REPORT AI SCIENCE

Giving your healthcare info to a chatbot is, unsurprisingly, a terrible idea / ChatGPT Health promises privacy. All you have is OpenAI's word.

by Robert Hart
Jan 23, 2026, 4:07 PM GMT

67 Comments (All New)

Image: Cath Virginia / The Verge, Getty Images

Robert Hart is a London-based reporter at *The Verge* covering all things AI and Senior Tarbell Fellow. Previously, he wrote about health, science and tech for *Forbes*.

Every week, more than 230 million people ask ChatGPT for health and wellness advice, according to OpenAI. The company says that many see the chatbot as an "ally" to help navigate the maze of insurance, file paperwork, and become better self-advocates. In exchange, it hopes you will trust its chatbot with details about your diagnoses, medications, test results, and other private medical information. But while talking to a chatbot may be starting to feel a bit like the doctor's office, it isn't one. Tech companies aren't bound by the same obligations as medical providers. Experts tell *The Verge* it would be wise to carefully consider whether you want to hand over your records.

Health and wellness is swiftly emerging as a key battleground for AI labs and a major test for how willing users are to welcome these systems into their lives. This month two of the industry's biggest players made overt pushes into medicine.

OpenAI released ChatGPT Health, a dedicated tab within ChatGPT designed for users to ask health-related questions in what it says is a more secure and personalized environment. Anthropic introduced Claude for Healthcare, a "HIPAA-ready" product it says can be used by hospitals, health providers, and consumers. (Notably absent is Google, whose Gemini chatbot is one of the world's most competent and widely used AI tools, though the company did announce an update to its MedGemma medical AI model for developers.)

OpenAI actively encourages users to share sensitive information like medical records, lab results, and health and wellness data from apps like Apple Health, Peloton, Weight Watchers, and MyFitnessPal with ChatGPT Health in exchange for deeper insights. It explicitly states that users' health data will be kept confidential and won't be used to train AI models, and that steps have been taken to keep data secure and private. OpenAI says ChatGPT Health conversations will also be held in a separate part of the app, with users able to view or delete Health "memories" at any time.

Related /

- **OpenAI launches ChatGPT Health, encouraging users to connect their medical records**
- **'Clinical-grade AI': a new buzzy AI word that means absolutely nothing**
- **Chatbots are struggling with suicide hotline numbers**

OpenAI's assurances that it will keep users' sensitive data safe have been helped in no small way by the company launching an identical-sounding product with tighter security protocols at almost the same time as ChatGPT Health. The tool, called ChatGPT for Healthcare, is part of a broader range of products sold to support businesses, hospitals, and clinicians working with patients directly. OpenAI's suggested uses include streamlining administrative work like drafting clinical letters and discharge summaries and helping physicians collate the latest medical evidence to improve patient care. Similar to other enterprise-grade products sold by the company, there are greater protections in place than offered to general consumers, especially free users, and OpenAI says the products are designed to comply with the privacy obligations required of the medical sector. Given the similar names and launch dates — ChatGPT for Healthcare was announced the day after ChatGPT Health — it is all too easy to confuse the two and presume the consumer-facing product has the same level of protection as the more clinically oriented one. Numerous people I spoke to when reporting this story did so.

Even if you trust a company's vow to safeguard your data... it might just change its mind.

Whichever security assurance we take, however, it is far from watertight. Users for tools like ChatGPT Health often have little safeguarding against breaches or unauthorized use beyond what's in the terms of use and privacy policies, experts tell *The Verge*. As most states haven't enacted comprehensive privacy laws — and there isn't a comprehensive federal privacy law — data protection for AI tools like ChatGPT Health "largely depends on what companies promise in their privacy policies and terms of use," says Sara Gerke, a law professor at the University of Illinois Urbana-Champaign.

Even if you trust a company's vow to safeguard your data — OpenAI says it encrypts Health data by default — it might just change its mind. "While ChatGPT does state in their current terms of use that they will keep this data confidential and not use them to train their models, you are not protected by law, and it is allowed to change terms of use over time," explains Hannah van Kolfschooten, a researcher in digital health law at the University of Basel in Switzerland. "You will have to trust that ChatGPT does not do so." Carmel Shachar, an assistant clinical professor of law at Harvard Law School, concurs: "There's very limited protection. Some of it is their word, but they could always go back and change their privacy practices."

Assurances that a product is compliant with data protection laws governing the healthcare sector like the Health Insurance Portability and Accountability Act, or HIPAA, shouldn't offer much comfort either, Shachar says. While great as a guide, there's little at stake if a company that voluntarily complies fails to do so, she explains. Voluntarily complying isn't the same as being bound. "The value of HIPAA is that if you mess up, there's enforcement."

