

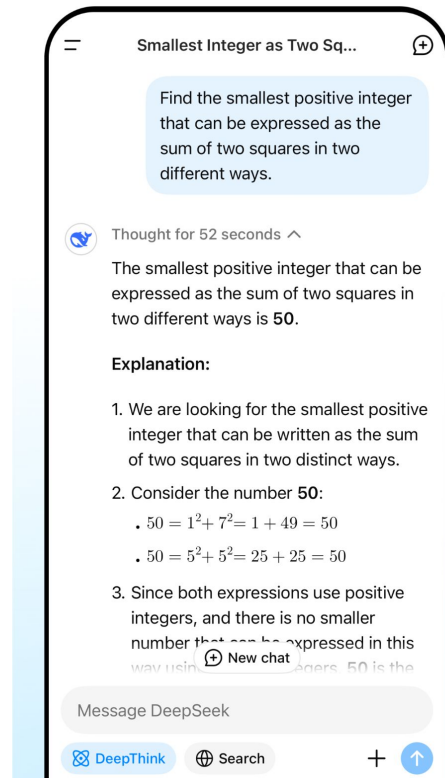
# How deepseek changes the LLM story

A quick and dirty talk by **Sasha Rush**



- < 200 employees
- Spin off of hedge-fund
- Consistent open-weights model releases


DeepThink first,  
solve puzzles smarter.



Published by  
Brian Lehrer Show Segments by The Brian Lehrer Show

## Why China's AI Startup 'DeepSeek' is a Sputnik Moment

Listen Queue

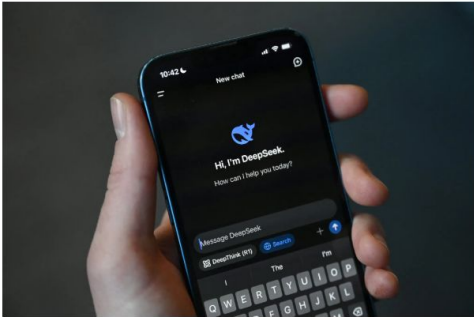


A video player interface showing a segment from the Brian Lehrer Show. The title is "Why China's AI Startup 'DeepSeek' is a Sputnik Moment". Below the title are buttons for "Listen", "Queue", and a menu icon. The main content area shows a blurred background with the OpenAI logo and a laptop in the foreground displaying the DeepSeek logo.

## What to Know About DeepSeek and How It Is Upending A.I.

How did a little-known Chinese start-up cause the markets and U.S. tech giants to quake? Here's what to know.

Listen to this article • 9:14 min [Learn more](#) Share full article



DeepSeek is a Chinese start-up run by a quantitative stock trading firm called HighFlyer. Greg Baker/Agence France-Presse — Getty Images

Leaders | Seek and ye shall find

## The real meaning of the DeepSeek drama

The Chinese model-maker has panicked investors. But it is good for the users of AI



ILLUSTRATION: ROSE WONG



lolzinventor · 1h

Would you consider releasing some model weights, and publishing some research?

... Reply 33



samaltman · 11m

OpenAI CEO Sam Altman | Verified

yes, we are discussing. i personally think we have been on the wrong side of history here and need to figure out a different open source strategy; not everyone at openai shares this view, and it's also not our current highest priority.

... Vote

AI EFFECT

## Nvidia sheds almost \$600 billion in market cap, biggest one-day loss in U.S. history

PUBLISHED MON, JAN 27 2025-4:08 PM EST | UPDATED MON, JAN 27 2025-5:26 PM EST

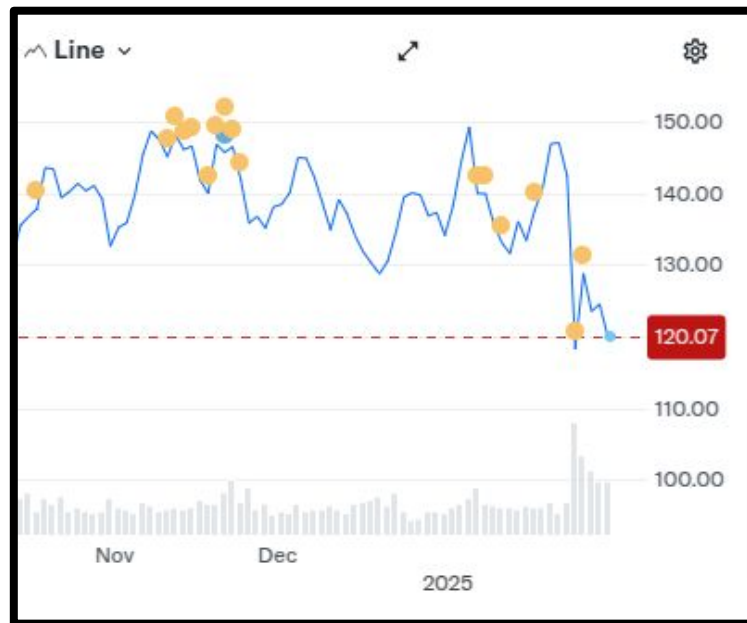


Samantha Subin  
@SAMANTHA\_SUBIN

WATCH LIVE

### KEY POINTS

- Nvidia shares plunged 17% on Monday, resulting in a market cap loss of close to \$600 billion, the biggest drop ever for a U.S. company.
- The sell-off, which hit much of the U.S. tech sector, was sparked by concerns about increased competition from Chinese AI lab DeepSeek.
- Data center companies that rely on Nvidia chips also plummeted, with Dell, Oracle and Super Micro Computer all falling by at least 8.7%.



[← Messages](#)

**Dad**

[Details](#)

Go Birds!

Yeah! Saquon is ridiculous.

What's up with Deep Seek?

# Hot Takes

Bloggers:

- <https://stratechery.com/2025/deepseek-faq/>
- <https://www.bloomberg.com/opinion/articles/2025-01-27/hedge-fund-ai-is-cheap-ai>
- <https://semianalysis.com/2025/01/31/deepseek-debates/>
- <https://www.interconnects.ai/p/deepseek-r1-recipe-for-o1>
- <https://darioamodei.com/on-deepseek-and-export-controls>

What can I add?

- *We're in this bizarre world where the best way to learn about llms... is to read papers by chinese companies* - Chris Manning

Review: What Happened?



## DeepSeek LLM

### Scaling Open-Source Language Models with Longtermism

Xiao Bi, Deli Chen, Guanting Chen, Shan  
Honghui Ding, Kai Dong, Qiushi Du, Zhe  
Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong  
Wenjie Hu, Panpan Huang, Erhang Li, Guowe  
Fangyun Lin, A.X. Liu, Bo Liu, Wen Liu, Xia  
Shanghao Lu, Fuli Luo, Shirong Ma, Xiaotao  
Tongzheng Ren, Zehui Ren, Chong Ruan, J  
Xuecheng Su, Jingxiang Sun, Yaofeng Sun, M  
Shiyu Wang, Yaohui Wang, Yongji Wang, Ton  
Yiliang Xiong, Hanwei Xu, R.X. Xu, Yanhong  
Xingkai Yu, B. Zhang, Haowei Zhang, Lecor  
Minghua Zhang, Wentao Zhang, Yicha  
Shangyan Zhou, Shunfeng Zh

## DeepSeek-Coder: When the Large Language Model Meets Programming - The Rise of Code Intelligence

Daya Guo<sup>\*1</sup>, Qihao Zhu<sup>\*1,2</sup>, Dejian Yang<sup>1</sup>, Zhenda Xie<sup>1</sup>, Kai Dong<sup>1</sup>, Wentao Zhang<sup>1</sup>  
Guanting Chen<sup>1</sup>, Xiao Bi<sup>1</sup>, Y. Wu<sup>1</sup>, Y.K. Li<sup>1</sup>, Fuli Luo<sup>1</sup>, Yingfei Xiong<sup>2</sup>, Wenfeng Liang<sup>1</sup>

<sup>1</sup>DeepSeek-AI

<sup>2</sup>Key Lab of HCST (PKU), MOE; SCS, Peking University  
{zhuqh, guodaya}@deepseek.com

<https://github.com/deepseek-ai/DeepSeek-Coder>

### Abstract

The rapid development of large language models has revolutionized code intelligence in software development. However, the predominance of closed-source models has restricted extensive research and development. To address this, we introduce the DeepSeek-Coder series

5 Jan 2024

26 Jan 2024

## DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models

Damai Dai<sup>1,2</sup>, Chengqi Deng<sup>1</sup>, Chenggang Zhao<sup>1,3</sup>, R.X. Xu<sup>1</sup>, Huazuo Gao<sup>1</sup>, Deli Chen<sup>1</sup>, Jiashi Li<sup>1</sup>, Wangding Zeng<sup>1</sup>, Xingkai Yu<sup>1,4</sup>, Y. Wu<sup>1</sup>, Zhenda Xie<sup>1</sup>, Y.K. Li<sup>1</sup>, Panpan Huang<sup>1</sup>, Fuli Luo<sup>1</sup>, Chong Ruan<sup>1</sup>, Zhifang Sui<sup>2</sup>, Wenfeng Liang<sup>†</sup>

