



# Have We Reached Peak AI?

EDWARD ZITRON – MARCH 18, 2024

Last week, *The Wall Street Journal* published a 10-minute-long interview with OpenAI CTO Mira Murati, with journalist Joanna Stern asking a series of thoughtful yet straightforward questions that Murati failed to satisfactorily answer. When asked about what data was used to train Sora, OpenAI's app for generating video with AI, Murati claimed it used publicly available data, and when Stern asked her whether it used videos from YouTube, Murati's face contorted in a mix of confusion and pain before saying she "actually wasn't sure about that." When Stern pushed a third time, asking about videos from Facebook or Instagram, Murati shook her head and said that if videos were "publicly available...to use, there might be the data, I'm not sure, I'm not confident about it."

Stern did well to get Murati to answer, but it's deeply concerning that the *Chief Technology Officer* of the most "important" AI company in the world can't answer a very basic question about training data. Even when asked about training on data from Shutterstock, a company that OpenAI has a partnership with, Murati stammered, shook her head, and said that she would not "go into the details of the data that was used, but it was publicly available **or** licensed data" (emphasis mine). Shortly after the interview, Stern adds that OpenAI shared that Shutterstock's data was used to train Sora's models.

**Side note:** One very useful thing that Stern does in this interview is break down exactly how Sora works for the masses, defining that as the "AI model analyzing lots of videos and learning to identify objects and actions." This may seem like a small gesture, one problem that has poisoned reporting on generative AI is an unwillingness to clearly describe **how** these things work, instead referring to its actions as some sort of magical process done on a big, scary computer.

This interview is important for a few reasons, but let's start with the most obvious: the Chief Technology Officer of OpenAI either can't or won't explain what materials its video-creating

generative AI was trained on. Throughout the interview she seemed lost, uneasy, unable to give many specifics about the product she was working on, describing everything in the broadest terms of “soon” and “eventually.” Had this interview not pulled out such a useful conversation around training data, Murati would have told the world very little. Though Murati’s vagueness might be at the request of her public relations and legal counsel, I actually think it’s more likely that she realized, in real time, that she either didn’t know these answers or that the truth would be far more underwhelming (or troubling) than the world might like.

“When asked whether it will be possible to fix Sora’s videos after they’ve been generated, Murati said “eventually,” and then couched that by saying “that’s what we’re trying to figure out...how to use this technology as a tool that people can edit and create with.” She promised that there would “eventually” be “more steerability, control and accuracy...and reflecting of intent of what you want.”

You’ll “eventually” be able to add audio to Sora videos, and when asked when Sora’s generative videos will be available to the public, she once again said “eventually,” and when pushed said that Sora’s launch would “definitely be this year, but could be a few months.”

Murati, living in a world of “eventualities,” provided no technical insights, no specifics, and very few details.

These are the kind of interviews that only the most popular tech companies get – relatively shallow back-and-forths where nobody seems to be able to say “hey, what does that actually mean?” Murati was astonishingly unprepared to answer specific questions about the technology behind the \$80 billion startup she briefly ran and seemed rattled whenever she had to expand on anything beyond the most general talking points.

While I generally loved how Stern approached this interview, I wanted to scream when she failed to push back against Murati’s claim that she sees Sora as “a tool for extending creativity,” and that OpenAI wants “people in the film industry, creators everywhere, to be part of informing” how OpenAI develops and deploys Sora. This is exactly the point at which you say, “what exactly does that mean?” and “how have you done that in the past when it comes to image generation?”

The answer, of course, is that OpenAI has zero interest in talking to anyone in the film industry or any creators anywhere, and Murati should’ve been verbally flayed for suggesting otherwise. While OpenAI has advertised positions for community specialists that would act as “ambassadors for OpenAI,” I see little evidence that OpenAI has any plans to help creators other than trying to convince them to use its tools.

This is, of course, all part of OpenAI’s flowery, hollow playbook. In a *Daily Show* interview from October 2022 (a full month before the launch of ChatGPT), Murati told Trevor Noah that OpenAI sees tools like ChatGPT and DALL-E 2 as “extensions of our creativity,” a direct copy-

paste of the messaging that OpenAI uses on DALL-E's website. In a seven-minute-long interview, Murati debuts the talking points that have underpinned OpenAI's entire messaging strategy – misinformation bad, OpenAI good, some jobs will be lost, but it's good, because that's happened before. In an interview with Bloomberg's Emily Chang from July 2023, Murati describes her job as “a combination of guiding the teams on the ground, thinking about long-term strategy, figuring out our gaps and making sure that the teams are well supported to succeed,” and says that one of the things she's most worried about is hallucinations (when a model authoritatively says something incorrect), a question that Chang fails to follow up with “so uh, how are you fixing that?”

