March 2026

Expert Insights
PERSPECTIVE ON A TIMELY POLICY ISSUE

OWEN J. DANIELS, HANNA DOHMEN

Open Models, Soft Power, and the Spectrum of U.S.-China Artificial Intelligence Competition

The competition between the United States and China in artificial intelligence (AI) is entering a new phase and is increasingly focused on influencing the global AI ecosystem, and open models are emerging as a critical aspect of soft power. *Soft power* refers to the ability to exert influence and shape other actors' behaviors and preferences through factors other than military strength, such as shared ideology, commercial opportunities, or cultural attractiveness. Our analysis informs the United States' open-source AI strategy by first examining the diverging frameworks of the United States' and China's AI action plans. We evaluate how AI developers in both countries are approaching open models to gain influence. We also offer recommendations for U.S. policymakers to ensure that the United States maintains its technological leadership. We argue that a strategy that emphasizes stronger support for an open-model ecosystem in the United States, a recalibration of export controls, and the creation of incentives for firms to adopt permissive licensing practices is needed to counter the influence that China hopes to gain through diffusing its technology and tech infrastructure. Beyond the pos-

RAND

The open-model ecosystem appears to be an increasingly important playing field for soft-power competition between the United States and China.

sibility that China might be seen as the primary champion of sharing AI's benefits globally, particularly in the developing world, the broad adoption of Chinese open models might collectively lead to AI developers and researchers being locked into using China's technology.

The United States' and the People's Republic of China's Artificial Intelligence Action Plans

As the releases of the United States' and China's new AI plans in 2025 suggest, market drivers and a battle for influence, in addition to national security concerns, appear to be central to the competition. National security is an important consideration in the U.S.-China AI rivalry, but it represents only one dimension of a broader and nuanced competition. The two action plans illustrate this point through their calls to develop open AI models, although they do so in differing ways and levels of detail.

Open models refer to AI systems whose *weights* (the parameters that determine the model's outputs for a given input), are freely available to users to download and adapt for specific use cases.[1] A model's level of openness can also vary depending on certain factors, such as the licensing restrictions governing the model's use.

The U.S. plan, *Winning the Race: America's AI Action Plan*, is postured as a domestic policy document to guide the federal government's execution of AI policy and recommends (1) ensuring access to large-scale computing power for researchers and start-ups, (2) creating public-private partnerships to increase widespread access to private-sector computing and other resources, and (3) encouraging small and medium-sized businesses to adopt open models.[2] The document cites the importance of creating open models "to ensure America has leading open models founded on American values" and because open models' potential to become global standards in commercial and academic research lends them "geostrategic value."[3] The plan notes that, "While the decision of whether and how to release an open or closed model is fundamentally up to the developer, the Federal government should create a supportive environment for open models."[4] The plan is the strongest indicator of the administration's stance on open-source models, but it should be read as part of a broader policy framework that promotes the exportation of the U.S. technology stack—the AI hardware, models, software, applications, and standards. National security concerns and economic influence, the latter of which is executed by U.S. firms, both play significant roles in the plan.

China's "Global AI Governance Action Plan" also embraces the open-model environment's geostrategic value, although the document couches the framing in

more globally beneficial terms because China's plan is primarily an instrument of global diplomacy.[5] Although the U.S. plan and China's plan have international scopes and policy implications, China's plan appears—at least in its framing—to focus more on establishing collaborations with other nations, particularly in the developing world. As the document states, China aims to "build transnational open-source communities and secure, reliable open-source platforms . . . promote the open sharing of basic resources, lower the barriers to technological innovation and application, avoid duplication of investment and waste of resources, and promote the inclusiveness and accessibility of AI technology services."[6] The plan identifies supporting countries in the Global South in developing, engaging with, and applying AI as critical steps to achieving more equitable and inclusive technology development. However, reading between the lines, we note that China's plan also appears set on entrenching its national technology in global stacks—note the language on avoiding wasteful "duplication of investment."[7] Still, language throughout the plan focuses on cooperating in international fora (such as the United Nations) to boost global AI development, create shared international standards and governance mechanisms, and, ultimately, "adhere to the spirit of openness and sharing."[8]

In short, the open-model ecosystem appears to be an increasingly important playing field for soft-power competition between the United States and China regarding AI. China arguably has a head start in developing its strategy. Since releasing its *Next Generation Artificial Intelligence Development Plan* in 2017,[9] China has called for the development of open AI models as a pathway to harnessing innovation in AI across sectors and domains. China's DeepSeek and Moonshot AI released the R1 and Kimi K2 models, respectively, in 2025, thus introducing highly customizable and performant Chinese models onto the market. Both models were capable of competing with leading U.S. closed and open models alike. Moreover, because of the generally open-source nature of China's leading AI models, the entire innovation ecosystem benefits from cross-pollination. Open-model developers can use their peer's model architectures, libraries, and mechanisms to build and improve their own models, which results in a stronger innovative ecosystem.

