tens of thousands of preventable deaths. That’s why health-care providers have software that notifies them if they, say, prescribe a deadly dose of a drug. Any software that could practice medicine autonomously would be similarly specialized: not generic ChatGPT, but a model trained in medical knowledge with the appropriate guardrails to minimize risks and mistakes. Such a model could check for drug interactions more consistently and comprehensively than physicians. And it would be able to check in with patients much more frequently than human health-care providers do.

Some AI models have already proved adept at such tasks. In one study published last year, expert physicians judged Google’s Articulate Medical Intelligence Explorer—a highly trained AI model for medical diagnosis and reasoning—as significantly better than human doctors at getting actors playing patients to open up. The “patients” revealed more about their complaints, along with their medical, family, and medication history, all of which can be crucial for an accurate diagnosis. In a 2023 [study](#), patients with type 2 diabetes were randomized to have their insulin dose adjusted by physicians or autonomous AI. The machines got patients to a stable dose faster than human physicians, and their patients took their insulin more reliably and had less distress related to their diabetes.

What about the physician-AI hybrids that medical associations are backing? Research on them is sparse, and more is needed. A 2025 review of 52 studies concluded that, on balance, such hybrids “neither outperformed medical AI alone nor surpassed the best of clinicians.” In at least a few of those studies, human-AI hybrids outperformed AI alone. But in those hybrid teams, the AI controlled the final clinical decision—the opposite of what the AMA and others have in mind.

Medical organizations are not entirely closed-minded to the possibility of change. For example, last week

the American College of Physicians published a [position paper](#) that allowed that “fully autonomous AI is possible,” but cautioned that its use should be infrequent and “limited to low-risk, low-complexity decisions, and always allow for bringing a clinician into the loop.” John Whyte, the CEO of the AMA, told *The Atlantic* that he was “disappointed” in the quality of the analysis in our *JAMA* commentary last month. He questioned the design of the studies we reviewed, many of which involved simulated situations with actors or a retrospective review of existing cases rather than real clinical outcomes, and did not follow patients for years or decades. We agree that longer-term, real-world studies are necessary; hopefully, once advanced AI medical tools have existed for more than a handful of years, results will begin to arrive. Whyte was enthusiastic about further study of physicians controlling AI tools, but only cautiously open to further exploration of autonomous AI in medicine, arguing that the latter would need to be highly regulated.

Reasons exist, of course, to be skeptical of AI-driven medical care. In nine of the 51 studies we reviewed, human physicians outperformed the machines. But some of these tested older AI models or general-use AI models that weren’t trained on medical information and specific guardrails for patient care. Prominent physicians have also argued that real-life doctors are harried by 15-minute appointments and administrative burdens such as prior authorization. That is certainly true, but it is also irrelevant to many of the experiments we reviewed, which examined relatively hassle-free simulations of medical appointments.

One of the most common objections to an AI-doctor future is that robots lack the human connection that patients need to build trust. The American College of Physicians’ recent position paper, for example, cited AI’s inability to make eye contact. But consider that doctors, too, can be robotic. In a Gallup [survey](#) conducted late last year, 21 percent of respondents said they had turned to AI because they had previously felt dismissed by a health-care provider. In the 2025 study of Google’s AMIE, patient actors reported that the AI was more emotionally supportive than the physicians were. In a systematic review of 15 studies comparing AI chatbots with physicians, 13 found that AI had significantly higher perceived empathy.

Like human physicians, AI is not and never will be perfect. But isn’t the possibility of reducing fatalities caused by physician error—even by just a little—worth exploring? Instead of writing off a possible source of progress, the medical community should help answer the key questions raised by AI: How should researchers devise ethical, real-world situations to test autonomous AI models? How should governments and professional organizations regulate them? And how should medical education be reformed so that human doctors can continue to do what they uniquely can?

Hundreds of years ago, doctors believed that disease was caused by an imbalance of “humors” in the body. Even after Louis Pasteur proved that bacteria caused infection and Joseph Lister introduced sterile surgical techniques, physicians cut patients to have them bleed out what ailed them, not bothering to clean the knife or their hands. But medicine progresses. Humans have learned from our past mistakes and created new drugs, procedures, and treatments to massively lengthen lifespan and improve quality of life.

Today need not be any different. Doctors have been presented with a powerful tool that has the potential to improve care for millions of people. It’s their duty not to stand in the way.