



# Two Loops

## How China's Open AI Strategy Reinforces Its Industrial Dominance

**Ngor Luong**, *Senior Policy Analyst*

MARCH 23, 2026

### Key Findings

- **China has opted to go all in on an open-source approach to AI.** Most Chinese labs publish model source code and weights. They also charge far less to use high-end products than their global competitors. **This has resulted in the acceleration of global uptake of Chinese AI and created a feedback loop where widespread adoption drives iteration, then further adoption.** As of publication, Alibaba's Qwen models accounted for the largest model ecosystem on Hugging Face, with over 100,000 derivatives.
- This open ecosystem enables China to innovate close to the frontier despite significant compute constraints. Chinese labs have narrowed performance gaps with top Western large language models. **They have also developed key architectural and training advances that are now industry standards.**
- Open model proliferation creates alternative pathways to AI leadership. China's strategy prioritizes data curation and refinement through the deployment of embodied AI in manufacturing, robotics, and research where **specialized, real-world data from widespread use may compound into advantages that proprietary U.S. models cannot easily replicate, even if they maintain technical superiority on benchmarks.**
- **China's open AI model strategy and its manufacturing dominance are mutually reinforcing.** As the Commission's 2025 Annual Report documented, China's industrial base generates "interlocking innovation flywheels" across adjacent sectors. Open models accelerate this dynamic by enabling low-cost AI deployment across factories, factories, logistics networks, and robotics—generating real world

data that feeds back into model improvement. Beijing has built the institutional infrastructure to exploit this advantage, designating data as a formal factor of production and permitting enterprises to carry data assets on their balance sheets.

- **U.S. export controls primarily target the digital loop—restricting access to advanced chips used for frontier model training—but are not well suited to addressing the physical loop of deployment-driven data creation and accumulation across China’s manufacturing base.** As open models reduce the compute required for effective deployment, China’s ability to generate proprietary industrial data at pace and scale becomes increasingly independent of access to cutting-edge hardware. This gap in the U.S. policy framework means that even successful controls on training compute may not prevent China from building AI advantages rooted in its physical economy.

## Introduction

The United States and China are locked in a contest over artificial intelligence (AI) that will shape the global balance of power for decades. In 2025, major U.S. and Chinese frontier AI labs released large language models (LLMs) that set new records for performance or efficiency on a regular basis. U.S.-based OpenAI released its long-anticipated ChatGPT-5, while China-based DeepSeek surprised the world with its efficient, near-frontier R1 model, which *Time Magazine* labeled as one of the “Best Inventions of 2025.”<sup>1</sup> New frontier model releases continue apace in 2026, with both countries now counting a number of organizations competing near the edge of AI capabilities.<sup>2</sup>

Behind this competition lie divergent national strategies. The U.S. AI ecosystem is diverse, but its leading labs have concentrated investment in compute-intensive frontier models, betting that superior hardware resources will yield transformative capabilities. China—constrained to some extent by U.S. export controls on advanced semiconductors but backed by sustained state support—has organized around open development and rapid deployment, integrating AI across its economy under a rubric of “general AI” (GAI).<sup>3</sup> Whether China’s approach reflects genuine strategic preference or adaptation to necessity is an open question, but the result is a fundamentally different theory of how AI leadership is won—not through a single breakthrough model but through ubiquitous adoption and iteration.

These approaches reflect respective national strengths. Bolstered by advantages in compute, top proprietary U.S. models remain a step ahead of Chinese competitors on key benchmarks, though the gap narrowed dramatically in 2025. With some exceptions, China’s frontier labs have openly published model details and offered access to their flagship products at a steep discount to U.S. peers.

China’s open approach has reshaped the competitive landscape. Permissive licensing, aggressive pricing, and an ecosystem that encourages collaboration are accelerating global uptake of Chinese AI and faster iteration among Chinese labs. While top U.S. models maintain a narrow lead in capabilities, they risk losing not only the race to a global user base but also the ability to set the technical standards and norms that will govern AI development for years to come.

## AI Literacy: Key AI Concepts

---

**Artificial General Intelligence (AGI):** The Commission’s 2024 Annual Report to Congress defined AGI as systems that are as good as or better than human capabilities across all cognitive domains and would surpass the sharpest human minds at every task.<sup>4</sup>

**General AI (GAI):** China’s term for AI systems with broad, cross-domain capabilities.\* Unlike AGI’s focus on rivaling human cognition, GAI focuses on developmental goals like increasing manufacturing productivity and prioritizes tailoring models toward industry-specific use cases.

**Open models:** Models that release their weights—the learned parameters determining how models process information. “Open” ranges from fully open-source models (which release code, training data, and weights) to open-weight models (which release only parameters).<sup>†</sup> <sup>5</sup> Most models described as “open” in this paper are only open-weight.<sup>6</sup>

**Closed models:** Proprietary models accessible only through the developer’s interface (e.g., ChatGPT) or an approved Application Programming Interface (API). Developers do not publish the weights of closed models.

**Large Language Models (LLMs):** Models that use deep neural networks to generate text by predicting the next token (a word or meaningful part of text) based on patterns learned from vast datasets.<sup>7</sup> LLMs excel at text generation because of their ability to capture meaningful relationships between distant tokens in a sequence.<sup>8</sup>

**Frontier models:** The most capable models available at a given time, achieving top performance across complex tasks like coding and math and processing multimodal data (e.g., text, images, and audio), often trained with the largest datasets and most compute available.

**Base models:** Pre-trained versions of LLMs or multimodal models before task-specific finetuning, alignment, or safety modifications; often used interchangeably with “foundation model.” Models that are modified from base models are known as derivatives.

**Benchmarks:** Standardized frameworks measuring model capabilities across domains, from general knowledge to scientific problem-solving to autonomous execution of multi-step processes.

\* The word order of the Chinese term 通用人工智能 places the adjective “general” before “artificial,” although it is commonly translated in English as either general artificial intelligence or artificial general intelligence. However, the term as used in Chinese documents refers to broadly capable AI and does not equate the capabilities of such systems to human intelligence (though they may be comparable in some or all domains). William Hannas et al., “China’s Advanced AI Research: Monitoring China’s Paths to ‘General’ Artificial Intelligence,” *Center for Strategic and Emerging Technology*, July 2022, 3. <https://cset.georgetown.edu/publication/chinas-advanced-ai-research/>.

<sup>†</sup> Open-weight models have varying degrees of licensing permission to allow developers to independently understand, audit, or recreate the model.

# U.S. Closed Models: Bigger is Better, Until It Isn't

U.S. frontier labs operate on the premise that bigger models trained on more data yield better performance—a belief grounded in “scaling laws,” the empirical finding that model performance improves predictably with increases in parameters, training compute, and dataset size.<sup>\* 9</sup> This logic reinforced a capital-intensive, proprietary approach to AI development.

Applying scaling laws requires complex and expensive training processes. Multi-million-dollar training runs are engineered to extract as much learning as possible from each unit of compute.<sup>10</sup> At the same time, running thousands of graphics processing units (GPUs)<sup>†</sup> in parallel for extended periods of time inevitably produces hardware failures, so training pipelines must incorporate enough redundancy that localized disruptions do not derail millions of dollars' worth of training compute.<sup>11</sup>

Data acquisition, curation, and refinement introduce another layer of competition.<sup>12</sup> Beyond publicly available data scraped from the internet, leading U.S. frontier labs rely on proprietary datasets, including enterprise data streams like search results and social media feeds as well as specialized collections like scientific publications.<sup>13</sup> Significant engineering expertise also goes into converting sources like microfilm and physical archives into usable training data. The frontier labs also buy labeled, curated, and refined data from third-party providers.<sup>14</sup>

Because competitive advantage under the scaling paradigm depends not just on model size but also on training knowhow, infrastructure design, and data composition, AI firms have strong incentives to guard these as trade secrets.<sup>15</sup> Maintaining proprietary weights helps protect these advantages and has reinforced the closed ecosystem among U.S. frontier labs.

## The Scaling Bet Hits Diminishing Returns

The differential between U.S. and Chinese investment in data center buildout reflects U.S. prioritization of scaling (see Table 1). For 2025, U.S. AI capital expenditure by Microsoft, Amazon, Meta, and Google was at least \$350 billion, according to Bloomberg Intelligence estimates, compared to less than \$40 billion for China's major cloud providers.<sup>‡</sup> <sup>16</sup> For 2026, Bloomberg Intelligence estimates the same four tech giants will spend more than \$400 billion, while China's spending will stay flat.<sup>17</sup>

\* A notable exception is Meta, which released the weights of its Llama family of models with some restrictions. Reports from mid to late 2025 suggested Meta may keep a new flagship model codenamed “Avocado” proprietary. The model is slated for early 2026. “Meta’s ‘Llama’ Coming to an End? New Large Model Code-Named Avocado Announced for Q1 2026 or Switching to Closed Source to Directly Compete with OpenAI,” *AI Base*, December 10, 2025. <https://news.aibase.com/news/23528>.

† GPUs are a type of semiconductor device widely used for training and running AI models. Originally developed to display video game graphics, GPUs are highly efficient in completing matrix operations used by deep neural networks.

‡ These include Alibaba, Tencent, Baidu, SenseTime, and Kingsoft Cloud. A June 2025 estimate from Bank of America placed China's annual capex higher, at \$98 billion. Vlad Savov, “China's AI Future Looks Like a Long Road to Small Profit,” *Bloomberg*, December 3, 2025. <https://www.bloomberg.com/news/newsletters/2025-12-03/china-s-ai-future-looks-like-a-long-road-to-small-profit>; Xinmei Shen, “China's AI Capital Spending Set to Reach up to US\$98 Billion in 2025 amid Rivalry with US,” *South China Morning Post*, June 25, 2025. <https://www.scmp.com/tech/tech-war/article/3315805/chinas-ai-capital-spending-set-reach-us98-billion-2025-amid-rivalry-us>.

**Table 1: Top U.S. and Chinese Models by Training Compute, Training Data, and Parameters, 2025**

Elements	Summary of Advantage	United States	China
<b>Training compute</b>	U.S. leads in frontier training compute capacity and scale, while China faces compute constraints.	xAI’s Grok 4 (5 x 10 <sup>26</sup> FLOPs)	Alibaba’s Qwen3-Max (1.5 x 10 <sup>25</sup> FLOPs)
<b>Training data size</b>	U.S. maintains a lead in large, high-quality training data, although China is rapidly closing the gap.	Nvidia’s Cosmos-1.0-Diffusion-14B Video2World (9,000 trillion tokens)*	Alibaba’s Qwen3-Max (36 trillion tokens)
<b>Model size</b>	U.S. continues to develop large-scale models, while China is catching up.	xAI’s Grok 4 (estimated 1.7–3 trillion parameters)	Alibaba’s Qwen3-Max (1 trillion parameters)

Note: Based on publicly available data in 2025 only.  
Source: Various.<sup>18</sup>

While U.S. firms are currently poised to maintain an advantage in compute, several factors are challenging strategies centered on scaling. First, successive generations of models require ever-more compute to yield incremental gains in performance, leading to diminishing returns with increased model size.<sup>19</sup> Second, returns from scaling models are not uniform across all tasks because improvements in prompting and task-specific optimization mean that small models can perform comparably to large models on some tasks.<sup>20</sup> Third, 2025 brought advances that delivered performance gains without relying on increased model size.<sup>21</sup> Novel architectures improved both training and inference efficiency, while improvements in post-training techniques and ability to handle long contexts enabled labs to strengthen performance with fewer parameters.<sup>22</sup> Reflecting these changes, OpenAI described its February 2025 GPT-4.5 as its largest model to date and its last to rely primarily on scaling pre-training rather than reasoning.<sup>23</sup>

## China’s AI Reaches the Frontier from Below

The release of DeepSeek’s R1 in January 2025 signaled that Chinese labs could approach frontier performance despite U.S. export controls on advanced semiconductors. R1 was followed by a succession of open models from Chinese labs that similarly did more with less, attaining scores nearing proprietary U.S. frontrunners on smaller budgets and offering competitive products at a fraction of the cost of top U.S. models (see Table 2). Updates to Alibaba’s Qwen family became widely used internationally. Beijing-based Moonshot AI set new bars for agentic workflows and context length—the amount of tokens a model can handle—though they were later overtaken by U.S. models.<sup>24</sup> DeepSeek went on to attain gold medal performance on questions from the 2025 International Math Olympiad (IMO) alongside OpenAI and Google DeepMind; DeepSeek-V3.2-Special, released in December 2025, was the first open model capable of IMO gold-level scores as of publication.<sup>25</sup>

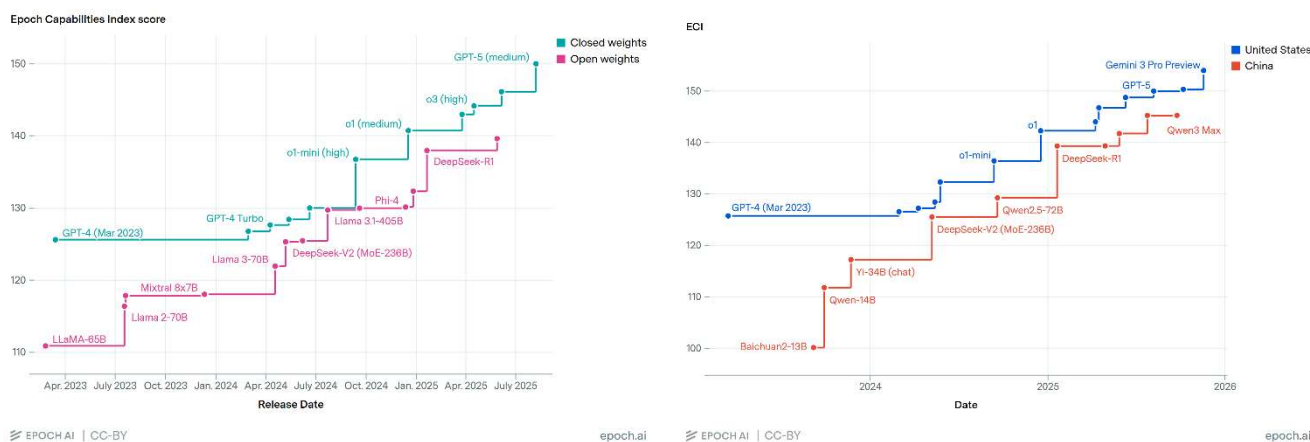
While U.S. closed models remain slightly ahead of Chinese open counterparts by most measures (see Figure 1), China’s successes in 2025 reflect a more fundamental challenge to U.S. AI supremacy than a narrowing lead: Chinese labs now play a central role in developing and refining innovations that drive frontier progress. Though DeepSeek

\* Nvidia’s model is a unique approach that handles high volumes of visual data to develop a model of the physical world. Its training data volume is an outlier even among other models trained on record-setting datasets.

R1 did not introduce the mixture-of-experts (MOE) architecture—an approach to improve efficiency by activating only the most relevant parameters—it popularized the use of large-scale MOE architectures in open frontier models.<sup>26</sup> It similarly helped boost the popularity of chain-of-thought reasoning in which models reason through intermediate steps before producing a final answer, improving accuracy on complex tasks. Both MOE and chain-of-thought are now standard for leading frontier models.<sup>27</sup>

The series of innovative solutions from China, along with the growing list of Chinese labs contributing near the frontier of AI capabilities, reflects not only a vibrant research and development (R&D) ecosystem but also necessity, as compute constraints drive experimentation with alternatives to scaling. Both are underpinned by sustained state support for open model development, promoting widespread adoption and reducing reliance on foreign technology.

