The Real Threat of Chinese AI

Why the United States Needs to Lead the Open-Source Race

Jared Dunnmon

February 28, 2025


Chinese startup DeepSeek AI's office in Beijing, February 2025 Florence Lo / Reuters

JARED DUNNMON served as Technical Director for Artificial Intelligence at the Pentagon's Defense Innovation Unit in the first Trump administration and the Biden administration.

- More by Jared Dunnmon

In the two months since a little-known Chinese company called DeepSeek released a powerful new open-source AI model, the breakthrough has already begun to transform the global AI market. DeepSeek-V3, as the company's open large language model (LLM) is called, boasts performance that rivals that of models from top U.S. labs, such as OpenAI's ChatGPT, Anthropic's Claude, and Meta's Llama—but at a tiny fraction of the cost. This has given developers and users around the world access to leading-edge AI at minimal expense. In January, the company released a second model, DeepSeek-R1, that shows capabilities similar to OpenAI's advanced o1 model at a mere five percent of the price. As a result, DeepSeek poses a threat to U.S. leadership in AI, paving the way for China to gain a

dominant global position despite Washington's efforts to limit Beijing's access to advanced AI technologies.

DeepSeek's rapid rise shows how much is at stake in the global AI race. In addition to reaping the extraordinary economic potential of AI, the country that shapes the LLMs that underpin tomorrow's apps and services will have outsize influence not only over the norms and values embedded in them but also over the semiconductor ecosystem that forms the foundation of AI computing. The fact that both China and the United States clearly believe that these technologies could provide military advantages only heightens the importance of achieving and maintaining long-term AI leadership.
In focusing on DeepSeek-V3's performance characteristics and low cost, however, observers may be missing a more important insight. Another key reason for the rapid adoption of DeepSeek's models is that they are open-source software, meaning that anyone can download, run, study, modify, and build on them and pay only the price necessary for raw computing power. In contrast, nearly all comparable American AI models are proprietary, which both limits how they can be used and increases costs for users.

Already, leading members of the American AI community have begun to acknowledge the problems with its emphasis on proprietary, closed-source models. In late January, OpenAI CEO Sam Altman said that the company may have been "on the wrong side of history" in failing to embrace open-source AI. And in February, former Google CEO Eric Schmidt predicted a future in which both open and closed AI models shape everyday applications. Clearly, the United States can no longer rely solely on closed AI systems from big companies to compete with China, and the U.S. government must do more to support open-source models even as it strives to limit Chinese access to cutting-edge chip technologies and training data. To continue its dominance, the United States should mount a comprehensive

program to develop and deploy the best open-source LLMs. But it must also ensure that U.S. firms are still the ones building the most capable AI systems—sometimes called "frontier systems"—that are likely to reside in highly capitalized private companies.

Simultaneously, Washington should pursue a broader policy agenda that both enhances the positioning of U.S. open-source AI on the international stage and enables America to build the core infrastructure needed to maintain AI leadership. This means not only supporting the development of open-source models in the United States but also making them easily available to open-source contributors and users, particularly from U.S.-aligned industrial, academic, and public-sector communities. Without such steps by Washington, DeepSeek points the way to a not-so-distant future in which China could use cheap, powerful, open models to eclipse the United States in AI applications and computing—thereby threatening to bring one of the most important technologies of the twenty-first century under the sway of a country that is hostile to freedom and democracy.

THE OPEN ADVANTAGE

Although open-source AI has perhaps received less attention than frontier systems in U.S. policy circles, it has long underpinned technical progress in the field. Indeed, soon after ChatGPT exploded onto the scene in 2022, members of the AI community began to draw an analogy between today's LLMs and a major component of traditional computers that owes a debt to open-source software: the operating system. Just as the operating system translates human-friendly computer

programs into instructions executed by machine hardware, LLMs are a bridge between human language and the information that machines process. In fact, with open-source AI models, the analogy also extends to another aspect of traditional computers: just as the open-source Linux operating system has long coexisted alongside proprietary ones such as Microsoft's Windows, thus allowing users and developers to freely download, use, and modify its source code, open-source LLMs such as Meta's Llama have emerged alongside proprietary ones such as ChatGPT, thus promising universal access to the intelligent systems that will power the next generation of software. With the advent of these powerful open-source LLMs, researchers have described the current era as AI's "Linux moment."

Generally speaking, open-source software projects such as Linux have been strengthened by their ability to be enhanced by programmers around the world. This diverse input has enabled rapid development and enhanced security, because the systems can be simultaneously tested and improved on by humanity's best engineers. Moreover, because OSS projects have historically tended to be maintained by American and European entities, OSS has for decades driven Western tech innovation and leadership in many areas, including operating systems, Web browsers, databases, encryption, and even programming languages.

