The AI Gold Rush

Entering a second stage

→ While frontier models remain strategically important, competitive advantage is increasingly shaped by access to compute, data centers, energy, distribution networks, financing capacity, and operational deployment.

→ Meanwhile, as Silicon Valley pours hundreds of billions into AI infrastructure, Europe risks remaining an AI consumer rather than an owner of the platforms and networks generating most of its economic value.

Introduction

For the last two years, the AI debate has focused on which model is best: OpenAI's GPT, Anthropic's Claude, Google's Gemini, Elon Musk's Grok, Meta's Llama, or China's DeepSeek. Investors, policymakers, and the media have treated AI as a technological race in which the best model would capture the greatest value.

Silicon Valley's AI race is increasingly being shaped by a broader ecosystem extending well beyond software

One of the most important current developments in Silicon Valley is AI's evolution from a predominantly software-driven business into a hybrid ecosystem where infrastructure, compute, energy, and deployment capabilities are becoming increasingly important.[1]

Technological revolutions do not automatically reward infrastructure owners. In the internet era, software and platform companies captured enormous value while telecommunications operators struggled despite owning critical infrastructure.[2] The broader lesson is that value tends to accrue to firms controlling scarce bottlenecks in the technological ecosystem—whether infrastructure, software platforms, distribution, data, or customer access.

In short, the AI gold rush is entering its second phase. The first rewarded innovation, while the next will likely reward firms that combine leading models with the infrastructure, financing capacity, distribution, and operational capabilities needed to deploy AI at scale.

Why OpenAI had to change

Nothing illustrates this transformation better than OpenAI. When it was founded in 2015, OpenAI presented itself as a research organization focused on developing AI for the benefit of humanity. Today, it is one of the most valuable private companies in the world and is widely expected to pursue a public listing.

OpenAI's evolution is not only philosophical. It also reflects the changing economics of frontier AI. As training and inference costs rise, access to compute, energy, cloud capacity, and financing becomes a strategic constraint.

Frontier AI has become extraordinarily expensive. The largest frontier model training runs now require enormous spending on chips, energy, cloud capacity, and engineering talent, with independent estimates suggesting future leading training runs could exceed $1 billion.[3]

In the next wave of AI, value shifts toward whoever controls the scarce bottlenecks of the AI stack

Developing the next generation of systems increasingly resembles financing a large industrial project rather than launching a software company. This is why access to public markets and financing capacity is becoming almost as important as access to innovation, and scale is becoming a competitive advantage.[4]

OpenAI is not alone. Anthropic, OpenAI's closest rival in the frontier model segment, has also seen its valuation soar and has recently filed for an IPO. That the two leading independent AI developers are both being discussed less as technology start-ups and more as future public-market giants illustrates how rapidly the sector is evolving.

Elon Musk's integrated ecosystem strategy

Musk's AI strategy extends beyond the model itself to control key parts of the AI ecosystem

This changing landscape also helps explain Elon Musk's strategy. Unlike most competitors, Musk appears to be pursuing a more vertically integrated AI strategy. Through xAI, X, Starlink, SpaceX, Tesla's autonomous driving capabilities, and robotics, he is assembling assets across data, distribution, connectivity, and physical infrastructure. This may offer strategic advantages through tighter integration across the AI stack, but it also carries significant execution, governance, and valuation risks. The model should therefore be presented as a high-risk integrated ecosystem strategy, not as a proven template.

SpaceX is an important part of that strategy. The company is no longer simply a space-launch provider. Through Starlink, it operates one of the world's largest satellite communications networks, providing connectivity to businesses, consumers, governments, and military users around the globe. Together with xAI, X, and Tesla, it forms part of an ecosystem that combines data, connectivity, computing, and distribution. This illustrates how the next phase of AI competition may depend less on individual models and more on control of the infrastructure that allows AI systems to operate at scale.