<sup>1</sup>DeepSeek-AI

<sup>2</sup>National Key Laboratory for Multimedia Information Processing, Peking University

<sup>3</sup>Institute for Interdisciplinary Information Sciences, Tsinghua University

<sup>4</sup>National Key Laboratory for Novel Software Technology, Nanjing University

{daidamai, szf}@pku.edu.cn, {wenfeng.liang}@deepseek.com

<https://github.com/deepseek-ai/DeepSeek-MoE>

Abstract

11 Jan 2024

## DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model

DeepSeek-AI


research@deepseek.com

### Abstract

We present DeepSeek-V2, a strong Mixture-of-Experts (MoE) language model characterized by economical training and efficient inference. It comprises 236B total parameters, of which 21B are activated for each token, and supports a context length of 128K tokens. DeepSeek-V2 adopts innovative architectures including Multi-head Latent Attention (MLA) and DeepSeekMoE. MLA guarantees efficient inference through significantly compressing the Key-Value (KV) cache into a latent vector, while DeepSeekMoE enables training strong models at an economical cost through sparse computation. Compared with DeepSeek 67B, DeepSeek-V2 achieves significantly stronger performance, and meanwhile saves 42.5% of training costs, reduces the KV cache by 93.3%, and boosts the maximum generation throughput to 5.76 times. We pretrain DeepSeek-V2 on a high-quality and multi-source corpus consisting of 8.1T tokens, and further perform Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL) to fully unlock its

[cs.CL] 19 Jun 2024



Sasha Rush 

@srush\_nlp



I don't know anything about production LLMs, but whenever I talk to someone who does they seem to be using DeepSeek. Seems to indicate evals haven't yet caught up to internal tests.

[Traducir post](#)

9:59 a. m. · 19 jul. 2024 · **10 mil** Visualizaciones



# DeepSeek-V3 Technical Report

DeepSeek-AI

research@deepseek.com

## Abstract

We present DeepSeek-V3, a strong Mixture-of-Experts (MoE) language model with 671B total parameters with 37B activated for each token. To achieve efficient inference and cost-effective training, DeepSeek-V3 adopts Multi-head Latent Attention (MLA) and DeepSeekMoE architectures, which were thoroughly validated in DeepSeek-V2. Furthermore, DeepSeek-V3 pioneers an auxiliary-loss-free strategy for load balancing and sets a multi-token prediction training objective for stronger performance. We pre-train DeepSeek-V3 on 14.8 trillion diverse and high-quality tokens, followed by Supervised Fine-Tuning and Reinforcement Learning stages to fully harness its capabilities. Comprehensive evaluations reveal that DeepSeek-V3 outperforms other open-source models and achieves performance comparable to leading closed-source models. Despite its excellent performance, DeepSeek-V3 requires only 2.788M H800 GPU hours for its full training. In addition, its training process is remarkably stable. Throughout the entire training process, we did not experience any irrecoverable loss spikes or perform any rollbacks.

[cs.CL] 27 Dec 2024

## DeepSeek-V3 Technical Report

DeepSeek-AI

research@deepseek.com

### Abstract

We present DeepSeek-V3, a strong Mixture-of-Experts (MoE) language model with 671B total parameters with 37B activated for each token. To achieve efficient inference and cost-effective training, DeepSeek-V3 adopts Multi-head Latent Attention (MLA) and DeepSeekMoE architectures, which were thoroughly validated in DeepSeek-V2. Furthermore, DeepSeek-V3 pioneers an auxiliary-loss-free strategy for load balancing and sets a multi-token prediction training objective for stronger performance. We pre-train DeepSeek-V3 on 14.8 trillion diverse and high-quality tokens, followed by Supervised Fine-Tuning and Reinforcement Learning stages to fully harness its capabilities. Comprehensive evaluations reveal that DeepSeek-V3 outperforms other open-source models and achieves performance comparable to leading closed-source models. Despite its excellent performance, DeepSeek-V3 requires only 2.788M H800 GPU hours for its full training. In addition, its training process is remarkably stable. Throughout the entire training process, we did not experience any irrecoverable loss spikes or perform any rollbacks.

[cs.CL] 27 Dec 2024

[cs.CL]

Training Costs	Pre-Training	Context Extension	Post-Training	Total
in H800 GPU Hours	2664K	119K	5K	2788K
in USD	\$5.328M	\$0.238M	\$0.01M	\$5.576M

Numbers for the *training run*.  
Not the whole process.



We believe they have access to around 50,000 Hopper GPUs, which is not the same as 50,000 H100, as some have claimed. There are different variations of the H100 that Nvidia made in compliance to different regulations (H800, H20), with only the H20 being currently available to Chinese model providers today. Note that H800s have the same computational power as H100s, but lower network bandwidth.

- Semianalysis



# Cost Complications

Bloggers:

- <https://stratechery.com/2025/deepseek-faq/>
- <https://www.bloomberg.com/opinion/articles/2025-01-27/hedge-fund-ai-is-cheap-ai>
- <https://semianalysis.com/2025/01/31/deepseek-debates/>
- <https://www.interconnects.ai/p/deepseek-r1-recipe-for-o1>

What can I add?

- Summary and technical context



# DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

DeepSeek-AI

research@deepseek.com

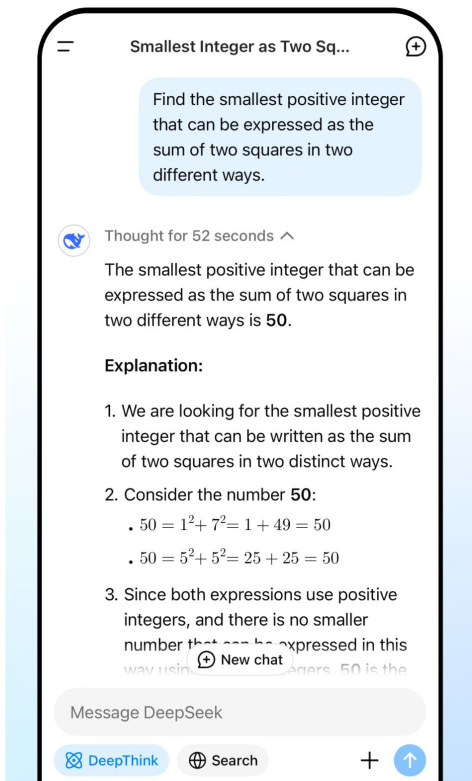
## Abstract

We introduce our first-generation reasoning models, DeepSeek-R1-Zero and DeepSeek-R1. DeepSeek-R1-Zero, a model trained via large-scale reinforcement learning (RL) without supervised fine-tuning (SFT) as a preliminary step, demonstrates remarkable reasoning capabilities. Through RL, DeepSeek-R1-Zero naturally emerges with numerous powerful and intriguing reasoning behaviors. However, it encounters challenges such as poor readability, and language mixing. To address these issues and further enhance reasoning performance, we introduce DeepSeek-R1, which incorporates multi-stage training and cold-start data before RL. DeepSeek-R1 achieves performance comparable to OpenAI-o1-1217 on reasoning tasks. To support the research community, we open-source DeepSeek-R1-Zero, DeepSeek-R1, and six dense models (1.5B, 7B, 8B, 14B, 32B, 70B) distilled from DeepSeek-R1 based on Qwen and Llama.

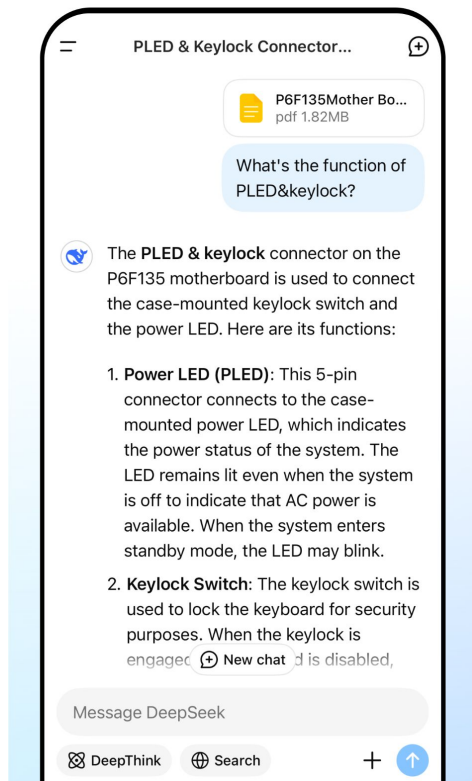
22 Jan 2025

[cs.CL]

## DeepThink first, solve puzzles smarter.



## Read documents, get quick summaries.



## DeepSeek v3

- Technical overview of LLM
- Primarily a *pretraining* innovation
- Think **OpenAI GPT-4o**

## DeepSeek R1

- Technical overview of reasoning
- Primarily a *post training* innovation
- Think **OpenAI o1**



We believe the most important thing right now is to participate global innovation. For years, Chinese companies have been accustomed to leveraging technological innovations developed elsewhere and monetizing them through applications. But this isn't sustainable. This time, our goal isn't quick profits but advancing the technological frontier to drive ecosystem growth.

