

DeepSeek

Pre-IPO opportunity in China's leading AI laboratory

\$10B

ROUND VALUATION

\$300M+

ROUND SIZE

2023

FOUNDED

China

JURISDICTION

SECTION 01

COMPANY SNAPSHOT

Company	DeepSeek (Hangzhou DeepSeek Artificial Intelligence Co., Ltd.)
Description	Chinese AI lab developing large language, reasoning, and code-generation models.
Founded	2023, Hangzhou, China. Founder: Liang Wenfeng (CEO of High-Flyer Capital Management, AUM ~\$14B).
Ownership	Liang Wenfeng holds 84%. This is the company's first external funding round since inception.

SECTION 02

DEEPSEEK AMONG LEADING AI COMPANIES

COMPANY	COUNTRY	VALUATION (MAY 2026)	REVENUE (ARR)
OpenAI	USA	\$852B	\$24B+
Anthropic	USA	\$380B	\$30B
xAI (Grok)	USA	\$50B+	~\$1B
Perplexity	USA	\$20B	~\$100M
Mistral	France	\$14B	N/A
DeepSeek	China	\$10B	Not disclosed

Source: TLDL AI Company Rankings 2026, Syntax Research, CNBC (April 2026)

KEY TAKEAWAY

DeepSeek is the only major AI player outside the US and EU with a valuation comparable to leading second-tier Western companies.

SECTION 03

INVESTMENT THESIS

DeepSeek has demonstrated that powerful AI models can be built and operated at a fraction of the cost incurred by competitors. The company is positioning itself as the primary supplier of affordable AI for enterprise use, automation, and government-scale AI infrastructure across China and aligned markets.

Core thesis	If OpenAI and Anthropic are premium models for cases where quality is non-negotiable, DeepSeek is the "good-enough and dramatically cheaper" option for the majority of real-world tasks.
Analogy	Just as Chinese smartphone makers captured the mid-tier market without displacing premium brands, DeepSeek can occupy a parallel position in AI.

SECTION 04

WHY NOW — MARKET SHIFT

The cost of running AI models is falling at an unprecedented rate:

- The price of a query to a GPT-3.5-class model fell **280×** between November 2022 and October 2024.
Source: Stanford AI Index 2025 — hai.stanford.edu/ai-index/2025-ai-index-report
- The price of a query to a GPT-4-class model fell from \$60 to under \$1.50 per million tokens — a **40×+ reduction** in two years.
*Source: arxiv.org/html/2511.23455v2 — *The Price of Progress* (2026)*
- The AI inference market: **\$106B in 2025** → **\$255B by 2030** (19% CAGR). Inference now accounts for 55% of cloud AI spend, having overtaken training costs.
Source: Grand View Research and Markets and MarketsandMarkets, cited in software.seni.com (March 2026)

IMPLICATION

In a world where inference cost falls roughly 10× every 1–2 years, the ability to build and serve cheap models becomes a durable competitive advantage.

SECTION 05

TECHNOLOGY ADVANTAGE — THE FACTS

Training cost

- V3 required 2.788M GPU-hours — approximately **\$5.6M** at H800 rental rates. GPT-4 training is estimated at \$50–100M, a **10–20× difference**.
Source: BentoML — bentoml.com/blog/the-complete-guide-to-deepseek-models
- V3 was trained on 2,048 H800 GPUs over roughly two months. Including R1, total combined cost was approximately **\$5.87M**.
Source: The Register — theregister.com/2025/09/19/deepseek_cost_train
- Reported training costs are 95% lower than industry leaders. The MoE architecture activates only 37B of 671B parameters per query.
Source: LocalAI Master — localaimaster.com/models/deepseek-v3-vs-v3-1-analysis

CAVEAT

According to The Register, the cost advantage may be overstated — Meta trained the comparable Llama 4 Maverick using a similar volume of GPU-hours but on a substantially larger dataset.

Source: theregister.com/2025/09/19/deepseek_cost_train

Gross margin

Structural gross margin on inference is reported at **85%**, higher than Anthropic (55%), Together AI (45%), and Groq (40%). This is achieved through MoE architecture rather than subsidization.

Source: SoftwareSeni — software.seni.com/the-ai-inference-market-in-2025

Latest models (April 2026)

DeepSeek-V4-Pro: 1.6 trillion parameters, 1M-token context window, MIT License. Costs **6–7× less** than GPT-5.5 and Claude Opus 4.7 with comparable benchmark performance.