Global researchers and businesses seek to reap the societal benefits of AI, as vaunted by model developers and technologists, in such areas as the economy, health care, and education—particularly in the developing world, where polling suggests optimism about AI's transformative potential is high.[10] However, China might have an advantage embedding its technology in AI stacks if the U.S. market does not produce viable alternatives. OpenAI released its gpt-oss-120b open model in early August 2025, the company's first open release since 2019, which is a seeming acknowledgment of open ecosystems' increasing influence in attracting AI users.[11]

Open models are not without downside risks compared with their closed counterparts, particularly from a security perspective.[12] Open models can be adapted to exploit model vulnerabilities, remove safety features, and misuse or maliciously deploy models. Furthermore, China has openly acknowledged that open models are a pathway to military-civil integration that promotes the adoption of civilian technologies for military ends.[13]

However, in the broader context of the U.S.-China rivalry, including the competition for global political and

economic influence and collaboration with other nations to uncover AI's full potential, open models will likely be an important tool in the countries' soft-power tool kits. As of early 2026, competition for the best open models is highly dynamic. Rankings on leaderboards that capture which models are most frequently downloaded and used are fluid, and many researchers gravitate toward using the best models available rather than orienting toward U.S. or Chinese models in particular. Data are limited and the market dynamics that drive researchers and businesses to adopt one model over another are not clear, especially because models' practical use cases remain largely nascent. However, perception that a powerful, adaptable model could be used to solve problems is clearly resonant: Researchers at the U.S. Center for AI Standards and Innovation found that downloads of DeepSeek models on model-sharing platforms increased nearly 1,000 percent since the release of DeepSeek's R1 model in January 2025.[14] Although China's flashy releases of models—such as those from DeepSeek or Moonshot AI—grab attention, Chinese dominance in the space is by no means preordained. That said, the United States needs to remain a key player in the open-model ecosystem, such as by providing funding for research and creating the conditions for companies to combat the possibility that researchers and commercial developers become locked in to using Chinese models over the long term. The U.S. government needs to closely monitor the dynamics and trends of the open-model ecosystem by collecting data that can inform future policies and exploring what incentives might lead more U.S. companies to make their model weights open access.

U.S. strategy should be geared toward facilitating a stronger U.S. ecosystem for open models in the long term; this way, developers can compete globally while taking advantage of the United States' leadership position in computing power.[15] Ways to strengthen the open-model ecosystem could include exploring policy mechanisms that encourage researchers and companies to use U.S. open models in exchange for access to greater computing resources, which would draw on the strength of the United States' cutting-edge AI infrastructure.[16] Such incentives could increase the likelihood that researchers and would-be commercial adopters of open models turn to U.S. developers and help create a flourishing open research and commercial environment in the United States. There is an important distinction between releasing powerful open models that can proliferate and having a vibrant research ecosystem that is best positioned to capitalize on such models. China has achieved some significant successes in the former, but the United States should strive for success in the latter. Maintaining the U.S. lead in the closed ecosystem frontier should continue to be a top priority, particularly for national security uses of AI, but the United States needs to complement its efforts by fostering a more-robust open-model ecosystem to meet China's soft-power challenge.

Open Models and Soft Power

To more fully understand the relationship between open models and soft-power influence, it is useful to understand the specific characteristics that differentiate open models from proprietary, closed models.

Benefits of Open Models

The benefits of open models lie in the models' capacity to allow a greater variety of actors to participate in AI research and development, including academic and commercial researchers who might not otherwise have access to the compute to pretrain their own models. Broadly speaking, the common users of open models can be split into two nonexclusive camps of researchers and commercial developers. Researchers are interested in accessing models for a variety of purposes, such as evaluating model performance and finding new use cases that are well suited to AI solutions by fine-tuning models. These researchers can visit open-model repositories (such as Hugging Face),[17] browse for models that could be used for specific tasks (such as multimodal content generation models, natural language–processing models, and computer vision models); download a model, then adapt or fine-tune the model to perform optimally with the researcher's data or in a specific application's context. Open models offer flexibility and performance with lower barriers to entry. Commercial developers—some of whom might be researchers seeking to profit from their academic work—might try to capitalize on open models for a less costly, less compute-intensive, and more adaptable way to build AI business applications.[18] Open-source models arguably also benefit from a broader user base's increased capacity to identify security issues, vulnerabilities, or other failure modalities within the models.

Thinking about open models' usefulness to researchers and commercial developers is helpful for framing policy approaches to the challenges and opportunities these models present. There certainly is not an iron curtain between research and commercial approaches to open-model development; researchers might be well positioned

The benefits of open models lie in the models' capacity to allow a greater variety of actors to participate in AI research and development.

to develop market applications of open models using their work, and they might favor the models that they used for research purposes when developing applications. Both researchers and commercial developers are likely to prize a powerful, low-latency model that suits their research or business tasks. Nonetheless, the research-commercial distinction highlights the varying uses of open models and how different actors might prize a particular model's unique characteristics.

Broadly speaking, researchers are often interested in accessing suites of the best available models and the original model's weights to test for factors (such as output replicability) in order to prod at models' limitations or to compare and rank the different models' performance. This type of research can help advance AI progress more broadly, occasionally without immediately obvious economic upsides.

Commercial developers—whether they are AI companies or other (non-AI) companies seeking to harness AI

Releasing powerful open models could help China position itself as the main diffuser of AI capabilities, particularly to the developing world.

for a specific product—might be more interested in accessing a powerful model that is well suited to a particular application. Developers can build products that are viable on the marketplace using an open model at a lower cost-to-performance ratio than using a proprietary model, thus achieving powerful performance at lower costs with open models that have near-frontier capabilities. Using a model to rapidly build prototypes, easily incorporate products into existing workflows, and attract the attention of financial backers or venture capitalists might be of paramount importance. Commercially speaking, a model's adaptability might be more significant for users than the ability to research a model's capabilities, so accessing the data that a model was trained on or its original weights might be less of a priority.