- Liang Wenfeng (DeepSeek Founder)

<https://thechinaacademy.org/interview-with-deepseek-founder-were-done-following-its-time-to-lead/>



**DeepSeek-V3** was actually the real innovation and what *should* have made people take notice a month ago (we certainly did). As a pretrained model, it appears to come close to the performance of<sup>4</sup> state of the art US models on some important tasks, while costing substantially less to train

– Dario Amodei

<https://darioamodei.com/on-deepseek-and-export-controls>



The reason why I spent so much time on V3 is that that was the model that actually demonstrated a lot of the dynamics that seem to be generating so much surprise and controversy. R1 is notable, however, because o1 stood alone as the only reasoning model on the market, and the clearest sign that OpenAI was the market leader.

— Ben Thompson

<https://stratechery.com/2025/deepseek-faq/>

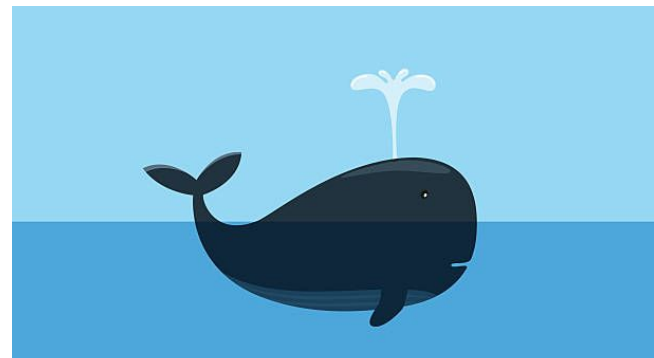
# Summary

DeepSeek v3



DeepSeek R1

- Long-running story in LLMs:  
Better base models enable new capabilities
- Long delayed insight into the state-of-the-art





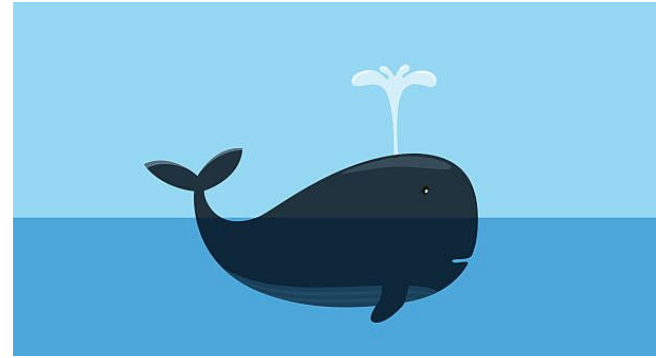
# Talk Plan

## DeepSeek v3

## DeepSeek R1

- Three key insights from both papers.
- What this tells us about where things are going

DeepSeek v3

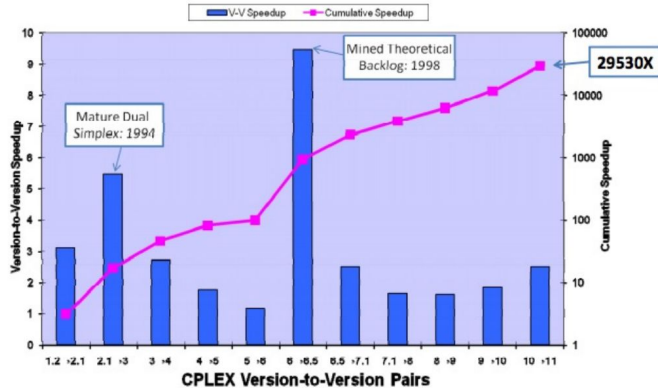




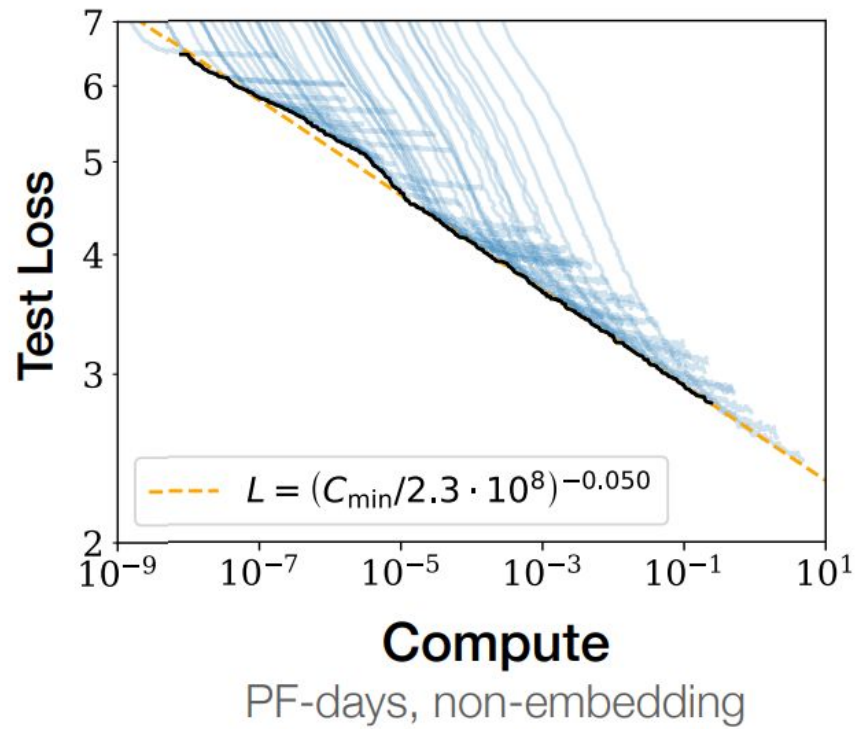
To understand how progress will proceed when intelligence becomes an algorithmic problem, we might do well to look at how existing algorithms improve. Several features of characteristic algorithmic improvement curves should interest us. How fast is progress? Does it tend to be smooth or lumpy? ...

– Katja Grace

### CPLEX MIP Performance Improvements



# Shifting the Curve



What this typically does is shift the curve: if the innovation is a 2x "compute multiplier" (CM), then it allows you to get 40% on a coding task for \$5M instead of \$10M; or 60% for \$50M instead of \$100M, etc.

