

AI infrastructure is becoming a new layer of soft power. While the AI stack splits into two ecosystems. – MAY 2026

U.S.

Consumer scale and frontier ceiling

CHINA

Efficiency and open-source dominance

AI RACE

SOFT POWER

NATION BRAND

When a student in Lagos trains on Qwen, a Berlin-based startup deploys DeepSeek, and an estimated 80% of open-source AI builders use Chinese base models, AI has become the most consequential geopolitical infrastructure since the internet.

RESULT

Bifurcated AI world. No clear winner

700M+

Monthly users of Chinese AI apps, inside China alone

900M

Weekly active users of ChatGPT globally (Feb 2026)

~35.7x

DeepSeek V3.2 costs \$0.14 per 1M input tokens; Claude Opus 4.7 costs \$5.

40%+

Qwen models make up 40% of new Hugging Face LLM derivatives.

Sources: QuestMobile / Tech in Asia (700M+ China AI app MAU, Sept 2025); OpenAI (ChatGPT 900M WAU, Feb 2026); DeepSeek & Anthropic API pricing (V3.2 vs Claude Opus 4.7); Hugging Face & arXiv "Measuring the Open Language Model Ecosystem" (Qwen ≈40% of new derivatives); a16z "Asserting American leadership in open-source AI" (~80% of open-source AI builders using Chinese open models).

CHART 1 OF 9 · AI COMPETITIVENESS OVERVIEW

CHINESE AI MODELS GAIN POPULARITY ON OPENROUTER

Weekly token usage – top models on OpenRouter (trillion tokens) · Jan 05–Apr 30 2026 · Source: FT chart data digitised; OpenRouter weekly estimates; live OpenRouter data on Apr 30 2026



Method note: combines FT chart data with OpenRouter data. Figures are reconstructed using an AI-assisted visual estimation method, so they should be treated as directional rather than exact audited figures.

CHART 1.2 · PERFORMANCE · SWE-BENCH %

THE CODING BENCHMARK GAP HAS NEARLY CLOSED

Best available model per country · 2023 Q4 – 2026 Q2



GPT-5.5: 88.7%

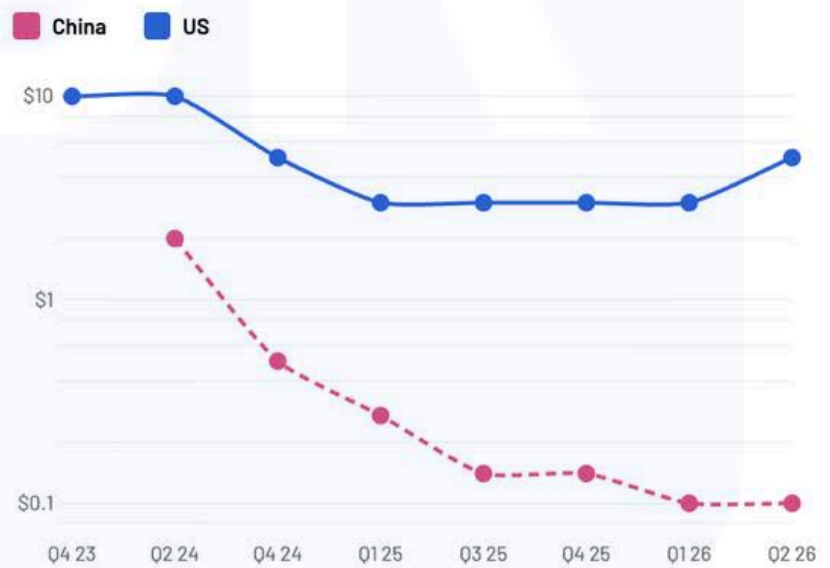
DeepSeek V4 Pro: 80.6%

Gap narrowed from 14pp to 8.1pp

CHART 1.3 · COST · \$ / 1M INPUT TOKENS

COST GAP REMAINS STRUCTURALLY WIDE

Best available model per country · 2023 Q4 – 2026 Q2



Step 3.5 Flash: \$0.10/M

GPT-5.5 / Claude Opus: \$5.00/M


Up to 50x input-price gap

CHART 2 OF 9 · CODING BENCHMARK

THE FRONTIER CEILING: US STILL LEADS, BUT BARELY

SWE-bench Verified scores — real GitHub issue resolution. The gap shrank from 20–30pp to roughly 5–8pp in 6 months.