**Figure 1: Chinese Labs and Open Models Close the Gap in Frontier Trends, 2023–2026**



Note: The “Epoch Capabilities Index” is a composite score on 39 different benchmarks.  
Source: Epoch AI.<sup>28</sup>

## Beijing’s Blueprint: State-Directed Open AI

Beijing views the diffusion of AI as a strategic enabler of multiple long-term goals, from boosting productivity as China confronts slowing growth and demographic decline to accelerating advances in other fields like synthetic biology. To support these objectives, Chinese policymakers have articulated a clear vision for how AI should evolve domestically: widespread societal adoption and deep economic integration spurred by an open R&D ecosystem that minimizes dependence on foreign technology and innovates within bounds set by the Party.<sup>29</sup> These pillars trace roots back to at least China’s sweeping 2017 AI Development Plan that established a broad set of objectives for AI development and encouraged AI integration across manufacturing, healthcare, transportation, and public administration as well as in military applications.<sup>30</sup> Similar objectives animate Beijing’s AI policy today, with the State Council’s August 2025 “AI+ Initiative” formalizing widespread adoption as the central organizing principle for China’s AI development. Among other goals, the initiative sets targets for AI integration into science and technology, industry, public services, governance, and international cooperation.<sup>31</sup>

Table 2: Top Chinese Labs and Open Models (Partial List)

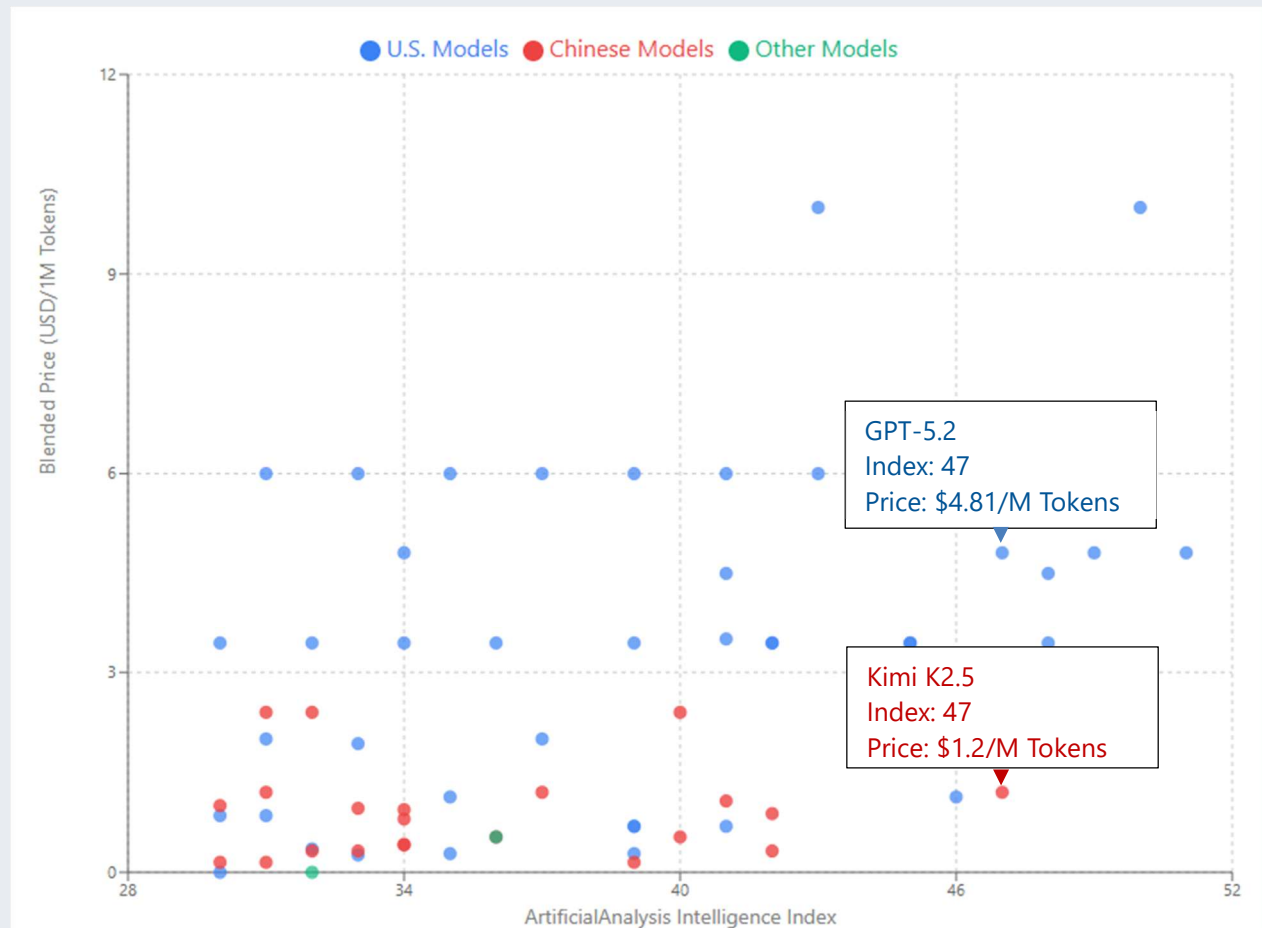
Lab	Model Family	Why the Model(s) Stands Out	Notes
<b>DeepSeek</b>	V3, R1, R2, R2 Speciale	<ul style="list-style-type: none"> <li>Focus on math and reasoning performance: High self-reported benchmark performance on competition-style math exams and complex reasoning tasks.</li> <li>Reasoning-first design: Trained to prioritize problem-solving, with outputs that expose chains-of-thought for transparency and distillation.</li> <li>Distillation scalability: R1's reasoning behavior was distilled into smaller, dense models (1.5B–70B parameters) while being cheaper to run.</li> </ul>	R1 was built on Qwen 2.5 and Llama 3 base architectures.
<b>Alibaba</b>	Tongyi Qianwen, Qwen 2, Qwen 2.5, Qwen3-Omni	<ul style="list-style-type: none"> <li>Focus on multilingual and multimodal performance: High self-reported scores, particularly in general knowledge tests, multilingual support (119 languages), and many audiovisual tasks.</li> <li>Cost-efficient architecture: Its MOE models achieve high performance while emphasizing efficient inference (e.g., the flagship model has a total of 235B parameters but only 22B activated ones per token).</li> </ul>	The Qwen family is the most downloaded on Hugging Face (see Figure 5).
<b>Moonshot AI</b>	Moonlight-A3B/Moonlight-16B-A3B, Kimi K2, K2 Thinking, K2.5 Thinking	<ul style="list-style-type: none"> <li>Focus on coding and agentic performance: High self-reported results in software engineering and tool-use tasks while also showing robust math reasoning.</li> <li>Cost-efficient architecture: Stabilized 1T-parameter training while activating only 32B parameters per token, aiming for fast, low-cost inference.</li> <li>Agentic capabilities: Strengthened for planning and tool use, making the model well suited for software engineering and complex problem solving.</li> <li>Dual-reasoning modes: Offers fast, cost-efficient outputs by skipping explicit reasoning traces when they are unnecessary.</li> </ul>	Kimi K2.5 offers the lowest cost for the highest performance capability (see Figure 2).
<b>MiniMax</b>	MiniMax M-1, M.2, M-2.1, M-2.5	<ul style="list-style-type: none"> <li>Focus on coding and agentic capabilities: Executes complex, long-horizon toolchains. Workflows include multi-file edits, compile-run-fix loops, tool calling, browsing, and long-horizon task execution.</li> <li>Cost-efficient architecture: Its models deliver lower latency, lower cost, and higher throughput for interactive agents and batched sampling while activating only ten billion activated parameters (230 billion in total).</li> <li>Dual-reasoning modes: The MiniMax family delivers fast, low-cost outputs that compete with larger closed models.</li> </ul>	Minimax went public on the Hong Kong Stock Exchange and raised \$619 million.
<b>Z.ai</b>	GLM-4, CodeGeeX4, GLM-Edge, GLM-4.6	<ul style="list-style-type: none"> <li>Focus on generalist performance: High self-reported scores across reasoning, coding, agentic, and visual benchmarks.</li> <li>Reliable agentic capabilities: Trained with synthetic tool-use and long-context data, aiming for more dependable function-calling and multi-step planning.</li> <li>Multi-stage expert training: Trained specialized expert models for reasoning, coding, and agentic tasks, then unified them through self-distillation to create a balanced generalist model.</li> </ul>	A spinoff of Tsinghua University, Z.ai was added to Department of Commerce's Entity List on January 16, 2025.

Source: Various.<sup>32</sup>

### Chinese Models Cost a Fraction of U.S. Prices

Chinese AI labs offer their models at increasingly competitive prices while narrowing the performance gap with U.S. counterparts. For example, Kimi K2.5 costs four times less than OpenAI’s GPT5.2 as of January 2026.\*<sup>33</sup> Both have an Intelligence Index of 47 on Artificial Analysis’s benchmark aggregator (see Figure 2). This aggressive pricing in part responds to a market reality in China: consumers are less willing to pay for software products.<sup>34</sup> To overcome this barrier, Beijing is subsidizing user access to existing models through APIs and the purchase of pre-trained model licenses outright.<sup>35</sup>

**Figure 2: U.S. vs. Chinese Inference Costs Relative to Capability**



Note: The models analyzed were released before February 1, 2026. The graph excludes OpenAI’s o1, which costs \$26.25 per million tokens but scores comparatively low, at 31 on the Intelligence Index.

Source: Artificial Analysis.<sup>36</sup>

\* Blended price on Artificial Analysis is a combination of output and input price token prices at a 3:1 ratio.

As with other sectors targeted by industrial policies, China's central and local governments promote AI development with a suite of supply-side and demand-side mechanisms. On the supply side, upstream measures include extensive support for energy infrastructure, such as subsidized electricity costs for data centers often powered by domestically made chips. For AI training and inference, Gansu, Guizhou, and Inner Mongolia have offered to slash cloud providers' power bills by as much as 50 percent.\*<sup>37</sup> Chinese companies also optimize the location of AI data centers in relation to power grids to allocate available power to these facilities and capitalize on cheap electricity.<sup>38</sup> On the demand side, the Chinese government is reducing barriers to AI adoption by subsidizing usage costs for already low-priced models.<sup>39</sup> Other policies to strengthen talent pipelines, fund startups, and increase Chinese labs' access to data are also contributing to China's AI development.

### From Vision to Reality: Building an Open Ecosystem

An open ecosystem is central to Beijing's policy design for competitive AI. As a latecomer facing compute constraints, China's government sought an approach that would maximize AI adoption while minimizing dependence on proprietary foreign technology.<sup>40</sup> Open models facilitate widespread use by lowering barriers for enterprises or developers to customize and deploy them without licensing fees or limits on adaptation. Support for open models also stems from China's broader, longstanding strategy to promote open-source technologies across its tech ecosystem—from operating systems to development platforms—to boost domestic innovation and reduce reliance on foreign proprietary software. (For more on China's open technology ecosystem, see Appendix I: "China's 'Open' Technology Long Game to Become the Industry Standard.")

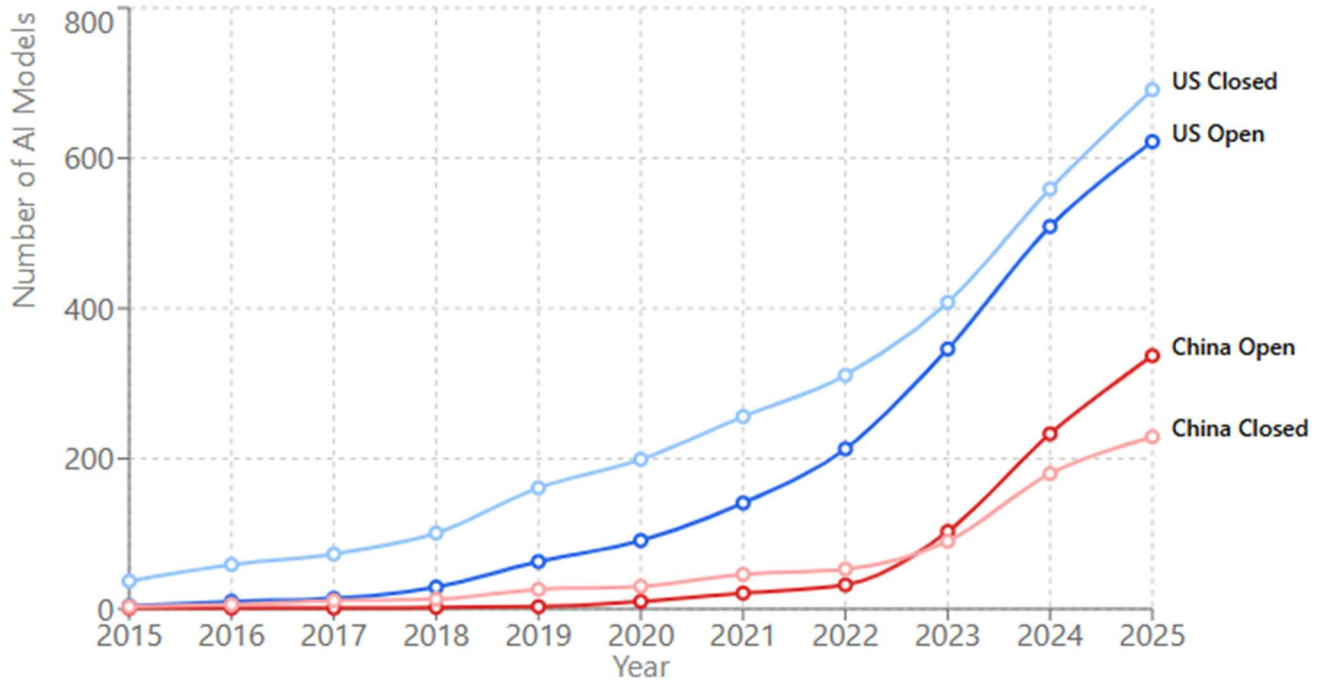
Recognizing these benefits, Beijing's 2017 AI Plan also sought to encourage an open AI ecosystem by fostering linkages between industry and academia and by promoting open-source platforms. Inspired by the success of Meta (then Facebook) and Google's open-source AI libraries, Chinese firms launched similar projects, such as Baidu's Apollo for developing autonomous vehicles.<sup>41</sup> This emphasis on open development carried forward in the 2021 14th Five-Year Plan, the 2025 AI+ Initiative mentioned above, as well as the draft 15th Five-Year Plan.<sup>42</sup> At the local level, governments in Beijing, Guangdong, and Hangzhou have all invested millions in developing open-source AI.<sup>43</sup>

By sheer count, state funding and policy support correspond with China's rapid introduction of open models. Between 2022—the year ChatGPT was released—and 2025, the number of Chinese open models expanded more than tenfold, from 32 to 337, according to data compiled by nonprofit Epoch AI (see Figure 3). The United States started from a much higher base, but its open models increased less than threefold in the same period, from 213 to 622.<sup>44</sup> In 2023, the number of Chinese open models surpassed the number of closed Chinese models; by contrast, there have always been more closed U.S. models than open U.S. models. Chinese government sources place the number of Chinese open models much higher than Epoch AI: in a July 2025 briefing, China's State Council Information Office claimed China had overtaken the United States in total number of open releases, at 1,509.<sup>†45</sup>

\* It is important to note that despite the government support, data center utilization in these regions was just 20–30 percent. Gigi Onag, "China to Set Up Cloud Service Selling Spare Data Center Capacity – Report," *Light Reading*, July 24, 2025. <https://www.lightreading.com/data-centers/china-to-set-up-cloud-service-selling-spare-data-center-capacity---report>.

† When comparing AI model counts across countries, methodology matters significantly: some trackers count finetuned variants and derivatives as distinct models, which can dramatically inflate totals for countries like China whose developers have been prolific finetuners of open Western base models, while counts limited to original foundation models and publicly known models tell a different story.

