

HEALTH

AI Is Taking Over Hospitals

This is health care's Uber moment.

By Benjamin Mazer




Illustration by Alicia Tatone. Source: Getty.

JUNE 17, 2026

SHARE AS GIFT 

DISCUSS  84

REMOVE 

Every knowledge-based profession may one day reach the point when AI outperforms the human experts. In medicine, that day appeared to come in April. A group of primarily Harvard and Stanford researchers announced the results of a study that pitted ChatGPT against hundreds of physicians in a diagnostic obstacle course involving written medical mysteries and information from real-world patients. The bot had won, and the humans weren't entirely happy about it.

“I get a little bit queasy about how some of these results might be used,” Adam Rodman, a lead author on the study, said at a press conference just ahead of its publication in the journal *Science*. The work had amounted to an academic exercise, he told reporters; as thorough as it may have been, it did not prove that ChatGPT or any other AI tool was ready to become a standard part of medical practice. His caution was in line with that of other experts, yet as Rodman knew, most people will ignore the warning. AI has already wormed its way into the U.S. health-care system, evidence and safeguards be damned.

Even as I was watching Rodman’s press conference, I got a message on my phone from the administrators at the medical center where I work as a pathologist. They’d emailed me to say that an “AI-powered clinical reasoning tool” was now available for me to use. This wasn’t the first time I’d gotten this sort of email; it wasn’t the second or third time either. In fact, I’ve lost count of how many generative-AI products have been rolled out to us in recent years, none of which has been approved for medical use by the FDA.

This enthusiasm feels unprecedented. Health care is typically among the last fields to adopt a new technology; I still use a pager, and I send faxes on a regular basis. (Younger readers can ask Claude to explain what these things are.) A tendency toward simple tech is in part a product of doctors’ safety-focused culture: We know that any ill-timed glitch has the potential to turn deadly. But these days, clinicians are allowed—encouraged, even—to run wild with the latest software, guided by a generic warning that “AI can make mistakes.”

Those mistakes can be consequential. Although Rodman’s research shows that generative AI can help diagnose rare diseases or make sense of unusual symptoms, a randomized trial that was published in *NEJM AI* just the week before found that intentionally erroneous output from an AI model can easily lead doctors astray. Nonprofessionals could be similarly misled. A recent study by Oxford scientists found that using AI did not significantly improve patients’ ability to diagnose themselves or others. Another one, led by researchers at Mount Sinai, suggested that chatbots may fail to alert users to potential medical emergencies.

Read: Drink whole milk, eat red meat, and use ChatGPT

Misdiagnosis is not the only concern. As AI permeates the health-care system, errors are cropping up in unexpected places. When I spoke with Rodman by phone after his press conference, he told me that he'd been surprised one day to find that his hospital, Beth Israel Deaconess Medical Center, had enlisted AI to draft messages to patients on his behalf—sometimes producing output for his review that he described as “completely absurd.” (Sarah Finlaw, a spokesperson for Beth Israel Lahey Health, told me that use of AI tools is voluntary and subject to hospital training and support. She also said that any output from AI tools must be approved by a physician.)

Part of the problem is that health-related AI products can be deployed without any vetting by officials at the FDA. If a software package that is intended for physicians is classified as a “clinical decision support tool,” and not a medical device, it usually avoids the agency's oversight. To be counted in this category, an AI-powered app generally must rely on the existing medical literature, avoid analyzing medical scans or images, explain its reasoning, and leave diagnosis and treatment up to a physician. Most of the generative-AI products that doctors use today seem to meet these criteria.

Consumer-wellness apps and devices may also bypass FDA review so long as they are intended for “maintaining or encouraging a healthy lifestyle” and not for diagnosing or treating specific conditions. With this in mind, Microsoft, OpenAI, Anthropic, and xAI all warn users that their health-related chatbots are not meant to provide medical care or issue diagnosis and treatment recommendations. In practice, though, the distinction isn't always clear. Elon Musk encourages people to use his Grok chatbot to generate second medical opinions and interpretations of their X-ray and MRI images; a marketing video for ChatGPT Health shows the app reassuring people that their lab results are in a healthy range and encouraging them to continue taking cholesterol medication.