There's a reason why medicine is a heavily regulated field

It's more than just privacy. There's a reason why medicine is a heavily regulated field — errors can be dangerous, even lethal. There are no shortage of examples showing chatbots confidently spouting false or misleading health information, such as when a man developed a rare condition after he asked ChatGPT about removing salt from his diet and the chatbot suggested he replace salt with the sodium bromide, which was historically used as a sedative. Or when Google's AI Overviews wrongly advised people with pancreatic cancer to avoid high-fat foods — the exact opposite of what they should be doing.

To address this, OpenAI explicitly states that their consumer-facing tool is designed to be used in close collaboration with physicians and is not intended for diagnosis and treatment. Tools designed for diagnosis and treatment are

designated as medical devices and are subject to much stricter regulations, such as clinical trials to prove they work and safety monitoring once deployed. Although OpenAI is fully and openly aware that one of the major use cases of ChatGPT is supporting users' health and well-being — recall the 230 million people asking for advice each week — the company's assertion that it is not intended as a medical device carries a lot of weight with regulators, Gerke explains. "The manufacturer's stated intended use is a key factor in the medical device classification," she says, meaning companies that say tools aren't for medical use will largely escape oversight even if products are being used for medical purposes. It underscores the regulatory challenges technology like chatbots are posing.

For now, at least, this disclaimer keeps ChatGPT Health out of the purview of regulators like the Food and Drug Administration, but van Kolfschooten says it's perfectly reasonable to ask whether or not tools like this should really be classified as a medical device and regulated as such. It's important to look at how it's being used, as well as what the company is saying, she explains. When announcing the product, OpenAI suggested people could use ChatGPT Health to interpret lab results, track health behavior, or help them reason through treatment decisions. If a product is doing this, one could reasonably argue it might fall under the US definition of a medical device, she says, suggesting that Europe's stronger regulatory framework may be the reason why it's not available in the region yet.

"When a system feels personalized and has this aura of authority, medical disclaimers will not necessarily challenge people's trust in the system."

Despite claiming ChatGPT is not to be used for diagnosis or treatment, OpenAI has gone through a great deal of effort to prove that ChatGPT is a pretty capable medic and encourage users to tap it for health queries. The company highlighted health as a major use case when launching GPT-5, and CEO Sam Altman even invited a cancer patient and her husband on stage to discuss how the tool helped her make sense of the diagnosis. The company says it assesses ChatGPT's medical prowess against a benchmark it developed itself with more than 260 physicians across dozens of specialties, HealthBench, that "tests how well AI models perform in realistic health scenarios," though critics note it is not very transparent. Other studies — often small, limited, or run by the company itself — hint at ChatGPT's medical potential too, showing that in some cases it can pass medical licensing exams, communicate better with patients, and outperform doctors at diagnosing illness, as well as help doctors make fewer mistakes when used as a tool.

OpenAI's efforts to present ChatGPT Health as an authoritative source of health information could also undermine any disclaimers it includes telling users not to utilize it for medical purposes, van Kolfschooten says. "When a system feels personalized and has this aura of authority, medical disclaimers will not necessarily challenge people's trust in the system."

Companies like OpenAI and Anthropic are hoping they have that trust as they jostle for prominence in what they see as the next big market for AI. The figures showing how many people are already using AI chatbots for health suggest they may be onto something, and given the stark health inequalities and difficulties many face in accessing even basic care, this could be a good thing. At least, it could be, if that trust is well-placed. We trust our private information with healthcare providers because the profession has earned that trust. It's not yet clear whether an industry with a reputation for moving fast and breaking things has earned the same.

67 COMMENTS

More in Report

Gaming's most fun sales data is full of old and obscure games
JAY PETERS JAN 24

Get ready for the AI ad-pocalypse
JESS WEATHERBED JAN 24

What TikTok's new owners mean for your feed
EMMA ROTH JAN 23

Epic and Google have a secret $800 million Unreal Engine and services deal
ADI ROBERTSON and SEAN HOLLISTER JAN 22

The state attorneys general are as mad as you are

SARAH JEONG JAN 22

Big games are getting bigger – and so are the stakes

JAY PETERS JAN 22

Top Stories

2:00 AM GMT

On the ground in Minneapolis after the killing of Alex Pretti

JAN 23

Why this winter storm will likely be a wild one

1:00 PM GMT

The great e-bike crackdown has begun

JAN 24

Get ready for the AI ad-pocalypse

JAN 24

ICE has killed another person in Minneapolis

JAN 24

Gemini with Personal Intelligence is awfully familiar