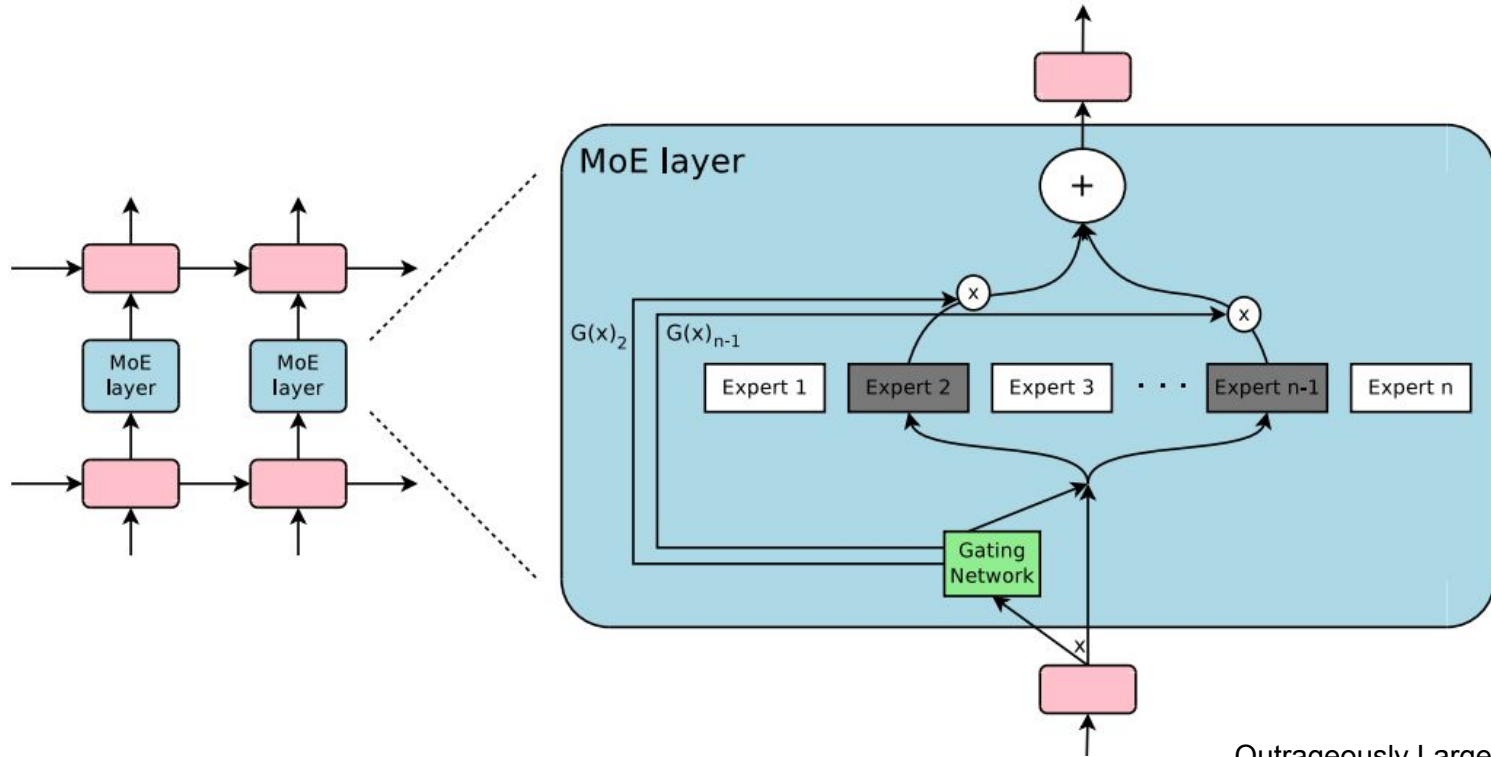
- Dario Amodei

# Three ways DeepSeek v3 - "Shifts the Curve"

- Architecture
- Parallelization
- Quantization

*Caveat:* Paper is detailed but not enough to split out individual contributions.

# Architecture: Mixture of Experts (Review)



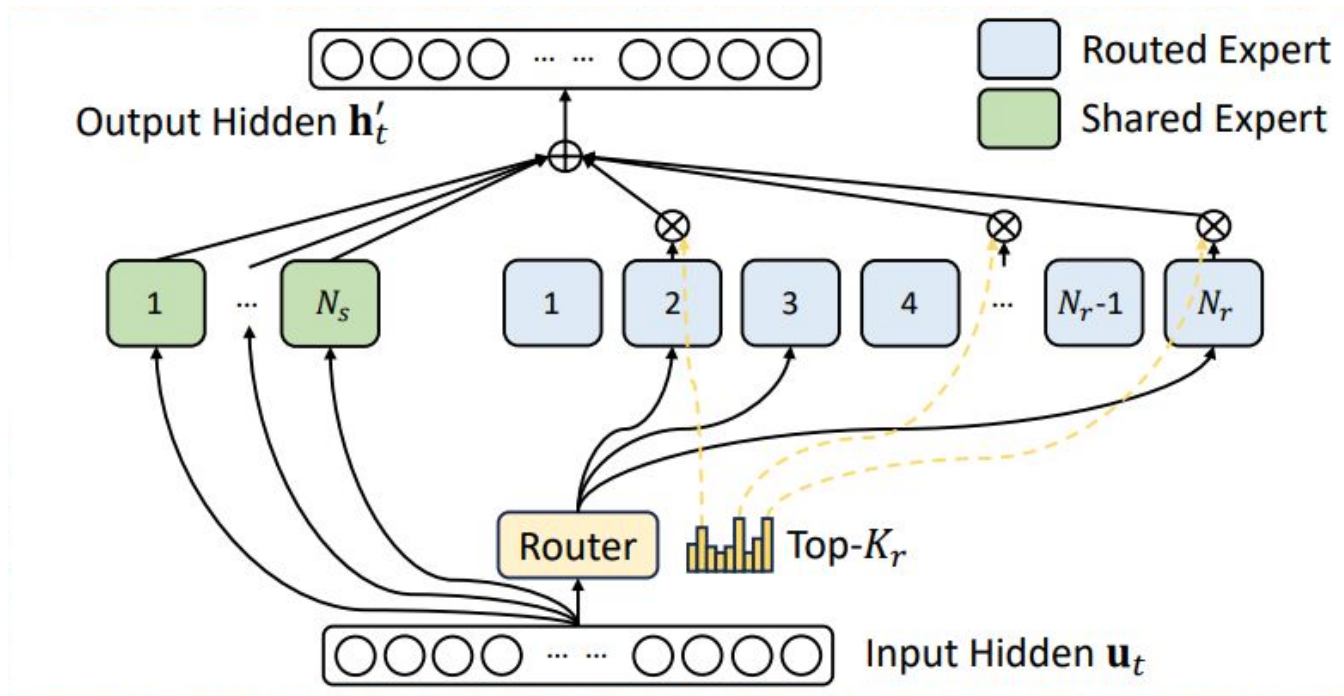
Outrageously Large Neural Networks:  
The Sparsely-Gated Mixture-of-Experts Layer  
(Shazeer et al. 2017)

# Architecture: Mixture of Experts

- GPT4 architecture has not been confirmed, long rumored to be MoE  
(Employed several MoE innovators)
- Llama-3 notably is not MoE but standard Transformer
- Mistral open-sourced MoE, but less training details



# Architecture: DeepSeek-MoE



# Architecture: DeepSeek-MoE

Mixture

$$\mathbf{h}_t^l = \sum_{i=1}^N \left( g_{i,t} \text{FFN}_i(\mathbf{u}_t^l) \right) + \mathbf{u}_t^l,$$

$$g_{i,t} = \begin{cases} s_{i,t}, & s_{i,t} \in \text{Topk}(\{s_{j,t} | 1 \leq j \leq N\}, K), \\ 0, & \text{otherwise,} \end{cases}$$