Source: VentureBeat — venturebeat.com/technology/deepseek-v4-arrives

Cost-reduction techniques

- | | |
|---------------------------|--|
| Mixture-of-Experts | Only 37B of 671B parameters activate per query — fewer computations while maintaining quality. |
| Sparse Attention | DeepSeek's attention mechanism skips low-value computations on long contexts. |
| RL training | The reinforcement learning approach behind R1 produced a reasoning model comparable to o1 at a fraction of the cost. |

PRICING ADVANTAGE

MODEL	INPUT (PER 1M TOKENS)	OUTPUT (PER 1M TOKENS)	LICENSE
DeepSeek-V3	\$0.28	\$1.10	MIT (open)
DeepSeek-R1	\$0.55	\$2.19	MIT (open)
GPT-5 (OpenAI)	\$1.25	\$10.00	Closed
OpenAI o1	\$15.00	\$60.00	Closed
Claude Sonnet	\$3.00	\$15.00	Closed
Gemini 2.5 Flash	\$0.15	\$0.60	Closed

Source: intuitionlabs.ai (March 2026), nxcodes.io/resources/news/deepseek-api-pricing-complete-guide-2026

- In a real-world scenario (1M input + 1M output): OpenAI and Google cost \$1.10–1.80, DeepSeek runs at **~\$0.07 — roughly 25× cheaper** for comparable quality.
- DeepSeek R1 vs OpenAI o1: \$0.55 / \$2.19 against \$15 / \$60 per million tokens — a **27× difference** on both ends.
- Open MIT license: developers download and run the model with no licensing fees. Result: **10M+ downloads of R1** on Hugging Face in early 2025.

SECTION 07

MARKET

Market size

Global AI inference market: **\$106B in 2025** → **\$255B by 2030**, growing 19% annually.

Source: [MarketsandMarkets](https://marketsandmarkets.com/Market-Reports/ai-inference-market-189921964.html) — marketsandmarkets.com/Market-Reports/ai-inference-market-189921964.html

Geopolitical distribution (key narrative point)

Per Microsoft data, DeepSeek's market share by country: **Belarus 56%, Russia 49%, Cuba 43%, China 89%**. It dominates in Iran, Syria, and several African countries, and is the default chatbot on Huawei smartphones.

Source: [Euronews](https://euronews.com/next/2026/01/09/deepseeks-ai-gains-traction-in-developing-nations-microsoft-report-says) — euronews.com/next/2026/01/09/deepseeks-ai-gains-traction-in-developing-nations-microsoft-report-says

IMPLICATION

DeepSeek has effectively become the default AI solution in markets outside the Western bloc. This creates a durable user base that is largely insulated from US geopolitical pressure.

Potential revenue streams

API usage	Pay-per-query access for developers and enterprises.
Enterprise deployments	Private installations for banks, government agencies, and large corporations.
Cloud partnerships	Distribution via Chinese and international cloud providers.
Model fine-tuning	Customization for specialized client tasks.
Government contracts	Strategic agreements within China and aligned markets.

RISK NOTE

Current revenue is not publicly disclosed — a key risk factor for valuation. The company has stated that profitability is not its current focus, which adds uncertainty around the monetization timeline.

SECTION 08

TRACTION AND BUSINESS METRICS

- **125M monthly active users** as of May 2025. **58%** of new AI startups use DeepSeek in their stack.
Source: SQ Magazine — sqmagazine.co.uk/deepseek-ai-statistics
- DeepSeek-Coder ranks **#2** among coding tools in the StackOverflow developer survey, behind only GitHub Copilot.
- Over **26,000 enterprise accounts**. R&D office opened in Zurich.

Source: Electro IQ — electroiq.com/stats/deepseek-ai-statistics

METRIC	VALUE	CONTEXT
MAU	125M+	As of May 2025
Enterprise accounts	26,000+	Confirmed enterprise
Startup adoption	58%	New AI startups using DeepSeek
Coding tool ranking	#2	StackOverflow developer survey
R&D offices	Zurich	International expansion

SECTION 09

COMPETITIVE LANDSCAPE

REGION	COMPETITORS	DEEPSEEK POSITIONING
USA	OpenAI, Anthropic, Google, Meta, xAI	10–30× cheaper, open source
China	Alibaba Qwen, Tencent, Baidu, MiniMax	Comparable quality, more transparent technology
Europe	Mistral, Llama (Meta)	Comparable pricing, more powerful model

PRIMARY COMPETITIVE RISK **MODERATE**

The cost advantage is replicable. Alibaba and Tencent have significantly larger resources. Historically, edges in AI models have lasted months, not years.

SECTION 10

DURABLE COMPETITIVE ADVANTAGES

Cost efficiency

Demonstrated ability to train and serve models at lower cost than competitors with comparable quality.

