INTERNET ARCHIVE Wayback Machine https://techcr 534 captures 14 Mar 2023 - 3 f
INTERNET ARCHIVE Wayback Machine https://techcrunch.com/2023/03/14/openai-releases-gpt-4-ai-that-it-claims-is-state-of-the-art/ Go 534 captures 14 Mar 2023 - 3 Feb 2026
APR AUG 17 NOV 2024 2025 2026 About this capture

Advertisement



IMAGE CREDITS: DANI PADGETT WATSON / STRICTLYVC

AI

OpenAI releases GPT-4, a multimodal AI that it claims is state-of-the-art

Kyle Wiggers 10:06 AM PDT · March 14, 2023

OpenAI has released a powerful new image- and text-understanding AI model, GPT-4, that the company calls "the latest milestone in its effort in scaling up deep learning."

GPT-4 is available today to OpenAI's paying users via ChatGPT Plus (with a usage cap), and developers can sign up on a waitlist to access the API.

Pricing is $0.03 per 1,000 "prompt" tokens (about 750 words) and $0.06 per 1,000 "completion" tokens (again, about 750 words). Tokens represent raw text; for example, the word "fantastic" would be split into the tokens "fan," "tas" and "tic." Prompt tokens are the parts of words fed into GPT-4 while completion tokens are the content *generated* by GPT-4.

GPT-4 has been hiding in plain sight, as it turns out. Microsoft confirmed today that Bing Chat, its chatbot tech co-developed with OpenAI, is running on GPT-4.

Other early adopters include Stripe, which is using GPT-4 to scan business websites and deliver a summary to customer support staff. Duolingo built GPT-4 into a new language learning subscription tier. Morgan Stanley is creating a GPT-4-powered system that'll retrieve info from company documents and serve it up to financial analysts. And Khan Academy is leveraging GPT-4 to build some sort of automated tutor.

GPT-4's new capabilities power a 'virtual volunteer' for the visually impaired

Tech and VC heavyweights join the Disrupt 2025 agenda

Netflix, ElevenLabs, Wayve, Sequoia Capital, Elad Gil — just a few of the heavy hitters joining the Disrupt 2025 agenda. They're here to deliver the insights that fuel startup growth and sharpen your edge. Don't miss the 20th anniversary of TechCrunch Disrupt, and a chance to learn from the top voices in tech — grab your ticket now and save up to $600+ before prices rise.

San Francisco | October 27-29, 2025

REGISTER NOW

Advertisement

DISRUPT
October 27-29, 2025 | San Francisco
Put your brand in front of 10,000+ tech and VC leaders across all three days of Disrupt 2025. Amplify your reach, spark real connections, and lead the innovation charge. Secure your exhibit space before your competitor does.
Book Your Table ›

Most Popular

- How your solar rooftop became a national security issue
- Co-founder of Elon Musk's xAI departs the company
- Google CEO adds a new calendar feature at Stripe co-founder's request
- NASA has sparked a race to develop the data pipeline to Mars
- iOS 26 beta 6 adds new ringtones, snappy app

GPT-4 can generate text and accept image and text inputs — an improvement over GPT-3.5, its predecessor, which only accepted text — and performs at "human level" on various professional and academic benchmarks. For example, GPT-4 passes a simulated bar exam with a score around the top 10% of test takers; in contrast, GPT-3.5's score was around the bottom 10%.

OpenAI spent six months "iteratively aligning" GPT-4 using lessons from an internal adversarial testing program as well as ChatGPT, resulting in "best-ever results" on factuality, steerability and refusing to go outside of guardrails, according to the company. Like previous GPT models, GPT-4 was trained using publicly available data, including from public webpages, as well as data that OpenAI licensed.

OpenAI worked with Microsoft to develop a "supercomputer" from the ground up in the Azure cloud, which was used to train GPT-4.