$$s_{i,t} = \text{Softmax}_i(\mathbf{u}_t^{lT} \mathbf{e}_i^l),$$

Hard selection

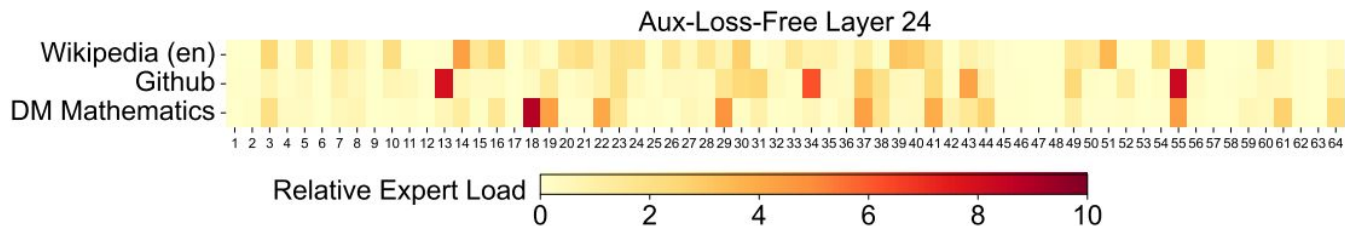
Router weights

# Architecture: DeepSeek-MoE

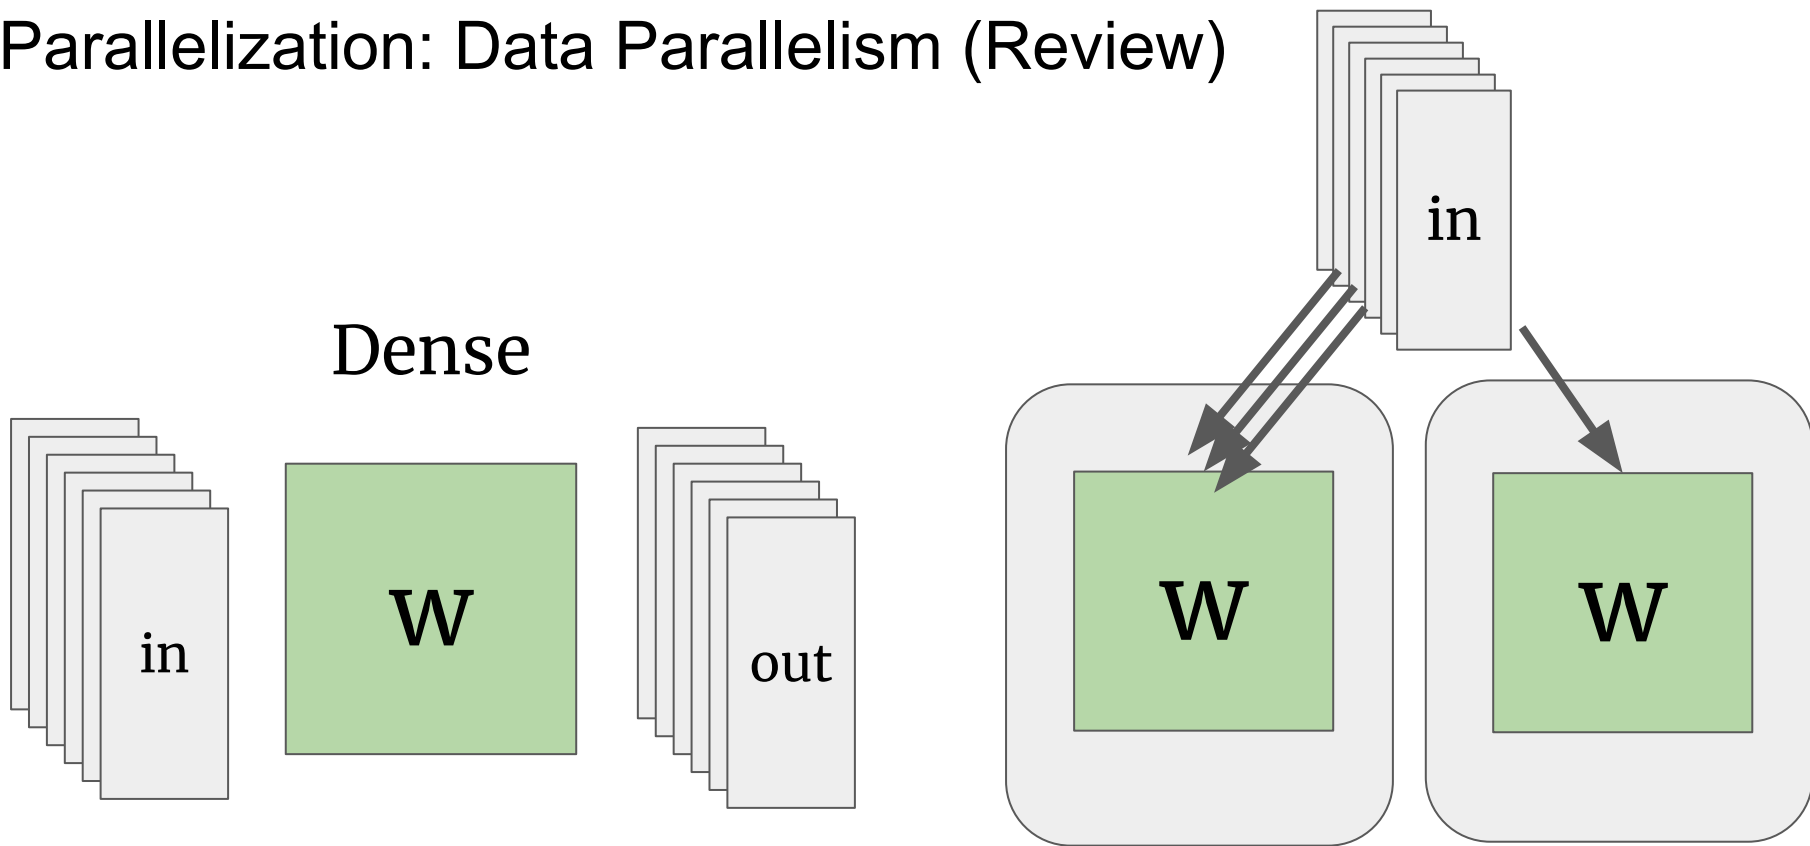
- Model setup

Model	#Total Params	#Activated Params
DeepSeek-V3	671B	37B

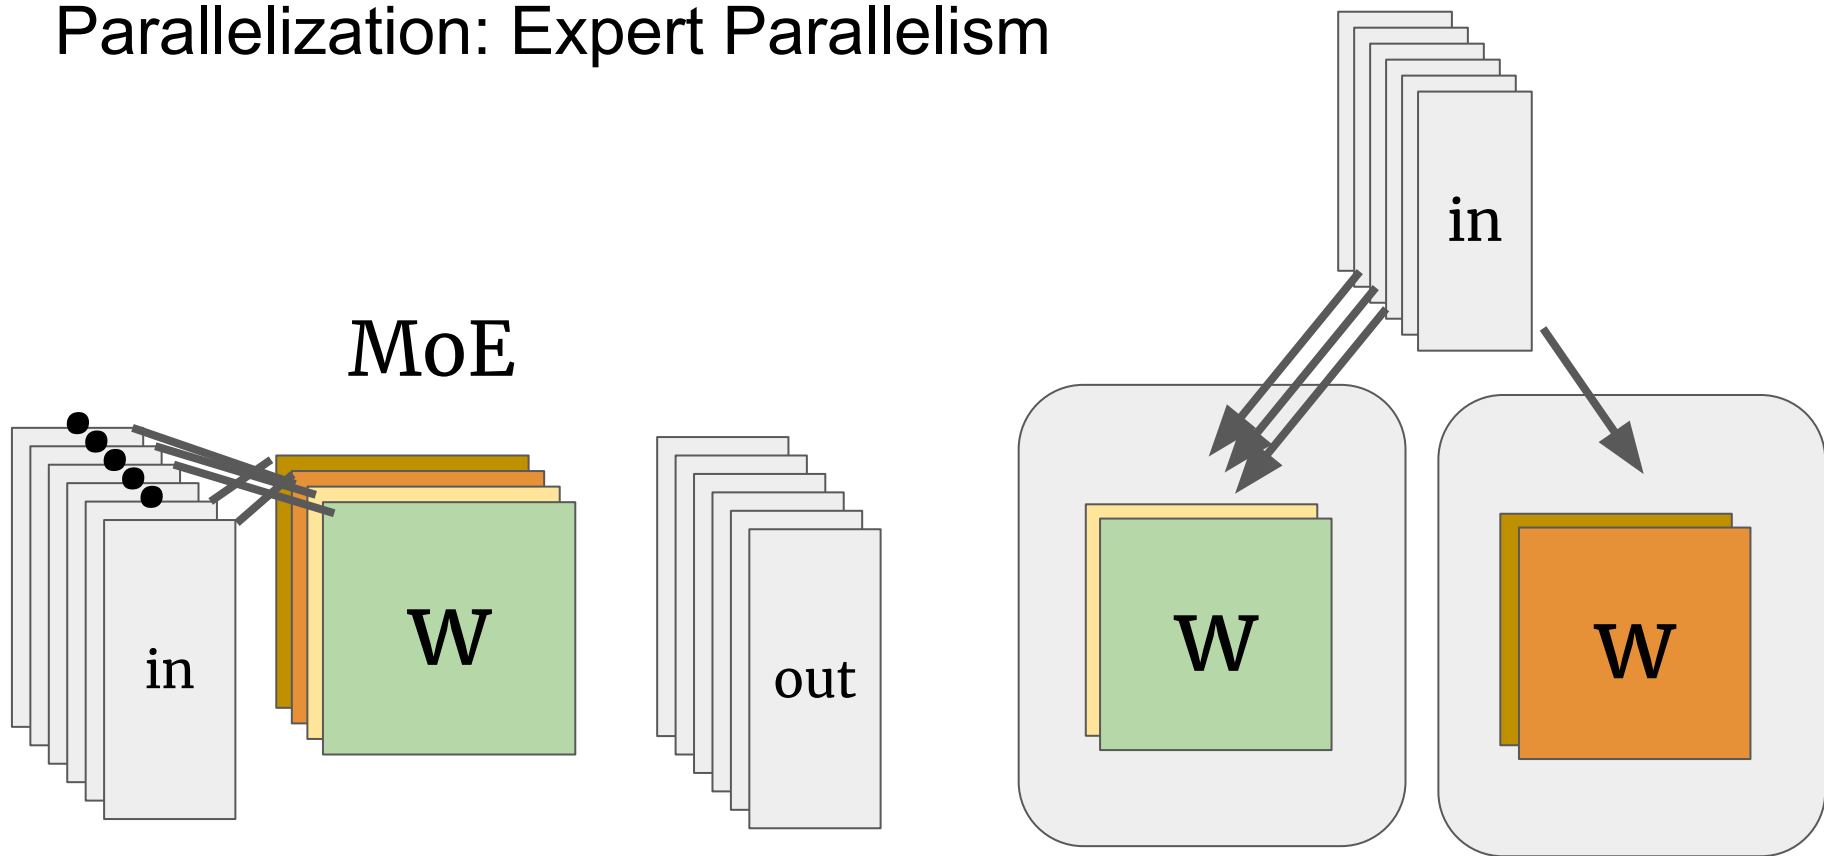
- Different layers exhibit expert specialization