SpaceX's IPO has reinforced this advantage. According to a June 3 filing, the company planned to raise up to $86 billion, the world's largest IPO, and would prioritize its proceeds toward expanding its AI infrastructure.[5] On June 12, the company raised approximately $75 billion, making it the largest IPO in history and valuing SpaceX at more than $2 trillion after its market debut.[6]

Whether this integrated approach ultimately proves superior remains uncertain. Recent volatility in AI-related equities suggests that investors are increasingly scrutinizing the profitability of large-scale infrastructure investments and the sustainability of AI-related valuations.[7]

The model wars are becoming less decisive

OpenAI remains the dominant brand in AI, but the gap between competitors has narrowed considerably. Benchmark rankings continue to change rapidly,

with leadership rotating across reasoning, coding, enterprise, and multimodal capabilities. This suggests that technological advantages may prove less durable than many investors initially assumed.

As frontier models converge, competitive advantage is increasingly shifting beyond the model itself to include complementary assets

This does not mean models are becoming irrelevant. Frontier models remain strategically important, as illustrated by national security debates and export control measures governing advanced AI capabilities and model weights. The U.S. government's June 12 issuance of an export control directive banning foreign nationals from accessing Anthropic's flagship Fable 5 and Mythos 5 AI models underscores that leading models remain important strategic assets.[8] In fact, the argument is narrower: as performance gaps compress, long-term advantage may increasingly depend on complementary assets surrounding the model.

Anthropic has established itself as a credible challenger. Google has recovered ground through Gemini. Meta continues to invest aggressively in open-source models. And Chinese firms are advancing rapidly despite restrictions on access to advanced chips.

As performance differences between models diminish, competitive advantage increasingly shifts elsewhere. Scale, distribution, computing capacity, and financing capacity become more important than marginal improvements in intelligence.

The industry increasingly resembles cloud computing: several players may succeed, but only those able to sustain massive investment can compete.

Can AI become a profitable business?

Yet perhaps the most important question facing industry is whether AI can become a sustainably profitable business.

While revenues are growing rapidly, costs remain exceptionally high. Model developers must finance training, inference, data acquisition, and talent while competing aggressively on pricing. Infrastructure providers, by contrast, generate revenue regardless of which model ultimately wins. This distinction may prove critical.

The real winners may be the hyperscalers

AI infrastructure may generate the greatest long-term value, but its profitability remains unproven

The hyperscalers may be among the largest beneficiaries, but the profitability of the AI infrastructure cycle is not yet proven. High capex creates both strategic barriers to entry and financial risk if AI revenues fail to scale fast enough.

Amazon, Microsoft, Google, and Meta are collectively committing hundreds of billions of dollars to AI-related capital expenditures in data centers, networking equipment, semiconductors, and energy systems, reinforcing their position as the indispensable backbone of the AI ecosystem.[9] Their capital expenditure plans increasingly resemble those of utilities or industrial conglomerates rather than traditional technology firms.

These firms possess something that many AI developers lack: the ability to finance massive infrastructure investments from their own balance sheets.

Amazon Web Services remains the dominant cloud platform. Microsoft benefits from its partnership with OpenAI. Google combines leading AI research with one of the world's largest cloud businesses. Meta is investing aggressively in both open-source models and computing capacity.

The relationship between model developers and infrastructure providers is therefore becoming increasingly symbiotic. OpenAI, Anthropic, and xAI may create applications that attract users and headlines, but their growth increasingly depends on the computing power, cloud capacity, and financing provided by hyperscalers.

As in the California gold rush, the most reliable profits accrue to those selling the picks and shovels.

U.S.-China competition

This broader framing is especially important when comparing the United States and China.

The U.S. approach remains heavily market driven. The state plays a role through export controls, defense procurement, research funding, and industrial policy, but the leading actors are private firms.

China may not need to dominate every frontier model category to become a formidable AI competitor

China is pursuing a different model. Its AI strategy is better understood as an ecosystem strategy.[10] It combines state support, domestic chip development, cloud infrastructure, robotics, industrial automation, open models, manufacturing deployment, and the integration of AI into strategic sectors. Even if Chinese firms remain constrained in accessing the most advanced chips, China’s industrial base gives it a distinct advantage: the ability to deploy AI at scale across factories, logistics systems, robotics platforms, electric vehicles, and manufacturing supply chains.

This matters because the future of AI will be determined not only by benchmark scores but also by where AI is used, how quickly it is embedded in production, and whether firms can turn intelligence into productivity. China may not need to dominate every frontier model category to become a formidable AI competitor. Instead, it may build an integrated industrial AI ecosystem that links software, hardware, automation, and state-backed deployment.