For reasons that make less sense after his removal from OpenAI's board, Chang also speaks with billionaire investor Reid Hoffman, who suggests that AI will be adopted “faster than iPhones” and that there will be a “co-pilot for every profession.” Chang weakly ripostes by laughing about her kids using ChatGPT to write papers, to which Hoffman retorts with his own version of “extending creativity,” saying that the hope would be that the interaction with AI will teach students to “create much more interesting papers,” a point at which Chang should have asked him *what that actually fucking means*.

While writing this piece, I took the time to watch several more interviews with Murati (and, indeed, OpenAI CEO Sam Altman), and for the most important company in Silicon Valley, there is very little fundamental explanation of why this technology matters, and what it actually does. In an interview with Joanna Stern at *The Wall Street Journal's* “Tech Live” event in October 2023, Sam Altman said that the thing that people really like about ChatGPT isn't that it “knows particular knowledge,” but that it has this “larval reasoning capacity that's going to get better and better,” and continues to mumble out an answer about how models will “set up all sorts of economic arrangements” that will have them explain how it will answer a question (?), before adding that “the fundamental thing about these models is not that they memorize a lot of data.” Stern fails to push back here on numerous fronts – that “larval reasoning” is a completely meaningless term, and that, in general, Altman has failed to actually explain what he means.

This is the problem with powerful people in tech. If you allow them to speak and fill in the gaps for them, they will happily do so. Murati and Altman continuously obfuscate how ChatGPT works, what it can do, what it could do, and profit handsomely from a complete lack of pushback from a press that routinely accepts AI executives' vague explanations at face value. OpenAI's messaging and explanations of what its technology can (or will) do have barely changed in the last few years, returning repeatedly to “eventually” and “in the future” and speaking in the vaguest ways about how businesses make money off of – let alone profit from – integrating generative AI.

Sam Altman is repeatedly given the ability to wax lyrical about the futuristic capabilities of artificial intelligence in a way that lets him paint a picture of a technology he is not actually building. Altman's fanciful claims include his kids “having more AI friends than human friends,” that human-level AI is “coming” without ever specifying when, that AI will replace

95% of tasks performed by marketing agencies, that ChatGPT will evolve in “uncomfortable ways,” that AI will kill us all, and that human beings are only separated from artificial intelligence because they “really care what others think.”

Every time Sam Altman speaks he almost immediately veers into the world of fan fiction, talking about both the general things that “AI” could do and non-specifically where ChatGPT might or might not fit into that without ever describing a real-world use case. And he’s done so in exactly the same way for years, failing to describe any industrial or societal need for artificial intelligence beyond a vague promise of automation and “models” that will be able to do stuff that humans can, even though OpenAI’s models continually prove themselves unable to match even the dumbest human beings alive.

Altman wants to talk about the big, sexy stories of Average General Intelligences that can take human jobs because the reality of OpenAI – and generative AI by extension – is far more boring, limited and expensive than he’d like you to know.

And I don’t believe things are likely to improve.

## Limited Intelligence

Last week, *The Information* published a story about Amazon and Google “tamping down generative AI expectations,” with these companies dousing their salespeople’s excitement about the capabilities of the tech they’re selling. A tech executive is quoted in the article saying that customers are beginning to struggle with questions like “is AI providing value?” and “How do I evaluate how AI is doing,” and a Gartner analyst told Amazon Web Services sales staff that the AI industry was “at the peak of the hype cycle around Large Language Models and other generative AI.”

The article confirms many of my suspicions – that, as *The Information* wrote, “other software companies that have touted generative AI as a boon to enterprises are still waiting for revenue to emerge,” citing the example of professional services firm KPMG buying 47,000 subscriptions to Microsoft’s co-pilot AI “at a significant discount on co-pilot’s \$30 per seat per month sticker price.” Confusingly, KPMG bought these subscriptions despite not having gauged how much value its employees actually get out of the software, but rather to “be familiar with any AI-related questions its customers might have.”