**Figure 3: Cumulative Count of U.S. and Chinese AI Models by Model Accessibility (2015-2025)**



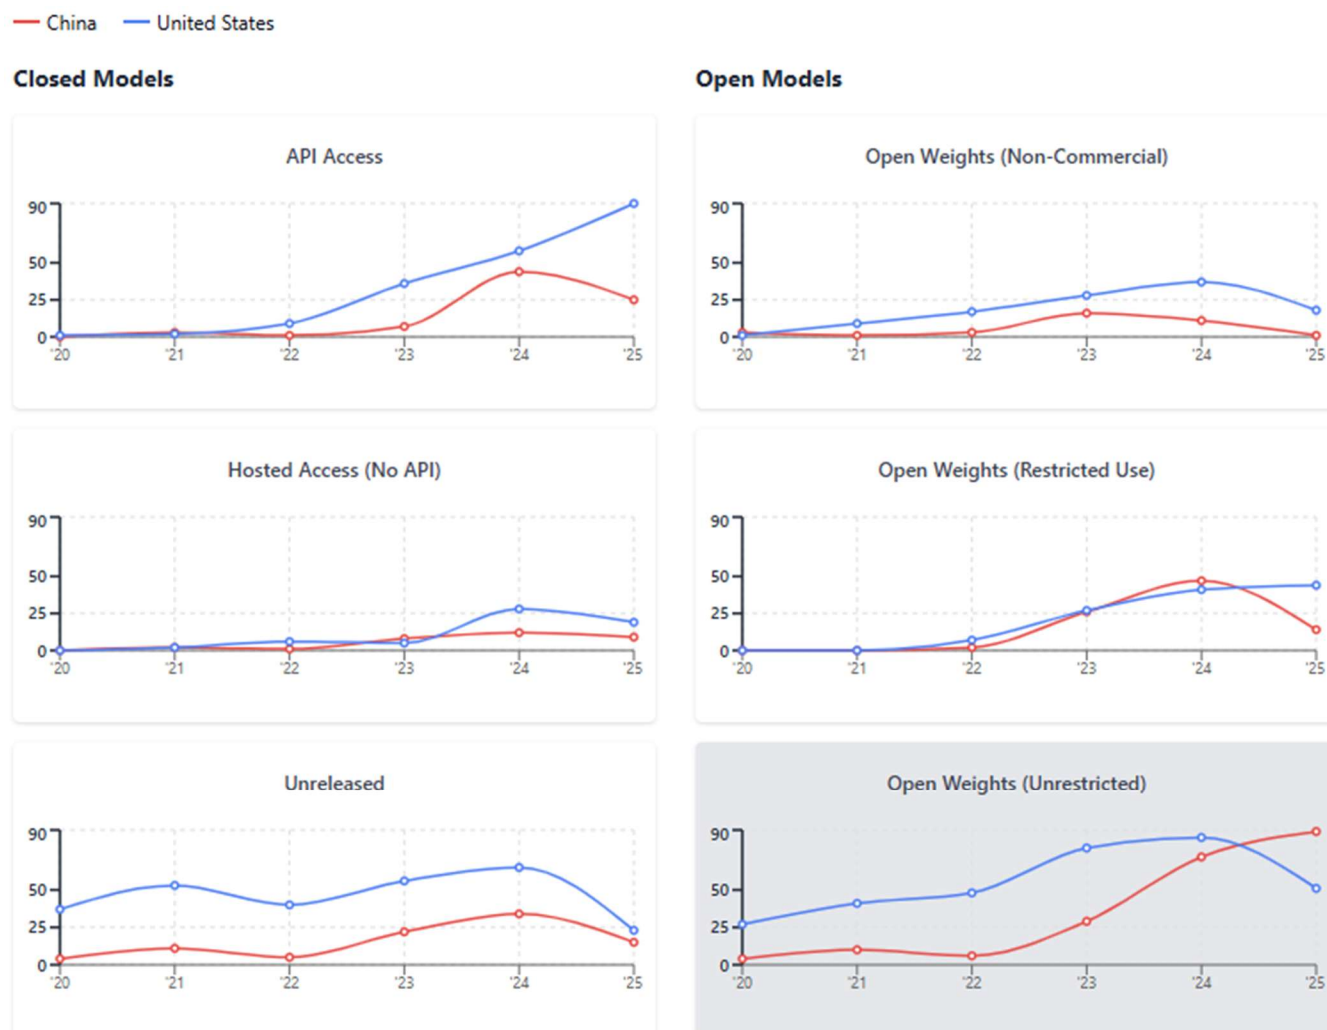
Note: The analysis assigns a country to each model if it has at least one organization based in the United States or China, and it excludes 101 models that are developed by both U.S. and Chinese organizations. Cumulative count begins in 1988 but only displays 2015–2025 and includes a handful of models other than LLMs or multimodal models.

Source: Commission staff analysis based on Epoch AI’s “AI Models” database, accessed January 23, 2026. <https://epoch.ai/data/ai-models>.

### China’s Open Ecosystem Accelerates Adoption and Innovation

More important than the number of releases, Chinese labs are accelerating adoption and iteration of their models far faster than U.S. labs, creating shared technical commons that lift the overall capabilities of China’s AI ecosystem. Though Chinese labs released fewer total open-weight models, they have recently made models more accessible than U.S. counterparts through permissive licensing, deploying a greater number of unrestricted models in 2025 (see Figure 4).<sup>46</sup> This strategy leverages what analysts describe as a diffusion-innovation feedback loop created by rapid rollout and experimentation: each model improvement becomes a foundation for subsequent innovations rather than remaining siloed in a single company. As Wang Zhe of Beihang University and Xue Lan of Tsinghua University noted, the spread of China’s open-source models enables exploration of diverse technological paths, easier capability sharing, and broader distribution of development costs.<sup>47</sup> These shared technical commons are the first of two feedback loops that compound China’s advantages—the second runs through the physical economy.

Figure 4: China vs. U.S. Models by Degree of Openness, 2020-2025



Note: The analysis assigns a country to each model if it has at least one organization based in the United States or China, and it excludes 101 models that are developed by both U.S. and Chinese organizations.

Source: Commission staff analysis based on Epoch AI’s “AI Models” database, accessed January 23, 2026. <https://epoch.ai/data/ai-models>.

## The First Loop

Data on open model downloads and iteration show this feedback loop in action. According to the American Truly Open Models (ATOM) project, in August 2025 Chinese models overtook U.S. models in total downloads on Hugging Face, a global open model collaboration platform.<sup>48</sup> By October 2025, Chinese models by Z.ai (formerly Zhipu AI), Alibaba, DeepSeek, Moonshot AI, and tech giant Meituan, best known as a food delivery platform, had displaced Meta’s models in popularity on LMArena (recently rebranded Arena), a leaderboard that ranks models based on human preference.<sup>49</sup> From November to December 2025, Chinese models accounted for seven of the top ten most downloaded models on Hugging Face. Developers uploaded derivative models back to Hugging Face at nearly twice the rate of U.S. models (see Figure 5).<sup>50</sup> By year’s end, derivatives of Alibaba’s Qwen family had become the largest

model ecosystem on Hugging Face, with over 100,000 derivatives—more than any Western counterpart, including Meta’s Llama.<sup>51</sup>

The iterative collaboration in China’s open ecosystem—the first of the two feedback loops driving China’s AI strategy—provides three structural advantages for accelerating AI advances and adoption:

- 1. Frontier labs can refine each other’s base models**, driving rapid capability improvements without costly investments in pre-training. For example, DeepSeek claims its Qwen derivative, R1-Distill-Qwen-32B, can outperform OpenAI’s o1-mini model.<sup>52</sup>
- 2. Firms can adapt high-performing base models to niche use cases.** In late 2025, the most downloaded model on Hugging Face was a small, specialized model for video captioning created by TikTok parent ByteDance. The short video giant had finetuned its Tarsier2-Recap-7b from Alibaba’s Qwen2-VL-7B-Instruct, which itself ranked fifth in downloads during that period (see Figure 5).<sup>53</sup> By providing ready-made foundations, open models allow developers to focus resources on specialized adaptations rather than recreating fundamental AI capabilities, particularly benefiting small and medium-sized enterprises lacking resources for proprietary alternatives.<sup>54</sup>
- 3. New entrants quickly gain a profile** through popular and widely downloaded open models, helping to attract talent, raise capital, and draw customers. Notable examples include Moonshot AI and Minimax, a Shanghai-based lab founded by former engineers from AI surveillance firm SenseTime in 2021. Aside from its widely downloaded eponymous foundational models, the startup’s Hailuo models are among the best free video generation models, rivaling OpenAI’s Sora in simulating physical environments.<sup>55</sup> In January 2026, Minimax was listed on the Hong Kong Stock Exchange, with its share price doubled on the first day of trading, underscoring investor confidence in the commercial potential of China’s open ecosystem.<sup>56</sup>

This growing dominance has even led to widespread adoption by U.S. companies.<sup>57</sup> One partner at U.S.-based venture capital firm Andreessen Horowitz provided a rough estimate that 80 percent of U.S. startups use Chinese base models to develop derivatives for their business.<sup>58</sup> Even larger companies like Airbnb use Alibaba’s Qwen for customer service chatbots due to its low cost and capabilities comparable with leading U.S. models.<sup>59</sup> However, this adoption raises security concerns: the U.S. National Institute of Standards and Technology (NIST) found that DeepSeek’s open models are more susceptible to cyber risks than comparable U.S. models.<sup>60</sup> Additionally, Chinese models may censor topics to comply with political pressures and domestic regulations, and using models hosted by Chinese entities could present data privacy risks.<sup>61</sup>

























The diffusion-iteration loop is not confined to software. As the Commission’s 2025 Annual Report documented, China’s state-directed industrial base generates “interlocking innovation flywheels” in which advances in one sector rapidly catalyze breakthroughs in adjacent industries, and production data feed back into R&D.\* Open AI models add a new dimension to this dynamic. By lowering the barrier for manufacturers, logistics firms, and robotics companies to deploy AI in operational settings, open models enable each deployment to generate proprietary, real-world data at scale that no amount of web scraping or synthetic generation can replicate. Beijing has long recognized data’s strategic importance: General Secretary of the Chinese Communist Party (CCP) Xi Jinping described data as “a new production factor,” a foundational and strategic resource as early as 2017, and the CCP Central Committee formally designated data as the fifth factor of production—alongside land, labor, capital, and

\* For more, see U.S.-China Economic and Security Review Commission, Chapter 6, “Interlocking Innovation Flywheels: China’s Manufacturing and Innovation Engine,” in *2025 Annual Report to Congress*, November 2025, 313–369.

technology—in 2020, followed by the “20 Data Measures” in 2022 and the creation of the National Data Administration (NDA) in 2023 to operationalize this vision.<sup>62</sup>

China’s advantages in translating this vision into practice are structural: an unmatched manufacturing base, extensive Internet of Things (IoT) and 5G infrastructure, and a policy apparatus that actively subsidizes AI adoption across the physical economy. The result is a second feedback loop, running in parallel with the first, in which widespread industrial deployment feeds data back into model improvement, which in turn enables more sophisticated deployment. It is the intersection of these two loops—one digital, one physical—that gives China’s open strategy its compounding force and poses the most serious long-term challenge to U.S. AI leadership.

**Figure 5: Top Ten Most Downloaded Models on Hugging Face (November–December 2025)**

Company	Model	Downloaded Nov 8 – Dec 8, 2025	Used as Base Model
<b>Chinese models</b>	<b>Sum of 7 models</b>	 51,064,545	 7,472
<b>U.S. models</b>	<b>Sum of 3 models</b>	 17,830,929	 4,808
Bytedance (China)	omni-research/Tarsier2-Recap-7b	 10,847,133	 –
Alibaba (China)	Qwen2.5-3B-Instruct	 8,780,419	 1,674
OpenAI (U.S.)	openai/gpt-oss-20b	 8,089,782	 672
Alibaba (China)	Qwen2.5-VL-3B-Instruct	 7,688,461	 699
Alibaba (China)	Qwen2.5-7B-Instruct	 7,245,333	 3,335
Alibaba (China)	Qwen3-4B-Instruct-2507	 6,289,445	 400
DeepSeek (China)	DeepSeek-OCR	 5,451,968	 125
Meta (U.S.)	Llama-3.1-8B-Instruct	 5,187,643	 3,975
Alibaba (China)	Qwen3-8B	 4,761,786	 1,239
OpenAI (U.S.)	gpt-oss-120b	 4,553,504	 161

Note: \* Data range from November 8 to December 8, 2025. Large-scale means more than three billion parameters. The number of models that use respective model as base model count those that are uploaded back to Hugging Face, which include finetunes, adapters (i.e., finetuning with extra trainable parameters), secondary finetunes, merges, and quantization (a compressing technique to reduce memory use and computing cost).

Source: Commission staff analysis based on Hugging Face, accessed December 8, 2025.

### The Second Loop: Deployment Generates Data

As internet-scale data become finite, the ability to gather high-quality, application-specific data through real-world deployment could become a critical source of AI advantage for China. Epoch AI estimates that leading U.S. AI companies may exhaust high-quality publicly available training data for AI language models between 2026 and 2032, reshaping the competitive landscape as the next frontier shifts toward companies’ proprietary data for specific use cases.<sup>63</sup>

Chinese AI scholars view data as a competitive opportunity. Weinan E, an academician at the Chinese Academy of Sciences and a machine learning scholar, argued in September 2025 that future progress will depend less on building larger models and more on securing high-quality, diverse data.<sup>64</sup> As a case in point, China’s connected factories and IoT infrastructure position the country to collect specialized industrial data at scale.<sup>65</sup> Data from 5G HD cameras in an intelligent factory in Guangdong, for instance, have enabled AI to improve quality inspection, reducing equipment repair rates by 20 percent and saving more than 1 million renminbi (RMB) (\$140,000) annually in costs.<sup>66</sup> This infrastructure creates millions of data collection points across factories, logistics networks, and smart cities. Leading Chinese robotics companies, such as AgiBot and Fourier, also have released open training datasets, expanding the pool available for development of embodied AI—systems that interact with the physical world.<sup>67</sup> This is the physical-economy feedback loop in practice.

However, Chinese labs also face significant data constraints. China’s data security law imposes strict rules on personal data, limiting data sharing between China’s super-apps like WeChat and Alipay and other developers.<sup>68</sup> Chinese webpages are also constantly being deleted by censors, while foreign data providers may block Chinese IP addresses.<sup>69</sup> In April 2025, Tang Jie, cofounder of Z.ai, wrote in the *People’s Daily* that China’s high-quality data resources remain “fragmented and scattered” and are now a leading bottleneck to improving domestic models, and he emphasized Xi Jinping’s call to integrate AI and big data with the real economy.<sup>70</sup> Recent policies to bolster China’s data labeling industry suggest the supply of high-quality annotated data remains an additional constraint.<sup>71</sup>

To address these gaps, Beijing constructed a layered institutional framework in March 2025: a National Public Data Resource Registration Platform under the ownership of the NDA.<sup>72</sup> The NDA reports that Chinese language data account for 60–80 percent of training data for most domestic Chinese models.<sup>73</sup> To address concerns over data protection laws, the Chinese government is also promoting open data-sharing platforms and funding large open datasets, such as Beijing Academy of Artificial Intelligence’s multilingual, multimodal FlagData.<sup>74</sup> To further address these concerns, the Chinese government is offering additional incentives. As the Commission noted in its 2024 Annual Report, China’s Ministry of Finance issued accounting standards in 2023, permitting enterprises to recognize qualifying data resources on their balance sheets as intangible assets or inventory—making China the first country to establish national data asset accounting standards.<sup>75</sup> Taken together, these measures reflect a systematic effort to build the institutional plumbing that converts China’s deployment advantage into a durable data asset.

## Where the Loops Converge: Small Models and China’s Structural Advantage

Two feedback loops—one digital, one physical—give China’s AI strategy its compounding force. The mechanism linking them is the emerging shift toward small, specialized AI models for real-world deployment.

The dominant narrative in U.S. AI strategy centers on frontier scale—ever-larger models trained on ever-more compute. But a growing body of evidence suggests that the models most consequential for economic and industrial application are not frontier LLMs but small language models (SLMs) finetuned for specific operational tasks. In a recent position paper, Nvidia researchers argued that small models “handle the bulk of operational subtasks” in agentic AI systems, delivering faster responses, fewer errors, and costs ten to 30 times lower than frontier alternatives.<sup>76</sup> The paper described a “heterogeneous” future in which large models are reserved for occasional complex reasoning, while small, specialized models power the majority of deployed AI applications.<sup>77</sup>

This trajectory has significant implications for U.S.-China competition because small models are precisely the kind of AI that open ecosystems produce most efficiently—and China dominates the global open ecosystem.