Licensing of Open Models

There are varying degrees of openness, which are partly determined by the licensing rules under which a model is released. For example, users might be able to download a model and modify its weights or fine-tune it for a specific purpose, but users cannot access the model's source code or distribute the derivative model or its outputs—this would make the model open weight, but the model would not meet open-source criteria.[19]

Generally speaking, more Chinese developers have released models under more-permissive licenses than U.S. developers. For instance, Meta's Llama models are distributed under a community license agreement,[20] which allows users to download the model weights and code but imposes limitations on distribution, including for businesses that provide AI services to more than 700 million monthly active users. Developers must attribute credit to Meta for models that are Llama offshoots. Google's Gemma models are also released under a proprietary license[21]—as opposed to an industry-standard license—which reserves the right to restrict use of Gemma services that the company believes are in violation of its prohibited use policy. These proprietary restrictions have led some proponents of open models to claim that Llama and Gemma models are not truly open source because the companies create obstacles for businesses to incorporate the models into products.[22]

Many of the latest and most-advanced open models released by Chinese developers adhere to standard industry licenses and are less restrictive to use.[23] DeepSeek's R1 was released under an MIT License, which allows anyone to use, modify, and distribute derivatives of the model and requires users only to include a copyright notice in distributions.[24] Kimi K2 models are released under a similar

but modified MIT license that maintains the permissive standard while requiring that commercial products with more than 100 million monthly active users or more than $20 million in monthly revenue must prominently display "Kimi K2" on their interfaces.[25] In a move that seemingly follows Chinese developers' more-permissive licensing approaches, OpenAI released its first open-source model, gpt-oss, under the Apache 2.0 License.[26] The Apache 2.0 License imposes no commercial-use limitations, instead allowing developers to release derivatives without having to adhere to company use policies.

A model's licensing details matter significantly for commercial developers. For commercial uses, permissive licenses (such as the MIT and Apache 2.0 licenses, under which the R1, Kimi K2, and gpt-oss models were released) encourage models' adoption by reducing the risk of legal uncertainty, regulatory burdens, and the possibility of model adopters profiting from technological applications or bearing responsibility for model misuses. Licensing agreements with use restrictions or distribution requirements often use ambiguous terms, which could lead to the commercial developer facing legal action from the model provider.[27] This is especially critical for start-ups with small teams and limited financial and legal resources. For researchers, however, the level of openness that such licenses provide does not guarantee replicability. Even the most permissively licensed models do not provide access to training data.

The choices that AI developers make in licensing their models—whether open sourcing models at all, using proprietary licenses, or using permissive licenses—are not only technical and legal decisions; they are also important factors in shaping commercial AI adoption and research. As OpenAI chief executive officer Sam Altman noted shortly after DeepSeek's R1 release, OpenAI's lack of an open-model strategy as of January 2025 put it "on the wrong side of history."[28] This idea links directly to the notion of soft-power influence with AI, and specifically in the context of the U.S.-China rivalry. Academic and commercial actors with constrained resources—in terms of either compute or financial resources—might prize a model's level of openness if its capabilities are equal to or comparable with a less-open competitor's model. The existing open-model space is fluid and competitive, and the degree to which a model is open could be a favorably differentiating factor. Models that can offer the right balance of performance and favorable licensing might be especially attractive to users, and the prominent marketing pushes of such companies as DeepSeek and Moonshot AI appear geared toward users seeking these qualities.

New Models, Same Strategy

Although the releases of powerful open models and the support for open ecosystems in the U.S. and Chinese AI action plans are notable, adopting an open strategy for emerging technology development is not new for China. Since the 2010s, China's AI technology developers and planners have advocated using open technology development as a strategic policy tool—including in the 2017 AI plan, the 2021 14th Five-Year Plan, and various other policy documents.[29] There appear to be several motivations for this focus. First, China hopes to broaden its influence abroad in global AI stacks, which is in line with other steps the country has taken to become a world-leading digital infrastructure provider in telecommunications,[30]

surveillance technologies,[31] and data center development.[32] Second, China aims to depend less on Western components of the AI supply chain, particularly for computing resources. This aim is especially salient as China contends with U.S. export controls restricting its ability to import most leading-edge semiconductors and to access other Western compute needed to develop leading AI models.[33] Third, as part of its well-established military-civil fusion efforts, China appears eager to harness advances in open models to seek advantages in its military rivalry with the United States.[34]

Releasing powerful open models could help China position itself as the main diffuser of AI capabilities, particularly to the developing world, which could have market and infrastructure implications. As mentioned previously, publics in developing countries appear more sanguine on AI's ability to provide benefits than those in their developed counterparts. For instance, Google and Ipsos polling found that more than 70 percent of people surveyed in emerging markets—such as Brazil, Mexico, South Africa, and the United Arab Emirates—said that AI would bring positive societal benefits across a variety of areas; for comparison, only half of U.S. respondents noted the same.[35] Through its willingness to export powerful open models, as well as its vocal and persistent advocacy for the Global South in conversations involving international AI governance at the United Nations and in its own multilateral global AI governance initiative,[36] China appears to be making a push to be seen as the primary exporter of AI's benefits to the non-Western world.