Salesforce CFO Amy Weaver said in its most recent earnings call that Salesforce was “not factoring in material contribution” from Salesforce’s numerous AI products in its Financial Year 2025 guidance. Software company Adobe’s shares slid as the company failed to generate meaningful revenue from its AI products, with analysts worried about its ability to actually monetize any of the generative products it’s proliferating. ServiceNow claimed in its earnings that generative AI was meaningfully contributing to its bottom line, yet *The*

*Information's* story quotes its Chief Financial Officer Gina Mastantuono as saying that “from a revenue contribution perspective, it’s not going to be huge.”

I believe a large part of the artificial intelligence boom is hot air, pumped through a combination of executive bullshitting and a compliant media that will gladly write stories *imagining* what AI can do rather than focus on what it’s actually doing. Notorious boss-advocate Chip Cutter of *The Wall Street Journal* wrote a piece last week about how AI is being integrated in the office, spending most of the article discussing how companies “might” use tech before digressing that every company he spoke to was using these tools experimentally and that they kept making mistakes. In an interview with Salesforce’s head of AI Clara Shih, *The New York Times* failed to get her to say much of anything about what its AI products do, other than how its “Einstein Trust Layer” handles data, to which Shih added that AI would “be transformational for jobs, the way the internet was.”

The media has been fooled, in the same way they were fooled by the metaverse, by the specious promises of AI and the executives that champion it. The half-truths and magical thinking have spread far faster due to the fact that AI actually exists, and it’s much easier to imagine how it might change our lives, even if the way it might do so is somewhere between improbable and impossible. It’s easy to think that tasks like data-entry, or “boring” work can be easily-automated, and when you use ChatGPT, you can almost kind-of-sort-of see how that might happen, even if ChatGPT really can’t do these things, all because ChatGPT was, at launch, able to do impressions of things that almost looked useful.

As we speak, there are few if any meaningful improvements to our lives as a result of the last year’s artificial intelligence boom. I just deleted a sentence where I talked about “the people I know who use ChatGPT,” and realized that in the last year, I have met exactly one person who has – a writer that used it for synonyms.

I can find no companies that have integrated generative AI in a way that has truly improved their bottom line other than Klarna, which claims its AI-powered support bot is “estimated to drive a \$40 million US in profit improvement in 2024,” which does not, as many have incorrectly stated, mean that it has “made Klarna \$40 million in profit.” Despite fears to the contrary, AI does not appear to be replacing a large number of workers, and when it has, the results have been pretty terrible. A study from Boston Consulting Group found that consultants that “solved business problems with OpenAI’s GPT-4” performed 23% worse than those who didn’t use it, even when the consultant was warned about the limitations of generative AI and the risk of hallucinations.

To be clear, I am not advocating for the replacement of workers with AI. I am, however, saying that if AI was actually capable of replacing the outputs of human beings – even if it was anywhere near doing so – any number of massive, scurrilous firms would be doing so at scale, and planning to do so more as models improved.

Unless, of course, it just wasn't possible. What if what we're seeing today isn't a glimpse of the future, but the new terms of the present? What if artificial intelligence isn't actually capable of doing much more than what we're seeing today, and what if there's no clear timeline when it'll be able to do more? What if this entire hype cycle has been built, goosed by a compliant media ready and willing to take career-embellishers at their word?

Every single time I've read about the "amazing" things that artificial intelligence can do, I see somebody attempting to add fuel to a fire that's close to going out. While Joanna Stern may have said that Sora's generative video clips "freaked her out," much of what makes them scary is the assumption that OpenAI will fix hallucinations, something that the company has categorically failed to do. AI hype is predicated on solving problems with AI models that are only getting worse, and OpenAI's only answers are a combination of "we'll work it out eventually, trust me" and "we need a technological breakthrough in chips and energy."

Generative AI's core problems – its hallucinations, its massive energy and unprofitable compute demands – are not close to being solved. Having now read and listened to a great deal of Murati and Altman's interviews, I can find few cases where they're even asked about these problems, let alone ones where they provide a cogent answer.

And I believe it's because there isn't one.

Generative AI models are expensive and compute-intensive without providing obvious, tangible mass-market use cases. Murati and Altman's futures depend heavily on keeping the world believing that development and improvement of their models' capabilities will continue a rapacious pace of progress that has unquestionably slowed, with OpenAI admitting that GPT-4 may be worse on some tasks.