The economics of small model development differ fundamentally from frontier training. Finetuning a small model for a specific industrial task requires a fraction of the compute needed to train a frontier LLM from scratch, making it accessible to a far wider range of firms. Open base models lower the barrier further: rather than building from zero, developers can adapt a high-performing foundation to their use case in days rather than months.

China's open ecosystem is already producing this dynamic at scale. As previously discussed, the most downloaded model on Hugging Face in late 2025 was not a frontier LLM but a small, specialized video captioning model—ByteDance's Tarsier2-Recap-7b, finetuned from Alibaba's Qwen2-VL-7B-Instruct.<sup>78</sup> Derivatives of Alibaba's Qwen family alone account for over 100,000 models on Hugging Face, many of them small, task-specific adaptations created by third-party developers.<sup>79</sup> This pattern—open base model, community-driven specialization, rapid iteration—is optimized for producing exactly the kind of AI the deployment shift demands.

Small models are the mechanism through which China's digital and physical feedback loops reinforce each other. The digital loop—open model diffusion and community iteration—generates an expanding library of capable, adaptable base models. The physical loop—deployment across China's manufacturing base, logistics networks, and robotics sector—provides the real-world environments where these models are put to work.

The connection runs in both directions. A quality inspection model deployed in a Guangdong factory does not require frontier-scale compute—it requires a small vision model finetuned on production-line data and running on edge hardware. That deployment generates proprietary operational data that feeds back into model improvement. The improved model enables more sophisticated deployment, generating richer data. Each cycle is faster and cheaper than frontier training and—critically—is not constrained by access to the most advanced semiconductors.

China's structural advantages compound at every stage of this process. Its open ecosystem provides the base models. Its manufacturing scale provides the deployment surface. Its IoT and 5G infrastructure provides the data collection architecture. And Beijing's policy apparatus—from the NDA to local AI adoption subsidies—actively lubricates the cycle. No other country possesses this combination at comparable scale.

If the center of gravity in AI value creation shifts from frontier training to deployed small models—as Nvidia and other industry leaders suggest—the competitive implications are significant.

The United States' AI advantages are concentrated in frontier scale: the largest training runs, the most advanced chips, and the most capital-intensive models. These remain important for research and for select consumer applications. But they may not determine which country captures value from AI in manufacturing, logistics, scientific research, and robotics—domains where deployment scale, iteration speed, and access to operational data matter more than benchmark performance.

U.S. export controls are calibrated to constrain frontier training by restricting access to advanced semiconductors. They do not address the small-model deployment cycle, which requires less advanced compute, draws on openly available base models, and generates advantage through application rather than pre-training. If the models that matter most for industrial AI are small, specialized, and open, the current U.S. policy framework could be targeting the wrong layer of the competition.

China's leadership appears to recognize this dynamic. Beijing's AI+ Initiative, its robotics industrial policy, and its institutional infrastructure for converting deployment data into a national asset all reflect a strategy organized around deployed AI rather than frontier scale. The convergence of China's open ecosystem dominance with its

manufacturing base could create a durable advantage that compounds independently of the frontier model race—and independently of the controls designed to slow it.

## Continued Momentum: From Labs to Factories

Even if China's leading models continue to lag behind U.S. closed models, China's open ecosystem, policy direction, and extensive manufacturing base—the foundation of the physical-economy feedback loop—could translate to advantage in adapting AI for two high-impact uses: research and robotics.

Scientific research is a central priority under China's AI+ Initiative. Early evidence compiled by Georgetown University's Center for Security and Emerging Technology suggests the adaptability of open models facilitates greater use by academic researchers.<sup>80</sup> This research extends to advancing AI, reinforcing the iterative cycle between industry and academia. Along these lines, computer vision startup Megvii, University of Chinese Academy of Sciences, and Huazhong University of Science and Technology claim they developed a smaller vision-language AI model that retains some capabilities of large systems while running efficiently on consumer-level GPUs by improving how visual information is encoded.\*<sup>81</sup>

Robotics and embodied AI represent a second domain where openness may confer advantage—and where the physical-economy feedback loop is directly operative. By lowering barriers to adaptation, open models allow robotics teams to integrate advanced language and multimodal capabilities without independently replicating costly pre-training, accelerating experiments with physical environments.<sup>82</sup> Shanghai AI Lab's Intern Robotics team, for instance, adapted Alibaba's Qwen2.5-3B-Instruct to build InternVLA-M1, which it describes as "a unified framework for spatial grounding and robot control that advances instruction-following robots toward scalable, general-purpose intelligence."<sup>83</sup>

Beyond near-term applications, this emphasis on practical deployment generates crucial advantages for embodied AI development: experience integrating sensors and actuators, data from real-world interactions, and iterative learning about how AI systems perform in unpredictable physical contexts.<sup>84</sup> Importantly, many Chinese researchers view deploying AI in the physical world as a likely technical pathway toward AGI, emphasizing the systems of multiple practical models rather than one singular transformative superintelligent model.<sup>† 85</sup>

China's leading scientists have advanced this view. Mu-ming Poo, scientific director of the Institute of Neuroscience, Chinese Academy of Sciences, argues that LLMs must interact with the physical world to develop an evolving "world model," or AI systems that understand the dynamics of the physical world.<sup>86</sup> Zhu Songchun, the president of Beijing's Institute of General Artificial Intelligence, insists that AI research should shift from reasoning and mathematical models toward cognition and value alignment.<sup>87</sup>

Academic scholars are searching for the means to put this view into practice. Research funded by China's National Natural Science Foundation and conducted at Tsinghua University, an embodied AI hub, explores the path from

\* Megvii was added to the Entity List in October 2019 for its role in China's surveillance state. U.S. Department of Commerce, *Addition of Certain Entities to the Entity List*, October 9, 2019. <https://www.federalregister.gov/documents/2019/10/09/2019-22210/addition-of-certain-entities-to-the-entity-list>.

† Not all Chinese AI developers are monolithic, as different companies such as DeepSeek, Alibaba, Moonshot AI, Zhipu, and Baidu pursue different technical paths. Afra Wang, "Topology of 'China AI,'" *Afra Wang's Substack*, September 11, 2025. <https://afraw.substack.com/p/topology-of-china-ai>.

LLMs and world models to embodied AI.\*<sup>88</sup> The researchers explained that embodied AI relies on LLMs to encode multimodal inputs such as text, images, and audio; perform semantic reasoning and task decomposition; and decode those interpretations into actionable outputs.<sup>89</sup> In other words, LLMs serve as the cognitive engine of embodied AI systems, translating high-level understanding into goal-directed interaction with the physical world. The feedback loop from physical interaction back to the LLM enables continuous refinement of the system's world model, potentially creating a pathway to AGI from accumulated real-world experience rather than scaled pre-training.

Chinese scholars' research agenda is directly mirrored in state policy. The draft 15th Five-Year Plan calls on China to "explore development pathways for general artificial intelligence" by promoting model innovation in areas such as multi-modalities, intelligent agents, and embodied AI, and to leverage "high-value scenarios to drive the practical implementation and iterative upgrading" of both general-purpose and industry-specific models. The Plan effectively cements the physical-economy feedback loop as a national strategy.<sup>90</sup> This is the logic of the physical-economy loop carried to its endpoint: deployment generates data, data refine models, and refined models enable more sophisticated deployment—a *cycle that compounds independently of access to frontier compute*.

## U.S. Models Still Lead for Key Applications

Available data suggest China's success in promoting open models has not yet clearly translated into market share gains for applications like chatbots or video generation, where many opportunities for monetizing technical advances lie. Use of AI models to power applications represents a further step in adoption, where the strategic value shifts beyond model performance itself to utility and integration for practical uses.

According to Microsoft's assessment of global AI adoption by country end users, adoption rates in China trail behind the United States, at only 16.3 percent versus 28.3 percent by the second half of 2025, though the study did not have equally robust coverage for China.<sup>91</sup> A RAND Corporation assessment of web traffic for leading U.S. and Chinese LLMs similarly found U.S. models captured 93 percent of global LLM site visits in August 2025.<sup>92</sup> Global web traffic data tracked by AICPB, a China-based global AI product ranking firm, likewise showed a U.S. lead across applications in data aggregated from multiple platforms.<sup>93</sup> For December 2025, U.S. AI tools topped charts for chatbots, code assistants, search engines, image generators, and video generators, while Chinese AI tools trailed (see Table 3).<sup>94</sup>

AI agents are a notable exception,<sup>†</sup> where the top three firms have China ties, though establishing domicile is challenging. China's Qihoo 360's Nano AI leads U.S.-headquartered GenSpark AI in monthly visits (see Table 3). However, GenSpark was cofounded by two former Baidu executives.<sup>95</sup> Manus, ranked third, is based in Singapore but was founded in China with funding from Tencent. The AI startup is reportedly under investigation by Chinese authorities

\* Tsinghua University plays a significant role in ushering China's advance in embodied AI, leading the nation in embodied AI research publications with 43 papers between 2022 and September 2025 and establishment of 18 companies by 2025. William C. Hannas et al., "China's Embodied AI: A Path to AGI," *Center for Security and Emerging Technology*, December 2025, 18–19. <https://cset.georgetown.edu/wp-content/uploads/CSET-China-Embodied-AI.pdf>.

† AI agents are modular, LLM-enabled systems designed for task-specific automation. They are distinct from agentic AI, or AI models performing functions for users (e.g., operating a web browser). Ranjan Sapkota, Konstantinos I. Roumeliotis, and Manoj Karkee, "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges," *arXiv*, May 15, 2025. <https://arxiv.org/abs/2505.10468>.

for potential violations of export restrictions and other regulations, following its acquisition by Meta in December 2025.\*

**Table 3: U.S. and Chinese Most Popular AI Web Tools by Category, December 2025**

Ranking Category	Country	Company	AI Product	December Monthly Visits (in millions)
<b>Overall and AI Chatbot</b>	U.S. ★	OpenAI	ChatGPT	5,700
	China	DeepSeek	DeepSeek	451.1
<b>AI Agent</b>	U.S.	GenSpark AI	GenSpark	13.5
	China ★	Qihoo360	Nano AI	189.41
<b>AI Code Assistant</b>	U.S. ★	GitHub	GitHub Copilot	304.1
	China	Baidu	Comate	1.85
<b>AI Search Engine</b>	U.S. ★	Microsoft	New Bing	1,330
	China	Qihoo360	Nano AI Super Search Agent	278.76
<b>AI Image Generator</b>	U.S. ★	SeaArt	SeaArt	25.54
	China	Shenzhen Liangmeng Tech	Jimeng AI	11.52
<b>AI Video Generator</b>	U.S. ★	OpenAI	Sora	35.14
	China	Kuaishou	Klingai	14.25

Note: ★ indicates category leader, based on the AI product's monthly visits in December 2025.

Source: Commission staff analysis based on AICPB's "Global AI Rankings by Users," accessed January 23, 2026. <https://www.aicpb.com>.

A few factors may explain Chinese models' lag relative to U.S. models' use in power applications. Even if Chinese models are nearly as good as top U.S. models for much lower costs, enterprise customers developing AI-base applications may prefer U.S. models for political and security reasons. Additionally, measuring use of AI models in applications is difficult compared to tracking model downloads and derivative uploads to gauge model diffusion

\* For more details on the Chinese Ministry of Commerce's investigation into the acquisition, see Graham Ayres et al., "February 2026 China Bulletin," *U.S.-China Economic and Security Review Commission*, February 4, 2026. <https://www.uscc.gov/trade-bulletins/china-bulletin-february-4-2026>.

(for more details on the challenges of measuring application use by nationality of models, see Appendix II: “Challenges Measuring AI Adoption in Applications”).

Despite its current lead, the United States should not get comfortable. Chinese AI firms have stressed the need to improve model integration in applications. Baidu’s CEO even sees the breadth of Chinese internet tech giants platforms as providing a leg up in developing models to power apps, noting “we [China] probably cannot match the investment of Google and OpenAI in training models, but we are closer to the applications.”<sup>96</sup> If China’s tech giants succeed in closing this gap, the compounding dynamics documented above could extend from model development into the application layer. Along these lines, Baidu, Alibaba, Tencent, and ByteDance are currently working to integrate their respective base models, developer infrastructure, enterprise productivity tools, marketing and creative platforms, and consumer applications.<sup>97</sup>

## U.S. Responses

Following DeepSeek’s success, major U.S. AI companies beyond Meta are entering the open ecosystem. In August 2025, OpenAI released two open-weight models—its first since GPT-2—after CEO Sam Altman admitted the company had been “on the wrong side of history” regarding open models.<sup>98</sup> In March 2026, Nvidia released Nemotron 3, an open-weight model designed to run on its hardware and support complex agentic AI systems, signaling the company’s commitment to open AI development alongside significant R&D investment to develop open-weight models.<sup>99</sup> Other efforts, such as the ATOM project, aim to foster U.S.-based labs for competitive open-source AI, but the overall health of the U.S. commercial open ecosystem remains uncertain.

Yet the trajectory of the U.S. open ecosystem’s most prominent champion is moving in the opposite direction. Meta released Llama 4 in April 2025 with two open-weight variants, Scout and Maverick, but withheld its most powerful model, Behemoth, from public release.<sup>100</sup> By late 2025, reports indicated that Meta’s next-generation model—codenamed Avocado—would adopt a closed, API-only approach, abandoning the weight downloads that defined the Llama series.<sup>101</sup> The shift was driven in part by concerns that Chinese labs, particularly DeepSeek, had leveraged Llama’s open architecture to accelerate their own capabilities—the very dynamic this paper documents.<sup>102</sup> CEO Mark Zuckerberg reportedly stated that Meta would not release models capable of superintelligence as open source.<sup>103</sup> The pivot contributed to leadership upheaval, including the departure of prominent AI scientist Yann LeCun from Meta’s Fundamental AI Research lab.<sup>104</sup> If sustained, Meta’s retreat from openness would leave the United States without a major frontier model developer anchoring its open AI ecosystem at precisely the moment China’s state-backed open development is accelerating.

U.S. policymakers have also recently rallied behind open models as a strategic lever in the global AI competition. The 2025 AI Action Plan explicitly positions them as critical to maintaining U.S. innovation leadership, reflecting growing expert consensus that open-source AI accelerates research, expands access, and provides alternatives to closed models.<sup>105</sup> The plan proposes national compute marketplaces to ease the bottlenecks from proprietary GPU cluster use by large AI companies.<sup>106</sup>

Beyond the private sector, the U.S. Department of Energy’s Genesis Mission mobilizes 17 national laboratories, industry, and academia to accelerate AI-driven scientific discovery, energy innovation, and national security—including investments in robotics, autonomous laboratories, and embodied AI.<sup>107</sup>

Responding to China’s focus on widespread adoption, both U.S. industry and policymakers have called for balancing pursuit of AGI with practical diffusion. Former CEO of Google Eric Schmidt called for both pursuit of superintelligence

models and broad deployment of less powerful models that impact everyday lives.<sup>108</sup> Jack Shanahan, retired U.S. Air Force lieutenant general and former director of the U.S. Department of Defense’s Joint Artificial Intelligence Center, similarly emphasized the need to “leverage AI as a practical tool—a driver of innovation and a measurable force multiplier for the American economy and society.”<sup>109</sup> These responses, however, primarily address the digital side of the competition—open model development and compute access—rather than the deployment-driven data advantages where China’s structural position is strongest.

## Areas for Additional Consideration

As of publication, Chinese labs have maintained momentum, with Tsinghua University spinoff Z.ai setting new records for open models with GLM-5, and Minimax’s high-profile Hong Kong initial product offering (IPO) has signaled investor confidence in China’s AI commercialization strategy. This trajectory presents a strategic dilemma: U.S. export controls limit China’s ability to match U.S. compute-intensive training, but they do not address the second feedback loop—deployment-driven data accumulation across China’s industrial base. The central policy challenge is that the United States has calibrated its response to one loop while China compounds advantages through both. Several gaps between current U.S. policy frameworks and the competitive dynamics documented in this analysis merit consideration by Congress.

