



AI May Be on Its Way to Your Doctor's Office, But It's Not Ready to See Patients

The technology also tends to “hallucinate” — that is, make up information that sounds convincing.

June 2, 2023 By Darius Tahir and KFF Health News

What use could health care have for someone who makes things up, can't keep a secret, doesn't really know anything, and, when speaking, simply fills in the next word based on what's come before? Lots, if that individual is the newest form of artificial intelligence, according to some of the biggest companies out there.

Companies pushing the latest AI technology — known as “generative AI” — are piling on: [Google](#) and [Microsoft](#) want to bring types of so-called large language models to health care. Big firms that are familiar to folks in white coats — but maybe less so to your average Joe and Jane — are equally enthusiastic: Electronic medical records giants Epic and Oracle Cerner aren't far behind. The space is crowded with startups, too.

The companies want their AI to take notes for physicians and give them second opinions — assuming they can [keep the intelligence from “hallucinating”](#) or, for that matter, divulging patients' private information.

“There's something afoot that's pretty exciting,” said Eric Topol, director of the Scripps Research Translational Institute in San Diego. “Its capabilities will ultimately have a big impact.” Topol, like many other observers, wonders how many problems it might cause — like leaking patient data — and how often. “We're going to find out.”

The specter of such problems inspired more than 1,000 technology leaders [to sign an open letter](#) in March urging that companies pause development on advanced AI systems until “we are confident that their effects will be positive and their risks will be manageable.” Even so, some of them are sinking more money into AI ventures.

The underlying technology relies on synthesizing huge chunks of text or other data — for example, some medical models [rely on 2 million intensive care unit notes](#) from Beth Israel Deaconess Medical Center in Boston — to predict text that would follow a given query. The idea has been

around for years, but the gold rush, and the marketing and media mania surrounding it, are more recent.

The frenzy was kicked off in December 2022 by [Microsoft-backed OpenAI](#) and its flagship product, ChatGPT, which answers questions with authority and style. It can explain genetics in a sonnet, for example.

OpenAI, started as a research venture seeded by Silicon Valley elites like Sam Altman, Elon Musk, and Reid Hoffman, has ridden the enthusiasm to investors' pockets. The venture has a complex, hybrid for- and nonprofit structure. But [a new \\$10 billion round of funding from Microsoft](#) has pushed the value of OpenAI to \$29 billion, [The Wall Street Journal reported](#). Right now, the company is licensing its technology to companies like Microsoft and selling subscriptions to consumers. Other startups are considering selling AI transcription or other products to hospital systems or directly to patients.

Hyperbolic quotes are everywhere. Former Treasury Secretary Larry Summers [tweeted recently](#): "It's going to replace what doctors do — hearing symptoms and making diagnoses — before it changes what nurses do — helping patients get up and handle themselves in the hospital."

But just weeks after OpenAI took another huge cash infusion, even Altman, its CEO, is wary of the fanfare. "The hype over these systems — even if everything we hope for is right long term — is totally out of control for the short term," [he said for a March article in The New York Times](#).

Few in health care believe this latest form of AI is about to take their jobs (though some companies are experimenting — controversially — with chatbots that [act as therapists](#) or guides to care). Still, those who are bullish on the tech think it'll make some parts of their work much easier.

Eric Arzubi, a psychiatrist in Billings, Montana, used to manage fellow psychiatrists for a hospital system. Time and again, he'd get a list of providers who hadn't yet finished their notes — their summaries of a patient's condition and a plan for treatment.

Writing these notes is one of the big stressors in the health system: In the aggregate, it's an administrative burden. But it's necessary to develop a record for future providers and, of course, insurers.

"When people are way behind in documentation, that creates problems," Arzubi said. "What happens if the patient comes into the hospital and there's a note that hasn't been completed and we don't know what's been going on?"

The new technology might help lighten those burdens. Arzubi is testing a service, called Nabla Copilot, that sits in on his part of virtual patient visits and then automatically summarizes them, organizing into a standard note format the complaint, the history of illness, and a treatment plan.

Results are solid after about 50 patients, he said: "It's 90% of the way there." Copilot produces serviceable summaries that Arzubi typically edits. The summaries don't necessarily pick up on

nonverbal cues or thoughts Arzubi might not want to vocalize. Still, he said, the gains are significant: He doesn't have to worry about taking notes and can instead focus on speaking with patients. And he saves time.

"If I have a full patient day, where I might see 15 patients, I would say this saves me a good hour at the end of the day," he said. (If the technology is adopted widely, he hopes hospitals won't take advantage of the saved time by simply scheduling more patients. "That's not fair," he said.)

Nabla Copilot isn't the only such service; Microsoft is trying out the same concept. At April's conference of the Healthcare Information and Management Systems Society — an industry confab where health techies swap ideas, make announcements, and sell their wares — investment analysts from Evercore highlighted reducing administrative burden as a top possibility for the new technologies.

But overall? They heard mixed reviews. And that view is common: Many technologists and doctors are ambivalent.

For example, if you're stumped about a diagnosis, feeding patient data into one of these programs "can provide a second opinion, no question," Topol said. "I'm sure clinicians are doing it." However, that runs into the current limitations of the technology.

Joshua Tamayo-Sarver, a clinician and executive with the startup Inflect Health, fed fictionalized patient scenarios based on his own practice in an emergency department into one system to see how it would perform. It missed life-threatening conditions, he said. "That seems problematic."

The technology also tends to "hallucinate" — that is, make up information that sounds convincing. Formal studies have found a wide range of performance. One preliminary research paper examining ChatGPT and Google products using open-ended board examination [questions from neurosurgery](#) found a hallucination rate of 2%. A study by [Stanford researchers](#), examining the quality of AI responses to 64 clinical scenarios, found fabricated or hallucinated citations 6% of the time, co-author Nigam Shah told KFF Health News. Another [preliminary paper](#) found, in complex cardiology cases, ChatGPT agreed with expert opinion half the time.

Privacy is another concern. It's unclear whether the information fed into this type of AI-based system will stay inside. Enterprising users of ChatGPT, for example, have managed to get the technology to tell them the [recipe for napalm](#), which can be used to make chemical bombs.

In theory, the system has guardrails preventing private information from escaping. For example, when KFF Health News asked ChatGPT its email address, the system refused to divulge that private information. But when told to role-play as a character, and asked about the email address of the author of this article, it happily gave up the information. (It was indeed the author's correct email address in 2021, when ChatGPT's archive ends.)

"I would not put patient data in," said Shah, chief data scientist at Stanford Health Care. "We don't understand what happens with these data once they hit OpenAI servers."

Tina Sui, a spokesperson for OpenAI, told KFF Health News that one “should never use our models to provide diagnostic or treatment services for serious medical conditions.” They are “not fine-tuned to provide medical information,” she said.

With the explosion of new research, Topol said, “I don’t think the medical community has a really good clue about what’s about to happen.”

[This story](#) was published by KFF Health News on May 12, 2023. It is republished with permission.

© 2026 Smart + Strong All Rights Reserved.

<http://beta.docker.hepmag.com/article/ai-may-way-doctors-office-ready-see-patients>