News and Perspectives

How Will the Food and Drug Administration Regulate Generative AI–Enabled Medical Devices?

Simon Spichak, JMIR Correspondent

Abstract

As the US Food and Drug Administration works to develop specific guidelines for the regulation of AI-enabled medical devices, could a previously cleared tool provide some clues as to what the new guidance will look like? In this *News and Perspectives* article, JMIR Correspondent Simon Spichak reports on and speaks with experts about the UpDoc diabetes management app and the future of AI regulation as AI-enabled devices continue to evolve.

Key Takeaways:

- The Food and Drug Administration (FDA) is seeking public comment to refine its approach to regulating generative AI–enabled devices but has not established any dedicated, binding guidance.
- A medical device advertised as a "clinical AI platform" nonetheless received FDA clearance to adjust insulin dosing and talk directly with patients.
- The precise architecture and role of the device's AI remain unclear, and its clearance documentation doesn't establish any obvious precedent for how other generative AI–enabled medical devices will be regulated.

Over 3 decades ago, Rita King,ound herself at a turning point in software regulation. As a founding member of the Underwriters Laboratories (now UL Solutions) team, King helped launch and manage the first US programs that assess safety-critical software—programs used in commercial and medical devices that can cause injury, death, and other damage if they falter.

"Back then, we did not have regulations, certifications, and understanding of how to go about to deal with safety-critical software," says King—now the CEO of MethodSense, a device regulatory consulting firm. Now, the Food and Drug Administration (FDA) is developing the first guidelines for regulating generative AI–enabled medical software, having sought public comment through October 19, 2026. This is another crossroads that feels both familiar and different for King; with generative AI, the big difference is exponentially more variability in user inputs and program outputs, making it particularly challenging to address the risks.

Could the so-touted "first FDA-cleared agentic clinical AI platform"—UpDoc's diabetes management app, which includes a conversational AI component and received FDA clearance in late 2025—provide a glimpse into the future of AI-enabled device regulation?

What's UpDoc?

Sharif Vakili, MD, MBA—cofounder and CEO of UpDoc—and his cofounder Ashwin Nayak, MD, MS, were fascinated by the concept of an agentic clinical AI. "AI software that does something that would have historically taken place with a licensed clinician," Vakili explains. "It was an itch we couldn't let go."

Their interest led them to develop their tool for type 2 diabetes—an app with both clinician- and patient-facing elements. Doctors can prescribe the tool to patients to provide closed-loop care delivery, including the autonomous management of dosing medications and a conversational AI layer that interacts with patients. In a randomized trial of 32 participants published in 2023, the tool helped participants reach optimal insulin dosing over 40 days faster and helped them achieve higher insulin adherence than the standard care group.

"If the physician has permitted this agent to be part of a care team member, it will call you and message you, and manage your medications," says Vakili. "So, all you have to do is, just pick up the phone and chat with it."

There remains some controversy and confusion, however, about exactly how UpDoc works and specifically the extent to which generative AI is used, a lot of which may have to do with discrepancies in the language used to describe the tool.

According to the FDA decision summary, the device received premarket clearance on the basis of substantial equivalence to a previously approved insulin calculator, with no mention made of "large language models" or "artificial intelligence." According to the 2023 study, the algorithm used to adjust insulin is rules-based and deterministic, constrained by the physician's treatment plan. Vakili himself confirmed that it operates similarly to an insulin pump, noting that it might best be described as "semiautonomous" but that the company otherwise "can't disclose the technical architecture," echoing previous reporting by STAT.

On the company's website and in their press release, however, the app's AI elements are emphasized; the latter

notes that "licensed clinicians for the first time can deploy AI agents that are FDA-cleared," giving the following example: if a patient's blood sugar dips or rises out of a healthy range, "UpDoc's platform identifies the trend, initiates a titration of the patient's insulin regimen within physician-approved parameters, triggers the appropriate follow up test to confirm safety, and documents the entire intervention in the EHR."

While the extent to which these "agents" are permitted to engage in independent clinical reasoning remains somewhat opaque, the regulatory distinction seems to be that the conversational AI layer appears to be separate from the treatment-adjusting component, which operates within physician-set parameters.

"We don't see generative AI around the safety-critical part of this device," says King. It's a minor component of the decision because the FDA may have thought that UpDoc didn't introduce substantial risks compared to other similar devices.

Yujan Shrestha, MD—founder of the regulatory consultancy firm Innolitics—wrote that "It appears that the FDA is comfortable with a protocolized clinical workflow that uses AI conversation around a deterministic core" and a "narrow and well-risk-controlled use case as far as agentic AI goes."

In other words, the approval of this device, especially given the limited public information around its architecture, doesn't reveal much about how or through which pathway generative AI–enabled medical devices will be regulated in the future. The wording of the clearance doesn't provide sufficient detail to gauge how the conversational AI layer was evaluated, nor does it seem to establish a clear precedent for other devices.

Regulating Generative AI

As the FDA seeks public comment to inform its regulatory framework, other agencies are funding the development of generative AI medical software—and UpDoc's technology has turned heads.

UpDoc is one of a few companies receiving a portion of US $62.7 million in funding from the Advanced Research Projects Agency for Health (ARPA-H) for a program the agency is calling ADVOCATE (Agentic AI-EnableD CardioVascular CAre TransfOrmation). The company will be developing a patient-facing clinical AI agent as part of the ADVOCATE program's overall aim to develop a disease-agnostic, open-source, agentic AI to manage heart failure.

> *"Humans can't review every single thing, but AI could. It's an extra checking gate."*
>
> Sharif Vakili, MD, MBA

"A third of the funding of this program is for the development of autonomous clinical AI agents that can deliver care for cardiovascular disease and particularly heart failure," says Vakili. "And the reason why they pick heart failure is because it is a big moonshot that is really hard to treat."

The hope is that researchers will be able to develop an open-source supervisory model that can act as an AI-in-the-loop for a range of medical tasks. "Humans can't review every single thing, but AI could," he says. "It's an extra checking gate."

Meanwhile, the FDA has also opened up the TEMPO pilot trial to test four generative AI–enabled medical devices in real-world settings.

How these models may receive approval and clearance is still up in the air while the FDA is actively seeking comment. Devices that integrate AI will still need to go through the same software development life cycle and validation programs that other medical software goes through, King says, but testing all the possible inputs and outputs with generative AI models will be challenging.

The FDA discussion document indicates that they're interested in a competence-based framework, which will require understanding what scenarios each device will need to be tested in and how these devices—conversational AI tools especially—interact with their users. It highlights that the risk calculation differs based on whether the device directs actions, takes medical actions itself, or simply provides information.

"What you do is you identify how a user may be using the device, and you identify the risks," says King. This, she thinks, will lead to a greater emphasis on human factors in the FDA's decision-making.

Keywords: artificial intelligence; medical device regulation; FDA 510(k) clearance; Food and Drug Administration; software as a medical device; clinical decision support systems; machine learning; artificial intelligence agents; regulatory guidance

Please cite as:
Spichak S
How Will the Food and Drug Administration Regulate Generative AI–Enabled Medical Devices?
J Med Internet Res 2026;28:e115030

URL: https://www.jmir.org/2026/1/e115030
doi: 10.2196/115030

© JMIR Publications. Originally published in the Journal of Medical Internet Research (https://www.jmir.org), 02.Nov.2026