- As this analysis documents, it is the intersection of these two loops—open model diffusion and industrial deployment—that gives China’s AI strategy its compounding force. Controls calibrated to only one loop could leave the more durable source of advantage unaddressed.
- There is no coordinated U.S. response to China’s open ecosystem dominance. The most prominent U.S. advocate for open AI development, Meta, has reportedly begun shifting its next-generation models toward a closed approach—driven in part by concerns that Chinese labs were leveraging Llama’s open architecture. While the 2025 AI Action Plan acknowledges open models’ importance, and efforts like the ATOM project and the industry-led AI Alliance attempt to foster a U.S.-based open ecosystem, no initiative matches the scale or coordination of China’s state-backed open development. U.S. researchers and companies increasingly rely on Chinese base models, creating long-term dependency on infrastructure with embedded censorship and potential security vulnerabilities.
- China’s open models could become the foundational infrastructure of global AI development. As documented in Appendix I, China has a long track record of leveraging open technology ecosystems to establish global technological dependencies. If Chinese base models become the default starting point for developers worldwide—as Hugging Face download data already suggest—China would shape the architecture, data formats, and security characteristics of AI systems far beyond its borders. Congress could consider directing NIST or the U.S. Office of Science and Technology Policy to engage proactively on AI interoperability standards, model evaluation frameworks, and deployment benchmarks before this infrastructure calcifies.
- U.S. domestic AI investment may be misaligned with the competitive dynamics this analysis documents. From CHIPS Act incentives to National Science Foundation AI research funding, policy mechanisms have primarily supported frontier model development rather than accelerating real-world deployment. There is no domestic equivalent to China’s systematic effort to convert industrial deployment into a data asset.
- An asymmetry in data access benefits Chinese models. Chinese open models hosted on global platforms like Hugging Face benefit from worldwide data contributions, community feedback, and derivative development.

Yet China's own data security law restricts outbound data sharing, limiting reciprocal access for foreign developers. This asymmetry—documented earlier in this analysis—means Chinese models improve through global openness while Chinese data remain largely closed. Congress could consider whether reciprocity requirements or transparency standards should apply to models hosted on U.S.-based platforms.

- Fundamental measurement gaps obscure competition in applications. Available metrics track model capabilities and training resources but provide limited visibility into how models are used to power applications like agents or video generators. Without standardized frameworks for measuring deployment, policymakers lack data to assess whether U.S. advantages in frontier models are translating into sustained leadership where strategic and economic value is captured—in manufacturing automation, logistics optimization, and research acceleration.
- Alternative pathways to AGI receive limited policy attention. Chinese researchers frame embodied AI as a viable route to AGI that bypasses compute-intensive scaling—a pathway China's policy framework explicitly supports through robotics initiatives, industrial AI deployment, and academic-industry collaboration on world models. U.S. investors and labs are betting heavily on AGI emerging from scaled pre-training of language models. If multiple technical pathways to advanced AI exist, U.S. industry and policy are optimized for only one, potentially ceding advantages in domains where physical deployment and iteration outweigh scale. As the preceding analysis suggests, the models most relevant to this physical pathway may be small, specialized, and openly available—a category where China already leads.
- The Genesis Mission offers a model that could extend beyond science. The Department of Energy's Genesis Mission, launched by executive order in November 2025, mobilizes 17 national laboratories, industry, and academia around AI-driven deployment for scientific discovery, energy innovation, and national security—including 14 robotics and autonomous laboratory projects. It represents the kind of deployment-focused, application-driven initiative this analysis suggests the United States needs. However, no equivalent effort exists for manufacturing, logistics, or industrial AI—the domains where China's physical-economy feedback loop is compounding most rapidly. Congress could consider whether the Genesis Mission model should be extended or replicated to accelerate AI deployment across the broader industrial base.
- Workforce and talent dynamics may be shifting. China's open ecosystem creates a visible pathway for AI talent—labs like Minimax and Moonshot AI have attracted engineers and capital through widely downloaded open models.<sup>110</sup> The United States' ability to attract and retain top AI researchers could be affected if the center of gravity in open AI development shifts toward Chinese platforms and institutions. At the same time, the shift toward deployed AI across the economy—from small, specialized models on factory floors to frontier systems in research labs—will require a workforce capable of developing, adapting, and maintaining these tools at every level. Congress could consider whether current education and training systems are preparing the U.S. workforce for an AI-integrated economy, including whether K-12 curricula, vocational programs, and post-secondary education are equipping workers with the skills to deploy and manage AI across industrial and commercial settings.

## Appendix I: China’s “Open” Technology Long Game to Become the Industry Standard

The Chinese government has long encouraged open-source technologies to drive domestic innovation and mitigate reliance on U.S. proprietary software.<sup>111</sup> In the early 2000s, China’s Academy of Sciences (CAS) launched the government-funded Red Flag Linux project to replace Microsoft Windows.<sup>112</sup> In recent years, the Chinese government has embraced and promoted open-source standards in many technologies. In June 2020, Chinese tech giants and state-affiliated entities founded China’s OpenAtom Foundation to promote collaboration on open-source technologies both domestically and internationally.<sup>113</sup> Huawei also contributed to OpenHarmony, an open-source, distributed operating system (OS) framework to displace Google’s Android OS.<sup>114</sup>

The earlier push for “open” technology development has laid a strong foundation for current open AI model development. China promoted Gitee as an alternative to GitHub to host source code.<sup>115</sup> In 2022, Alibaba also launched ModelScope, a platform for open-source AI models, as an alternative to Hugging Face since China restricted access to this global platform.<sup>116</sup> To compete with U.S. AI developing tools such as Meta’s PyTorch and Google’s TensorFlow, which developers use in their workflow to complete code, debug, and test their systems, China launched its own versions, such as Baidu’s PaddlePaddle and Huawei’s MindSpore.<sup>117</sup> With government and tech giant backing, more than 90 percent of Chinese enterprises used open-source technologies as of 2023.<sup>118</sup> As part of the global community, Chinese open-source projects account for 17 percent of the top 100 most active open-source software projects globally as of 2023, trailing only the United States, according to the Cloud Computing Standard and Open Source Promotion Committee of China.<sup>119</sup>

Since the release of DeepSeek’s open models in early 2025, the Chinese government has promoted open AI models with the goal to frame China as the global provider of open AI development and propel Chinese models into industry standard.<sup>120</sup> The Chinese leadership has already positioned China as a global champion of open-source AI development. In the summer of 2025, Premier Li Qiang pronounced at the World Economic Forum that “China’s innovation is open and open-source [and China is] willing to share indigenous technologies and innovation scenarios with the world,” and he continued to promote China’s open AI development at the World Artificial Intelligence Conference.<sup>121</sup> The Chinese government also puts open model development at the forefront of its domestic AI plan, announcing the AI+ Initiative that calls for government support of AI products and prioritizes the development of open-source AI ecosystems.<sup>122</sup> In September, the Cyberspace Administration of China, the country’s most influential AI regulator, released a framework for AI technical standards with emphasis on risks from open-source models.<sup>123</sup> China’s domestic technical standards-setting for open models will have greater ramifications for international AI development as Chinese models are gaining popularity at home and abroad.

## Appendix II: Challenges Measuring AI Adoption in Applications

For applications, the distinction between open and closed model strategies and benefits is less readily apparent than for models because similar products and services can be built on top of either type of model.<sup>124</sup> Both approaches have the potential to accelerate widespread adoption. Closed models offer plug-and-play API access that enables application development, while open models allow finetuning and deployment for specialized use cases.<sup>125</sup> In theory, China’s success in diffusing open models and the United States’ lead in frontier model development could translate into distinct patterns of use in AI-powered applications. In practice, assessing relative

diffusion in this layer remains premature due to fundamental data limitations, market fragmentation, and the nascent state of many application categories.

Measuring AI adoption in the application layer is complicated for several reasons.

- Market structures differ fundamentally. U.S.-developed applications typically target global audiences, while Chinese applications primarily serve the domestic market that benefits from a large internet population, advanced mobile ecosystem, lower regulatory requirements, and early commercialization through government procurement and enterprise partnerships.<sup>126</sup> Only a minority of successful Chinese AI app developers have attempted to reach the global market to generate revenue.<sup>127</sup> Therefore, raw user traffic metrics conflate these different user bases, skewing significantly toward AI apps and tools that serve the global market.
- Even with plausible user traffic data as a measurement of diffusion, the lack of sufficient data makes it difficult for cross-country comparison of adoption rates. As Jeff Ding, a professor at George Washington University, notes, assessing the widespread uptake of a technology requires numerous metrics that are often inconsistent or incomparable across contexts.<sup>128</sup>
- Different methodologies attempt to address this challenge with varying degrees of success. Microsoft measures AI diffusion across countries by estimating the number of AI users, but the data lack China's coverage.<sup>129</sup> Similarly, Epoch AI's analysis of AI adoption rates also relies on estimating active users but focuses more specifically on open models.<sup>130</sup> RAND's methodology relies on SimilarWeb, which collects anonymized web traffic data from thousands of global devices and tallies total monthly website visits by country but is limited by the data's exclusion of application-based use for granular observations.<sup>131</sup> Chinese government-reported token usage data showing an increase from 100 billion daily tokens in early 2024 to over 30 trillion in late 2025 offer evidence of growth at a systemic level but do not account for different applications' token intensity and are not comparable to data on token usage in the United States.<sup>132</sup> In this paper, an early assessment of U.S. and Chinese AI diffusion at the application layer relies on AICPB, a China-based global AI product ranking that tracks web traffic from multiple platforms and third-party datasets.<sup>133</sup> While it provides cross-country comparisons of AI tool popularity by application-specific use, its methodology is not entirely transparent, including category definition and classification.
- Visible consumer-facing applications represent only part of the diffusion story. Enterprise deployments, AI features embedded within existing platforms such as Baidu search or Google's Gemini agent, and vertical-specific tools often operate below the radar of web traffic measurements via API access.<sup>134</sup> Future definitive assessment should require more comprehensive data collection, clearer definitional standards, better methods for capturing invisible deployments, and sufficient time for applications to mature and markets to stabilize.

**Disclaimer:** This paper is the product of professional research performed by staff of the U.S.-China Economic and Security Review Commission, and was prepared at the request of the Commission to support its deliberations. Posting of the report to the Commission's website is intended to promote greater public understanding of the issues addressed by the Commission in its ongoing assessment of U.S.-China economic relations and their implications for U.S. security, as mandated by Public Law 106-398 and Public Law 113-291. However, the public release of this document does not necessarily imply an endorsement by the Commission, any individual Commissioner, or the Commission's other professional staff, of the views or conclusions expressed in this staff research report.

# Endnotes

- <sup>1</sup> Kali Hays, "The Best Inventions of 2025: DeepSeek R1," *TIME*, October 9, 2025. <https://time.com/collections/best-inventions-2025/7318231/vera-c-rubin-observatory/>; "Introducing GPT-5," *OpenAI*, August 7, 2025. <https://openai.com/index/introducing-gpt-5/>.
- <sup>2</sup> "Gemini 3.1 Pro: A Smarter Model for Your Most Complex Tasks," *Google*, February 19, 2026. <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-1-pro>; Eduardo Baptista, "A Year On from DeepSeek Shock, Get Set for Flurry of Low-Cost Chinese AI Models," *Reuters*, February 12, 2026. <https://www.reuters.com/world/china/year-deepseek-shock-get-set-flurry-low-cost-chinese-ai-models-2026-02-12/>; "Introducing Claude Opus 4.6," *Anthropic*, February 5, 2026. <https://www.anthropic.com/news/claude-opus-4-6>.
- <sup>3</sup> Michael Froman, "China, the United States, and the AI Race," *Council on Foreign Relations*, October 10, 2025. <https://www.cfr.org/article/china-united-states-and-ai-race>; China's State Council, "Opinions of the State Council on Deepening the Implementation of the 'Artificial Intelligence+' Initiative," *Center for Security and Emerging Technology* (translation), August 21, 2025, 1–2. [https://cset.georgetown.edu/wp-content/uploads/t0652\\_AI\\_plus\\_opinions\\_EN.pdf](https://cset.georgetown.edu/wp-content/uploads/t0652_AI_plus_opinions_EN.pdf); William Hannas et al., "China's Advanced AI Research: Monitoring China's Paths to 'General' Artificial Intelligence," *Center for Strategic and Emerging Technology*, July 2022, v. 3. <https://cset.georgetown.edu/publication/chinas-advanced-ai-research/>.
- <sup>4</sup> U.S.-China Economic and Security Review Commission, *2024 Annual Report to Congress*, November 2024, 733.
- <sup>5</sup> Sunil Ramlochan, "Openness in Language Models: Open Source vs. Open Weights vs. Restricted Weights," *Prompt Engineering and AI Institute*, accessed March 6, 2026. <https://promptengineering.org/llm-open-source-vs-open-weights-vs-restricted-weights/>; Justin Riddiough, "Understanding Open Weights and Parameters in Open Source AI," *AI Models*, December 10, 2023. <https://aimodels.org/open-source-ai/open-weights/>.
- <sup>6</sup> Ben Cottier et al., "How Far Behind Are Open Models?" *Epoch AI*, November 4, 2024. <https://epoch.ai/blog/open-models-report>.
- <sup>7</sup> James Dunham, "Large Language Models (LLMs): An Explainer," *Center for Security and Emerging Technology*, August 1, 2023. <https://cset.georgetown.edu/article/large-language-models-llms-an-explainer/>.
- <sup>8</sup> Sebastian Raschka, "Understanding and Coding the Self-Attention Mechanism of Large Language Models from Scratch," February 9, 2023. <https://sebastianraschka.com/blog/2023/self-attention-from-scratch.html>.
- <sup>9</sup> "Artificial Power: 2025 Landscape Report," *AI Now Institute*, June 3, 2025. <https://ainowinstitute.org/publications/research/1-1-the-agi-mythology-the-argument-to-end-all-arguments>; Kari Briski, "How Scaling Laws Drive Smarter, More Powerful AI," *Nvidia*, February 12, 2025. <https://blogs.nvidia.com/blog/ai-scaling-laws>; Sam Altman, "Three Observations," *Sam Altman's Personal Blog*, February 9, 2025. <https://blog.samaltman.com/three-observations>; Casey Newton, "AI Companies Hit a Scaling Wall," *Platformer*, November 14, 2024. <https://www.platformer.news/openai-google-scaling-laws-anthropic-ai/>; Billy Perrigo, "Anthropic CEO Dario Amodei on Being an Underdog, AI Safety, and Economic Inequality," *Time*, June 23, 2024. <https://time.com/6990386/anthropic-dario-amodei-interview/>; Leopold Aschenbrenner, "Situational Awareness," *Situational Awareness*, June 2024. <https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>.
- <sup>10</sup> Jie Ding, *GenAI Course: Resource Index*, Stat 8105: GenAI (Fall 2025), School of Statistics, University of Minnesota, accessed February 23, 2026. <https://genai-course.jding.org/resource/index.html>.
- <sup>11</sup> Lex Fridman, Nathan Lambert, and Sebastian Raschka, "State of AI in 2026: LLMs, Coding, Scaling Laws, China, Agents, GPUs, AGI," *Lex Fridman Podcast #490* (podcast), February 1, 2026, 00:58:36. <https://lexfridman.com/ai-sota-2026-transcript>.
- <sup>12</sup> Leopold Aschenbrenner, "Situational Awareness," *Situational Awareness*, June 2024, 26–28. <https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>; Blaise Agüera y Arcas and Peter Norvig, "Artificial General Intelligence Is Already Here," *Noema Magazine*, October 10, 2023. <https://www.noemamag.com/artificial-general-intelligence-is-already-here/>.
- <sup>13</sup> Sarah Perez, "One of Google's Biggest AI Advantages Is What It Already Knows about You," *TechCrunch*, December 1, 2025. <https://techcrunch.com/2025/12/01/one-of-googles-biggest-ai-advantages-is-what-it-already-knows-about-you/>; Alex Reisner, "The Company Quietly Funneling Paywalled Articles to AI Developers," *Atlantic*, November 4, 2025. <https://www.theatlantic.com/technology/2025/11/mon-crawl-ai-training-data/684567/>.
- <sup>14</sup> Rachel Hovde, "Foreign Influence Risks in the AI Data Annotation Sector," *Americans for Responsible Innovation*, September 24, 2025, 3. <https://ari.us/wp-content/uploads/2025/09/White-Paper-AI-Data-Annotation-Risks-9-24.pdf>.
- <sup>15</sup> OpenAI, "GPT-4 Technical Report," *arXiv*, March 2, 2023, 2. <https://arxiv.org/pdf/2303.08774>.
- <sup>16</sup> Matt Day and Annie Bang, "Big Tech to Spend \$650 Billion This Year as AI Race Intensifies," *Bloomberg*, February 6, 2026. <https://www.bloomberg.com/news/articles/2026-02-06/how-much-is-big-tech-spending-on-ai-computing-a-staggering-650-billion-in-2026>; Vlad Savov, "China's AI Future Looks Like a Long Road to Small Profit," *Bloomberg*, December 3, 2025. <https://www.bloomberg.com/news/newsletters/2025-12-03/china-s-ai-future-looks-like-a-long-road-to-small-profit>.
- <sup>17</sup> Vlad Savov, "China's AI Future Looks Like a Long Road to Small Profit," *Bloomberg*, December 3, 2025. <https://www.bloomberg.com/news/newsletters/2025-12-03/china-s-ai-future-looks-like-a-long-road-to-small-profit>.
- <sup>18</sup> Ji Zhou et al., "A Comparative Evaluation of Large Vision-Language Models for 2D Object Detection under SOTIF Conditions," *arXiv*, January 30, 2026; "AI Models," *Epoch AI*, accessed January 23, 2026. <https://epoch.ai/data/ai-models?view=table>; Ege Erdil, "Inference Economics of Language Models," *Epoch AI*, June 17, 2025. <https://epoch.ai/blog/inference-economics-of-language-models>; Josh You, "How Much Energy Does ChatGPT Use?" *Epoch AI*, February 7, 2025. <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use/>; Ron Barron and Elon Musk, "Fireside Chat: Starman Elon Musk and Ron Baron Discuss the Future," *Baron Capital*, November 14, 2024, 35:20. <https://www.baroncapital-group.com/conference-2025/ron-baron-elon-musk-discuss-the-future>.