China US



GPT-5.5 holds ceiling

DeepSeek V4 Pro: 83.7%


Gap: 20pp → 5pp in 6 months

CHART 3 OF 9 · API PRICING

THE PRICE WAR HAS ONE WINNER — AND IT ISN'T CLOSE

Input token pricing (\$/million tokens), as of early 2026 list prices. US frontier models remain materially more expensive, while Chinese models compete aggressively on price.

China US



DeepSeek V3.2: \$0.14/M input

Claude Opus 4.7: \$5/M input, \$25/M output


≈36× input gap · ≈89× output gap

CHART 4 OF 9 · OPEN-SOURCE STATUS

CHINA OWNS THE OPEN-WEIGHT STACK

License permissiveness vs SWE-bench score. Open weights at frontier performance — a combination US labs have not matched.

China · open-weight China · closed US · closed US · open-weight



MIT + Apache dominate open stack

No US · closed model below \$2/M

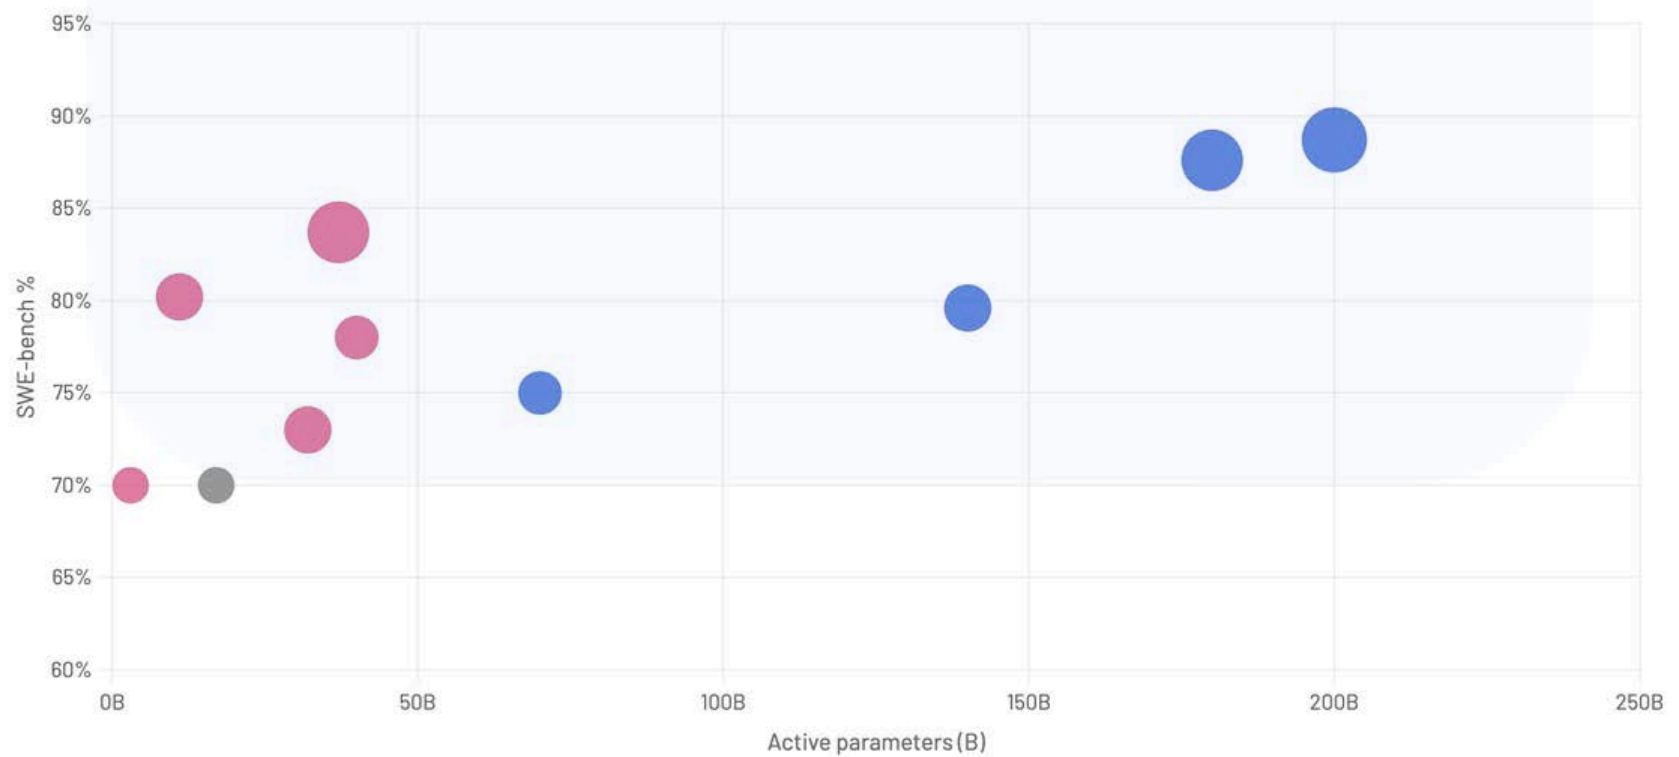
Qwen overtook Llama on HuggingFace

CHART 5 OF 9 · ARCHITECTURE EFFICIENCY

MOE: DOING MORE WITH LESS IS CHINA'S PLAYBOOK

Active parameters vs SWE-bench score. China models punch above their active parameter count.

China US US · open-weight



Kimi K2.6: 11B active → 80.2% SWE

Qwen3.6: 3B active → laptop-deployable

MoE = efficiency moat

CHART 6 OF 9 · CONVERGENCE TIMELINE

THE REASONING GAP NARROWED FROM 15PP TO 5.9PP

Best US vs best China GPQA % over time. The frontier reasoning gap compresses into single digits by 2026.



CHART 7 OF 9 · TRAINING ECONOMICS

DEEPSEEK BROKE THE "\$100M TO PLAY" ASSUMPTION

Estimated training costs. The efficiency gap is not catching up – it is a structural re-architecture of what frontier AI costs.

