ABUNDANCE INSTITUTE

DEEPSEEK AND THE FUTURE OF AI: A POLICYMAKER'S GUIDE

Over the last year, the big commercial US artificial intelligence labs have started locking down their research, choosing to keep their methods confidential. This reflects heightened market competition, an increasing desire to be the first to reach the milestone of "artificial general intelligence," and concerns, held with differing degrees of sincerity, about "AI safety."

It was a shock when a Chinese lab, DeepSeek, recently published a pair of papers that openly advanced the frontier of AI technology. The lab also publicly released an advanced model, R1, that is nearly as good as the top American models while costing a fraction as much to train and run. They also released a smartphone app that uses their model.

Like Meta, but unlike the other American AI companies, DeepSeek fully released the model's weights. This means anyone can download everything needed to run the model on their own computing infrastructure. This game-changing release marks China's arrival at the cutting edge of the AI technology scene.

Why does DeepSeek's advance matter?

DeepSeek R1 demonstrates that China has caught up to US AI development capabilities despite US export controls on advanced chips. China has world-class AI researchers. We can no longer pretend that the US has an insurmountable lead in AI.

It also accelerates the pace of increased efficiency in training and operating AI. US researchers have been impressed by R1's performance, particularly considering the relatively low compute cost claimed by the company.

DeepSeek wants to represent a new model of Chinese innovation. Past Chinese innovation has focused on copying and scaling the deployment of foreign inventions. Arguably, DeepSeek continues this tradition – but DeepSeek is deliberately positioning itself as an original, open innovator. In an interview in July, DeepSeek CEO Liang Wenfeng emphasized his desire to shift the perception of China-style innovation, saying, "The real gap [between US and Chinese AI innovation] is the difference between originality and imitation. If this doesn't change, China will always be only a follower..."

DeepSeek's commitment to openness, originality, and performing at the bleeding edge challenges the US approach of restricting access to AI technology in the name of premature commercial dominance and/or "AI safety."

Did export controls on advanced chips fail?

It's more accurate to say that they didn't slow China down as much as expected. DeepSeek improved the efficiency of its models out of necessity, partly motivated by export limits on cutting-edge chips like Nvidia H100s. DeepSeek claims to have trained V3, the core model underlying its pioneering R1 model, using only 2.788 million GPU hours using Nvidia H800s, a permitted chip. This outcome raises questions about the effectiveness of restricting access to advanced chips as a strategy for maintaining technological advantage.

Did DeepSeek really spend so little money on their models?

DeepSeek claims it spent the equivalent of $5.6 million USD in training costs for the base V3 model, on which the R1 model is based. While this amount cannot be independently verified, experts consider it plausible, given the model size and the details in the research papers. However, this is just the bare cost of the compute time for the model—DeepSeek notes that they are not counting staff or research costs, which could easily be $10s of millions of USD.

The cost to operate R1 (known as “inference cost”) appears to be significantly less than near-equivalent U.S. models, but the magnitude of the efficiency gain is unknown. DeepSeek is charging for API access about 1/20th of what OpenAI charges for API access to its o1 model, but this doesn’t prove that inference costs are actually that low. Microsoft is offering R1 API access for free

Does DeepSeek present security risks to users?

Some. For security reasons, there is an important difference between using the DeepSeek models and using the DeepSeek app. The DeepSeek model is essentially a very large spreadsheet of numbers that can be used on any suitable hardware. Running a DeepSeek model on personal or American-owned hardware presents minimal security risks. For example, American companies like Microsoft are running DeepSeek models on their hardware.

On the other hand, using the DeepSeek app presents a greater risk. Any data you upload into the app or prompts you provide are accessible by the DeepSeek company. Because DeepSeek is a Chinese company, that data could be exposed to the Chinese government.

However, the model itself poses another security risk. DeepSeek’s choice to release its models publicly and for free (known as “open weights”) has instantly made it one of the top tools for AI researchers and developers. However, DeepSeek is imbued with CCP values. For example, it will not accurately respond to questions of political sensitivity to the Chinese government.

It is in the US’s cultural, economic, and national security interest that the world’s most popular models be imbued with American values, not those of the CCP.

How should the US respond to China’s rise as an AI peer?

The US should shift its focus from hobbling China’s AI industry to accelerating our own innovation. President Trump set the correct tone by rescinding former President Biden’s AI executive order and replacing it with one that puts federal agencies in a posture of supporting American AI innovation. Defense-focused policies like export controls may have limited effectiveness. If AI is not developed rapidly and openly in America, it will be developed in China and adopted by the world. The US federal and state governments must focus on preventing or removing regulatory barriers that would reduce competition and slow deployments. American companies likewise should recalibrate their risk tolerance, dialing back unnecessary “trust and safety” practices that hinder model effectiveness and slow their release.

The US is known for its culture of open innovation, which DeepSeek is now adopting. However, if we abandon our values and adopt a European precautionary approach to AI regulation, China could completely eclipse us as the epicenter of AI development.

Glossary

AI SAFETY

An approach to AI development and regulation that focuses on potential speculative harms from AI, up to and including extinction of the human race.

ARTIFICIAL GENERAL INTELLIGENCE

A level of AI in which the model can do virtually any task as well as a typical human.

MODEL WEIGHTS

The numerical representation of the AI model that is generated during the training of the model.

OPEN WEIGHTS

An approach to AI development where the model weights are released to the public. Meta and DeepSeek use open weights, while labs like OpenAI, Google, and Anthropic use closed weights.

PRETRAINING

The earliest, most computationally intensive step in training an AI model, during which the model learns how to predict subsequent tokens in a way that is recognizable as language as well as improves in usefulness.

GET IN TOUCH

abundance.institute
@abundanceinst
@abundanceinstute

Chris Koopman
CEO
chris@abundance.institute

Neil Chilson
Head of AI Policy
neil@abundance.institute