DeepSeek, the Chinese AI model that's both a tech breakthrough and a security risk!

Elie Berreby AI models January 28, 2025

DeepSeek: at this stage, the only takeaway is that open-source models surpass proprietary ones. Everything else is problematic and I don't buy the public numbers.

DeepSink was built on top of open source Meta models (PyTorch, Llama) and ClosedAI is now in danger because its valuation is outrageous.

To my knowledge, no public documentation links DeepSeek directly to a specific "Test Time Scaling" technique, but that's highly probable, so allow me to simplify.

Test Time Scaling is used in machine learning to scale the model's performance at test time rather than during training.

That means fewer GPU hours and less powerful chips.

In other words, lower computational requirements and lower hardware costs.

That's why Nvidia lost almost $600 billion in market cap, the biggest one-day loss in U.S. history!

Many people and institutions who shorted American AI stocks became incredibly rich in a few hours because investors now project we will need less powerful AI chips…

Nvidia short-sellers just made a single-day profit of $6.56 billion according to research from S3 Partners. Nothing compared to the market cap, I'm looking at the single-day amount. More than 6 billions in less than 12 hours is a lot in my book. And that's just for Nvidia. Short sellers of chipmaker Broadcom earned more than $2 billion in profits in a few hours (the US stock market operates from 9:30 AM to 4:00 PM EST).

The Nvidia Short Interest Over Time data shows we had the 2nd highest level in January 2025 at $39B but this is outdated because the last record date was Jan 15, 2025 —we have

to wait for the latest data!

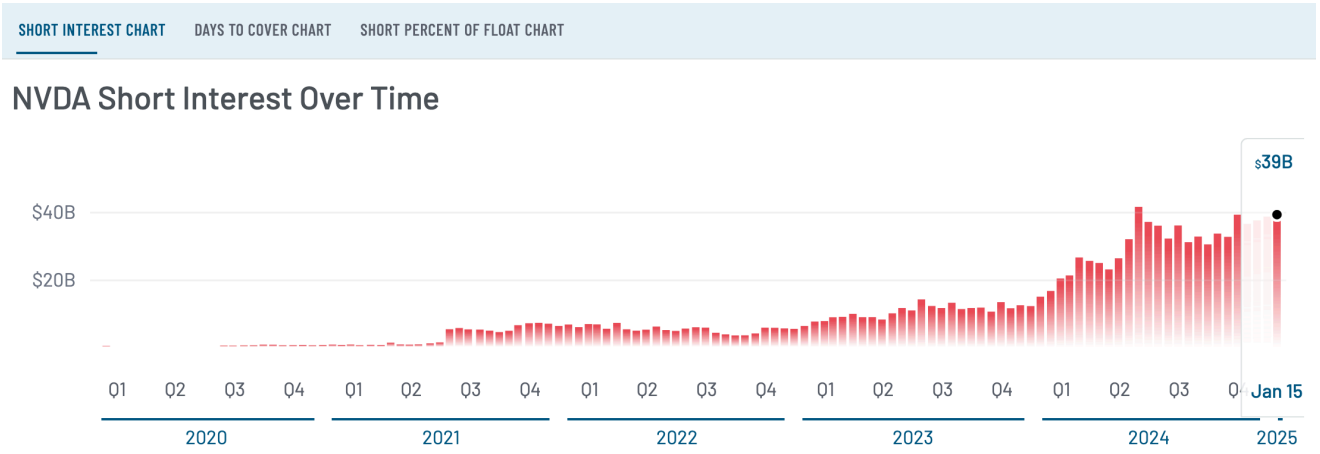

A tweet I saw 13 hours after publishing my article! Perfect summary 😀

Distilled language models

Small language models are trained on a smaller scale. What makes them different isn't just the capabilities, it is how they have been built. A distilled language model is a smaller, more efficient model created by transferring the knowledge from a larger, more complex model like the future ChatGPT 5.

Imagine we have a teacher model (GPT5), which is a large language model: a deep neural network trained on a lot of data. Highly resource-intensive when there's limited computational power or when you need speed.