This influence extends beyond the models themselves and includes the broader open-model ecosystem's infrastructure. For example, support for Chinese model architectures (e.g., DeepSeek's Multi-Head Latent Attention) within such industry frameworks as vLLM helps ensure that these models become developer defaults.[37] This locks users and developers into using China's architecture and can solidify China's diffusion of the AI ecosystem.

However, although Chinese open models (such as DeepSeek's R1) have fewer licensing restrictions, they might be restricted in other important ways. Researchers have raised concerns that R1 might effectively censor information about topics that are considered politically sensitive in China. Some research indicates that publicly released, fine-tuned versions of DeepSeek's models deflect or refuse to answer questions about certain topics, such as the Tiananmen Square massacre, and this censorship remained even in the models' distilled versions.[38] In September 2025, researchers at cybersecurity firm CrowdStrike suggested that DeepSeek was providing unhelpful answers or computer code with security flaws in response to user queries from groups that the Chinese government opposes;[39] whether intentional or not, the finding provides anecdotal support of the possibility that Chinese models might prove unreliable to China's political opponents.

There are several implications for the United States' and China's soft-power competition. Barring other alternatives, researchers and companies around the globe might seek to use powerful open Chinese models (such as DeepSeek's R1) that have fewer potential hurdles to make new findings or achieve profitability, respectively. The ability to fully manipulate models to meet localized needs might be especially attractive in the developing world, benefitting Chinese model developers' reputations and potentially making it more difficult for U.S. companies to break into emerging markets. For example, it is not dif-

ficult to imagine the good will—not to mention the potential lock-in effects—that could result from a researcher in a country other than the United States or China using a Chinese open model to help solve a national disease outbreak, which could beget broader use of Chinese models compared with U.S. ones.

This kind of outcome is in line with the objectives detailed in the Chinese AI action plan, which are explicitly aimed at cultivating cooperation in AI with the Global South. The tone of the document also contrasts markedly with the U.S. plan, which emphasizes the goal of achieving U.S. global dominance in AI. In addition, the uptake of Chinese open models might propagate censorship or disinformation that aligns with PRC objectives, including the censorship of Chinese history and information about China's troubling human rights practices. Yet the possibility that Chinese models could be considered unsafe or unreliable for use by certain actors and groups adds yet another nuance.

Formulating an Open-Model Strategy for the United States

As the U.S.-Chinese AI competition evolves to incorporate a greater focus on open models, diffusion, and technology adoption, there are several considerations for developing a new U.S. AI strategy. The United States' 2025 AI action plan represents an evolution of the U.S. position on open models and specifically advocates encouraging open-model development because the administration recognizes open-source AI's value for spurring innovation and the broader adoption of AI.[40] Michael Kratsios, director of the White House's Office of Science and Technology Policy, emphasized this point on social media when he stated, "A year ago, policymakers in DC were debating a ban on open weight models . . . from day one, the Trump administration changed course and made clear America needs to offer the best open models for innovators to build on and for America to export."[41]

> To compete effectively, the United States should focus on promoting multiple pathways to AI progress.

These developments represent a potential inflection point for the global open-source AI ecosystem in which the U.S. public and private sectors have recognized the importance of competing with China by using open models. However, whether this truly is an inflection point will depend on how the U.S. government implements policies to support the open AI ecosystem for both academic researchers and commercial developers, as well as how and whether Meta, OpenAI, and their competitors adapt their approaches. As the Trump administration begins to implement its AI action plan, it should focus on the following areas.

Export controls could slow the spread of Chinese models in global AI stacks, while an invigorated U.S. open ecosystem could counter China's strategy.

Promoting Open-Source Artificial Intelligence Development

To compete effectively, the United States should focus on promoting multiple pathways to AI progress, including both open-model ecosystem growth and closed model development. Open models offer distinctive benefits for commercial developers and researchers alike. In the private sector, the U.S. government should encourage leading firms to shift toward open-source models. Some U.S. companies have already been pursuing open-model strategies; Meta has been the leading U.S. open-model player, but others exist, such as the Allen Institute for AI's Tulu models, Microsoft's Phi models, and NVIDIA's GR00T N1 model. OpenAI's release of gpt-oss is the latest and most notable indicator of this emerging trend.[42] As the Trump administration pursues closer ties to AI firms,[43] it should explore incentives it could use to encourage these companies to develop viable alternatives to Chinese models in terms of capacity and openness.

Model performance alone is unlikely to be sufficient in securing global adoption of U.S. open models. Costs, licensing, and levels of model access all play roles. Commercial developers and researchers consider costs, reliability, transparency, and legal requirements.[44] According to a 2025 survey of developers across 41 countries conducted by McKinsey & Company, the Mozilla Foundation, and the Patrick J. McGovern Foundation, 63 percent of respondents prefer open-source AI because of lower costs.[45] Moreover, developers also indicated that their preference was attributable to the full visibility into model weights, which allows them to adapt models to specific use cases. Markets might provide firms with incentives to provide less-restrictive licensing requirements in the push to gain users.