Europe’s risk of missing the second phase

Europe has world-class researchers and industrial firms, but it remains dependent on foreign cloud providers, chips, and platforms.[11] At the same time, the United States dominates hyperscale cloud computing, while China builds a parallel ecosystem.

This matters because technological revolutions rarely reward only innovators. They often reward owners of the underlying infrastructure.

If Europe remains primarily a consumer rather than an owner of AI infrastructure, it risks capturing only a fraction of the economic value generated by technology.

The challenge is technological and strategic. Europe's objective should not aim to copy OpenAI. Its more credible opportunity lies in anchoring AI value in industrial deployment, sovereign compute, applied data, energy systems, and sector-specific operational AI.

Although Europe is unlikely to outspend the United States or China on frontier infrastructure, it has stronger opportunities in industrial AI, given its strengths in manufacturing, aerospace, energy, automotive, logistics, and engineering.

Although Europe is unlikely to outspend the United States or China on frontier infrastructure, it has a more credible opening in industrial AI, where its strengths still matter

Conclusion

Between models and applications sits an emerging operational layer that connects AI to real-world organizations, workflows, and decision-making. For example, Palantir is building systems that connect AI to operational processes, governance, and enterprise decision-making.[12] As model performance converges, this layer could become increasingly important.

The first phase of AI rewarded innovation. The second phase is likely to reward firms that combine strong models with the infrastructure, financing capacity, operational systems, distribution networks, and execution capabilities needed to deploy AI at scale. The key question is not whether infrastructure will replace software, but which firms will control the most valuable bottlenecks across the AI value chain.

The next phase of AI is likely to reward firms that combine innovation with scale, capital, and execution

EXHIBIT 2.0 – CUMULATIVE VENTURE CAPITAL INVESTMENT IN AI (MILLIONS OF USD, 2012-2024)

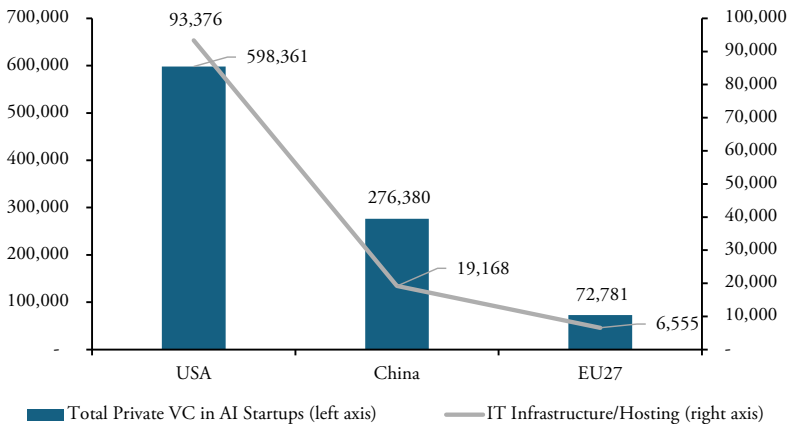

Source: U.S. Federal Reserve.

Notes

[1] https://www.iea.org/reports/energy-and-ai/executive-summary

[2] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-critical-bets-on-the-future-of-telco-value-creation

[3] https://hai.stanford.edu/ai-index/2026-ai-index-report/research-and-development

[4] https://openai.com/index/five-new-stargate-sites/

[5] https://www.ft.com/content/95286aa1-ee55-4a45-86b8-951cc0fb23a8?syn-25a6b1a6=1

[6] https://www.reuters.com/world/china/big-techs-635-billion-ai-spending-faces-energy-shocktest-sp-global-says-2026-03-31/

[7] https://www.wsj.com/finance/stocks/tech-stocks-retreat-on-concerns-over-ai-spending-valuations-9776df5a?eafs_enabled=false

[8] https://www.anthropic.com/news/fable-mythos-access

[9] https://www.reuters.com/world/china/big-techs-635-billion-ai-spending-faces-energyshocktest-sp-global-says-2026-03-31

[10] https://merics.org/en/report/chinas-drive-toward-self-reliance-artificial-intelligencechipslarge-language-models

[11] https://commission.europa.eu/topics/competitiveness/ai-continent_en

[12] https://www.palantir.com/docs/foundry/aip/overview