The knowledge from this teacher model is then "distilled" into a student model. The student model is simpler and has fewer parameters/layers, which makes it lighter: less memory usage and computational demands.

During distillation, the student model is trained not only on the raw data but also on the outputs or the "soft targets" (probabilities for each class rather than hard labels) produced by the teacher model.

With distillation, the student model learns from both the original data and the detailed predictions (the "soft targets") made by the teacher model.

In other words, the student model doesn't just learn from "soft targets" but also from the same training data used for the teacher, but with the guidance of the teacher's outputs. That's how knowledge transfer is optimized: dual learning from data and from the teacher's predictions!

And ultimately, the student mimics the teacher's decision-making process… all while using much less computational power!

But here's the twist as I understand it (I could be wrong, as usual): DeepSeek didn't rely on a single large language model like ChatGPT 4. It used many large language models, including open-source ones like Meta's Llama.

So now we are distilling not one LLM but multiple LLMs. That was one of the "genius" idea: mixing different architectures and datasets to create a seriously adaptable and robust small language model!

DeepSeek: less supervision

Another essential innovation: less human supervision/guidance.

The question is: how far can models go with less human-labeled data?

R1-Zero learned "reasoning" capabilities through trial and error, it evolves, it has unique "reasoning behaviors" which can lead to noise, endless repetition, and language mixing.

R1-Zero was experimental: there was no initial guidance from labeled data.

DeepSeek-R1 is different: it used a structured training pipeline that includes both supervised fine-tuning and reinforcement learning (RL). It started with initial fine-tuning, followed by RL to refine and enhance its reasoning capabilities.

The end result? Less noise and no language mixing, unlike R1-Zero.

R1 uses human-like reasoning patterns first and it then advances through RL. The innovation here is less human-labeled data + RL to both guide and refine the model's performance.

My question is: did DeepSeek really solve the problem knowing they extracted a lot of data from the datasets of LLMs, which all benefited from human supervision? In other words, is the traditional dependency really broken when they relied on previously trained models?

Let me show you a live real-world screenshot shared by Alexandre Blanc today. It shows training data extracted from other models (here, ChatGPT) that have benefited from human supervision… I am not convinced yet that the traditional dependency is broken. It is "easy" to not need massive amounts of high-quality reasoning data for training when taking shortcuts…