To facilitate the growth of the U.S. open-source model ecosystem for academic research that could be beneficial for the globe, the administration and Congress could also consider strengthening support for the National AI Research Resource,[46] an important pilot program providing computational resources and development infrastructure to under-resourced organizations.[47] Congress can codify and fully fund the program to bolster its compute resources and strengthen the U.S. open-source model ecosystem.[48]

Apart from the academic community, the U.S. government could also encourage open-model adoption by partnering with major cloud service providers, such as Amazon Web Services, Microsoft Azure, or Google Cloud Platform. Cloud service providers already offer compute credits for start-ups.[49] Similarly, the administration could partner with these providers to offer subsidized or free

compute credits that are designated for open-model development and adoption. The U.S. government could also consider differentiating between trusted and nontrusted open models, rewarding developers and adopters who align with certain criteria of openness. Federal funding could be conditioned on meeting licensing, transparency, and reproducibility standards.

Taken together, these efforts could help align adoption incentives with the U.S. government's strategic objectives. In the long term, the efforts could create conditions in which researchers and commercial users alike develop preferences for U.S. models, thus increasing the likelihood that the United States remains globally competitive in open models even amid Chinese companies' flashy new model releases. The potential efforts offer positive incentives for researchers and firms that could pay future dividends in soft-power competition.

Rightsizing Export Controls for the Inference Paradigm

Complementary to positive policy incentives, the United States needs to calibrate its export control strategy to new developments in the field. Put simply, most computing power in advanced AI models is devoted to two processes: training, which is used to teach and tune a model before its deployment, and inference, which enables a model to respond to user needs once deployed. Chinese open-source AI developers, such as DeepSeek, are capitalizing on an emerging paradigm shift in AI computing resources from training to inference capabilities. Traditionally, AI models were optimized to provide immediate responses to user queries. However, the existing trend in AI development has shifted toward test-time inference scaling, which allows the models to allocate additional computational resources during the deployment phase to improve the model's performance and output. A primary method within test-time inference scaling is multistep reasoning, often implemented through chain-of-thought structures in which the model "thinks" before it provides an output to the user's query. This shift has notable implications for the computational demands for both training and deploying state-of-the-art models.

As reasoning-capable models grow in prominence and are used more widely, inference compute becomes more critical both for capability uplift and widespread deployment. For instance, DeepSeek's R1 allegedly used far less compute for training using the company's development techniques,[50] but it relied more heavily on inference to "reason" by decomposing a problem into component parts and effectively showing its work to users. This trend underscores the demand for hardware that is optimized for inference workloads rather than training alone. Although many AI chips can be used in both the training and inference stages of AI development, certain qualities of AI hardware—such as high memory capacity[51]—make some chips more suitable for inference than for training.

Previously, U.S. export control policy regarding AI primarily focused on inhibiting China's ability to train AI models,[52] seemingly imposing prohibitive entry costs to Chinese companies looking to develop frontier models from scratch. However, as demonstrated by R1, Kimi K2, and many other Chinese models that have been released recently (as of February 2026), export controls may not substantially slow China's model development. DeepSeek's success demonstrates that algorithmic optimizations and more-efficient

The U.S. government needs to help present emerging U.S. open models as attractive alternatives to Chinese competitors.

harnessing of the performance of larger models to improve smaller models can reduce export controls' effectiveness on limiting training compute. The AI paradigm shift toward an increasing demand for inference-optimized compute suggests that export controls that target such compute could constrain Chinese companies' abilities to achieve capability uplift by focusing on inference and deploying models at scale for users both inside China and abroad.

If the U.S. government's objectives remain focused on slowing China's ability to develop and deploy frontier AI models in service of U.S. national security and foreign policy interests, export controls could remain a critical policy lever for limiting China's inference capacity in China and abroad as part of an overarching AI strategy. A risk of allowing PRC firms to access AI chips for both training and inference is that Chinese compute infrastructure providers can scale some of their services, including services for users in other countries who hope to access and tailor powerful Chinese open models for their needs. Although we acknowledge that lifting some U.S. export restrictions can provide a commercial benefit for U.S. companies, the combination of right-sized controls and shifting global compute demands could hamper a key element of China's AI diffusion strategy and present a narrow window of opportunity for the United States to expand its AI soft power.

Countering China's Narratives While Protecting U.S. Interests

Promoting open-model development and using calibrated export restrictions might appear to be contradictory policy levers. However, they are both needed to effectively export U.S. AI, protect U.S. advantages at the technological frontier, and compete for AI soft power using open models. Export controls could slow the spread of Chinese models in global AI stacks, while an invigorated U.S. open ecosystem (one that offers alternatives to such models as R1, Kimi K2, and Alibaba's Qwen) could counter China's strategy to establish itself as the main diffuser of AI's benefits. But to avoid being branded globally as the world's policeman of AI diffusion, the U.S. market needs to produce viable alternatives that users perceive to be as open as China's models. By offering performant and accessible open models, U.S. companies give AI developers across the globe reliable alternatives that reduce the risk that they and researchers become entrenched in the Chinese AI technology stack. For commercial developers in particular, when an application is built on one specific model and optimized for that model's unique behavior, it could become burdensome and costly to switch models if the product is not developed with interoperability in mind.

Over the long term, the U.S. government will need to understand the incentive structures that motivate the leading U.S. AI firms to adopt permissive licenses for their open models. OpenAI's recent choice to use a permissive license—potentially driven by DeepSeek's success—stands out among actions that leading labs have taken, and the greater proclivity of Chinese firms to use less restrictive licenses is notable. To some extent, this is a soft-power problem because it involves differing elements of perception and practicality: Llama's licensing restrictions for businesses providing services to more than 700 million monthly active users is a measure that few companies might actually expect to reach, and developers gladly used Llama when it was considered to be the premier open model available to users. But, coupled with Llama's restrictions on certain uses and the improved performance of alternative models, commercial developers might opt to use models in which navigating licensing restrictions is most seamless. The United States might need more large firms to loosen their licensing restrictions so that the firms can compete reputationally with Chinese open-model companies over time.

