

Have we crossed the AI event horizon?

Despite Sam Altman's claims that we've crossed the AI "event horizon," the reality is that today's models remain far behind human intelligence in scale, flexibility, efficiency, and learning, says Lewis Liu

Sam Altman is currently peddling the "gentle singularity", claiming we are already "past the event horizon" for AI – at the point where the technology will be infused in every part of our lives and, in his words, "wonders become routine, and then table stakes". But despite his claims that we've crossed this threshold, the math tells a different story.

I want to take a cold, hard look at what this actually means and examine the technical realities behind Altman's vision, especially given how the rise of AGI (Artificial General Intelligence – AI genuinely smarter than humans) will shape human civilisation in the years to come.

There are obviously two camps: those who believe we are already at or super close to achieving AGI (Altman is the high priest of this movement), and the naysayers who argue that LLMs are nothing more than stochastic parrots with some useful applications (Gary Marcus is an intellectual ringleader for this camp). Some fall in between, depending on definition and context – full disclosure, that's me.

However, from the perspective of true, genuine conscious AGI, I think we are far, far off. Let's examine the technical parameters that tell the real story.

Scale Gap

Let's first talk about the size of these models. The human adult brain is estimated to contain 100-500 trillion synapses, the connections between neurons. The best analogy for LLMs are parameters, where the largest current model, GPT-4, has roughly 1.8 trillion parameters. So just looking at pure size alone, the most modern LLMs are roughly one per cent the size of the adult human brain. For context, the number of parameters of GPT-4 is roughly the same order of magnitude as the number of synapses of a honeybee.

Another way of thinking about this: Broca's area of the human brain, which handles "LLM-type language functions," represents less than one per cent of total brain volume. Perhaps we can say our most powerful LLM models have reached some level of parity with just one per cent of our brain – the part responsible for generating language.

Flexibility Gap

This rough comparison ignores crucial differences in how these systems actually work. Neurons and synapses are extraordinarily dynamic: they can strengthen, weaken, form new connections and even change their fundamental properties based on experience. A single neuron can participate in multiple circuits simultaneously and modify its behaviour in real-time based on context.

LLM parameters, by contrast, are static numerical numbers (called weights), that remain fixed once training is complete. While neurons have evolved over millions of years to create flexible, adaptive networks that can rewire themselves, LLM architectures are rigid mathematical constructs. It's like comparing a Lego set for a car vs a plastic model of a car: superficially similar, but fundamentally different in capability.

Efficiency Gap

The energy comparison is staggering. The human brain operates on roughly 20 watts, less power than a light bulb, while processing sensory input, maintaining consciousness, controlling motor functions, and generating thoughts simultaneously. GPT-4, meanwhile, requires 100 to 1,000 times more power to produce a fraction of the brain's output, and that's just for language generation, no vision, no motor control, no real-time learning.

This isn't just about energy costs. It reveals something fundamental about the underlying architectures. The brain achieves its remarkable capabilities through billions of years of evolutionary optimisation, while current AI systems are essentially brute-forcing their way to "intelligence" through massive computational power.

Learning Gap

Perhaps most critically, the brain continuously updates its "model" at virtually zero additional energy cost. Every conversation you have, every experience you encounter, seamlessly integrates into your existing knowledge without requiring a complete system rebuild. You learn your friend's new phone number, adjust to

a changed commute route, or pick up a new skill, all while maintaining everything you already knew.

LLMs, by contrast, require hundreds of billions of dollars and the equivalent of nuclear power plants to incorporate new information through retraining. They cannot learn from individual conversations or adapt to new information without starting the entire training process from scratch. This inability to learn continuously isn't just a technical limitation: it represents a fundamental gap between how biological and artificial intelligence currently operate. Until we solve this problem, claims of approaching human-level AGI remain premature.

From a raw numbers perspective, we are orders of magnitude off from the size, scale, efficiency and flexibility needed for AI models to reach human parity. However, there is an argument that model size is growing exponentially and it may only take a few years for models to catch up to human scale (ignoring the flexibility gap between parameters and neurons and ignoring efficiency).

But we are already seeing significant plateauing of model performance simply from adding more parameters or data. The amount of [high-quality data in the world is finite](#), and furthermore, data quality is deteriorating as the vast majority of new internet content is now AI-generated, leading to a risk of progressively worse models through what's known as "model collapse", a concept [I wrote about previously](#). As such, there is no clear path to reaching the next levels of intelligence through current approaches.

I know many brilliant well-funded people are working on these problems (and know many of these folks myself), and I'm willing to bet there will be major breakthroughs in the coming years. But reaching human-scale AGI will still take considerable time, and we have certainly not "crossed the event horizon."

I think it's extremely dangerous for people like Altman to make such claims. This rhetoric leads to both irrational fear and over-inflated expectations of what AI can currently achieve, which, to be clear, is already quite a lot. Obviously, these statements serve strategic purposes: influencing public policy, claiming "winner status" in a hyper-competitive race, and generating headlines. But they distract from the real work at hand.

From an AI founder, scientist, and investor perspective, here's my view: AI is already transforming the world and OpenAI, despite what I consider dubious IP

and copyright practices, has made significant contributions to that transformation. But does it really matter what exact benchmark score the latest model achieves? What matters more is how existing technology can improve our lives.

The question we should be asking isn't whether we've crossed some arbitrary, marketing-driven threshold of "event horizon." Instead, we should focus on how to use this powerful technology responsibly: to solve real problems rather than create new ones.

Let's stay grounded. Focus on scientific facts, not hype. Focus on practical use cases that AI can genuinely achieve today. Focus on AI safety, trust and governance. Most importantly, focus on building AI that betters humanity rather than generating more clickbait and confusion.

Dr Lewis Z. Liu is co-founder and CEO of Eigen Technologies