As I've written before, hallucinations are a feature not a bug. These models do not "know" anything. They are mathematical behemoths generating a best guess based on training data and labeling, and thus do not "know" what you are asking it to do. You simply cannot fix them. Hallucinations are not going away.

Every bit of excitement for this technology is based on the idea what it might do, which quickly becomes conflated with what it could do, allowing Altman – who is far more a marketing person than an engineer – to sell the dream of OpenAI based off of the least-specific promises since Mark Zuckerberg said we'd live in our Oculus headsets.

## **Altman, Freed**

I believe that Sam Altman has been tapdancing this entire time, hoping that he could amass enough power and revenue that his success would be inevitable. Yet his ultra-successful hype campaign was deeply specious, and he – along with the rest of the AI industry – has found himself suddenly having to deliver a future he's not even close to developing.

What I fear isn't automation taking our jobs, but the bottom falling out of generative AI as companies realize that the best they're going to see is a few digits of profit growth. Companies like Nvidia, Google, Amazon, Snowflake and Microsoft have hundreds of billions of dollars of market capitalization – as well as expected revenue growth – tied into the idea that everybody will be integrating AI into everything, and that they will be doing so more than they are today.

If the AI bubble pops, the entire tech industry will suffer as venture capitalists are once again washed out through chasing an unprofitable, barely-substantiated trend. And again the entire industry suffers because people don't want to build new things or try new ideas, but fund the same people doing similar things again and again because it feels good to be part of a consensus, even if you're wrong. Silicon Valley will continually fail to innovate at scale until it learns to build real things again – things that people use because the things in question actually do something.

Altman needs us to build more efficient chips and “energy breakthroughs” because he knows, at his heart, that generative AI can neither fix its own problems nor develop much further without technology that doesn't (and may never) exist. Murati's mealy-mouthed answers around “publicly-available data” heavily suggest that OpenAI's models are trained on YouTube and Facebook videos, meaning that any public launch of Sora will be one that's immediately met with an apocalyptic legal fight. How apocalyptic? Well, a study from last week revealed that every single model produced copyrighted material, with OpenAI's GPT-4 producing it on 44% of the prompts constructed for the study, and Nvidia is being sued by authors claiming its NeMo language model violates their copyright.

Eventually, one of these companies will lose a copyright lawsuit, causing a brutal reckoning on model use across any industry that's integrated AI. These models can't really “forget,” possibly necessitating a costly industry-wide retraining and licensing deals that will centralize power in the larger AI companies that can afford them. And in the event that Sora and other video models are actually trained on copyrighted material from YouTube and Instagram, there is simply no way to square that circle legally without effectively restarting training the model.

## **Artificial Hype**

Sam Altman desperately needs you to believe that generative AI will be essential, inevitable and intractable, because if you don't, you'll suddenly realize that trillions of dollars of market capitalization and revenue are being blown on something remarkably mediocre. If you focus on the present – what OpenAI's technology can do today, and will likely do for some time – you see in terrifying clarity that generative AI isn't a society-altering technology, but another

form of efficiency-driving cloud computing software that benefits a relatively small niche of people.

If you stop saying things like “AI could do” or “AI will do,” you have to start asking what AI can do, and the answer is...not that much, and not much more in the future. Sora is not going to generate movies. It’s going to continue making horrifying human-adjacent creatures that walk like the AT-ATs from Star Wars, and cartoons that look remarkably like copyrighted material from YouTube.

I believe that artificial intelligence has three quarters to prove itself before the apocalypse comes, and when it does, it will be that much worse, savaging the revenues of the biggest companies in tech. Once usage drops, so will the remarkable amounts of revenue that have flowed into big tech, and so will acres of data centers sit unused, the cloud equivalent of the massive over hiring we saw in post-lockdown Silicon Valley.

I fear that the result could be a far worse year for the tech industry than we saw in 2023, one where the majority of the pain hits workers rather than the ghouls who inflated this perilous bubble.

---

**Ed Zitron** is the CEO of **EZPR**, a national tech and business public-relations agency. He is also the author of the tech and culture newsletter ***Where’s Your Ed At*** and the host of the ***15 Minutes in Hell*** podcast.