Beyond licensing, encouraging and incentivizing the further use of U.S. compute infrastructure globally should also play a key role in the U.S. government's strategy. When global developers train and deploy a model—even if the model is a derivative of a Chinese open model—it is beneficial for the United States to encourage the use of U.S. hardware. Overly limiting access to U.S. compute outside China for inference could yet further incentivize Chinese chipmakers' pursuit of leading-edge capabilities and drive global capital toward China to fund the pursuit of alternatives to U.S. hardware.

Soft power in AI is as much about carrots as it is about sticks, and the U.S. government needs to help present emerging U.S. open models as attractive alternatives to Chinese competitors. The United States should embrace the benefits that open models can provide to many countries while highlighting the continuing importance of export controls for national security, especially given concerns about China's development and deployment of AI. Embracing the development of a stronger U.S. open-model ecosystem can provide allies and partners in more- and less-developed countries with attractive alternatives to China's models, foster global collaboration and AI innovation, and promote technological influence abroad while bolstering innovation at home. Combined with attractive U.S. alternatives to Chinese open models, export controls to limit China's compute capacity both inside and outside China could hinder China's ability to develop and diffuse technology that might propagate narratives and disinformation that run counter to U.S. interests.

More broadly, it is worth remembering that technology is but one variable in the more complex formula that makes up a country's soft power. Being seen as a willing diffuser of AI's benefits could certainly benefit the United States' global image, but AI policy will also interact with other political, economic, and security factors in the larger equation, many of which could lessen allies' and partners' willingness to embrace a U.S. vision for AI.

Notes

[1] Daniels and Hanna, "China's Overlooked AI Strategy."

[2] White House, *Winning the Race: America's AI Action Plan*.

[3] White House, *Winning the Race: America's AI Action Plan*, p. 4.

[4] White House, *Winning the Race: America's AI Action Plan*, p. 4.

[5] Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan."

[6] Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan."

[7] Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan."

[8] Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan."

[9] Webster et al., "DigiChina: China's 'New Generation Artificial Intelligence Development Plan' (2017)."

[10] Edelman Trust Institute, *2025 Edelman Trust Barometer*.

[11] OpenAI, "gpt-oss-120b & gpt-oss-20b Model Card"; Miller, "Open Foundation Models: Implications of Contemporary Artificial Intelligence."

[12] Miller, "Open Foundation Models: Implications of Contemporary Artificial Intelligence."

[13] China's State Council, *A Next Generation Artificial Intelligence Development Plan*.

[14] Center for AI Standards and Innovation, *Evaluation of DeepSeek AI Models*.

[15] McGuire, *China's AI Chip Deficit: Why Huawei Can't Catch Nvidia and U.S. Export Controls Should Remain*.

[16] Synergy Research Group, "Cloud Market Share Trends—Big Three Together Hold 63% While Oracle and the Neoclouds Inch Higher."

[17] Hugging Face, homepage.

[18] Miller, Hoffmann, and Gelles, *The Use of Open Models in Research*.

[19] Open Source Initiative, "The Open Source Definition (Annotated)."

[20] Meta, "Llama 4 Community License Agreement."

[21] Google AI for Developers, "Gemma Terms of Use."

[22] Wiggers, "'Open' AI Model Licenses Often Carry Concerning Restrictions."

[23] Soos, "Apache vs. MIT License."

[24] DeepSeek-AI, "MIT License."

[25] Moonshot AI, "Kimi K2."

[26] OpenAI, "Introducing gpt-oss."

[27] Wiggers, "'Open' AI Model Licenses Often Carry Concerning Restrictions."

[28] Wiggers, "Sam Altman: OpenAI Has Been on the 'Wrong Side of History' Concerning Open Source."

[29] Arcesati and Meinhardt, *China's Open-Source Tech Development: Insights into a Growing Ecosystem*.

[30] Lee, "Navigating the US-China 5G Competition."

[31] Beraja et al., "Exporting the Surveillance State via Trade in AI."

[32] Shilov, "China Is Developing Nation-Spanning Network to Sell Surplus Data Center Compute Power—Latency, Disparate Hardware Are Key Hurdles."

[33] Bureau of Industry and Security, Department of Commerce, "Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections."

[34] Pomfret and Pang, "Exclusive: Chinese Researchers Develop AI Model for Military Use on Back of Meta's Llama."

[35] Google and Ipsos, *Our Life with AI: The Reality of Today and the Promise of Tomorrow*.

[36] Daniels, Narayanan, and McFaul, "Analyzing the Global AI Governance Ecosystem."

[37] vLLM, "vLLM Large Scale Serving: DeepSeek @ 2.2k tok/s/H200 with Wide-EP."

[38] Rannaberg, "Deep Dive into Censorship of DeepSeek R1 Based Models."

[39] Menn, "AI Firm DeepSeek Writes Less-Secure Code for Groups China Disfavors."

[40] White House, *Winning the Race: America's AI Action Plan*.