- <sup>19</sup> Sara Hooker, "On the Slow Death of Scaling," *SSRN*, last revised January 6, 2026, 5. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5877662](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5877662); Moritz Hain, "When Bigger Isn't Better: The Diminishing Returns of Scaling AI Models," *Sapien Blog*, October 31, 2025. <https://www.sapien.io/blog/when-bigger-isnt-better-the-diminishing-returns-of-scaling-ai-models>.
- <sup>20</sup> Annie Wong et al., "Reasoning Capabilities of Large Language Models on Dynamic Tasks," *arXiv preprint*, August 10, 2025. <https://doi.org/10.48550/arXiv.2505.10543>; Olivia Shone, "3 Key Features and Benefits of Small Language Models," *Microsoft*, September 25, 2024. <https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/09/25/3-key-features-and-benefits-of-small-language-models>; Sher Badshah and Hassan Sajjad, "Quantifying the Capabilities of LLMs across Scale and Precision," *arXiv preprint*, May 8, 2024. <https://doi.org/10.48550/arXiv.2405.03146>.
- <sup>21</sup> Cedric Clyburn, "The State of Open Source AI Models in 2025," *Red Hat*, January 7, 2026. <https://developers.redhat.com/articles/2026/01/07/state-open-source-ai-models-2025>; Ege Erdil, "Frontier Language Models Have Become Much Smaller," *Epoch AI*, December 13, 2024. <https://epoch.ai/gradient-updates/frontier-language-models-have-become-much-smaller>.
- <sup>22</sup> Sebastian Raschka, "The State of LLMs 2025: Progress, Problems, and Predictions," *Ahead of AI* (blog), December 30, 2025. <https://magazine.sebastianraschka.com/p/state-of-llms-2025>.
- <sup>23</sup> "Introducing GPT-4.5," *OpenAI*, February 27, 2025. <https://openai.com/index/introducing-gpt-4-5>.
- <sup>24</sup> "Gemini 3 Pro Preview (High) vs. Kimi K2 Thinking: Model Comparison," *Artificial Analysis*, accessed March 3, 2026. <https://artificialanalysis.ai/models/comparisons/gemini-3-pro-vs-kimi-k2-thinking>; "China AI Startup Moonshot Targets \$10 Billion Valuation," *Bloomberg*, February 17, 2026. <https://www.bloomberg.com/news/articles/2026-02-17/china-ai-startup-moonshot-seeks-10-billion-value-in-new-funding>; DeepMind, "Gemini 3 Pro Model Card," November 2025, 1, 4. <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-Pro-Model-Card.pdf>; Alex Volkov, "ThursdAI – Nov 6, 2025 – Kimi's 1T Thinking Model Shakes Up Open Source, Apple Bets \$1B on Gemini for Siri, and Amazon vs. Perplexity!" *ThursdAI: The Top AI News from the Past Week* (podcast page), November 7, 2025. <https://sub.thursdai.news/p/thursdai-nov-6-2025-kimis-1t-thinking>.
- <sup>25</sup> "DeepSeek-V3.2 Release," *DeepSeek*, December 1, 2025. <https://api-docs.deepseek.com/news/news251201>; Vincent Chow, "DeepSeek Releases First Open AI Model with Gold-Level Scores at Maths Olympiad," *South China Morning Post*, November 29, 2025. <https://www.scmp.com/tech/tech-trends/article/3334553/deepseek-releases-first-open-ai-model-gold-level-scores-maths-olympiad>; Emily Riehl, "AI Took on the Math Olympiad—But Mathematicians Aren't Impressed," *Scientific American*, August 7, 2025. <https://www.scientificamerican.com/article/mathematicians-question-ai-performance-at-international-math-olympiad/>.
- <sup>26</sup> DeepSeek-AI, "DeepSeek-R1 Technical Overview: Model Architecture," *DeepWiki*. <https://deepwiki.com/deepseek-ai/DeepSeek-R1/2-model-architecture>; William Fedus, Barret Zoph, and Noam Shazeer, "Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity," *arXiv*, June 16, 2022. <https://arxiv.org/abs/2101.03961>.
- <sup>27</sup> Shruti Koparkar, "Mixture of Experts Powers the Most Intelligent Frontier AI Models, Runs 10x Faster to Deliver 1/10 the Token Cost on NVIDIA Blackwell NVL72," *Nvidia*, December 3, 2025. <https://blogs.nvidia.com/blog/mixture-of-experts-frontier-models/>; "The Struggle to Get Inside How AI Models Really Work," *Financial Times*, June 4, 2025. <https://www.ft.com/content/b349f590-de84-455d-914a-cc5d9eef04a6?>
- <sup>28</sup> Luke Emberson, "Chinese AI Models Have Lagged the US Frontier by 7 Months on Average since 2023," *Epoch AI*, January 2, 2026. <https://epoch.ai/data-insights/us-vs-china-eci>; Luke Emberson, "Open-Weight Models Lag State-of-the-Art by around 3 Months on Average," *Epoch AI*, October 30, 2025. <https://epoch.ai/data-insights/open-weights-vs-closed-weights-models>.
- <sup>29</sup> Xin Wen, "Challenges and Recommendations for Building Open Source Innovation Ecosystem for Large-Models in China," *Bulletin of Chinese Academy of Sciences* 39, no. 8 (August 20, 2024). <https://bulletinofcas.researchcommons.org/cgi/viewcontent.cgi?article=2664&context=journal>.
- <sup>30</sup> Graham Webster et al., "Full Translation: China's 'New Generation Artificial Intelligence Development Plan,'" *DigiChina*, August 1, 2017. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
- <sup>31</sup> China's State Council, "Opinions of the State Council on Deepening the Implementation of the 'Artificial Intelligence+' Initiative," *Center for Security and Emerging Technology* (translation), August 21, 2025, 1–2. [https://cset.georgetown.edu/wp-content/uploads/t0652\\_AI\\_plus\\_opinions\\_EN.pdf](https://cset.georgetown.edu/wp-content/uploads/t0652_AI_plus_opinions_EN.pdf).
- <sup>32</sup> "Kimi (API Docs)," accessed March 12, 2026. <https://platform.moonshot.ai/docs/overview>; "MiniMaxAI/MiniMax-M2.5," *Hugging Face*, accessed March 12, 2026. <https://huggingface.co/MiniMaxAI/MiniMax-M2.5>; Julian Horsey, "Minimax M2.5 Benchmarks: Targets \$1 per Hour for 100 Tokens per Second," *Geeky Gadgets*, February 17, 2026. <https://www.geeky-gadgets.com/minimax-m2-5-benchmarks/>; "China's AI Startup MiniMax Group Raises \$619 Million in Hong Kong IPO," *Reuters*, January 8, 2026. <https://www.reuters.com/world/asia-pacific/chinas-ai-startup-minimax-group-raises-619-million-hong-kong-ipo-2026-01-08/>; Caroline Meinhardt et al., "Beyond DeepSeek: China's Diverse Open-Weight AI Ecosystem and Its Policy Implications," *Stanford University's Human-Centered Artificial Intelligence Institute*, December 2025, 11. <https://hai.stanford.edu/assets/files/hai-digichina-issue-brief-beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-policy-implications.pdf>; U.S. Department of Commerce, "Addition of Entities to and Revision of Entry on the Entity List," Fed. Reg. 250108-0010 (January 16, 2025). <https://www.federalregister.gov/documents/2025/01/16/2025-00704/addition-of-entities-to-and-revision-of-entry-on-the-entity-list>.
- <sup>33</sup> "LLM Leaderboard - Comparison of over 100 AI Models from OpenAI, Google, DeepSeek & Others," *Artificial Analysis*, accessed February 1. <https://artificialanalysis.ai/leaderboards/models?deprecation=all>.
- <sup>34</sup> Tao Tian Yu and Guo Hong Yun, "2025 AI创业真相" [The Realities of AI Start-Ups in 2025], *tmtpost*, August 2, 2025. <https://www.less-wrong.com/posts/jcj5hXoDDdNcAgRna/translation-the-realities-of-ai-start-ups-in-2025>.
- <sup>35</sup> Shannon Vaughn, "From Vouchers to Visas: China's Innovative Plan for AI Dominance," *Foreign Policy Research Institute*, September 10, 2025. <https://www.fpri.org/article/2025/09/from-vouchers-to-visas-chinas-innovative-plan-for-ai-dominance/>.
- <sup>36</sup> "LLM Leaderboard - Comparison of over 100 AI Models from OpenAI, Google, DeepSeek & Others," *Artificial Analysis*, accessed February 1. <https://artificialanalysis.ai/leaderboards/models?deprecation=all>.

- <sup>37</sup> Zijiang Wu and Eleanor Olcott, "China Offers Tech Giants Cheap Power to Boost Domestic AI Chips," *Financial Times*, November 4, 2025. <https://www.ft.com/content/cad2cdd6-7cce-4de3-8710-977de667378c>.
- <sup>38</sup> Andrew Stokols, "Energy and AI Coordination in the 'Eastern Data Western Computing' Plan," *Jamestown Foundation*, March 8, 2025. <https://jamestown.org/program/energy-and-ai-coordination-in-the-eastern-data-western-computing-plan/>; Sara Hsu, "China's Overlooked AI Energy Edge over the US: Cheaper Energy," *Diplomat*, February 5, 2025. <https://thediplomat.com/2025/02/chinas-overlooked-ai-energy-edge-over-the-us-cheaper-energy/>.
- <sup>39</sup> Shannon Vaughn, "From Vouchers to Visas: China's Innovative Plan for AI Dominance," *Foreign Policy Research Institute*, September 10, 2025. <https://www.fpri.org/article/2025/09/from-vouchers-to-visas-chinas-innovative-plan-for-ai-dominance/>.
- <sup>40</sup> Xin Wen, "Challenges and Recommendations for Building Open Source Innovation Ecosystem for Large-Models in China," *Bulletin of Chinese Academy of Sciences* 39, no. 8 (August 20, 2024). <https://bulletinofcas.researchcommons.org/cgi/viewcontent.cgi?article=2664&context=journ>.
- <sup>41</sup> "中华人民共和国国民经济和社会发展第十四个五年规划和 2035 年远景目标纲要" [Outline of the People's Republic of China 14th Five-Year Plan for National Economic and Social Development and Long-Range Objectives for 2035], *Xinhua*, March 12, 2021, 40. CSET Translation. [https://cset.georgetown.edu/wp-content/uploads/t0284\\_14th\\_Five\\_Year\\_Plan\\_EN.pdf](https://cset.georgetown.edu/wp-content/uploads/t0284_14th_Five_Year_Plan_EN.pdf); Junxiao Han et al., "An Empirical Study of the Landscape of Open Source Projects in Baidu, Alibaba, and Tencent," *arXiv*, March 2, 2021. <https://arxiv.org/pdf/2103.01590>; Graham Webster et al., "Full Translation: China's 'New Generation Artificial Intelligence Development Plan,'" *DigiChina*, August 1, 2017. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
- <sup>42</sup> "中华人民共和国国民经济和社会发展第十五个五年规划纲要" [Outline of the People's Republic of China 15th Five-Year Plan for National Economic and Social Development], *Xinhua*, March 13, 2026. <https://www.news.cn/politics/20260313/085af5de5a4b4268aa7d87d90817df2f/c.html>; "李强主持召开国务院常务会议" [Li Qiang Presided over the State Council Executive Meeting], *People's Daily*, August 1, 2025. <http://politics.people.com.cn/n1/2025/0801/c1024-40534026.html>; "中华人民共和国国民经济和社会发展第十四个五年规划和 2035 年远景目标纲要" [Outline of the People's Republic of China 14th Five-Year Plan for National Economic and Social Development and Long-Range Objectives for 2035], *Xinhua*, March 12, 2021, 40. CSET Translation. [https://cset.georgetown.edu/wp-content/uploads/t0284\\_14th\\_Five\\_Year\\_Plan\\_EN.pdf](https://cset.georgetown.edu/wp-content/uploads/t0284_14th_Five_Year_Plan_EN.pdf); Graham Webster et al., "Full Translation: China's 'New Generation Artificial Intelligence Development Plan,'" *DigiChina*, August 1, 2017. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
- <sup>43</sup> Ly Yating, Du Zhihang, and Denise Jia, "China's Regions Power Up AI and Robotics Innovation with Financial Incentives," *Caixin*, March 12, 2025. <https://www.caixinglobal.com/2025-03-12/chinas-regions-power-up-ai-and-robotics-innovation-with-financial-incentives-102297178.html>; Beijing Municipal Government, 北京市人民政府办公厅关于印发《北京市促进通用人工智能创新发展的若干措施》的通 [Notice from the General Office of the Beijing Municipal People's Government on Issuing the "Several Measures of Beijing Municipality for Promoting the Innovative Development of General Artificial Intelligence"], May 23, 2025. [https://www.beijing.gov.cn/zhengce/zhengcefaqiu/202305/t20230530\\_3116869.html](https://www.beijing.gov.cn/zhengce/zhengcefaqiu/202305/t20230530_3116869.html); Hangzhou City Government, 杭州市人民政府办公厅关于印发支持人工智能全产业链高质量发展若干措施的通知 [Notice from the General Office of the Hangzhou Municipal People's Government on Issuing Several Measures to Support the High-Quality Development of the Entire Artificial Intelligence Industry Chain], July 18, 2024. <https://zj87.jxt.zj.gov.cn/zlq/citySite/views/article/news/detail.html?id=250064>.
- <sup>44</sup> "AI Models," *Epoch AI*, accessed January 23, 2026. <https://epoch.ai/data/ai-models>.
- <sup>45</sup> "China Tops Global AI Model Count with over 1,500 Large Models Released," *Xinhua*, July 28, 2025. [https://english.scio.gov.cn/china-voices/2025-07/28/content\\_117999955.html](https://english.scio.gov.cn/china-voices/2025-07/28/content_117999955.html).
- <sup>46</sup> "AI Models," *Epoch AI*, accessed January 23, 2026. <https://epoch.ai/data/ai-models>; <https://epoch.ai/blog/open-models-report>.
- <sup>47</sup> Wang Zhe and Xue Lan, "大模型开源创新公地：历史演进、价值逻辑与中国叙事" [Open-Source Innovation Commons for Large Language Models: Historical Evolution, Value Logic, and the Chinese Narrative], *Tsinghua University's Center for International Security and Strategy*, September 24, 2025. <https://ciss.tsinghua.edu.cn/info/zlyaq/8618>.
- <sup>48</sup> "Reinvigorating AI Research in the U.S. by Building Leading, Open Models at Home," *ATOM Project*, accessed January 23, 2026. <https://atom-project.ai/>.
- <sup>49</sup> Kevin Schaul, "China Now Leads the U.S. in This Key Part of the AI Race," *Washington Post*, October 13, 2025. <https://www.washingtonpost.com/technology/2025/10/13/china-us-open-source-ai/>.
- <sup>50</sup> Commission staff analysis based on Hugging Face, accessed December 8, 2025. <https://huggingface.co/>.
- <sup>51</sup> "Reinvigorating AI Research in the U.S. by Building Leading, Open Models at Home," *ATOM Project*, accessed January 23, 2026. <https://atom-project.ai/>; Adina Yakefu and Irene Solaiman, "The Future of the Global Open-Source AI Ecosystem: From DeepSeek to AI+," *Hugging Face*, February 3, 2026. <https://huggingface.co/blog/huggingface/one-year-since-the-deepseek-moment-blog-3>.
- <sup>52</sup> "deepseek-ai/DeepSeek-R1-Distill-Qwen-32B," *Hugging Face*, accessed December 8, 2025. <https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-32B>.
- <sup>53</sup> "omni-research/Tarsier2-Recap-7b," *Hugging Face*, accessed December 8, 2025. <https://huggingface.co/omni-research/Tarsier2-Recap-7b>.
- <sup>54</sup> Anna Hermansen and Cailean Osborne, "The Economic and Workforce Impacts of Open Source AI," *Linux Foundation*, May 2025, 9. [https://www.linuxfoundation.org/hubfs/Research%20Reports/lfr\\_marketimpacts25\\_052725a.pdf?hsLang=en](https://www.linuxfoundation.org/hubfs/Research%20Reports/lfr_marketimpacts25_052725a.pdf?hsLang=en).
- <sup>55</sup> "MiniMax Hailuo 2.3: A New Level of Complex Video Performance & Media Agent," *Minimax*, October 28, 2025. <https://www.minimax.io/news/minimax-hailuo-23>; Eric Hal Schwartz, "Forget Virtual Pets – The Next AI Video Craze Is Cats Doing Olympic Diving, and It's All Thanks to This New Google Veo 3 Rival," *Tech Radar*, June 20, 2025. <https://www.techradar.com/computing/artificial-intelligence/forget-virtual-pets-the-next-ai-video-craze-is-cats-doing-olympic-diving-and-its-all-thanks-to-this-new-google-veo-3-rival>.