Talent concentration

One of the most prestigious AI labs in China for attracting top researchers.

Developer trust

Open-weight models drive broad adoption and an independent ecosystem.

Strategic importance

Potentially a cornerstone of China's AI sovereignty strategy.

Hardware integration

Optimization for Huawei Ascend creates deep alignment with Chinese infrastructure.

Global brand

Rare for a Chinese AI company — DeepSeek is widely recognized outside China.

SECTION 11

RISKS

VALUATION RISK MODERATE

Strong demand from large strategic investors — including reports that Tencent is considering acquiring up to 20% of the company, alongside interest from Alibaba — may drive the round valuation higher. Result: entry at a higher valuation than initially indicated.

OVERSUBSCRIPTION RISK MODERATE

High demand from major Chinese tech companies and state-backed funds may make the round inaccessible despite a soft commit. Result: full refund of committed capital within up to 2 months.

MARKET RISK MODERATE

The company may not reach IPO, or the timeline may be substantially longer than expected. Result: inability to realize gains in a liquid market within a foreseeable timeframe.

LIQUIDITY RISK MODERATE

Chinese private companies are not traded by most international brokers. Several jurisdictions restrict investment in Chinese assets. Result: extended search for buyers on the secondary market and potentially additional transaction costs.

BANKRUPTCY RISK MODERATE

Intense competition in domestic, Asian, and international AI markets creates operational risk. Result: partial or total capital loss.

All risks listed above are currently assessed as moderate by the team.

SECTION 12

EXIT SCENARIOS

IPO on the Hong Kong Stock Exchange (HKEx)

The most realistic scenario given current geopolitics. HKEx is actively attracting Chinese tech listings and has experience with AI assets. Expected horizon: **2–3 years** from the date of investment.

M&A / strategic acquisition

Less likely — Liang Wenfeng has historically resisted ceding control. However, in a substantial market shift, major players (Alibaba, Tencent, ByteDance) could make an acquisition offer.

Secondary sale

Prior to a liquidity event, investors may sell their position on the secondary market. Given the limited number of brokers handling Chinese private assets and jurisdictional restrictions, finding a buyer may take time and incur additional costs.

Buyback / tender offer

When raising new financing rounds, the company may organize a partial buyback from existing investors. Historically, this mechanism has been used by major tech companies to provide liquidity to early backers and retain key employees.

BASE CASE

The team views an HKEx IPO as the most likely exit on a 3-year horizon. Secondary sale is the fallback if early liquidity is required.

REFERENCES

SOURCES

1. Stanford HAI AI Index 2025 — hai.stanford.edu/ai-index/2025-ai-index-report
2. TLDL AI Company Rankings 2026 — tldl.io/resources/ai-companies-landscape-2026
3. Syntax Research: Anthropic Peer Analysis — syntaxdata.com (Feb 2026)
4. CNBC: Anthropic \$900B valuation — cnbc.com (Apr 2026)
5. BentoML: DeepSeek models guide — bentoml.com/blog/the-complete-guide-to-deepseek-models
6. The Register: DeepSeek training cost — theregister.com/2025/09/19/deepseek_cost_train
7. LocalAI Master: DeepSeek V3 analysis — localaimaster.com/models/deepseek-v3-vs-v3-1-analysis
8. SoftwareSeni: AI Inference Market 2025 — softwareseni.com/the-ai-inference-market-in-2025
9. VentureBeat: DeepSeek V4-Pro — venturebeat.com/technology/deepseek-v4-arrives
10. IntuitionLabs: LLM pricing — intuitionlabs.ai/articles/llm-api-pricing-comparison-2025
11. Nxcode: DeepSeek API pricing 2026 — nxcode.io/resources/news/deepseek-api-pricing-complete-guide-2026
12. MarketsandMarkets: AI Inference Market — marketsandmarkets.com/Market-Reports/ai-inference-market-189921964.html
13. Euronews: DeepSeek in developing nations — euronews.com/next/2026/01/09/deepseeks-ai-gains-traction
14. SQ Magazine: DeepSeek statistics — sqmagazine.co.uk/deepseek-ai-statistics
15. Electro IQ: DeepSeek enterprise stats — electroi.com/stats/deepseek-ai-statistics
16. Price of Progress 2026 — arxiv.org/html/2511.23455v2
17. SemiAnalysis: DeepSeek true costs (cited in VentureBeat)

Confidential. This document does not constitute an investment recommendation or an offer to sell securities. Investments in private companies and Chinese assets carry a high degree of risk, including the potential for total loss of capital. Prospective investors should consult their own financial, legal, and tax advisors before making any investment decision.