Most of these apps also invite users to connect their medical records and wearable health devices. AI companies wouldn't need to gobble up all of these data just to offer generic health information. A new product from the medical start-up Hims & Hers, called Labs AI, goes so far as to help interpret the results from “up to 130 biomarker tests” for its users and then provide a “deep, personalized, and actionable

analysis on whole body health, risks, and patterns.” I, too, analyze a patient’s lab results and then give personalized, actionable advice. What’s the difference?

When I reached out to the makers of these products, they reaffirmed that no actual medical advice is being given out to users. Dominic King, the vice president of health at Microsoft AI, told me in an emailed statement that its Copilot app provides “helpful information and support for conversations with clinicians” and not “a single, firm diagnosis.” Patrick Carroll, the chief medical officer of Hims & Hers, told me that Labs AI does not diagnose or recommend treatment: “That responsibility belongs to clinicians, and Labs is designed to reinforce that boundary.” Anthropic and xAI did not respond to my inquiries. OpenAI declined to comment for this article.

Perhaps that boundary—between doctor and algorithm—is somewhat artificial to begin with. One idea kicking around the medical literature is to stop treating AI products as if they were merely standard medical devices. Given their humanlike ability to learn new information and tailor answers to individual patients, medical AIs may function more like doctors than defibrillators—so perhaps they should be evaluated in the same way that physicians are. Instead of requiring FDA approval for each and every function it can perform, a chatbot might be asked to pass a medical-licensing exam and undergo a period of supervision akin to a medical residency.

Read: The ChatGPT symptom spiral

For the moment, though, that idea remains on the fringe. Haider Warraich, a cardiologist and program manager at the Advanced Research Projects Agency for Health, the U.S. government’s program for developing advanced health technology, is leading a major effort to get medical chatbots approved in the traditional way. His agency is providing funding for the development of an AI tool that is tailor-made for heart conditions, and then to send it through a full FDA-authorization process. Warraich’s hope is that by undergoing such a rigorous evaluation, the chatbot will be able to safely evaluate and treat patients without the involvement of a doctor. Rodman praised this approach but warned that the process is going to take years, during which time a plethora of new health AIs will have slipped into the market

with little scrutiny.

In this way, the emergence of today's AI health products remind me of the rise, in the 2010s, of ride-sharing services such as Uber and Lyft. The taxi industry is heavily regulated, making it difficult for new players to enter the market. Yet by skirting and at times ignoring those rules, ride-sharing companies were able to acquire a critical mass of users in a short period of time. Pretty soon, governments had little choice but to adjust their laws to match what had by then become the status quo. The same pattern could end up playing out in medicine. Will regulations meant to ensure that medical products are safe and effective remain in force? Or will they instead be weakened or removed to clear the path for tools that everyone is already using?

We'll soon find out. The health-care system is not going to "slow down and wait for the evidence to accrue," Rodman told me. Eighty percent of doctors are already using AI tools in their job, according to a 2026 survey by the American Medical Association. Patients aren't far behind. The benefits of AI may remain uncertain, but they're already too enticing to pass up.

[View Discussion](#) 84

ABOUT THE AUTHOR

Benjamin Mazer

Benjamin Mazer is a physician specializing in pathology and laboratory medicine.

Explore More Topics

[Scientific Research](#)

About

Contact

Podcasts

Subscription

Follow

[Privacy Policy](#) [Your Privacy Choices](#) [Advertising Guidelines](#) [Terms & Conditions](#) [Terms of Sale](#)
[Responsible Disclosure](#) [Site Map](#)

TheAtlantic.com © 2026 The Atlantic Monthly Group. All Rights Reserved.

This site is protected by reCAPTCHA and the Google Privacy Policy and Terms of Service apply