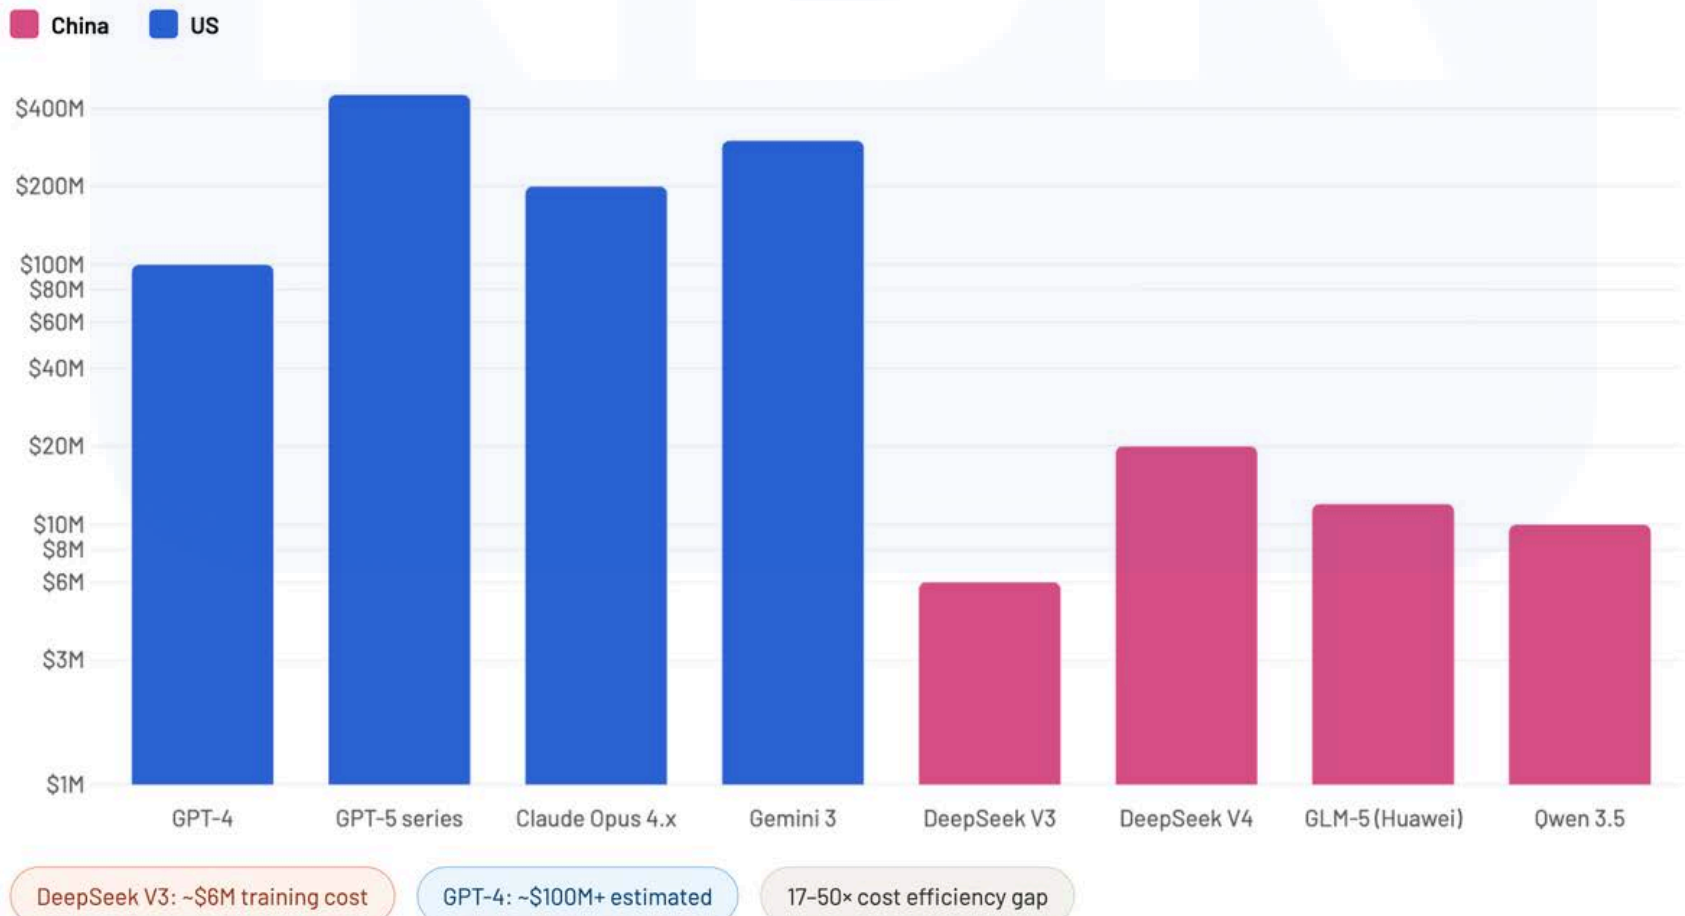



CHART 8 OF 9 · OPEN-SOURCE ECOSYSTEM

QWEN SURPASSED LLAMA. THE INFRASTRUCTURE LAYER SHIFTED.

Share of new HuggingFace LLM derivative models by base model family, using only the Excel values: Qwen 40%+ vs Llama 15%. Remaining share is calculated as other / mixed.

