

[View in browser](#)

Bloomberg

Tech Daily | AI

Hi, hello, it's Rachel in San Francisco. There's been so much talk about AI hallucinating that it's making me feel like I'm hallucinating. But first...

Help us make this newsletter better by [filling out this survey](#)

Today's must-reads:

- China [hit Micron with a chips review](#)
- Twitter users [balked at paying for blue check marks](#)
- Italian regulators [launched a probe into OpenAI](#)

Choice of words

Somehow the idea that an artificial intelligence model can “hallucinate” has become the default explanation anytime a chatbot messes up.

It's an easy-to-understand metaphor. We humans can at times hallucinate: We may see, hear, feel, smell or taste things that aren't truly there. It can happen for all sorts of reasons (illness, exhaustion, drugs).

Companies across the industry have applied this concept to the new batch of extremely powerful but still flawed chatbots. Hallucination is listed as a limitation on the [product page](#) for OpenAI's latest AI model, GPT-4. Google, which [opened access to its Bard chatbot in March](#), reportedly brought up AI's propensity to hallucinate in [a recent interview](#).

Even skeptics of the technology are embracing the idea of AI hallucination. A couple of the signatories on a petition that went out last week urging a six-month

halt to training powerful AI models mentioned it along with [concerns about the emerging power of AI](#). Yann LeCun, Meta Platforms Inc.'s chief scientist, has [talked about it repeatedly on Twitter](#).

Granting a chatbot the ability to hallucinate — even if it's just in our own minds — is problematic. It's nonsense. People hallucinate. Maybe some animals do. Computers do not. They use math to make things up.

Humans have a tendency to anthropomorphize machines. (I have a robot vacuum named Randy.) But while ChatGPT and its ilk can produce convincing-sounding text, they don't actually understand what they're saying.

In this case, the term “hallucinate” obscures what's really going on. It also serves to absolve the systems' creators from taking responsibility for their products. (Oh, it's not our fault, it's just hallucinating!)

Saying that a language model is hallucinating makes it sound as if it has a mind of its own that sometimes derails, said Giada Pistilli, principal ethicist at Hugging Face, which makes and hosts AI models.

“Language models do not dream, they do not hallucinate, they do not do psychedelics,” she wrote in an email. “It is also interesting to note that the word ‘hallucination’ hides something almost mystical, like mirages in the desert, and does not necessarily have a negative meaning as ‘mistake’ might.”

As a rapidly growing number of people access these chatbots, the language used when referring to them matters. The discussions about how they work are no longer exclusive to academics or computer scientists in research labs. It has seeped into everyday life, informing our expectations of how these AI systems perform and what they're capable of.

Tech companies bear responsibility for the problems they're now trying to explain away. Microsoft Corp., a major OpenAI investor and a user of its technology in Bing, and Google rushed to bring out new chatbots, regardless of the risks of spreading misinformation or hate speech.

ChatGPT reached a million users in the days following its release, and people have conducted over 100 million chats with Microsoft's Bing chatbot. Things are going so well that Microsoft is even [trying out ads](#) within the answers Bing spits out; you might see one the next time you ask it about buying a house or a car.

But even OpenAI, which started the current chatbot craze, appears to agree that hallucination is not a great metaphor for AI. A footnote in one of its technical papers ([PDF](#)) reads, “We use the term ‘hallucinations,’ though we recognize ways this framing may suggest anthropomorphization, which in turn can lead to harms or incorrect mental models of how the model learns.” Even still, variations of the word appear 35 times in that paper. —[Rachel Metz](#)

The big story

Tech giants from Microsoft to Meta axing jobs are also shedding real estate, [leading to a glut of empty offices](#) in major American cities and a sea of struggling landlords.

Get fully charged

Microsoft is trying to make every last drop of its [\\$1 billion climate fund count](#).

Apple won a legal challenge against the UK antitrust watchdog into its dominance over the mobile phone market [due to a procedural technicality](#).

Lemon8, a new app that’s a sort of mishmash of Instagram and Pinterest, is surging in popularity in the US and [drawing attention due to its owner](#): Beijing-based ByteDance, which also owns TikTok.

The former Grubhub driver only won \$65, but the settlement of an eight-year federal court case [may have far-reaching implications](#).

Watch: After Huawei posted its first annual profit decline in more than a decade, the company’s USA chief security officer [talked about what happened](#) in a TV interview on *Bloomberg Technology*.