"In a casual conversation, the distinction between GPT-3.5 and GPT-4 can be subtle," OpenAI wrote in a blog post announcing GPT-4. "The difference comes out when the complexity of the task reaches a sufficient threshold — GPT-4 is more reliable, creative and able to handle much more nuanced instructions than GPT-3.5."

Without a doubt, one of GPT-4's more interesting aspects is its ability to understand images as well as text. GPT-4 can caption — and even interpret — relatively complex images, for example identifying a Lightning Cable adapter from a picture of a plugged-in iPhone.

The image understanding capability isn't available to all OpenAI customers just yet — OpenAI's testing it with a single partner, Be My Eyes, to start with. Powered by GPT-4, Be My Eyes' new Virtual Volunteer feature can answer questions about images sent to it. The company explains how it works in a blog post:

"For example, if a user sends a picture of the inside of their refrigerator, the Virtual Volunteer will not only be able to correctly identify what's in it, but also extrapolate and analyze what can be prepared with those ingredients. The tool can also then offer a number of recipes for those ingredients and send a step-by-step guide on how to make them."

A more meaningful improvement in GPT-4, potentially, is the aforementioned steerability tooling. With GPT-4, OpenAI is introducing a new API capability, "system" messages, that allow developers to prescribe style and task by describing specific directions. System messages, which will also come to ChatGPT in the future, are essentially instructions that set the tone — and establish boundaries — for the AI's next interactions.

For example, a system message might read: "You are a tutor that always responds in the Socratic style. You *never* give the student the answer, but always try to ask just the right question to help them learn to think for themselves. You should always tune your question to the interest and knowledge of the student, breaking down the problem into simpler parts until it's at just the right level for them."

Even with system messages and the other upgrades, however, OpenAI acknowledges that GPT-4 is far from perfect. It still "hallucinates" facts and makes reasoning errors, sometimes with great confidence. In one example cited by OpenAI, GPT-4 described Elvis Presley as the "son of an actor" — an obvious misstep.

"GPT-4 generally lacks knowledge of events that have occurred after the vast majority of its data cuts off (September 2021), and does not learn from its experience," OpenAI wrote. "It can sometimes make simple reasoning errors which do not seem to comport with competence across so many domains, or be overly gullible in accepting obvious false statements from a user. And sometimes it can fail at hard problems the same way humans do, such as introducing security vulnerabilities into code it produces."

OpenAI does note, though, that it made improvements in particular areas; GPT-4 is less likely to refuse requests on how to synthesize dangerous chemicals, for one. The company says that GPT-4 is 82% less likely overall to respond to requests for "disallowed" content compared to GPT-3.5 and responds to sensitive requests — e.g. medical advice and anything pertaining to self-harm — in accordance with OpenAI's policies 29% more often.



OpenAI GPT-4

ringtones, snappy app launches, and more

- Ford throws out Henry Ford's assembly line to make low-cost EVs in America
- The computer science dream has become a nightmare

Advertisement

IMAGE CREDITS: OPENAI

There's clearly a lot to unpack with GPT-4. But OpenAI, for its part, is forging full steam ahead — evidently confident in the enhancements it's made.

"We look forward to GPT-4 becoming a valuable tool in improving people's lives by powering many applications," OpenAI wrote. "There's still a lot of work to do, and we look forward to improving this model through the collective efforts of the community building on top of, exploring, and contributing to the model."

Topics: AI AI Artificial Intelligence (AI) Enterprise gpt-4 OpenAI

Advertisement



Kyle Wiggers
AI Editor

Kyle Wiggers was TechCrunch's AI Editor until June 2025. His writing has appeared in VentureBeat and Digital Trends, as well as a range of gadget blogs including Android Police, Android Authority, Droid-Life, and XDA-Developers. He...

View Bio ›

Loading the next article

Error loading the next article

Advertisement

© 2025 TechCrunch Media LLC.

Could not connect to the reCAPTCHA service. Please check your internet connection and reload to get a reCAPTCHA challenge.