■ Qwen / Chinese base ■ Llama / US base ■ Other / mixed



■ Qwen / Chinese base ■ Llama / US base ■ Other / mixed

Qwen-derived: 40%+ of new derivatives

Llama: 15%

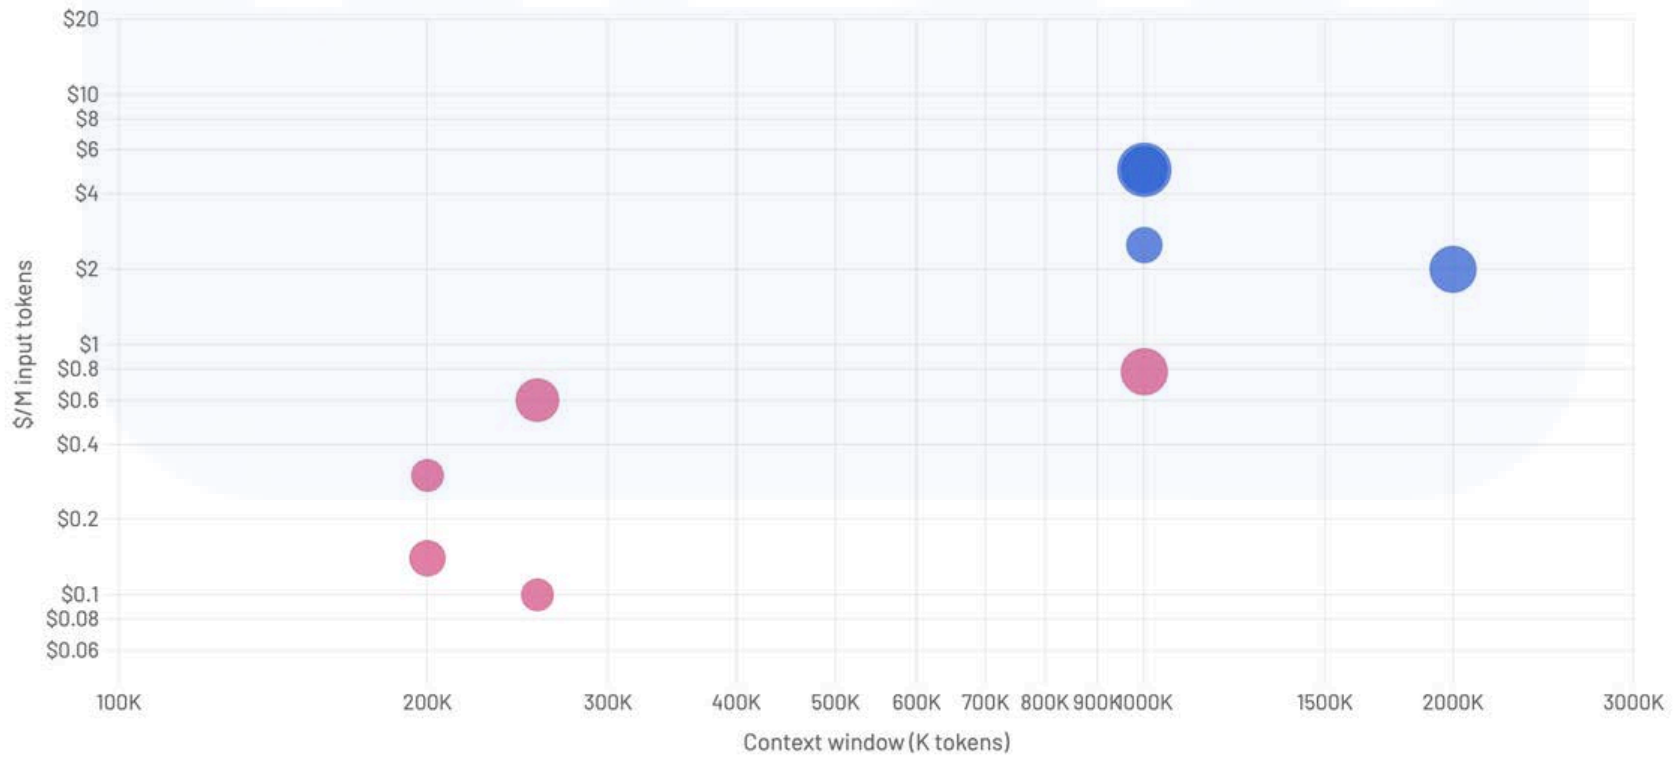
Other / mixed: ~45% remainder

CHART 9 OF 9 · CONTEXT WINDOW

CONTEXT WINDOW RACE: CHINA SCALES, US ANCHORS ON QUALITY

Context window size (K tokens) vs list input price (\$/million tokens), as of early 2026. Larger context at lower cost remains a structural Chinese advantage.