[41] Kratsios, "A year ago, policymakers in DC were debating a ban on open weight models. From day one, the Trump administration changed course and made clear America needs to offer the best open models for innovators to build on and for America to export. Today, our companies are leading."

[42] OpenAI, "Introducing gpt-oss."

[43] Schiffer and Knight, "OpenAI Announces Massive US Government Partnership."

[44] Bisht et al., "Open Source in the Age of AI."

[45] Bisht et al., "Open Source in the Age of AI."

[46] National Artificial Intelligence Research Resource Pilot, homepage.

[47] Miller and Gelles, "The NAIRR Pilot: Estimating Compute."

[48] Hourihan, *Eight Ways Congress Can Help U.S. Researchers Drive AI-Powered Scientific Breakthroughs*.

[49] Google Cloud, "The Next Generation of Startups Are Building on Google Cloud"; AWS Startups, "Apply for Activate Credits."

[50] DeepSeek-AI, *DeepSeek-R1: Incentivizing Reasoning Capability in LLMs Via Reinforcement Learning*.

[51] Somala, "Ban the H20: Competing in the Inference Age."

[52] Bureau of Industry and Security, Department of Commerce," Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification"; Bureau of Industry and Security, Department of Commerce, "Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections."

References

Arcesati, Rebecca, and Caroline Meinhardt, *China's Open-Source Tech Development: Insights into a Growing Ecosystem*, Mercator Institute for China Studies, May 2021.

AWS Startups, "Apply for Activate Credits," webpage, undated. As of November 6, 2025:
https://aws.amazon.com/startups/credits

Beraja, Martin, Andrew Kao, David Y. Yang, and Noam Yuchtman, "Exporting the Surveillance State via Trade in AI," National Bureau of Economic Research, Working Paper 31676, September 2023.

Bisht, Ankit, Roger Roberts, Brittany Presten, and Katherine Ottenbreit, "Open Source in the Age of AI," McKinsey & Company, February 11, 2025.

Bureau of Industry and Security, Department of Commerce, "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification," *Federal Register*, Vol. 87, No. 197, October 13, 2022.

Bureau of Industry and Security, Department of Commerce, "Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections," *Federal Register*, Vol. 88, No. 205, October 25, 2023.

Center for AI Standards and Innovation, *Evaluation of DeepSeek AI Models*, National Institute of Standards and Technology, September 30, 2025.

China's State Council, *A Next Generation Artificial Intelligence Development Plan*, trans. by Rogier Creemers, Graham Webster, Paul Triolo, and Elsa Kania, New America, July 20, 2017.

Daniels, Owen J., and Hanna Dohmen, "China's Overlooked AI Strategy," *Foreign Affairs*, July 25, 2025.

Daniels, Owen J., Mina Narayanan, and Cole McFaul, "Analyzing the Global AI Governance Ecosystem," *2025 IEEE Conference on Artificial Intelligence*, Institute of Electrical and Electronics Engineers, May 2025.

DeepSeek-AI, *DeepSeek-R1: Incentivizing Reasoning Capability in LLMs Via Reinforcement Learning*, arXiv, arXiv:2501.12948v1, January 20, 2025.

DeepSeek-AI, "MIT License," software license, HuggingFace, January 20, 2025. As of November 5, 2025:
https://huggingface.co/deepseek-ai/DeepSeek-R1/blob/main/LICENSE

Edelman Trust Institute, *2025 Edelman Trust Barometer: Trust and the Crisis of Grievance with Insights for the Technology Sector*, undated.

Google, Ipsos, *Our Life with AI: The Reality of Today and the Promise of Tomorrow*, January 2024.

Google AI for Developers, "Gemma Terms of Use," webpage, last updated March 24, 2025. As of November 5, 2025:
https://ai.google.dev/gemma/terms

Google Cloud, "The Next Generation of Startups Are Building on Google Cloud," webpage, undated. As of November 6, 2025:
https://cloud.google.com/startup

Hourihan, Matt, *Eight Ways Congress Can Help U.S. Researchers Drive AI-Powered Scientific Breakthroughs*, Association of American Universities, July 11, 2025.

Hugging Face, homepage, undated. As of November 5, 2025:
https://huggingface.co/

Kratsios, Michael [@mkratsios47], "A year ago, policymakers in DC were debating a ban on open weight models. From day one, the Trump administration changed course and made clear America needs to offer the best open models for innovators to build on and for America to export. Today, our companies are leading," X post, August 5, 2025. As of November 6, 2025:
https://x.com/mkratsios47/status/1952860184721932344

Lee, Nicol Turner, "Navigating the US-China 5G Competition," Brookings Institution, April 2020.

McGuire, Chris, *China's AI Chip Deficit: Why Huawei Can't Catch Nvidia and U.S. Export Controls Should Remain*, Council on Foreign Relations, December 15, 2025.

Menn, Joseph, "AI Firm DeepSeek Writes Less-Secure Code for Groups China Disfavors," *Washington Post*, September 16, 2025.

Meta, "Llama 4 Community License Agreement," webpage, April 5, 2025. As of November 5, 2025:
https://www.llama.com/llama4/license/

Miller, Kyle, "Open Foundation Models: Implications of Contemporary Artificial Intelligence," Center for Security and Emerging Technology, March 12, 2024.

Miller, Kyle, and Rebecca Gelles, "The NAIRR Pilot: Estimating Compute," Center for Security and Emerging Technology, May 8, 2024.

