

ChatGPT Moves Into Health Records and Voice-Directed Work

AI News Digest

2026-07-24

ChatGPT Moves Into Health Records and Voice-Directed Work

By AI News Digest • July 24, 2026

OpenAI expands ChatGPT into connected health records and voice-directed desktop work, while Microsoft launches new in-house image and voice models. Also: multimodal model releases, an open on-device AI coworker, and a new NVIDIA–KAIST research lab for agentic AI.

ChatGPT moves deeper into personal data and desktop workflows

OpenAI begins U.S. rollout of Health in ChatGPT

OpenAI is rolling out **Health in ChatGPT** to U.S. users, letting them connect Apple Health and supported medical records to review information in context, track changes, and support health conversations. With permission, connected context can be used across conversations—for example, to compare a new result with prior tests or summarize changes since an appointment—and OpenAI says this data is not used to train foundation models or target ads. [1, 2]

OpenAI says more than 300 million people ask ChatGPT health-related questions weekly. *Why it matters:* the product extends ChatGPT from answering general health questions toward working with a user’s own longitudinal records, making privacy handling and contextual accuracy central to the experience. [3]

Voice control arrives for ChatGPT’s desktop agents

ChatGPT Voice is rolling out globally on macOS and Windows for Plus, Pro, Business, Edu, and Enterprise plans. OpenAI says users can control their computer and direct multiple agents in ChatGPT Work or Codex by voice, with

GPT-Live simultaneously speaking, listening, and coordinating work in the app. [4]

Why it matters: this is a shift from voice as a conversational interface to voice as an input layer for multi-agent desktop work.

Microsoft expands its in-house image and voice lineup

Microsoft AI has launched **MAI-Image-2.5-Pro** in Foundry preview, describing it as its highest-fidelity professional image model for high-quality generation, detailed editing, and precise in-image text. It joins Microsoft’s image family so builders can choose among quality, speed, and cost trade-offs. [5]

Separately, **MAI-Voice-2-Flash** is in public preview; Microsoft says it is twice as fast as MAI-Voice-2 and 32% cheaper at \$15 per million characters, and powers Dynamics 365 Contact Center. Microsoft also reports up to 89% lower GPU costs for that deployment. [6]

Why it matters: Microsoft is pairing frontier-facing model releases with product deployment claims—including an 84% image-model cost reduction in Power-Point and lower latency in OneDrive—signaling a focus on operating economics alongside capability. [7]

Multimodal competition broadens from content to action

xAI released **Grok 4.5** across grok.com, X, and iOS and Android, calling it its most capable model yet. [8]

Black Forest Labs introduced **FLUX 3**, a unified multimodal architecture spanning image, video, audio, and action prediction; the company says the model can be extended for robotics action prediction, while FLUX 3 Video is available in early access. [9]

Why it matters: the releases show two different routes to broader AI interfaces: a general-purpose model expanding across consumer platforms, and a single architecture designed to connect media generation with robotic action prediction.

Andrew Ng launches an open, on-device AI coworker

Andrew Ng and Rohit Prasad announced **OpenWorker**, an open-source agent intended to produce completed work such as customer briefs, reports, calendar updates, and Slack triage rather than simply chat. It works across files and everyday tools, and checks in before consequential actions, according to the announcement. [10]

OpenWorker runs on Mac, supports user-selected API-based or local models such as Ollama, and keeps data on-device except when a user elects to use an LLM provider or integration; Windows support is planned. The project is available at openworker.com, with source code on GitHub. [10]

Why it matters: it offers a model-independent, privacy-oriented alternative in the emerging AI-coworker category, where control over data and model choice is becoming a primary design distinction.

NVIDIA and KAIST create a joint agentic-AI research lab

NVIDIA and the Korea Advanced Institute of Science and Technology announced a Seoul-based lab dedicated to advancing agentic AI for South Korea. NVIDIA describes it as the first joint AI research lab between a Korean university and a global technology company; the collaboration combines KAIST researchers with NVIDIA's full-stack expertise, Nemotron open models, and AI Cloud partner computing. [11]

The announcement came alongside discussions between South Korean leaders, NVIDIA, and ecosystem partners on expanding the country's AI infrastructure and expertise. *Why it matters:* it links university research, open models, and compute infrastructure in a national AI-development effort. [11]

Sources

1. X post by @OpenAI
2. X post by @OpenAI
3. X post by @OpenAI
4. X post by @OpenAI
5. X post by @mustafasuleyman
6. X post by @mustafasuleyman
7. X post by @mustafasuleyman
8. X post by @grok
9. X post by @bfl_ai
10. X post by @AndrewYNg
11. At AI Summit, South Korea Outlines Its AI Future With NVIDIA and Partners