Alberta Doctors' Digest

The AI is in

A 60-year-old man arrives at the emergency department insisting his neighbour is poisoning him. Over the next 24 hours, he experiences worsening hallucinations and attempts to escape the hospital. The culprit? Sodium bromide – an industrial cleaning salt he purchased online after ChatGPT suggested it as a substitute for table salt.

Welcome to the era of ChatGPT Health

Launched in January 2026, ChatGPT Health is OpenAI's dedicated "sandbox" for personal health management. Unlike the general chatbot interface, this compartmentalized workspace allows users to securely integrate medical records, sync biometric data from Apple Health and MyFitnessPal, and receive AI-generated insights grounded in their actual clinical history. OpenAI claims over 230 million people already ask health questions on ChatGPT weekly. They claim the new platform makes those conversations more personalized and, theoretically, more private.

What patients are actually doing

The use cases are pragmatic. Patients upload lab reports for plain-language explanations, generate data-driven question lists before appointments, and analyze how biometric patterns correlate with activity levels.

They're using it for everything from comparing insurance plans, drafting letters, to translating discharge instructions for non-English-speaking family members. In an overstretched health care system with rising out-of-pocket costs, ChatGPT Health serves as a 24/7 "clinical processor" – helping patients organize fragmented information.

The promise: A more literate patient

The platform shines brightest in bridging communication gaps. For multilingual communities and health literacy challenges, the impact could be profound. ChatGPT Health translates complex discharge instructions into a patient's native language or rephrases dense pathology reports for elderly caregivers in simple terms.

This shifts the clinic conversation from "What does this mean?" to "Given this, what should we do?" That's a subtle but profound upgrade in patient engagement. In an ideal world, patients arrive better prepared with organized symptom timelines, specific questions grounded in their own lab values, and a clearer understanding of what they're being asked to monitor between visits.

The problem: Confident hallucinations

But the early evidence is sobering. When *The Washington Post* tested the platform with a decade of Apple Watch data, ChatGPT Health gave a healthy reporter a failing grade

for heart health – a conclusion his cardiologist dismissed as nonsense. When asked the same question in different sessions, the AI's assessment swung wildly.

Experts warn that because the AI sounds confident and personalized, users fail to recognize where helpful tips end and dangerous misinformation begins. Hallucination rates ranged from 50% to 82.7% across six popular medical AI chatbots tested, with these models not only accepting false information but often expanding on it, producing confident explanations for non-existent conditions.

Unlike regulated medical devices, ChatGPT Health has no mandatory safety controls, no post-market surveillance, and operates outside FDA jurisdiction.

What physicians should do

The folder of Google/WebMD printouts has been replaced by AI-generated health briefs.

Physicians need to embrace patient agency while guarding against algorithmic overconfidence.

When patients arrive with ChatGPT summaries, physicians should ask what data was uploaded/connected, what might be missing, and what questions drove the analysis.

In this new landscape, the physician's most important tool may simply be skepticism applied equally to both the patient's data and the algorithm's conclusions.

Editor’s note: The views, perspectives and opinions in this article are solely the author’s and do not necessarily represent those of the AMA.