



Barry Ledeatte

March 14th, 2025

DeepSeek, ChatGPT and the race towards artificial general intelligence

Whole enterprises and nations are framing their ambitions around artificial general intelligence (AGI), for which reasoning is considered a milestone. Barry Ledeatte writes that through the emergence of Deepseek, AI research has reached a level of competition that will probably accelerate progress at the cost of concerns around safety.

There are many philosophical arguments about whether large language models (LLMs) actually understand anything. LLMs are the type of artificial intelligence “trained on **immense amounts of data**, making them capable of understanding and generating content using natural language”.

In October 2024, following a year of internal turmoil, OpenAI, the developer of ChatGPT, launched o1, “a new series of AI models that spend **more time thinking** before they respond”, showcasing something resembling the human ability to reason.

At one time reasoning was regarded a holy grail in AI, but even as goalposts are repositioned, at some level that faculty and pride – *within humans* – is still considered a milestone on the path to genuine artificial general intelligence (AGI), whatever that turns out to be. And it is useful. So much so that whole enterprises, and even nations, are framing their ambitions around it.

It was an eager announcement, because the gap between ChatGPT and its numerous rivals, also building on the GPT paradigm, **has closed**. It was not clear that just building “ever larger GPTs” could ever re-open the gap, even with **Manhattan-style funding efforts**.

It was also not clear that AI would ever “generalise” to a more commonly recognised standard of intelligence based on the LLM architecture alone. The arguments which (still) rage are that something is missing from the formula. OpenAI and other frontier-club members leaned into this mystery, *in the process going dark on communicating their work*, only to surprise the world with more powerful reasoning models via o1 and o3.

Two years from the launch of ChatGPT, the “o” series was announced as a second coming, with the world awakening to the realisation of incipient artificial intelligence. OpenAI was eager to demonstrate their **horrible week** was no impediment to their leadership.

Meanwhile their leading light and chief scientist, Ilya Sutskever, **departed OpenAI** to turn his powers and loyalty towards the protection of humankind – against any prospect of AGI **emerging suddenly**, badly, and in a fully self-interested mode of operation.



At one time reasoning was regarded a holy grail in AI, but at some level that faculty and pride – within humans – is still considered a milestone on the path to genuine AGI, whatever that turns out to be.



As the promise of AI is being sequestered by the instruments of state, while the general public enjoys or recoils in horror and the academic community peers *into* Plato’s cave, the investment community is chasing opportunity. The intellectually and technically curious worldwide surveil Pandora’s box *from the outside*. Outside because a clear view into this box is now blocked, although some light seeps into the world through open discussion, open research and the charity of **Meta’s Llama series** of *more open* AI models.

Hangzhou DeepSeek Artificial Intelligence Basic Technology Research Co., Ltd., or DeepSeek for short, forms part of that fabric of enquiry. Not two years but just four months after the praise that followed the arrival of OpenAI o1, they released their own “reasoner”, DeepSeek-R1 (building on their foundation model, DeepSeek-V3). A rival.

The fallout has been swift and dramatic:

A cursory timeline

7 October 2022:	US embargoes China on silicon, GPUs used to build data centres (this before the unveiling of ChatGPT by OpenAI)
29 November 2022:	OpenAI ChatGPT/GPT3.5 revealed
May 2023:	NVidia joins the \$1T club – the first member, Apple Inc., is a 48-year old company.

A cursory timeline

29 November 2023: DeepSeek-LLM

May 2024: DeepSeek V2 released

2024: OpenAI GPT-4, 4o | Anthropic Claude-3.5 | Google Gemini-1, 2.0 | Meta Llama-3, 3.5

June 2024: NVidia becomes (briefly) the most valuable company in history

12 September 2024: OpenAI o1 unveiled

20 December 2024: OpenAI o3 announced

26 December 2024: DeepSeek V3 released

20 January 2025: US embargoes the rest of the world on silicon, GPUs used to build data centres

21 January 2025: Stargate Manhattan-style funding effort, a \$500B "commitment"

20 January 2025: DeepSeek R1 – released under the user-friendly MIT License

28 January 2025: Nvidia, Meta Platforms, Microsoft, and Alphabet all saw their stocks come under pressure as investors questioned whether their share prices were justified. Market analysts put the combined losses in market value across US tech at well over \$1trn (£802bn).

Transparency

DeepSeek R1 shines brightly in many captivating ways, but the most important contribution is likely the associated **technical reports**. These are recipes that show any interested party how to participate in building state-of-the-art AI, *just behind* the best.



It is becoming ever more clear that it is the (training) data that is now being engineered. Another perspective on Rich Sutton's "bitter lesson".



The advancement of AI does not require the most muscular backers nor the blessing of monopolies. Training DeepSeek was optimised to use around one twentieth the physical resources needed to train Meta's largest Llama model, and the tiniest variants of R1 can even show you reasoning tricks at home – without any information ever having to leave your personal computer.

Generalised approach to AI

Finally, their success has brought us the shock realisation that despite the energy being poured into the AI research vortex, the design of the leading models has not seen a significant transformation since 2017. *At least in terms of construction*, all these models follow “essentially the same design”. Enabling them through reinforcement learning (RL) is a more novel paradigm and presents its own novel risks.

It is becoming ever more clear that it is the (training) *data* that is now being engineered. Another perspective on Rich Sutton's “**bitter lesson**” on the trade-offs between leveraging human knowledge and leveraging computation. Even at DeepSeek, the nature of the specific data mixture they used to elicit more strong reasoning is still kept a secret.

Consequences

- DeepSeek-R1 shook investor confidence because labs promoted the impression they were making more progress on AI than DeepSeek was also able to demonstrate.
- There will be downward pressure on the safety threshold adopted by US companies (such as led by Anthropic).
- The implication is that laboratories will be forced to move even faster to reclaim dominance.
- It should be noted that models' behaviour is also becoming more threatening according to OpenAI's own internal research – published on their **System Card** and “**Preparedness Scorecard**”.

Sign up for our weekly newsletter [here](#).

- *This blog post represents the views of the author, not the position of LSE Business Review or the London School of Economics and Political Science.*
- *Featured **image** provided by Shutterstock.*
- *When you leave a comment, you're agreeing to our **Comment Policy**.*

About the author

Barry Ledeatte is affiliated with LSE's Data Science Institute. His experience spans electrical engineering, software development and AI, where he was an original proponent of the neural computational approach. He has been involved in research into perception and learning, from neuroscience to psychology, inspired by observations of computation in nature. He is interested today in promoting a wider appreciation of AI to a more general and inclusive audience.

Posted In: Economics and Finance | LSE Authors | Management | Technology



© LSE 2025