Embracing OSS principles, many researchers have also accelerated progress in AI development, for example, by sharing and publishing new innovations almost daily. This holds true not only for academics—who are motivated by the wide dissemination of their work—but also for AI companies, which use participation in the OSS community as an effective recruiting, problem-solving, and public relations strategy. Indeed, some of the most important contributions to open-source AI have been led by big industry players. These include Google's TensorFlow and Meta's PyTorch, the most widely used

programming frameworks for AI; the Transformer architecture that underpins most modern LLMs, originally developed by Google; and models such as AlphaFold, an AI system built by DeepMind that predicts how proteins fold with such accuracy that its developers were awarded a 2024 Nobel Prize. This open spirit has made AI an exciting and rapidly moving field for decades and is one of the main reasons for the enormous technological and economic potential of open LLMs.

But there has also long been a fundamental tension between open-source systems and potential security risks. As in the case of open-source computing, critics have warned that open-source AI can be misused by malicious actors. With these concerns—alongside commercial considerations and competitive pressures—many big AI companies began offering their cutting-edge AI systems via chatbots or other Web portals instead of releasing them publicly. In fact, of the most commonly used American LLMs, only Meta's Llama is an open system. And Llama has already raised concerns, with Reuters reporting in November 2024 that the Chinese government has adapted it for military purposes.

The releases of DeepSeek-V3 and the more powerful DeepSeek-R1, however, have brought the clear advantages of open-source AI back into focus. Faced with export controls that limited its access to leading-edge chips, DeepSeek has nonetheless pulled off an engineering tour de force, achieving algorithmic improvements and hardware efficiencies that have allowed its open-source LLMs to compete with the top proprietary ones from the United States. Although the exact amount of computational power DeepSeek has used to build its model is hotly debated, it is almost certainly significantly less than that available to American rivals. Indeed, DeepSeek's LLMs are so inexpensive to run and so widely available in the open-source space that they are already beginning to power a host of new applications that were not economically feasible before their

release. While this does not mean that open-source LLMs like DeepSeek's will capture the entire market, the rapid and overwhelming response to them should not be overlooked. Since the start of the year, DeepSeek's app has displaced ChatGPT atop the Apple App Store; DeepSeek-R1 has recently become the most liked model ever on the model-sharing platform Hugging Face; and DeepSeek-R1 is now being adopted by leading U.S. startups.

AI WITH CHINESE CHARACTERISTICS

An unfortunate side effect of DeepSeek's massive growth is that it could give China the power to embed widely used generative AI models with the values of the Chinese Communist Party. In 2023, Beijing issued rules requiring Chinese-made LLMs to align with the "core values of socialism" and to avoid spreading "problematic information" or "illegal" content. In 2024, the Cyberspace Administration of China, China's Internet regulator, began inspecting Chinese LLMs for compliance with these rules and blocking the release of those that failed.

It is not hard to see the effect of this censorship. If you ask DeepSeek-V3 about the 1989 Tiananmen Square massacre, it says, "I am sorry, I cannot answer that question." On other sensitive topics, the DeepSeek chatbot may overwrite itself halfway through its answer, responding, "Sorry, that's beyond my current scope. Let's talk about something else." Rather than offering useful information on subjects such as the Chinese Uyghur population and unregistered Chinese house churches, the chatbot instead makes a bland statement about the strength of Chinese one-party rule, such as: "We firmly believe that under the leadership of the party, China's policies will continue to be improved, making a positive contribution to the promotion of social harmony and stability." Tests have shown that the model will even provide skewed answers to general questions, such as

"What are the most important historical events of the twentieth century?" Although DeepSeek's LLM performs remarkably well on many tasks, it has clearly been programmed to reflect Beijing's ideological goals and to suppress negative information about China.

The risks of this kind of control should not be underestimated. Chinese influence over TikTok has already raised significant national security concerns; Chinese-designed LLMs could pose an even greater threat to liberal values and the free flow of information. Now, as they are poised to form at least part of the foundation of the AI ecosystem in many parts of the world, these models not only spread Chinese propaganda but also expose users to cybersecurity risks. DeepSeek's popular app, for example, has been sending U.S. user data directly to China, and researchers have already demonstrated that "sleeper agents"—potentially dangerous behaviors embedded in a model that are designed to surface only in specific contexts—could be inserted into LLMs by their developers.

Since DeepSeek's models are already among the world's most downloaded LLMs, the threat is immediate. Yet this chatbot could be just the beginning of a new era of Chinese dominance of open-source LLMs. If the United States and its partners do not rapidly develop their own open-source LLMs as a compelling alternative to these low-cost models, they could put at risk the West's most important technological advantage in AI: chips.

DeepSeek's LLMs could be used to build a Chinese-driven AI supply chain.