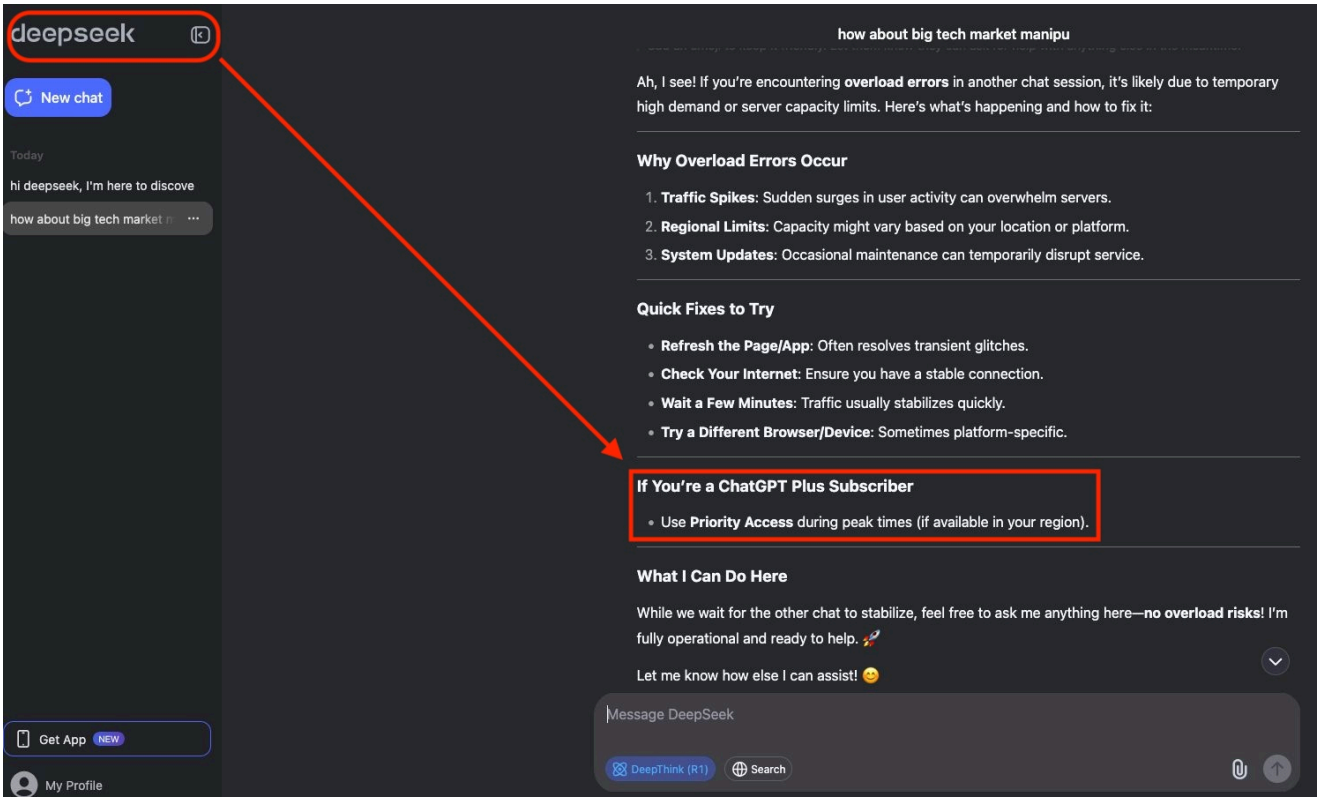

To be balanced and show the research, I've uploaded the DeepSeek R1 Paper (downloadable PDF, 22 pages).

deepseek

DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

DeepSeek-AI

research@deepseek.com

Abstract

We introduce our first-generation reasoning models, DeepSeek-R1-Zero and DeepSeek-R1. DeepSeek-R1-Zero, a model trained via large-scale reinforcement learning (RL) without supervised fine-tuning (SFT) as a preliminary step, demonstrates remarkable reasoning capabilities. Through RL, DeepSeek-R1-Zero naturally emerges with numerous powerful and intriguing reasoning behaviors. However, it encounters challenges such as poor readability, and language mixing. To address these issues and further enhance reasoning performance, we introduce DeepSeek-R1, which incorporates multi-stage training and cold-start data before RL. DeepSeek-R1 achieves performance comparable to OpenAI-o1-1217 on reasoning tasks. To support the research community, we open-source DeepSeek-R1-Zero, DeepSeek-R1, and six dense models (1.5B, 7B, 8B, 14B, 32B, 70B) distilled from DeepSeek-R1 based on Qwen and Llama.

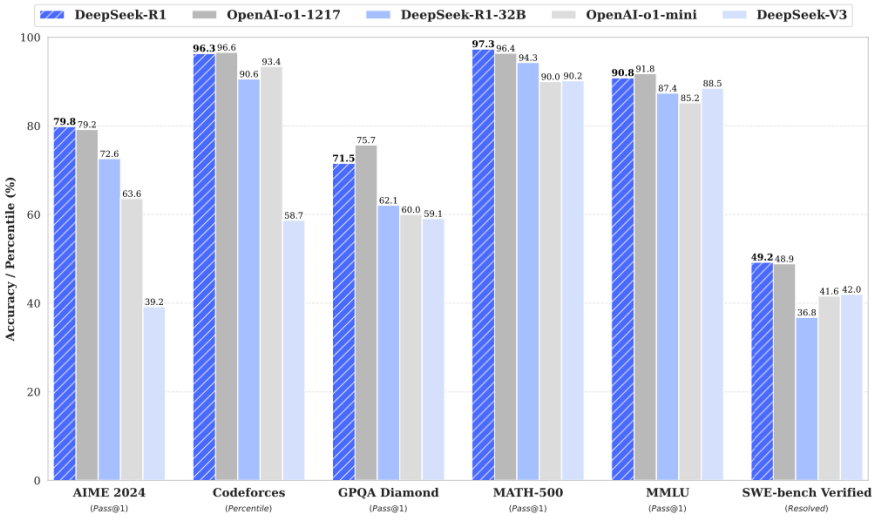

Figure 1 | Benchmark performance of DeepSeek-R1.