■ China ■ US



Qwen 3.6 Plus: 1M context @ \$0.78/M input

Claude Opus 4.7: 1M context @ \$5/M input

Grok 4.20: 2M context @ ~\$2/M input

Five ways the AI race reshapes global influence, none appear in benchmark tables

Performance gaps close. Pricing gaps widen. But the most durable geopolitical effects operate on a slower frequency: who trains the next cohort of developers, whose governance frameworks get adopted at the UN, and which model's defaults get embedded in a government ministry's AI stack in Nairobi or Jakarta.

01 – DISTRIBUTION

Open-weight models travel without visas

MIT and Apache 2.0 licenses remove every traditional barrier to geopolitical influence: no trade agreement, no diplomatic presence, no enterprise contract needed. A government ministry building on Qwen inherits its training data, its fine-tuning choices, and its content defaults – regardless of intent on either side.

40%+ of new Hugging Face LLM derivatives built on Chinese base models. Among developers building with open-source AI tools, 80% are using Chinese open-source models.

HuggingFace model metadata; a16z State of AI 2026

02 – ECONOMIC DEPENDENCY

Cost structures determine adoption – especially where capital is scarce

At \$0.14/M tokens versus \$5–15/M for US frontier models, the price differential is not a preference; it is a structural constraint for most of the world. DeepSeek adoption in Africa runs 2–4× higher than in Western markets, driven by cost elimination and existing Huawei telecom distribution infrastructure already in place across the continent.

DeepSeek Africa adoption is strong in the developing markets. China announced roughly \$50.7B Africa in 2024. Digital infrastructure and tech connectivity were part of the FOCAC 2024 cooperation agenda.

Microsoft AI Diffusion Report H2 2025; China-Africa FOCAC 2024

03 – REGULATORY CAPTURE

The governance race is being run at the UN – and China is showing up

In July 2024, China's AI capacity-building resolution passed the UN General Assembly unanimously with 143 co-sponsors. The US led a parallel resolution in March 2024, then declined to sign the Paris AI Summit declaration in February 2025. Governance frameworks adopted early become the template nations build on – especially those without the capacity to write their own.

143 nations co-sponsored China's UN AI resolution (Jul 2024). The US did not sign Paris AI Summit declaration (Feb 2025).

UNGA Resolution A/78/L.75; Lawfare Media, Jan 2026

04 – TALENT PIPELINE

The developer cohort being trained today will ship products for a decade

Huawei ICT Academy works with 3,500+ colleges in 110+ countries and has trained 1.8M+ students. UNESCO's G77+China project supports the development of digital and AI competencies across member states. The US equivalent – export-controlled, premium-priced – reaches fewer institutions in these regions.

1.3M students trained via Huawei ICT Academy globally. 50+ countries co-designing AI curricula in Chinese-anchored frameworks.

Huawei ICT Academy 2025 report; UNESCO G77+China initiative, Apr 2025

05 – GEOPOLITICAL STRATEGIC POSITIONING

AI is the new layer of the Digital Silk Road – and it compounds existing infrastructure dependencies

China's 15th Five-Year Plan (2026–2030) explicitly frames AI as an instrument of “normative leadership abroad,” coupling domestic capability build-out with standard-setting and market access across developing markets. Digital Silk Road extends through Belt-and-Road-linked markets – AI models now layer on top of existing infrastructure dependencies. The strategic logic is compounding: a country that runs on Huawei telecom, trains developers on Huawei ICT Academy tools, and deploys Qwen-based applications has built a stack where switching costs rise with every layer added.

The US position is structurally different: strong on consumer scale (900M ChatGPT WAU), strong on frontier capability, and strong on enterprise trust in regulated markets. But its engagement model – export-controlled hardware, premium API pricing, and an increasingly transactional approach to Global South partnerships – limits its footprint at the infrastructure and education layers where long-term influence is accumulated. The Paris 2025 abstention was read by many non-aligned nations not as a technical disagreement, but as a signal of disengagement from multilateral AI governance.