# Parallelization: Data Parallelism (Review)



# Parallelization: Expert Parallelism

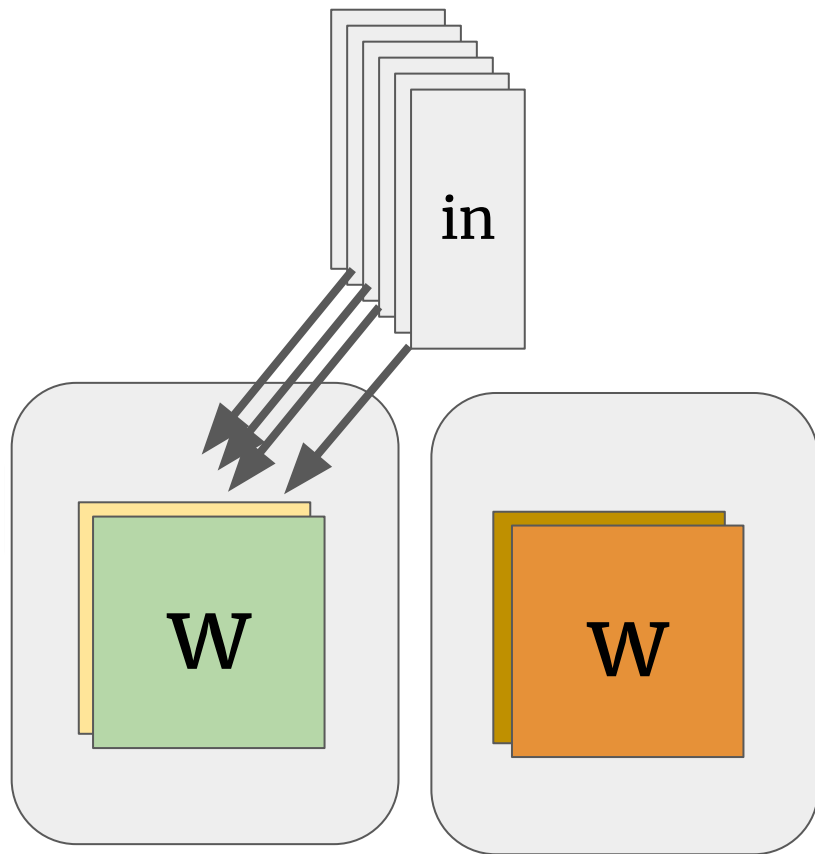


# Parallelization: Balance

- Need to ensure all compute is being used

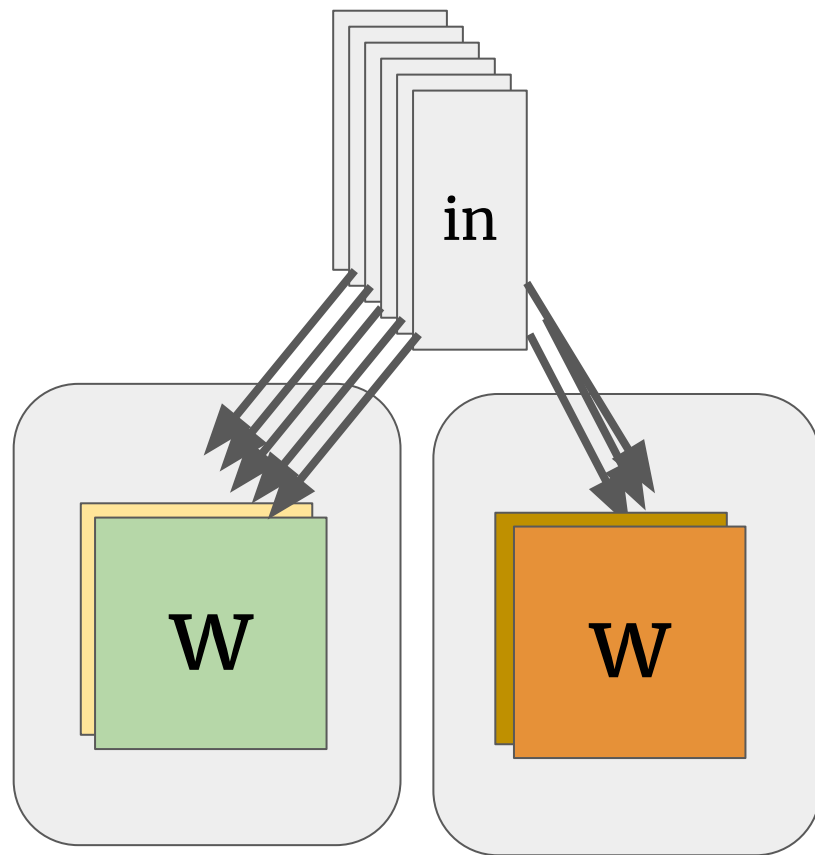
$$g'_{i,t} = \begin{cases} s_{i,t}, & s_{i,t} + b_i \in \text{Topk}(\{s_{j,t} + b_j | 1 \leq j \leq N_r\}, K_r), \\ 0, & \text{otherwise.} \end{cases}$$

- Adjust routing based on overload



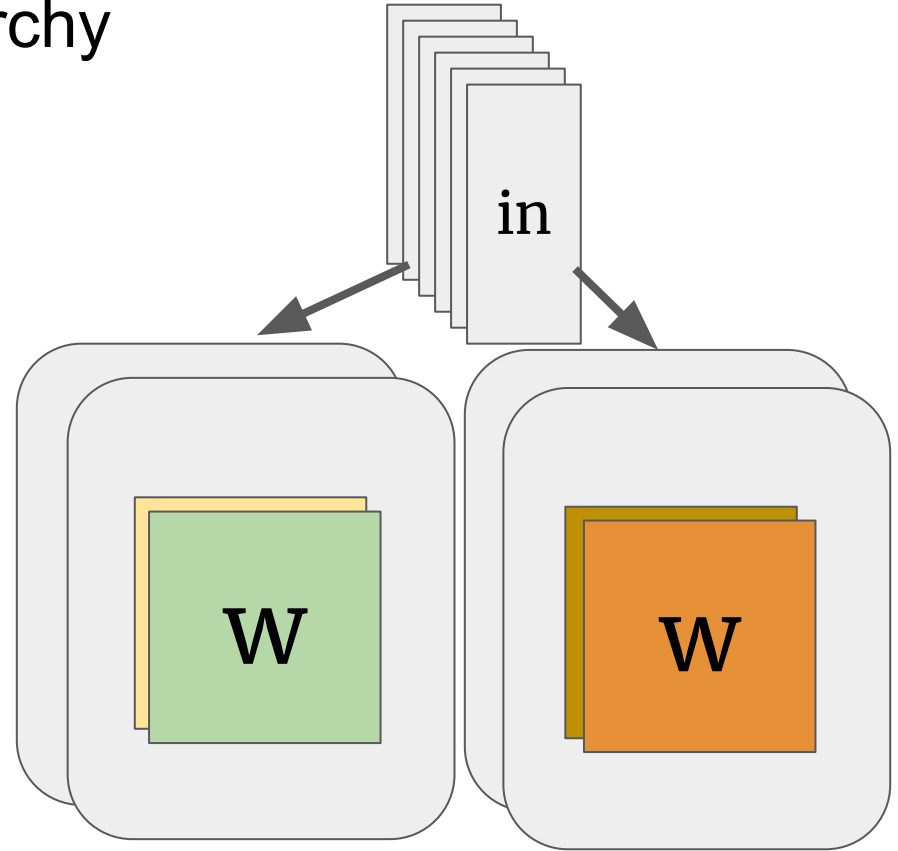
# Parallelization: Traffic

- Too many active experts can overload network



# Parallelization: Network Hierarchy

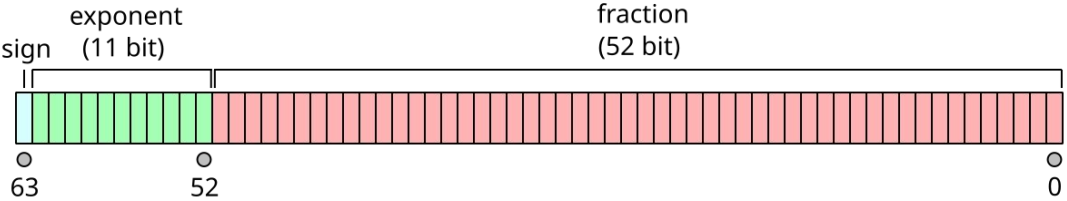
- Intra (NVLink): 8 GPUs
- Inter (IB): Cluster