- <sup>56</sup> Yantoultra Ngui, Kane Wu, and Liam Mo, "MiniMax, China's Second 'AI Tiger' to Go Public, Doubles in Value in Hong Kong Debut," *Reuters*, January 9, 2026. <https://www.reuters.com/world/asia-pacific/china-ai-firm-minimax-set-surge-hong-kong-debut-2026-01-09/>.
- <sup>57</sup> "China Is Quietly Upstaging America with Its Open Models," *Economist*, August 21, 2025. <https://www.economist.com/business/2025/08/21/china-is-quietly-upstaging-america-with-its-open-models>.
- <sup>58</sup> "China Is Quietly Upstaging America with Its Open Models," *Economist*, August 21, 2025. <https://www.economist.com/business/2025/08/21/china-is-quietly-upstaging-america-with-its-open-models>.
- <sup>59</sup> Rachel Cheung, "Cheap and Open Source, Chinese AI Models Are Taking Off," *Wire China*, November 9, 2025. <https://www.thewirechina.com/2025/11/09/cheap-and-open-source-chinese-ai-models-are-taking-off/>.
- <sup>60</sup> U.S. National Institute of Standards and Technology, *Evaluation of DeepSeek AI Models*, September 30, 2025. [https://www.nist.gov/system/files/documents/2025/09/30/CAISI\\_Evaluation\\_of\\_DeepSeek\\_AI\\_Models.pdf](https://www.nist.gov/system/files/documents/2025/09/30/CAISI_Evaluation_of_DeepSeek_AI_Models.pdf).
- <sup>61</sup> Zeyi Yang, "How Chinese AI Chatbots Censor Themselves," *Wired*, February 26, 2026. <https://www.wired.com/story/made-in-china-how-chinese-ai-chatbots-censor-themselves/>; Bobby Allyn, "International Regulators Probe How DeepSeek Is Using Data. Is the App Safe to Use?" *NPR*, January 31, 2025. <https://www.npr.org/2025/01/31/nx-s1-5277440/deepseek-data-safety>.
- <sup>62</sup> Xi Jinping, "不断做强做优做大我国数字经济" [Unceasingly Make Our Country's Digital Economy Stronger, Better, and Bigger], *DigiChina*, January 15, 2022. <https://digichina.stanford.edu/work/translation-xi-jinpings-speech-to-the-politburo-study-session-on-the-digital-economy-oct-2021/>; China's State Council and CCP Central Committee, 中共中央 国务院关于构建更加完善的要素市场化配置体制机制的意见 [Opinions of the CCP Central Committee and the State Council on Improving the Systems and Mechanisms for Market-Based Allocation of Factors of Production], April 9, 2020. [https://www.gov.cn/zhengce/2020-04/09/content\\_5500622.htm](https://www.gov.cn/zhengce/2020-04/09/content_5500622.htm).
- <sup>63</sup> Pablo Villalobos et al., "Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data," *Epoch AI*, June 6, 2024. <https://epoch.ai/blog/will-we-run-out-of-data-limits-of-llm-scaling-based-on-human-generated-data>.
- <sup>64</sup> "我国正式全面启动'人工智能+'新时代" [China Officially Launches the New Era of "Artificial Intelligence Plus"], *China Development and Reform Daily*, September 1, 2025. [https://www.ndrc.gov.cn/wsdwhfz/202509/t20250901\\_1400201.html](https://www.ndrc.gov.cn/wsdwhfz/202509/t20250901_1400201.html).
- <sup>65</sup> Kevin Xu, "China's Structural Advantage in Open Source AI," *Interconnected* (blog), June 25, 2025. <https://interconnect.substack.com/p/china-structural-advantage-in-open>; "China Establishes over 30,000 Smart Factories: Ministry," *Xinhua*, February 12, 2025. [https://english.www.gov.cn/news/202502/12/content\\_WS67ac9475c6d0868f4e8ef978.html](https://english.www.gov.cn/news/202502/12/content_WS67ac9475c6d0868f4e8ef978.html).
- <sup>66</sup> "5G-A Powers All-Scenario IoT to Enable Intelligent Connections for All," *Huawei*, June 19, 2025. <https://www.huawei.com/en/news/2025/6/mwvsh-5ga-all-scenario-aiot>; "AgiBot Releases Humanoid Manipulation Dataset to Enable Large-Scale Learning," *Robot Report*, December 30, 2024. <https://www.therobotreport.com/agibot-releases-humanoid-manipulation-dataset-to-enable-large-scale-learning/>.
- <sup>67</sup> Qiao Xinyi, "Chinese Robotics Pioneer Fourier Makes Humanoid Robot Dataset Open Source," *Yicai Global*, March 18, 2025. <https://www.yicaiglobal.com/news/chinas-fourier-makes-humanoid-robot-dataset-open-source>.
- <sup>68</sup> Ran Guo, "Assetizing, Trading, Franchising: China's Strategy for Building a National Data Economy," *Asia Society*, February 13, 2026. <https://asiasociety.org/policy-institute/assetizing-trading-franchising-chinas-strategy-building-national-data-economy>; Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," *RAND*, June 26, 2025. <https://www.rand.org/pubs/perspectives/PEA4012-1.html>, 131; Rogier Creemers, "China's Emerging Data Protection Framework," *Journal of Cybersecurity* 8, no.1 (2022). <https://doi.org/10.1093/cybsec/tyac011>.
- <sup>69</sup> Mengying Tao, "Chasing Artificial General Intelligence: China between Breakthroughs and Bottlenecks," *Sinolytics*, July 31, 2025, 11. <https://sinolytics.de/global-business-news/blog/technology/agi-in-china/>; Irene Zheng, "AI Proposals at 'Two Sessions': AGI as 'Two Bombs, One Satellite'?" *China Talk* (blog), March 8, 2023. <https://www.chinatalk.media/p/ai-proposals-at-two-sessions-agi>.
- <sup>70</sup> Tang Jie, "大力推动我国人工智能大模型发展 〔深入学习贯彻习近平新时代中国特色社会主义思想·学习《习近平经济文选》第一卷专家谈〕 [Vigorously Promoting the Development of Large-Scale Artificial Intelligence Models in My Country (In-Depth Study and Implementation of Xi Jinping Thought on Socialism with Chinese Characteristics for a New Era: Experts' Discussions on Volume 1 of "Selected Works of Xi Jinping on the Economy")], *People's Daily*, May 8, 2025. [https://web.archive.org/web/20250509065808/http://paper.people.com.cn/rmrb/pc/content/202505/08/content\\_30071810.html](https://web.archive.org/web/20250509065808/http://paper.people.com.cn/rmrb/pc/content/202505/08/content_30071810.html).
- <sup>71</sup> China's National Development and Reform Commission, National Data Administration, Ministry of Finance, and Ministry of Human Resources and Social Security, 国家发展改革委等部门关于促进数据标注产业高质量发展的实施意见, [Implementation Opinions of the National Development and Reform Commission and Other Ministries on Promoting the High-Quality Development of the Data Labeling Industry], March 25, 2025. <https://cset.georgetown.edu/publication/china-ndrc-data-labeling-opinions/>.
- <sup>72</sup> China's National Data Bureau, 国家公共数据资源登记平台上线运行 [National Public Data Resource Registration Platform Launched], March 1, 2025. [https://www.gov.cn/lianbo/bumen/202503/content\\_7009412.htm](https://www.gov.cn/lianbo/bumen/202503/content_7009412.htm).
- <sup>73</sup> "国家数据局: 国内多数模型训练使用中文数据占比超60%" [National Data Administration: The Majority of Domestic Model Training Uses Chinese Data, Accounting for over 60%], *People's Daily*, August 19, 2025. [https://web.archive.org/web/20251209201309/https://www.gov.cn/lianbo/bumen/202508/content\\_7037033.htm](https://web.archive.org/web/20251209201309/https://www.gov.cn/lianbo/bumen/202508/content_7037033.htm).
- <sup>74</sup> Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," *RAND*, June 26, 2025. <https://www.rand.org/pubs/perspectives/PEA4012-1.html>; China's National Development and Reform Commission, National Data Administration, Ministry of Finance, and Ministry of Human Resources and Social Security, 国家发展改革委等部门关于促进数据标注产业高质量发展的实施意见, [Implementation Opinions of the National Development and Reform Commission and Other Ministries on Promoting the High-Quality Development of the Data Labeling Industry], March 25, 2025. <https://cset.georgetown.edu/publication/china-ndrc-data-labeling-opinions/>.
- <sup>75</sup> U.S.-China Economic and Security Review Commission, Chapter 1, "U.S.-China Economic and Trade Relations (Year in Review)," in *2024 Annual Report to Congress*, November 2024, 43.