China's 15th FYP (2026–30) explicitly targets AI standard-setting and market access across developing markets as a policy goal. Digital Silk Road active in ~150 countries – AI models now layer on top of existing infrastructure dependency.

CRS Analysis of China's 15th Five-Year Plan, Apr 2026; East Asia Forum Digital Silk Road analysis, Apr 2025

LOOKING AHEAD

The United States continues to lead in frontier model performance and private investment, with tech giants committing hundreds of billions to AI development. But scaling is increasingly constrained by a “power wall.” With data center electricity demand projected to surge by 2030, the challenge is shifting from building better models to sustaining them at scale. This is not a permanent constraint, but it is a real near-term bottleneck as energy infrastructure, permitting, and grid modernization struggle to keep pace with AI growth.

At the same time, access to advanced semiconductors is becoming the decisive choke point. Cutting-edge AI depends on a small number of highly specialized chips, and control over their design and supply chains is now as strategic as oil once was. The U.S. maintains a clear advantage here through its ecosystem and export controls, but this also globalizes the competition by forcing alternative pathways.

China’s response is not simply to catch up, it is to reconfigure the game. Constrained in high-end chips, it is scaling through energy availability, industrial policy, and system-level efficiency. By rapidly expanding renewable and nuclear capacity, China can deploy compute at lower marginal cost, allowing it to compensate for hardware gaps while embedding AI deeply into manufacturing, logistics, and robotics. This is less about brute force alone and more about optimization at scale.

A parallel divergence is emerging in ecosystem strategy. The U.S. largely captures value through proprietary, high-performance models and elite research networks. China, by contrast, is pushing more open and cost-efficient systems outward, particularly across developing markets, to drive

adoption, standards, and long-term dependency. With a large and increasingly self-sufficient talent pipeline, its focus is tilting toward applied and embodied AI, especially in physical systems.

What is emerging is a layered competition. The U.S. is likely to remain dominant in frontier “digital brains”: high-end intelligence, advanced models, and breakthrough research. China is positioning itself to control more of the “digital plumbing”: the infrastructure, industrial integration, and deployment layers through which AI actually scales in the real world.

Other players, including India and blocs like the European Union, will not fully align with either system, instead hedging across both while building selective autonomy.

In developing markets, AI-driven global influence will be determined primarily by control over deployment ecosystems rather than leadership in frontier innovation. Influence increases to the extent that AI systems are embedded into core functions such as public services, finance, logistics, and education, creating structural dependency through everyday use. In this context, affordability, adaptability, and ease of integration will matter more than absolute model performance.

About NationBrandResearch (NBR)

NationBrandResearch is a global research initiative dedicated to understanding and advancing how countries build, maintain, and enhance their competitive positioning and soft power influence in an increasingly interconnected world.

We provide evidence-based research, data-driven insights, and strategic recommendations to governments, institutions, and organizations seeking to understand and shape global perceptions of their nations.

LEARN MORE: WWW.NATIONBRANDRESEARCH.COM

SLIDE #1

700M+ Chinese AI mobile apps reached 700M+ monthly active users (Sept 2025)

Source: Tech in Asia / QuestMobile, SCMP

900M ChatGPT weekly active users globally (Feb 2026)

Source: OpenAI official, TechCrunch confirmed

DeepSeek up to ~90× cheaper on output, ~36× cheaper on input vs Claude Opus 4.7

Source: DeepSeek API docs vs Anthropic official pricing

40%+ Qwen-based derivatives now 40%+ of new HuggingFace models (Mar 2026)

Source: arXiv: "Measuring the Open Language Model Ecosystem"

80% a16z: 80% of open-source AI builders use Chinese open models

Source: Andreessen Horowitz, Martin Casado

CHART 1 – AI COMPETITIVENESS OVERVIEW

Weekly OpenRouter token usage, SWE-bench performance trend, and input-price trend, Jan 05–Apr 30 2026.

Source: Nation Brand Research; OpenRouter weekly API usage data; FT chart data digitised; SWE-bench benchmark trend data; DeepSeek, Anthropic, and OpenAI pricing.