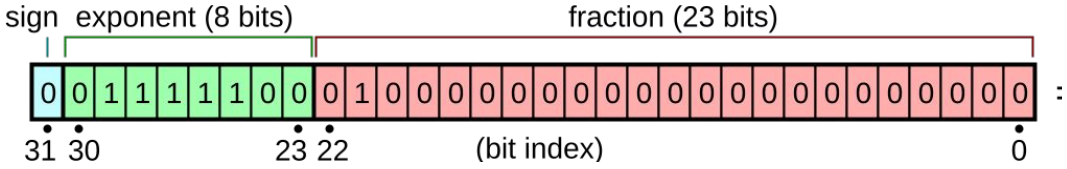


# Quantization: Float (Review)

fp64



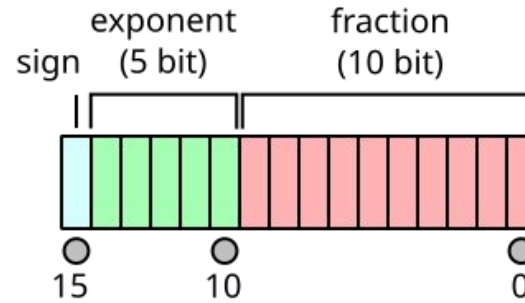
fp32



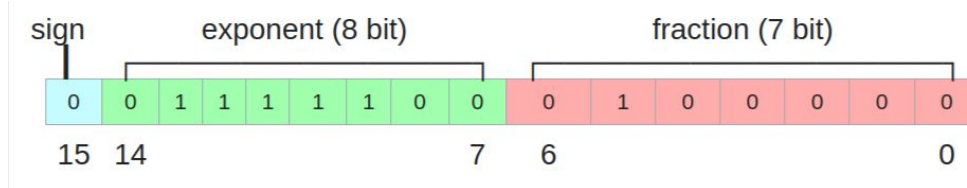
$$(-1)^{b_{31}} \times 2^{(b_{30}b_{29}\dots b_{23})_2 - 127} \times (1.b_{22}b_{21}\dots b_0)_2$$

# Quantization: Float (Review)

fp16



bf16



# Quantization: Tensor Cores

	H100 SXM	H100 PCIe
FP64	34 teraFLOPS	26 teraFLOPS
FP64 Tensor Core	67 teraFLOPS	51 teraFLOPS
FP32	67 teraFLOPS	51 teraFLOPS
TF32 Tensor Core	989 teraFLOPS*	756 teraFLOPS*
BFLOAT16 Tensor Core	1,979 teraFLOPS*	1,513 teraFLOPS*
FP16 Tensor Core	1,979 teraFLOPS*	1,513 teraFLOPS*
FP8 Tensor Core	3,958 teraFLOPS*	3,026 teraFLOPS*

wgmma - instruction

"Warpgroup level Matrix Multiply Accumulate"

$D(m \times n) = A(m \times k) \times B(k \times n) + D(m \times n)$

For  $m=64$   $k=16$ ,  $n = \text{power of } 2$

fp8

E4M3

sign	exponent					significand		
0	0	0	0	0	0	0	0	0

E5M2

sign	exponent						significand	
0	0	0	0	0	0	0	0	

$E_4M_3$

$M$

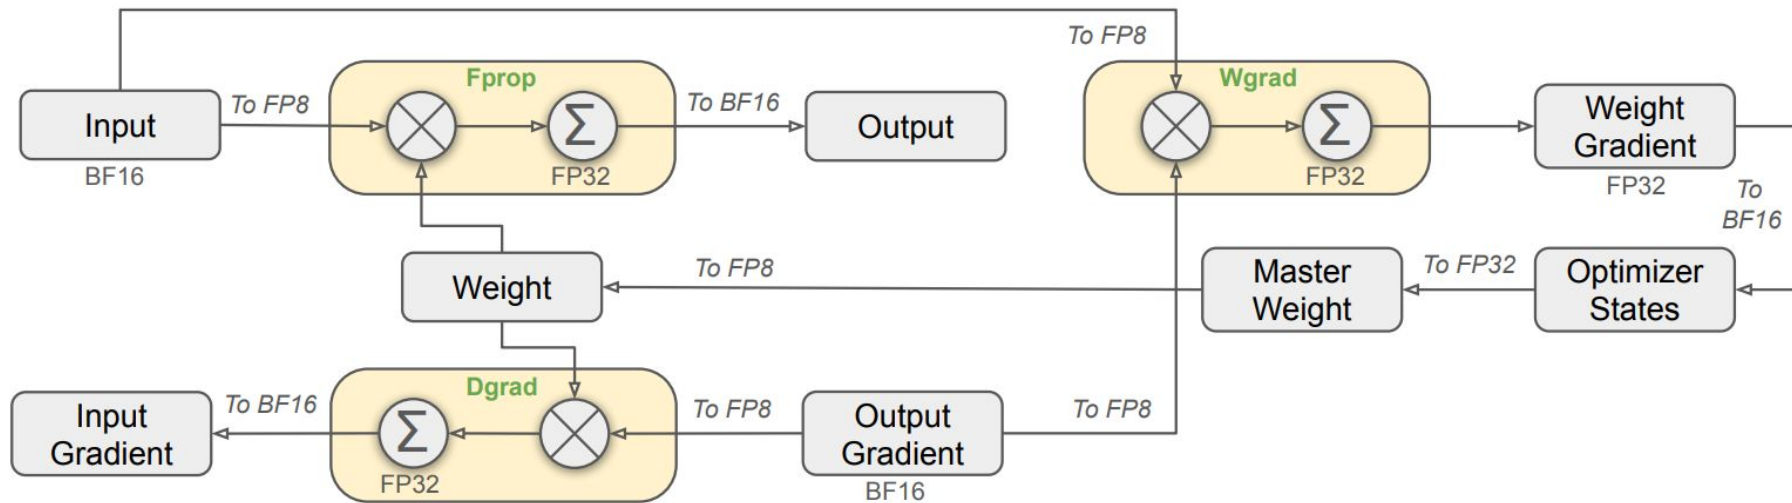


$E$



<b>0</b> <b>0111</b> ...	1	1.125	1.25	1.375
<b>0</b> <b>1000</b> ...	2	2.25	2.5	2.75
<b>0</b> <b>1001</b> ...	4	4.5	5	5.5
<b>0</b> <b>1010</b> ...	8	9	10	11
<b>0</b> <b>1011</b> ...	16	18	20	22

# Quantization: Full Pipeline



# What Don't We Know?

- V3 Paper gives relatively little information on data!
- Kernels are described but not provided (open-weights)

# OpenAI Allegations

- OpenAI insinuated that DeepSeek had access to large amount of outputs
- Well known "distillation" properties from small amounts of data

## OpenAI says DeepSeek may have 'inappropriately' used its data

OpenAI itself has been accused of building ChatGPT by inappropriately accessing content it didn't have the rights to.



### Imitation Attacks and Defenses for Black-box Machine Translation Systems

Eric Wallace   Mitchell Stern   Dawn Song  
UC Berkeley  
{ericwallace, mitchell, dawnson}@berkeley.edu

#### Abstract

Adversaries may look to *steal* or *attack* black-box NLP systems, either for financial gain or to exploit model errors. One setting of particu-

tion model allows an adversary to avoid long-term API costs or launch a competitor service. Moreover, attacking a system using adversarial examples (Szegedy et al., 2014) allows an adversary to