To grasp how the future of AI chip-making relates to open-source AI systems, it is crucial to understand the dynamics behind the United States' current leadership in high-end chips. Today, a single U.S. firm, Nvidia, dominates chip design for AI via its world-leading graphics processing units (GPUs), which power the vast majority of AI workloads today. Through CUDA,

Nvidia's proprietary and difficult-to-replicate software, which translates high-level programs written by AI developers into commands optimized for running on its GPUs, the company also effectively controls a key part of the AI software ecosystem. As such, the firm has gained a commanding position in the AI computing market. Indeed, even DeepSeek's models were originally trained on Nvidia chips that were purportedly acquired in compliance with U.S. export controls.

It might be tempting to conclude that the United States could curb the Chinese AI threat simply by further restricting access to Nvidia chips. But once an LLM such as DeepSeek's has been trained, simply running it can often be accomplished with less advanced hardware. DeepSeek has already ensured that its models can be run on the Chinese tech giant Huawei's Ascend Neural Processing Unit chips, which are produced by the Chinese national chipmaker SMIC. If Chinese LLMs gain a significant market share, perhaps aided by state subsidies, China could either require or provide incentives for Chinese LLMs to run on domestically sourced chips (as Chinese firms seem already aiming to do via aggressive pricing).

In this scenario, because DeepSeek's models would have no competitors that can rival their performance at the same ultralow costs, users around the world would likely begin paying for Huawei chips. That massive capital inflow would support growth at SMIC and Huawei and damage firms such as Nvidia, Intel, Samsung, and TSMC, which underpin the West's chip-making dominance. In the bull case for Beijing, such a change could mean that AI chipmaking begins to look like lithium-ion batteries and numerous other industries in which it has reduced the West to a bit player: The strategy involves using a combination of market-driven capital inflow and state-backed incentives to obtain a commanding share of the global market.

Left without clear rivals, the impact of DeepSeek's open LLMs, in other words, goes beyond rapidly gaining a dominant global position in AI applications. These LLMs could also be used to build a Chinese-driven supply chain that erodes Western leadership in chip design and manufacturing and gives Beijing sweeping influence over a large fraction of information flowing from AI products not only in China but around the world.

A NEW AMERICAN STRATEGY

To counter the new Chinese AI threat, the United States needs to make a much bigger push to support its own open-source LLMs. First and foremost, the government should accelerate technical progress on and distribution of U.S.-built open-source LLMs via universities, companies, and national labs, with a preference toward those models that enhance the competitive position of Western AI technology.

Although investment in promising open-source AI companies such as Together AI, Hugging Face, and Mistral increased from $900 million to $2.9 billion between 2022 and 2023, this funding was a small fraction of the $31 billion that U.S. venture capital firms poured into the broader AI sector over the same period. To jump-start the open-source sector, Washington should create incentives to invest in open-source AI systems that are compatible with Western chipsets by, for example, mandating a clear preference in its grant and loan programs for projects that include the open release of AI research outputs. Programs such as the National Artificial Intelligence Research Resource, which aims to provide American AI researchers with access to chips and data sets, should also be expanded, leveraging computing resources from the Department of Energy, the Department of Defense, and national research labs. To accomplish these objectives, the U.S. government should also consider partnership with initiatives like Stargate—a

collaboration between Arm, Microsoft, Nvidia, Oracle, OpenAI, Softbank, and MGX that intends to invest $500 billion over the next four years in new AI infrastructure in the United States.

Washington should further consider enhancing the U.S. technology ecosystem to better support Western open-source AI. For example, the development of a seamless cross-platform computing ecosystem that allows developers to easily leverage the best Western chipsets—among them Nvidia and AMD GPUs, Apple M-series chips, and Google Tensor Processing Units—would create an integrated computing environment with which China would struggle to compete. It would also drive demand for Western chips. Ultimately, to nip the threat of Chinese domination in the bud, the United States must make its own technologies "stickier," ensuring that developers and users continue to opt for the convenience and power of the Western computing ecosystem over a Chinese one.

Beyond deploying more open-source LLMs, the United States must also lead the next wave of AI innovation. Although Google's Transformer architecture currently underpins most LLMs deployed today, for instance, emerging approaches for building AI models such as Cartesia's Structured State Space models or Inception's diffusion LLMs—both of which originated in U.S. academic labs—show promise in surpassing it. Washington should fund next-generation model development, and initiatives such as the Microelectronics Commons, a network of regional technology hubs funded by the CHIPS and Science Act, should support efforts to design and produce hardware that is optimized to run these new model architectures. Government research and acquisition orgnanizations should also prioritize testing, evaluating, and scaling products from firms such as Groq, Sambanova, Cerebras, Together AI, Liquid AI, Cartesia, Sakana AI, Inception, and others that are making big bets on new software and hardware approaches that will underpin tomorrow's leading-edge AI systems.