Miller, Kyle, Mia Hoffmann, and Rebecca Gelles, *The Use of Open Models in Research*, Center for Security and Emerging Technology, October 2025.

Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan," July 26, 2025. As of November 5, 2025: https://perma.cc/5PDL-N4WZ

Moonshot AI, "Kimi K2," webpage, July 2025. As of November 5, 2025: https://www.kimi.com/preview/1980c532-5631-8c55-885f-517d380005e7

National Artificial Intelligence Research Resource Pilot, homepage, undated. As of November 6, 2025: https://nairrpilot.org/

OpenAI, "gpt-oss-120b & gpt-oss-20b Model Card," August 5, 2025.

OpenAI, "Introducing gpt-oss," August 5, 2025.

Open Source Initiative, "The Open Source Definition (Annotated)," webpage, last updated February 16, 2024. As of November 5, 2025: https://opensource.org/definition-annotated

Pomfret, James, and Jessie Pang, "Exclusive: Chinese Researchers Develop AI Model for Military Use on Back of Meta's Llama," Reuters, November 1, 2024.

Rannaberg, Carl, "Deep Dive into Censorship of DeepSeek R1 Based Models," *Medium*, March 24, 2025.

Schiffer, Zoë, and Will Knight, "OpenAI Announces Massive US Government Partnership," *Wired*, August 6, 2025.

Shilov, Anton, "China Is Developing Nation-Spanning Network to Sell Surplus Data Center Compute Power—Latency, Disparate Hardware Are Key Hurdles," Tom's Hardware, July 24, 2025.

Somala, Venkat, "Ban the H20: Competing in the Inference Age: How Inference Scaling Should Change US AI Strategy," *ChinaTalk*, March 7, 2025. As of November 6, 2025: https://www.chinatalk.media/p/ban-the-h20-competing-in-the-inference

Soos, "Apache vs. MIT License," February 16, 2022.

Synergy Research Group, "Cloud Market Share Trends—Big Three Together Hold 63% While Oracle and the Neoclouds Inch Higher," November 19, 2025.

vLLM, "vLLM Large Scale Serving: DeepSeek @ 2.2k tok/s/H200 with Wide-EP," blog post, December 17, 2025. As of February 20, 2026: https://blog.vllm.ai/2025/12/17/large-scale-serving.html

Webster, Graham, Rogier Creemers, Elsa Kania, and Paul Triolo, "DigiChina: China's 'New Generation Artificial Intelligence Development Plan' (2017)," webpage, August 1, 2017. As of November 5, 2025: https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/

White House, *Winning the Race: America's AI Action Plan*, July 2025.

Wiggers, Kyle, "'Open' AI Model Licenses Often Carry Concerning Restrictions," *TechCrunch*, March 14, 2025.

Wiggers, Kyle, "Sam Altman: OpenAI Has Been on the 'Wrong Side of History' Concerning Open Source," *TechCrunch*, January 31, 2025.

Acknowledgments

We would like to thank Michael J. Mazarr, Anca Agachi, Helen Toner, Igor Mikolic-Torreira, Kyle Miller, and Mia Hoffmann for their feedback and assistance. We also thank our reviewers, William Marcellino and Jennifer Staats.

About the Authors

Owen J. Daniels is the associate director of analysis and Andrew W. Marshall Fellow at Georgetown University's Center for Security and Emerging Technology and a Nonresident Senior Fellow at the Atlantic Council. His research focuses on the governance of AI and the integration of AI into militaries, among other topics. He holds a B.A. in International Politics with minors in Arabic language and near eastern studies from Princeton University.

Hanna Dohmen is a senior research analyst at Georgetown University's Center for Security and Emerging Technology and a Nonresident Fellow at the Atlantic Council. Her research focuses on U.S. competitiveness in emerging technologies, economic statecraft, and U.S. innovation policy, among other topics. She holds a B.S. in regional and comparative studies of Europe and Asia.

About This Paper

In this paper, we aim to inform the United States' open-source artificial intelligence (AI) strategy and make recommendations to advance U.S. competitiveness in the context of competition with China. We analyze the United States' and China's respective AI action plans, how AI companies in both countries approach open-source AI models, and the implications for U.S. soft power. We conclude by recommending that U.S. policymakers show stronger support for U.S. open-source ecosystems, recalibrate export controls that account for inference demand, and incentivize firms to adopt more permissive licensing practices.

Center for the Geopolitics of AGI

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was commissioned by the division's Center for the Geopolitics of Artificial General Intelligence (AGI), which is committed to helping decisionmakers understand, anticipate, and prepare to navigate the national security and geopolitical implications of AGI. The center convenes leading technologists, strategists, economists, political scientists, and outside experts to consider the feasibility and effectiveness of prospective AGI-enabled capabilities; the domestic and international implications of their use; and the strategies and policies that governments, businesses, and civil society could adopt to respond to new realities. For more information, visit www.rand.org/geopolitics-of-agi.

Funding

This effort was independently initiated and conducted within the Center for the Geopolitics of Artificial General Intelligence using income from operations and gifts from RAND supporters, including philanthropic gifts made or recommended by DALHAP Investments Ltd., Ergo Impact, Founders Pledge, Charlottes och Fredriks Stiftelse, Good Ventures, Longview, and Coefficient Giving. RAND donors and grantors have no influence over research findings or recommendations.

www.rand.org

PE-A4686-1