CHART 2 – THE FRONTIER CEILING: US STILL LEADS, BUT BARELY

SWE-bench Verified scores for selected leading US and Chinese models. Independent benchmark runs show the US frontier still ahead, with the gap now in single-digit percentage points.

Source: Nation Brand Research, SWE-bench benchmark data (Q3 2024 – Q1 2026)

CHART 3 – THE PRICE WAR HAS ONE WINNER

Input token pricing (\$/million tokens), based on early-2026 list prices.

Source: Nation Brand Research; DeepSeek API pricing; Anthropic pricing; OpenAI pricing; xAI/Grok pricing trackers where official/current pricing is unavailable.

CHART 4 – CHINA OWNS THE OPEN-WEIGHT STACK

License permissiveness vs SWE-bench score. Open weights at frontier performance—a combination US labs have not matched.

Source: Nation Brand Research, HuggingFace model registry analysis

CHART 5 – MOE: DOING MORE WITH LESS IS CHINA'S PLAYBOOK

Active parameters vs SWE-bench score. China models punch above their active parameter count.
Source: Nation Brand Research, model architecture & benchmark data analysis (Q1 2026)

CHART 6 – THE REASONING GAP NARROWED FROM 15PP TO 5.9PP

Best US vs best China GPQA % over time.
Source: Nation Brand Research, GPQA benchmark trend data from Excel table.

CHART 7 – DEEPSEEK BROKE THE "\$100M TO PLAY" ASSUMPTION

Estimated training costs. The efficiency gap is not catching up—it is a structural re-architecture of what frontier AI costs.
Source: Nation Brand Research, estimated training cost analysis (DeepSeek V3: ~\$6M, GPT-4: ~\$100M+ estimated); 17-50x efficiency gap

CHART 8 – QWEN SURPASSED LLAMA. THE INFRASTRUCTURE LAYER SHIFTED.

Share of new Hugging Face LLM derivative models by base model family, using only the Excel values: Qwen-derived 40%+, Llama 15%, and other/mixed as the calculated remainder.
Source: Nation Brand Research; Excel model-ecosystem table; arXiv "Measuring the Open Language Model Ecosystem."

CHART 9 – CONTEXT WINDOW RACE: CHINA SCALES, US ANCHORS ON QUALITY

Context window size (K tokens) vs API input price. Larger context at lower cost is another Chinese structural advantage.
Source: Nation Brand Research, official API documentation comparison (DeepSeek, Claude, Grok pricing as of Feb 2026)

DISTRIBUTION

- "Qwen derivatives account for 40%+ of new HuggingFace language-model derivatives." (HuggingFace State of Open Source Spring 2026; arXiv "Measuring the Open Language Model Ecosystem")
- "a16z estimates that among developers building with open-source AI tools, 80% are using Chinese open-source models." (Andreessen Horowitz, 2026)

ECONOMIC DEPENDENCY

- "DeepSeek V3.2 input pricing is \$0.14/M tokens, compared with Claude Opus 4.7 at \$5/M input and \$25/M output. GPT-5.5 input pricing varies by tier, with long-context input around \$5/M."
- "DeepSeek and other low-cost Chinese models are positioned strongly in developing markets because lower API costs reduce adoption barriers, especially where capital is scarce."

REGULATORY CAPTURE

- "China's AI capacity-building resolution adopted by UN General Assembly July 2024 with 143 co-sponsors." (UNGA Resolution A/78/L.79)
- "US and UK did not sign Paris AI Summit declaration on inclusive AI, February 2025." (Reuters, Feb 2025)

TALENT PIPELINE

- "Huawei ICT Academy works with 3,500+ colleges in 110+ countries, training 1.8M+ students." (Huawei official, 2025-26)
- "UNESCO's G77+China project supports digital and AI competency development across member states." (UNESCO official)

GEOPOLITICAL POSITIONING

- "China's 15th Five-Year Plan positions AI, digital infrastructure, and standards as strategic priorities. The Digital Silk Road extends through Belt-and-Road-linked markets, creating channels for AI infrastructure influence."
- "Chinese AI, including low-cost and open-weight models such as DeepSeek and Qwen, may extend the Digital Silk Road into AI infrastructure and standard-setting."