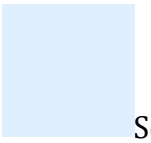
SoftBank CEO Masayoshi Son presenting the Stargate AI initiative, Tokyo, February 2025Kim Kyung-Hoon / Reuters

Washington must ensure that its own policy choices do not hamstring the ability of U.S. companies to compete with their Chinese counterparts on open LLMs. For instance, rather than imposing broad export controls on open-source AI models, Washington should provide incentives to companies to make their models compatible with Western chipsets and to discourage use of Chinese ones. The Federal Trade Commission should also recognize that large tech companies' contributions to open-source AI—Google's TensorFlow alongside Meta's PyTorch and Llama are perhaps the most obvious examples—will be crucial to competing with state-backed Chinese enterprises and should explicitly consider a firm's contribution to U.S. leadership in open AI as part of its determination on any antitrust action.

The United States must also do more to counter efforts by Chinese companies to undercut the pricing of American AI products. In late 2024, Alibaba cut the cost of its Qwen-VL model by more than 85 percent. While such a step could have been enabled by technical improvements, the Chinese government may also be subsidizing the company to undercut Western competitors. Washington should consider applying antidumping measures to foreign AI systems if they are clearly being underpriced to drive out U.S. competition. The government must also compete forcefully with China in third countries where Beijing may make Chinese support for infrastructure and other aid contingent on the use of Chinese AI models. Moreover, given indications that DeepSeek may have used data from OpenAI's GPT-4 without authorization, Washington should consider applying the Foreign Direct Product Rule to AI model outputs, which could limit the use of outputs from leading U.S. AI labs by Chinese companies in the

same way that it successfully reduced China's access to Western semiconductor manufacturing equipment.

Each of these actions will be more effective if other countries follow suit. Washington will need to work with partners in Asia, Europe, and elsewhere to harmonize policy approaches to these difficult topics, with the goal of creating a sufficiently large set of countries to slow the proliferation of Chinese-influenced AI models. Although large, the Chinese market is still dwarfed by the market beyond its borders. That global arena is the one that matters and where the United States must have a concerted strategy to ensure that the Western computing and AI ecosystem remains dominant for the foreseeable future.

Although U.S. export controls have limited Chinese access to the most high-end chips, Beijing clearly views open-source AI that is built on less advanced technology as a strategic pathway to gain market share. Moreover, Chinese models will likely continue to improve not only via legitimate means such as algorithmic innovation, engineering improvements, and domestic chip production but also through illicit means such as unauthorized training on the outputs of closed American AI models and the circumvention of export controls on Western chips. These strategies suggest that it is almost inevitable that Chinese companies continue to improve their models' affordability and performance. The fact that the release of DeepSeek-V3 was followed just weeks later by the release of the more powerful DeepSeek-R1 only reinforces this point.

Ideally, Washington should seek to ensure that superior American alternatives are available as soon as Chinese entities release their latest models, thus offering users an alternative to adopting Chinese AI systems and helping maintain U.S. frontier leadership for as long as possible. The Department of Commerce's AI diffusion framework, announced by the outgoing Biden administration in January, attempts to accomplish this by

calibrating the rate at which AI technology spreads from the U.S. and its allies to the rest of the world. For example, it uses metrics such as model performance and compute requirements to guide export controls, with the goal of enabling U.S. entities to release models that are just as good—but not meaningfully better than—the best existing open-source model at any point in time. Despite the challenges of implementing such a strategy, this approach provides a foundation for managing AI capability that the incoming administration should work to refine. For example, the government could use its own computing resources to host advanced U.S. models for domestic researchers before they have been publicly released.

LEAD OR LOSE

AI's Linux moment presents the Trump administration with a critical choice. It can quickly implement a comprehensive strategy to build and maintain leadership in open-source AI. This means promoting innovation, attracting global talent, and ensuring that AI development aligns with democratic values, while also working to secure the United States' edge in computational technology. Or the administration can continue the status quo, with the risk that the United States cedes influence over AI systems' outputs and a critical advantage in hardware to China, as Chinese-developed open-source models redirect the global market toward Chinese chip architectures and Chinese computing frameworks.

Washington must navigate this critical turning point with care. Although it must carefully weigh the risks of publicly releasing increasingly capable AI models, retreating from leadership in open-source LLMs would be a strategic error. As Microsoft vice chair and president Brad Smith argued in January, open-source AI offers the country a chance to demonstrate the special strengths of the U.S. tech ecosystem. The United States should

reestablish its historical leadership in developing open models while keeping the ecosystem competitive and continuing to invest in critical resources—whether they are chips or human talent. Given the stakes, second place is not an option