The Silicon Valley sceptic warning tech's new bubble is about to burst

Uber veteran Gary Marcus believes the AI revolution is delivering diminishing returns

Andrew Orlowski

Most interviews take place over leisurely meals at fancy restaurants or in hotel lobbies, but we're on to our third circuit of Russell Square. Gary Marcus, the Cassandra of artificial intelligence, says walking helps him think – even if it is a freezing morning.

"I think the day of reckoning may have come – the signs are there," Marcus says. "AI is reaching a point of diminishing results."

Marcus, formerly the head of AI for Uber, has become one of the most compelling commentators on AI, and also one of the most vilified. He is in London to spread an idea that is subversive within tech: the current AI "revolution" is trumped up and misguided.

Three years ago Marcus, who is a professor emeritus of psychology and neural science at New York University, warned that the industry's strategy was flawed. In a piece published by the science magazine Nautilus, Marcus argued that AI was "hitting a wall".

He predicted that devoting more resources, like vast new data centres, would cease to yield impressive results. Artificial intelligence needed a whole new blueprint, he argued.

Diminishing returns

"I was arguing that if the scaling of pure LLMs [large language models, the current AI systems] would run out, that some problems like

‘hallucinations’ [where the bots make up facts] would be very hard to solve,” Marcus says.

“Then you would expect to have multiple teams competing and reaching a point of diminishing returns, with essentially no competitive moat, and a lot of competition over price.”

This view, which he set out just as ChatGPT was sweeping the world, was seen as heresy by the AI true believers. Only a year ago, the British Nobel Prize winner Sir Demis Hassabis, whom Marcus admires greatly, dismissed his concerns.
But the latest large language model, the best-known application of AI today, has been underwhelming.

“Instead of refining and enhancing the model, the quality seems to worsen,” one ChatGPT user wrote last year. “Disappointed with the new models, waste of time,” another complained.

OpenAI withdrew its latest model, GPT-4o, after a few days because it was returning “misguided, incorrect and downright harmful ideas”, reported VentureBeat. It was also too sycophantic, the company acknowledged.

OpenAI, co-founded by Sam Altman, has withdrawn its latest ChatGPT model Credit: SeongJoon Cho/Bloomberg

Real breakthroughs are still possible but the industry needs to rethink how to achieve them, Marcus argues. He wants to anchor the language models in symbolic logic, which was the approach of the first generation of AI, rather than language analysis.

Investors are starting to share his concerns. China's debut on the world AI stage spooked the market after DeepSeek, a Chinese AI model developed by a largely unknown team, upended the thinking about how AI should work.

It wasn't the best in the field, and it still isn't today. But DeepSeek could be built and run for a fraction of the cost of power-hungry American models. It has even been demonstrated running on a Dell office PC from 1998. It turned out that simple tuning and sensible design yielded huge wins – something anyone could have discovered.

While it seemed obvious after the fact, the debut upended the industry assumption that AI would be dominated by just a handful of companies with the deepest pockets.

The era of diminishing returns envisioned by Marcus "has clearly been reached", he believes.

DeepSeek changes the picture dramatically, in good and bad ways: we may have more AI everywhere, with large language models running in your earbud or microwave oven. Think of how spell-checking software is everywhere.

But that also means the business case for capital expenditure on huge computer infrastructure now looks fraught. If it's everywhere, and no one is paying, why burn the capital? Nobody selling spell checkers in the Eighties ever bought a Learjet.

"The economics just don't make a lot of sense," says Marcus, who is also a popular science writer and an entrepreneur. "Too many companies already

know that secret formula; nobody holds an advantage any more for more than a month or two.

“There’s a real data centre bubble, something that Alibaba’s chair Joseph Tsai talked about recently. Tsai argued that people are investing ahead of demand that they’re seeing today. He also claimed that data centres are being built on spec without binding agreements from big tenants like Google and Microsoft.”

Microsoft is now reportedly pulling out of data centre contracts, the analyst firm TD Cowen has reported.

RECOMMENDED

Microsoft's £2.5bn investment in Britain at risk from creaking power grid

Read more

The elephant in the room is reliability. While generative AI excels at novelty – old movies cast with the characters you choose, for example – it still struggles with solid work, such as planning, reasoning and general purpose of abstraction.

“What it does well is mimicry,” Marcus explains.

He terms the current crop of AI BSI, or “broad shallow intelligence”. He told an audience the night before we met at the London School of Economics that he chose the B and the S for a reason. “It can BS its way through a chess game,” he says, jokingly.

Automation must be reliable if it is to boost productivity and then, even more hopefully, GDP growth. The Government commissioned a report estimating that £45bn of annual savings could accrue to the UK from AI – but that didn’t explain how. If Marcus is only half right, they may well be disappointed.

Still, today coders are happily using AI to generate their work, undergraduates are using it to write essays and depressed teenage boys are getting therapy from these bots. Surely if costs fall, as DeepSeek has demonstrated, all that won’t go away?

“Long term, I am somewhat optimistic; I think a much better AI could be built than we have now,” Marcus says. “In the short term, I think we are already starting to see a correction.

“The current technology just isn’t that reliable; LLMs won’t disappear, but they aren’t magic and they aren’t trustworthy, which really limits their applications.”

Marcus made a bet with the AI researcher Miles Brundage at the end of last year. The pair drew up a list of 10 tasks that AI should be able to perform by the end of 2027. For example, it should be able to watch a completely new mainstream movie, follow plot twists, and accurately describe characters’ conflicts and motivations. Or write good legal briefs free of errors.

The logical reasoning required for these sorts of tasks is beyond the AI of today. Marcus, as you’d expect, bet that computers wouldn’t be able to achieve the checklist of tasks.

“The end game here is not that large language models will vanish,” Marcus predicts, “but that they will be virtually free. And eventually we will have much better technology and laugh at how absurd it was that people took them so seriously.”