- <sup>76</sup> Peter Belcak et al., "Small Language Models Are the Future of Agentic AI," *arXiv*, June 2025. [arxiv.org/abs/2506.02153](https://arxiv.org/abs/2506.02153).
- <sup>77</sup> Peter Belcak et al., "Small Language Models Are the Future of Agentic AI," *arXiv*, June 2025. [arxiv.org/abs/2506.02153](https://arxiv.org/abs/2506.02153).
- <sup>78</sup> "omni-research/Tarsier2-Recap-7b," *Hugging Face*, accessed December 8, 2025. <https://huggingface.co/omni-research/Tarsier2-Recap-7b>.
- <sup>79</sup> Adina Yakefu and Irene Solaiman, "The Future of the Global Open-Source AI Ecosystem: From DeepSeek to AI+," *Hugging Face*, February 3, 2026. <https://huggingface.co/blog/huggingface/one-year-since-the-deepseek-moment-blog-3>.
- <sup>80</sup> Kyle Miller, Mia Hoffmann, and Rebecca Gelles, "The Use of Open Models in Research," *CSET*, October 2025. <https://cset.georgetown.edu/wp-content/uploads/CSET-The-Use-of-Open-Models-in-Research.pdf>; "georgetown-cset/open-models," *GitHub*, accessed December 8, 2025. <https://github.com/georgetown-cset/open-models>.
- <sup>81</sup> Haoran Wei et al., "Small Language Model Meets with Reinforced Vision Vocabulary," *arXiv*, January 23, 2024. <https://arxiv.org/abs/2401.12503>.
- <sup>82</sup> Tongtong Feng et al., "Embodied AI: From LLMs to World Models," *arXiv*, September 24, 2025, 10. <https://arxiv.org/pdf/2509.20021>.
- <sup>83</sup> "InternRobotics/InternVLA-M1," *Hugging Face*, accessed December 8, 2025. <https://huggingface.co/InternRobotics/InternVLA-M1>; "Intern Robotics," *GitHub*, accessed December 8, 2025. <https://github.com/InternRobotics>; "Intern Robotics," *Intern Robotics*, accessed December 8, 2025. <https://archive.ph/biGwB>; "InternVLA-M1," *GitHub*, accessed December 8, 2025. <https://internrobotics.github.io/internvla-m1.github.io/>.
- <sup>84</sup> William Hannas et al., "China's Embodied AI: A Path to AGI," *Center for Strategic and Emerging Technologies*, December 2025, 4–5. <https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-Embodied-AI.pdf>.
- <sup>85</sup> William Hannas et al., "China's Embodied AI: A Path to AGI," *Center for Strategic and Emerging Technologies*, December 2025, 4–5. <https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-Embodied-AI.pdf>; William Hannas and Huey-Meei Chang, "China's Artificial General Intelligence," *Center for Security and Emerging Technology*, August 29, 2025. <https://cset.georgetown.edu/article/chinas-artificial-general-intelligence/>.
- <sup>86</sup> Huang Haihua, "视频|蒲慕明：目前最需要的科学家，不是只顾“钓鱼之乐”，而是“要钓大鱼” [Mu-ming Poo: The Scientists Most Needed Today Are Those Who Are "Driven to Catch Big Fish," Not Those Who Only Care about "Fishing for the Joy of It"], *Liberation Daily*, March 26, 2025, 6. [https://cset.georgetown.edu/wp-content/uploads/t0633\\_mu-ming\\_poo\\_interview\\_EN.pdf](https://cset.georgetown.edu/wp-content/uploads/t0633_mu-ming_poo_interview_EN.pdf).
- <sup>87</sup> "通用人工智能关键在立'心'" [The Key to General Artificial Intelligence Lies in Establishing a "Mind"], *S&T Daily*, August 5, 2024. [https://www.ncsti.gov.cn/kjdt/kjrd/rqzn\\_kjrd/202408/t20240805\\_174149.html](https://www.ncsti.gov.cn/kjdt/kjrd/rqzn_kjrd/202408/t20240805_174149.html).
- <sup>88</sup> "World Model," *Nvidia*, accessed March 12, 2026. <https://www.nvidia.com/en-us/glossary/world-models/>; Tongtong Feng et al., "Embodied AI: From LLMs to World Models," *arXiv*, September 24, 2025, 9. <https://arxiv.org/pdf/2509.20021>.
- <sup>89</sup> Tongtong Feng et al., "Embodied AI: From LLMs to World Models," *arXiv*, September 24, 2025, 9. <https://arxiv.org/pdf/2509.20021>.
- <sup>90</sup> "中华人民共和国国民经济和社会发展第十五个五年规划纲要" [Outline of the People's Republic of China 15th Five-Year Plan for National Economic and Social Development], *Xinhua*, March 13, 2026. <https://www.news.cn/politics/20260313/085af5de5a4b4268aa7d87d90817df2f/c.html>.
- <sup>91</sup> "Global AI Adoption in 2025—A Widening Digital Divide," *Microsoft AI Economy Institute*, January 8, 2026, 5. <https://www.microsoft.com/en-us/corporate-responsibility/topics/ai-economy-institute/reports/global-ai-adoption-2025/>.
- <sup>92</sup> Austin Horng-En Wang and Kyle Siler-Evans, "U.S.-China Competition for Artificial Intelligence Markets," *RAND*, January 14, 2026. [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RRA4300/RRA4355-1/RAND\\_RRA4355-1.pdf](https://www.rand.org/content/dam/rand/pubs/research_reports/RRA4300/RRA4355-1/RAND_RRA4355-1.pdf).
- <sup>93</sup> "AICPB Rankings Methodology," *AICPB*, accessed March 12, 2026. <https://www.aicpb.com/methodology>.
- <sup>94</sup> Commission staff analysis based on AICPB's "Global AI Rankings by Users," accessed January 23, 2026. <https://www.aicpb.com>.
- <sup>95</sup> "Former Baidu Executives Launch Startup, Become Unicorn in One and a Half Years," *36Kr Europe*, November 23, 2025. <https://eu.36kr.com/de/p/3566544496180358>.
- <sup>96</sup> Charlie Campbell, "'We're Not That Far Behind.' Baidu's Robin Li on China's Push to Diffuse AI throughout Society," *TIME*, January 25, 2026. <https://time.com/7357630/robin-li-baidu-interview/>.
- <sup>97</sup> "The State of Chinese AI Apps 2025," *Tech Buzz China Insider*, October 17, 2025. <https://techbuzzchina.substack.com/p/the-state-of-chinese-ai-apps-2025>.
- <sup>98</sup> Rachel Metz, "OpenAI Releases Two 'Open' AI Models after DeepSeek's Success," *Bloomberg*, August 5, 2025. <https://www.bloomberg.com/news/articles/2025-08-05/openai-releases-open-weight-models-after-deepseek-s-success>; Victor Tangermann, "Sam Altman Regrets Ditching Open Source, Says He's Been on the 'Wrong Side of History,'" *Futurism*, February 3, 2025. <https://futurism.com/sam-altman-open-source-wrong-side-history>.
- <sup>99</sup> Kari Briski, "New NVIDIA Nemotron 3 Super Delivers 5x Higher Throughput for Agentic AI," *Nvidia*, March 11, 2026. <https://blogs.nvidia.com/blog/nemotron-3-super-agentic-ai/>; U.S. Securities and Exchange Commission, *Nvidia Corporation*, accessed March 12, 2026. <https://www.sec.gov/ix?doc=/Archives/edgar/data/1045810/000104581025000230/nvda-20251026.htm#fact-identifier-948>.
- <sup>100</sup> "The Llama 4 Herd: The Beginning of a New Era of Natively Multimodal AI Innovation," *Meta AI Blog*, April 5, 2025. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>.
- <sup>101</sup> "Meta Working on New AI Model Codenamed Avocado," *Bloomberg*, December 2025. <https://www.bloomberg.com/news/articles/2025-12-10/inside-meta-s-pivot-from-open-source-to-money-making-ai-model>.
- <sup>102</sup> Jordan Novet, "From Llamas to Avocados: Meta's Shifting AI Strategy Is Causing Internal Confusion," *CNBC*, December 9, 2025. <https://www.cnbcm.com/2025/12/09/meta-avocado-ai-strategy-issues.html>.
- <sup>103</sup> Kif Leswing, "Zuckerberg Says Meta Likely Won't Open Source All of Its 'Superintelligence' AI Models," *TechCrunch*, July 30, 2025. <https://techcrunch.com/2025/07/30/zuckerberg-says-meta-likely-wont-open-source-all-of-its-superintelligence-ai-models/>.
- <sup>104</sup> Jordan Novet, "From Llamas to Avocados: Meta's Shifting AI Strategy Is Causing Internal Confusion," *CNBC*, December 9, 2025. <https://www.cnbcm.com/2025/12/09/meta-avocado-ai-strategy-issues.html>.

- <sup>105</sup> Andy Haupt et al., "Inside Trump's Ambitious AI Action Plan," *Stanford HAI Institute*, July 24, 2025. <https://hai.stanford.edu/news/inside-trumps-ambitious-ai-action-plan>.
- <sup>106</sup> Andy Haupt et al., "Inside Trump's Ambitious AI Action Plan," *Stanford HAI Institute*, July 24, 2025. <https://hai.stanford.edu/news/inside-trumps-ambitious-ai-action-plan>; Kevin Xu, "Five Interesting Things about the White House AI Action Plan," *Interconnected* (blog), July 23, 2025. <https://interconnected.blog/five-interesting-things-about-the-white-house-ai-action-plan/>.
- <sup>107</sup> U.S. Department of Energy, *Energy Department Advances Investments in AI for Science*, December 10, 2025. <https://www.energy.gov/articles/energy-department-advances-investments-ai-science>; U.S. White House, *Launching the Genesis Mission*, November 24, 2025. <https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/>.
- <sup>108</sup> Eric Schmidt and Selina Xu, "Silicon Valley Is Drifting Out of Touch with the Rest of America," *New York Times*, August 19, 2025. <https://www.nytimes.com/2025/08/19/opinion/artificial-general-intelligence-superintelligence.html>.
- <sup>109</sup> Jack Shanahan and Kevin Frazier, "Stop Obsessing over AGI," *Center Square*, July 30, 2025. [https://www.thecentersquare.com/opinion/article\\_9e8cf0a6-4af1-4065-8f99-ea0bce2064ba.html](https://www.thecentersquare.com/opinion/article_9e8cf0a6-4af1-4065-8f99-ea0bce2064ba.html); U.S. Department of Defense, *Establishing of the Joint Artificial Intelligence Center*, June 27, 2018. [https://admin.govexec.com/media/establishment\\_of\\_the\\_joint\\_artificial\\_intelligence\\_center\\_osd008412-18\\_r....pdf](https://admin.govexec.com/media/establishment_of_the_joint_artificial_intelligence_center_osd008412-18_r....pdf).
- <sup>110</sup> "China's Minimax Reports Strong Revenue Growth, Charts Broader AI Ambitions," *Reuters*, March 2, 2026. <https://www.reuters.com/world/china/chinas-minimax-reports-strong-revenue-growth-charts-broader-ai-ambitions-2026-03-02/>; "China AI Startup Moonshot Targets \$10 Billion Valuation," *Bloomberg*, February 17, 2026. <https://www.bloomberg.com/news/articles/2026-02-17/china-ai-startup-moonshot-seeks-10-billion-value-in-new-funding>; Caroline Meinhardt et al., "Beyond DeepSeek: China's Diverse Open-Weight AI Ecosystem and Its Policy Implications," *Stanford University's Human-Centered Artificial Intelligence Institute*, December 2025, 14. <https://hai.stanford.edu/assets/files/hai-digichina-issue-brief-beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-policy-implications.pdf>; Esther Shittu, "China Startup Moonshot AI Rivals U.S. with Cheap Open Model," *Tech Target*, July 14, 2025. <https://www.techtarget.com/searchcenter-priseai/news/366627679/China-startup-Moonshot-AI-rivals-US-with-cheap-open-model>.
- <sup>111</sup> Rebecca Arcesati and Caroline Meinhardt, "China Bets on Open-Source Technologies to Boost Domestic Innovation," *Mercator Institute for China Studies*, May 19, 2021. <https://merics.org/en/report/china-bets-open-source-technologies-boost-domestic-innovation>.
- <sup>112</sup> Rebecca Arcesati and Caroline Meinhardt, "China Bets on Open-Source Technologies to Boost Domestic Innovation," *Mercator Institute for China Studies*, May 19, 2021. <https://merics.org/en/report/china-bets-open-source-technologies-boost-domestic-innovation>; "Microsoft in China: Clash of Titans," *CNN*, February 23, 2000. <https://web.archive.org/web/20081022160400/http://archives.cnn.com/2000/TECH/computing/02/23/microsoft.china.idg/>.
- <sup>113</sup> Sunny Cheung, "Open-Source Technology and PRC National Strategy: Part I," *Jamestown Foundation*, May 10, 2025. <https://jamestown.org/program/open-source-technology-and-prc-national-strategy-part-i/>.
- <sup>114</sup> Sunny Cheung, "Open-Source Technology and PRC National Strategy: Part I," *Jamestown Foundation*, May 10, 2025. <https://jamestown.org/program/open-source-technology-and-prc-national-strategy-part-i/>; Itsuro Fujino, "Huawei Breaks Free from Google Ecosystem with Homegrown OS," *Nikkei Asia*, May 8, 2024. <https://asia.nikkei.com/business/china-tech/huawei-breaks-free-from-google-ecosystem-with-homegrown-os>.
- <sup>115</sup> Manuel Torres, "China Is Building Its GitHub Alternative Gitee," *TechCrunch*, August 21, 2020. <https://techcrunch.com/2020/08/21/china-is-building-its-github-alternative-gitee/>.
- <sup>116</sup> Wendy Chang, Rebecca Arcesati, and Antonia Hmaid, "China's Drive toward Self-Reliance in Artificial Intelligence: From Chips to Large Language Models," *MERICs*, July 22, 2025. <https://merics.org/en/report/chinas-drive-toward-self-reliance-artificial-intelligence-chips-large-language-models>.
- <sup>117</sup> Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," *RAND*, June 26, 2025. <https://www.rand.org/pubs/perspectives/PEA4012-1.html#:~:text=and%20state%20agencies>.
- <sup>118</sup> Sunny Cheung, "Open-Source Technology and PRC National Strategy: Part I," *Jamestown Foundation*, May 10, 2025. <https://jamestown.org/program/open-source-technology-and-prc-national-strategy-part-i/>.
- <sup>119</sup> Mao Li, "从小红书破圈到DeepSeek崛起：中国互联网创新的国际突围路径" [From Xiaohongshu to the Rise of DeepSeek: The International Breakthrough Path of China's Internet Innovation], *S&T Daily*, February 21, 2025. [https://www.stdaily.com/web/gdxw/2025-02/21/content\\_299631.html](https://www.stdaily.com/web/gdxw/2025-02/21/content_299631.html).
- <sup>120</sup> Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," *RAND*, June 26, 2025. <https://www.rand.org/pubs/perspectives/PEA4012-1.html>.
- <sup>121</sup> "China Calls for Global AI Cooperation Days after Trump Administration Unveils Low-Regulation Strategy," *Guardian*, July 26, 2025. <https://www.theguardian.com/technology/2025/jul/26/china-calls-for-global-ai-cooperation-days-after-trump-administration-unveils-low-regulation-strategy>; "李强出席2025世界人工智能大会暨人工智能全球治理高级别会议开幕式并致辞" [Li Qiang Attended the Opening Ceremony of the 2025 World Artificial Intelligence Conference and the High-Level Conference on Global Governance of Artificial Intelligence and Delivered a Speech], *Xinhua*, July 26, 2025. [https://www.gov.cn/yaowen/liebiao/202507/content\\_7033942.htm](https://www.gov.cn/yaowen/liebiao/202507/content_7033942.htm); Li Qiang, "Address by China Premier Li Qiang to the Annual Meeting of New Champions 2025," *World Economic Forum*, June 25, 2025. <https://www.weforum.org/stories/2025/06/li-qiang-transcript-china-chinese-premier-summer-davos-amnc-world-economic-forum/>.
- <sup>122</sup> "李强主持召开国务院常务会议" [Li Qiang Presided over the State Council Executive Meeting], *People's Daily*, August 1, 2025. [https://www.gov.cn/yaowen/liebiao/202507/content\\_7033942.htm](https://www.gov.cn/yaowen/liebiao/202507/content_7033942.htm).
- <sup>123</sup> Matt Sheehan and Scott Singer, "How China Views AI Risks and What to Do ABOUT Them," *Carnegie Endowment for International Peace*, October 16, 2025. <https://carnegieendowment.org/research/2025/10/how-china-views-ai-risks-and-what-to-do-about-them>.
- <sup>124</sup> Insikt Group, "Measuring the US-China AI Gap," *Recorded Future*, May 8, 2025, 22. <https://assets.recordedfuture.com/insikt-report-pdfs/2025/ta-2025-0508.pdf>; Matthew S. Smith, "Morgan Stanley and Bank of America Are Focusing AI Power on Tools to Make Employees

More Efficient," *Business Insider*, May 2, 2025. <https://www.businessinsider.com/internal-artificial-intelligence-solutions-banking-companies-employee-training-2025-5>.

<sup>125</sup> "Open-Source vs. Closed-Source LLM Software: Unveiling the Pros and Cons," *Charter Global*, May 20, 2025. <https://www.charter-global.com/open-source-vs-closed-source-llm-software-pros-and-cons/>.

<sup>126</sup> "The State of Chinese AI Apps 2025," *Tech Buzz China Insider*, October 17, 2025. <https://techbuzzchina.substack.com/p/the-state-of-chinese-ai-apps-2025>.

<sup>127</sup> "The State of Chinese AI Apps 2025," *Tech Buzz China Insider*, October 17, 2025. <https://techbuzzchina.substack.com/p/the-state-of-chinese-ai-apps-2025>.

<sup>128</sup> Jeffrey Ding, "The Diffusion Deficit in Scientific and Technological Power: Re-Assessing China's Rise," working paper, August 2022, 32–34. <https://jeffreyding.github.io/documents/Diffusion%20Deficit%20working%20paper%20August%202022.pdf>.

<sup>129</sup> Amit Misra et al., "Measuring AI Diffusion: A Population-Normalized Metric for <sup>Tracking</sup> Global AI Usage," *arXiv*, November 4, 2020, 2, 5. <https://arxiv.org/pdf/2511.02781>.

<sup>130</sup> Arden Berg and Anson Ho, "After the ChatGPT Moment: Measuring AI's Adoption," *Epoch AI*, July 17, 2025. <https://epochai.substack.com/p/after-the-chatgpt-moment-measuring>.

<sup>131</sup> Austin Horng-En Wang and Kyle Siler-Evans, "U.S.-China Competition for Artificial Intelligence Markets," *RAND*, January 14, 2026, 3. [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RRA4300/RRA4355-1/RAND\\_RRA4355-1.pdf](https://www.rand.org/content/dam/rand/pubs/research_reports/RRA4300/RRA4355-1/RAND_RRA4355-1.pdf).

<sup>132</sup> "大模型产业：全链条突破 全场景落地" [Large Language Model Industry: Breakthroughs across the Entire Value Chain and Implementation in All Scenarios], *S&T Daily*, October 10, 2025. <https://archive.ph/L4rYh>; "AI Model Rankings," *Open Router*, accessed December 8, 2025. <https://openrouter.ai/rankings?view=trending>.

<sup>133</sup> "AICPB Rankings Methodology," *AICPB*, accessed March 12, 2026. <https://www.aicpb.com/methodology>.

<sup>134</sup> Austin Horng-En Wang and Kyle Siler-Evans, "U.S.-China Competition for Artificial Intelligence Markets," *RAND*, January 14, 2026, 4–5. [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RRA4300/RRA4355-1/RAND\\_RRA4355-1.pdf](https://www.rand.org/content/dam/rand/pubs/research_reports/RRA4300/RRA4355-1/RAND_RRA4355-1.pdf).