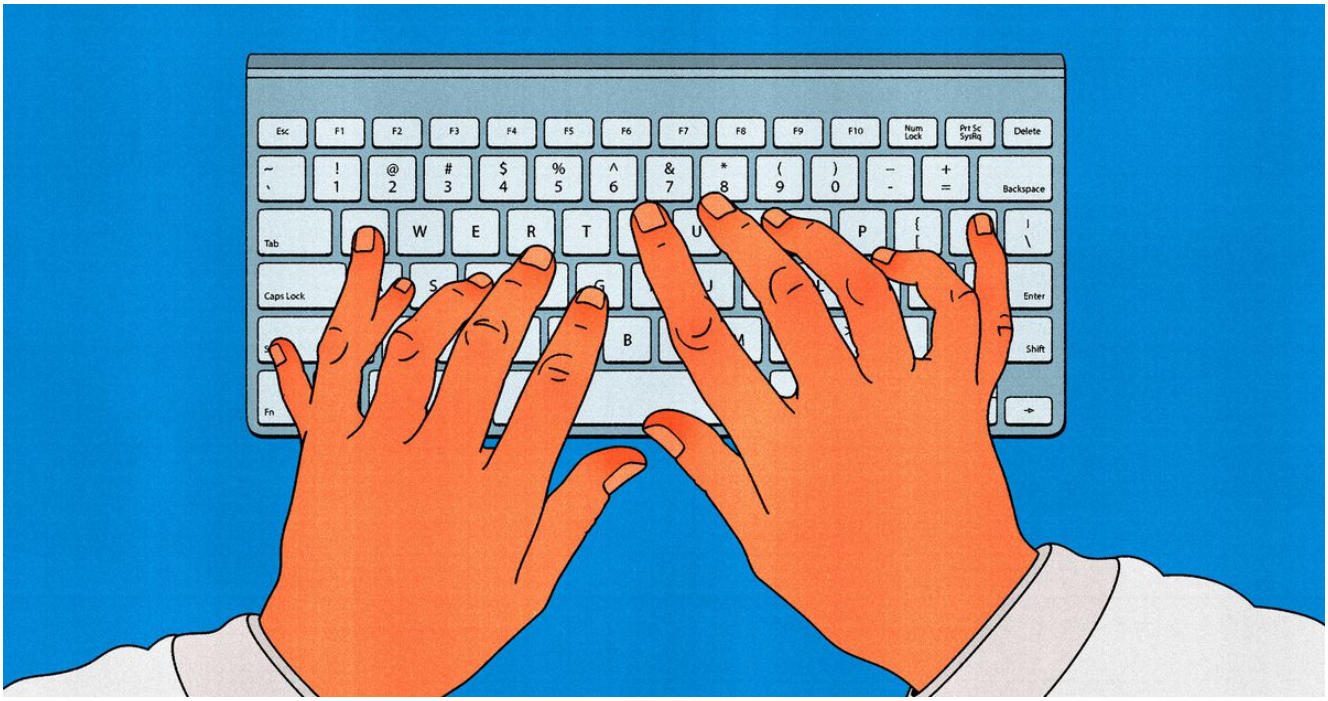
Google scrambles to manually remove weird AI answers in search



Social media is abuzz with examples of Google's new AI Overview product saying weird stuff, from telling users to put glue on their pizza to suggesting they eat rocks. The messy rollout means Google is racing to manually disable AI Overviews for specific searches as various memes get posted, which is why users are seeing so many of them disappear shortly after being posted to social networks.

It's an odd situation, since Google has been testing AI Overviews for a year now — the feature launched in beta in May 2023 as the Search Generative Experience — and CEO Sundar Pichai has said the company served over a billion queries in that time.

But Pichai has also said that Google's brought the cost of delivering AI answers down by 80 percent over that same time, "driven by hardware, engineering and technical breakthroughs." It appears that kind of optimization might have happened too

early, before the tech was ready.

“A company once known for being at the cutting edge and shipping high-quality stuff is now known for low-quality output that’s getting meme’d,” one AI founder, who wished to remain anonymous, told *The Verge*.

Google continues to say that its AI Overview product largely outputs “high quality information” to users. “Many of the examples we’ve seen have been uncommon queries, and we’ve also seen examples that were doctored or that we couldn’t reproduce,” Google spokesperson Meghann Farnsworth said in an email to *The Verge*. Farnsworth also confirmed that the company is “taking swift action” to remove AI Overviews on certain queries “where appropriate under our content policies, and using these examples to develop broader improvements to our systems, some of which have already started to roll out.”

Gary Marcus, an AI expert and an emeritus professor of neural science at New York University, told *The Verge* that a lot of AI companies are “selling dreams” that this tech will go from 80 percent correct to 100 percent. Achieving the initial 80 percent is relatively straightforward since it involves approximating a large amount of human data, Marcus said, but the final 20 percent is extremely challenging. In fact, Marcus thinks that last 20 percent might be the hardest thing of all.

“You actually need to do some reasoning to decide: is this thing plausible? Is this source legitimate? You have to do things like a human fact checker might do, that actually might require artificial general intelligence,” Marcus said. And Marcus and Meta’s AI chief Yann LeCun both agree that the large language models that power current AI systems like Google’s Gemini and OpenAI’s GPT-4 will not be what creates AGI.

Look, it’s a tough spot for Google to be in. Bing went big on AI before Google did with Satya Nadella’s famous “we made them

dance” quote, OpenAI is reportedly working on its own search engine, a fresh AI search startup is already worth $1 billion, and a younger generation of users who just want the best experience are switching to TikTok. The company is clearly feeling the pressure to compete, and pressure is what makes for messy AI releases. Marcus points out that in 2022, Meta released an AI system called Galactica that had to be taken down shortly after its launch because, among other things, it told people to eat glass. Sounds familiar.

Google has grand plans for AI Overviews — the feature as it exists today is just a tiny slice of what the company announced last week. Multistep reasoning for complex queries, the ability to generate an AI-organized results page, video search in Google Lens — there’s a lot of ambition here. But right now, the company’s reputation hinges on just getting the basics right, and it’s not looking great.

“[These models] are constitutionally incapable of doing sanity checking on their own work, and that’s what’s come to bite this industry in the behind,” Marcus said.

https://www.theverge.com/2024/5/24/24164119/google-ai-overview-mistakes-search-race-openai