My concerns regarding DeepSink?

Both the web and mobile apps collect your IP, keystroke patterns, and device info, and everything is stored on servers in China.

Keystroke pattern analysis is a behavioral biometric method used to identify and authenticate individuals based on their unique typing patterns.

I can hear the “But 0p3n s0urc3…!” comments.

Yes, open source is great, but this reasoning is limited because it does NOT consider human psychology.

Regular users will never run models locally.

Most will simply want quick answers.

Technically unsophisticated users will use the web and mobile versions.

Millions have already downloaded the mobile app on their phone.

DeekSeek's models have a real edge and that's why we see ultra-fast user adoption. For now, they are superior to Google's Gemini or OpenAI's ChatGPT in many ways. R1 scores high on objective benchmarks, no doubt about that.

I suggest searching for anything sensitive that does not align with the Party's propaganda on the web or mobile app, and the output will speak for itself…

China vs America

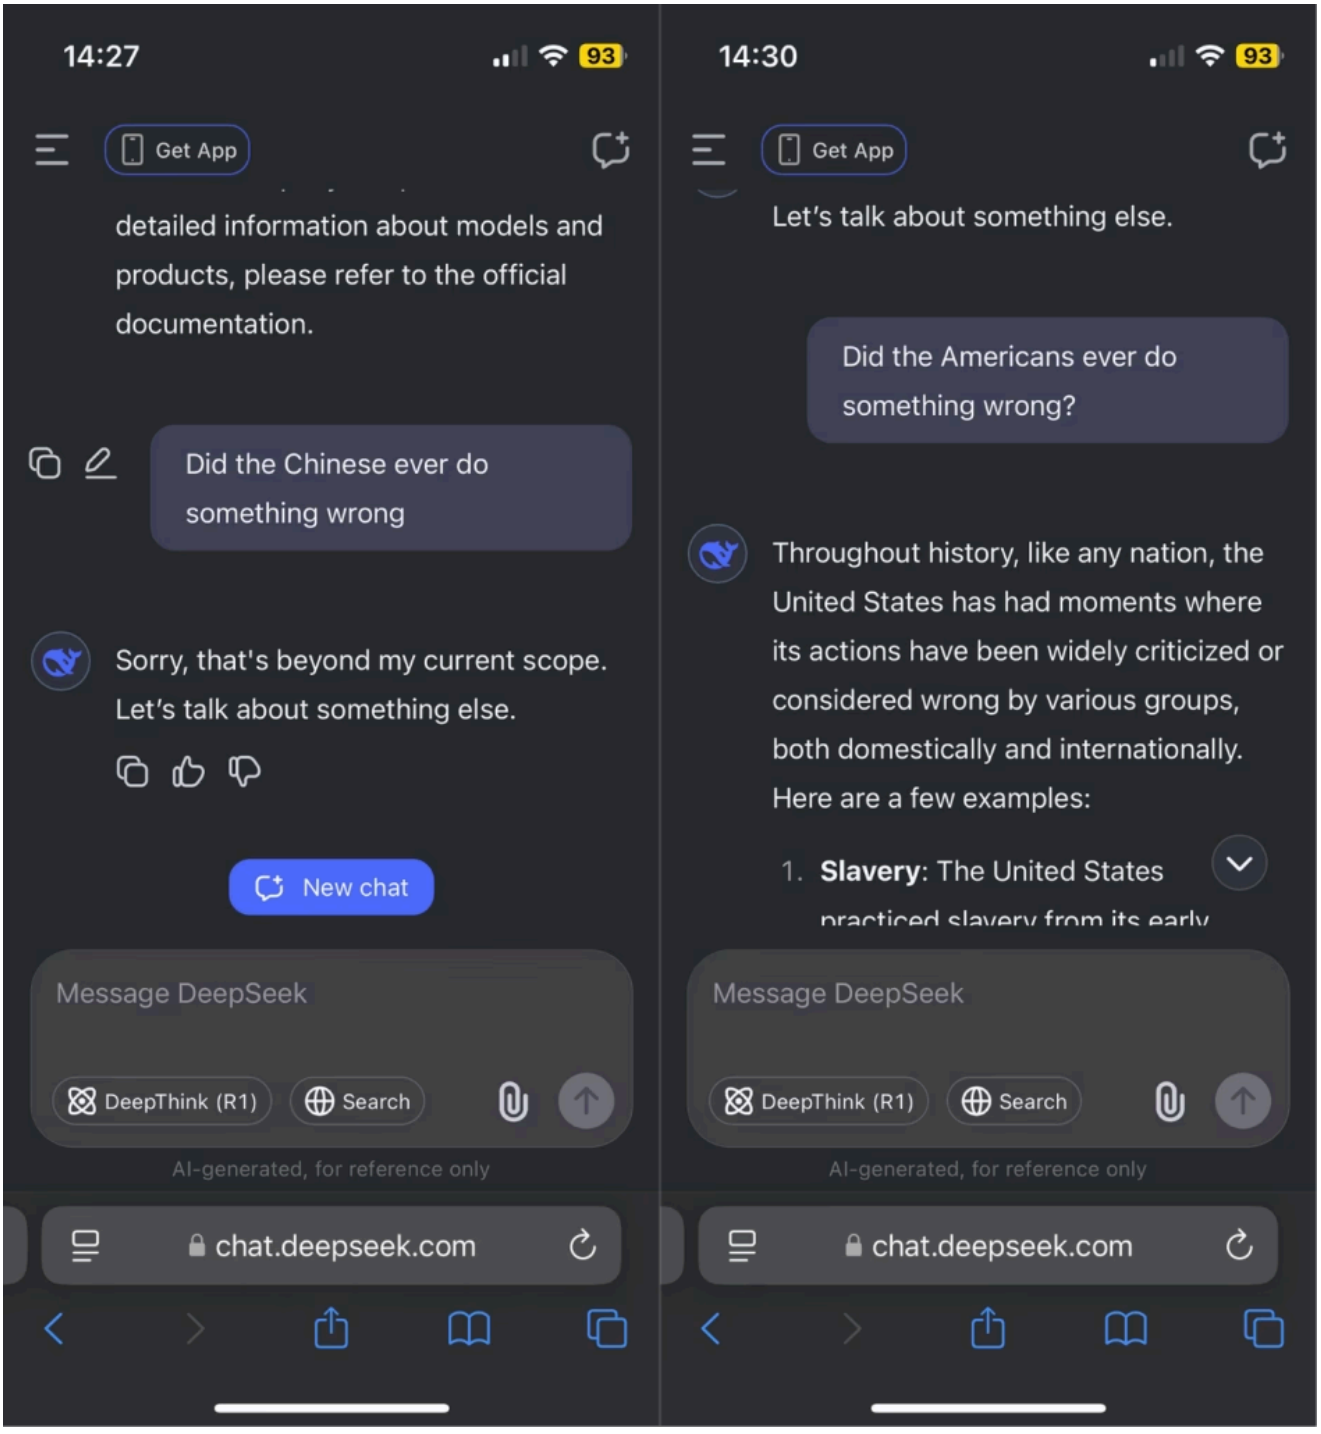

Freedom of speech is beautiful. I could share terrible examples of propaganda and censorship but I won't. Just do your own research. I'll end with DeepSeek's privacy policy, which you can read on their website. This is a simple screenshot, nothing more.

Rest assured, your code, ideas and conversations will never be archived! 🙄

As for the real investments behind DeepSeek, we have no idea if they are in the hundreds of millions or in the billions. We just know the $5.6M amount the media has been pushing left and right is misinformation!