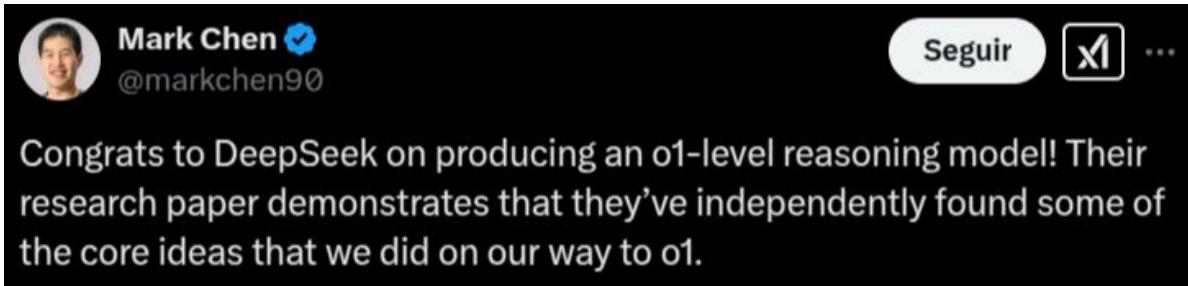


# DeepSeek is quite ambitious



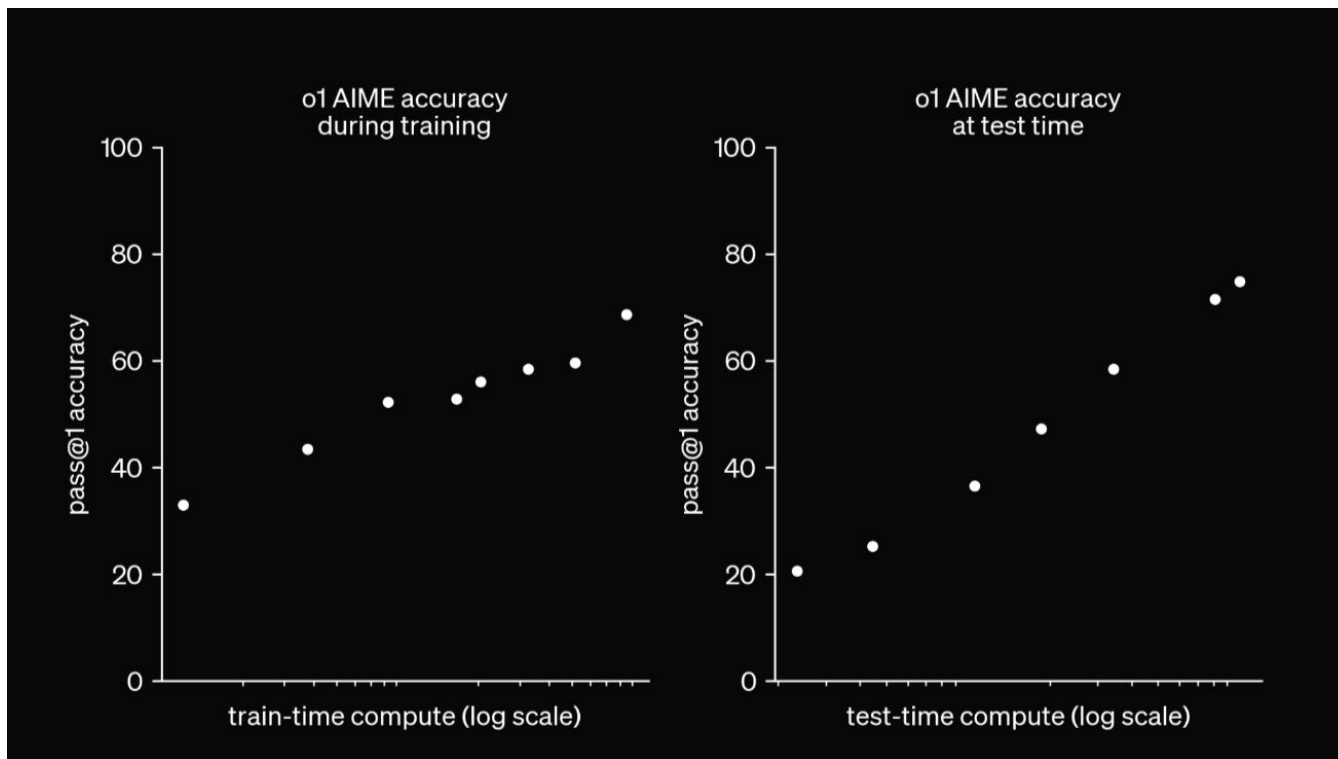
*We will consistently study and refine our model architectures, aiming to further improve both the training and inference efficiency, striving to approach efficient support for infinite context length. Additionally, we will try to break through the **architectural limitations of Transformer**, thereby pushing the boundaries of its modeling capabilities.*

DeepSeek R1



(Chief Research Officer at OpenAI)

# Reminder: OpenAI o1





o1 learns to hone its chain of thought and refine the strategies it uses. It learns to recognize and correct its mistakes. It learns to break down tricky steps into simpler ones. It learns to try a different approach when the current one isn't working.

- OpenAI

## Example: Model "Thinking"

First, the cipher seems connected to the plaintext.

Given the time constraints, perhaps the easiest way is to try to see patterns.

Option 1: Try to find mapping from letters to letters.

Do any letters match?

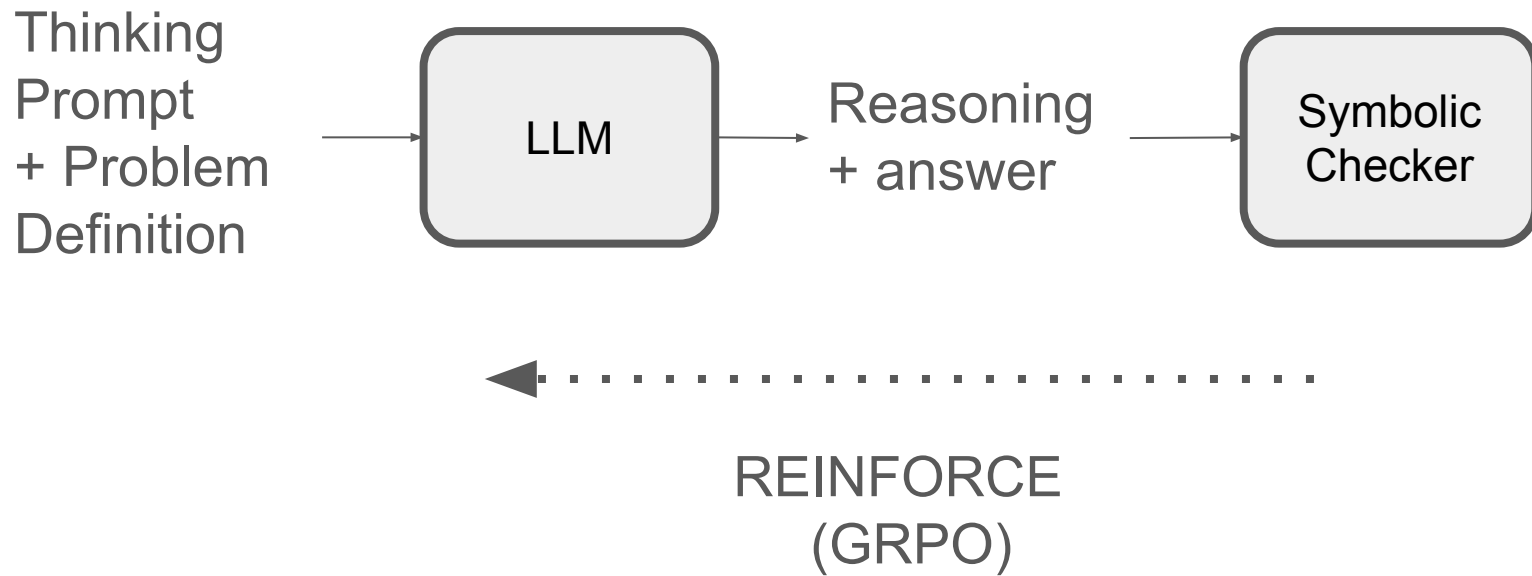
First, let's write down the ciphertext and plaintext letters on top of each other.

# Speculations on o1

After the release of o1, unsure of method

- Massive annotation effort?
- Process Reward Model?
- MCTS?
- Classic AI search?

# R1-Zero: Simple, Scalable Method





# Where does the reward come from?

- Hard-coded verifiers for math and coding puzzles
- Hard-coded linters for style

# The "Aha Moment"

- The RL process increases length of responses
- Responses resemble "thinking" without prodding

---

Quest

Respo

To sol

$(\sqrt{a-x})$

Rearrange to isolate the inner square root term:

$$(a-x^2)^2 = a+x \implies a^2 - 2ax^2 + (x^2)^2 = a+x \implies x^4 - 2ax^2 - x + (a^2 - a) = 0$$

...

**Wait, wait. Wait. That's an aha moment I can flag here.**

Let's reevaluate this step-by-step to identify if the correct sum can be ...

We started with the equation:

$$\sqrt{a-x} - \sqrt{a+x} = x$$

First, let's square both sides:

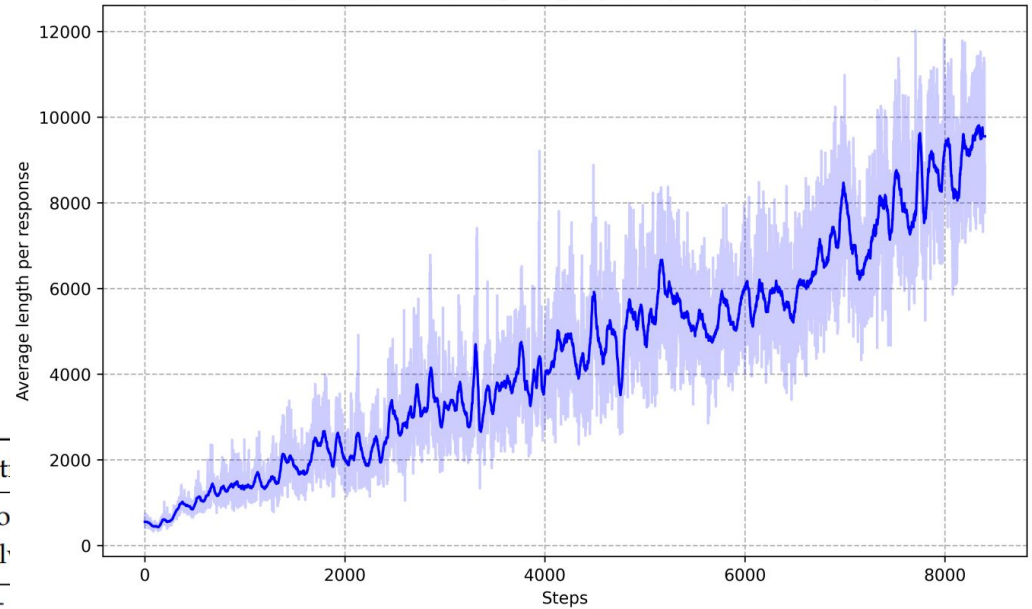
$$a - \sqrt{a+x} = x^2 \implies \sqrt{a+x} = a - x^2$$

Next, I could square both sides again, treating the equation: ...

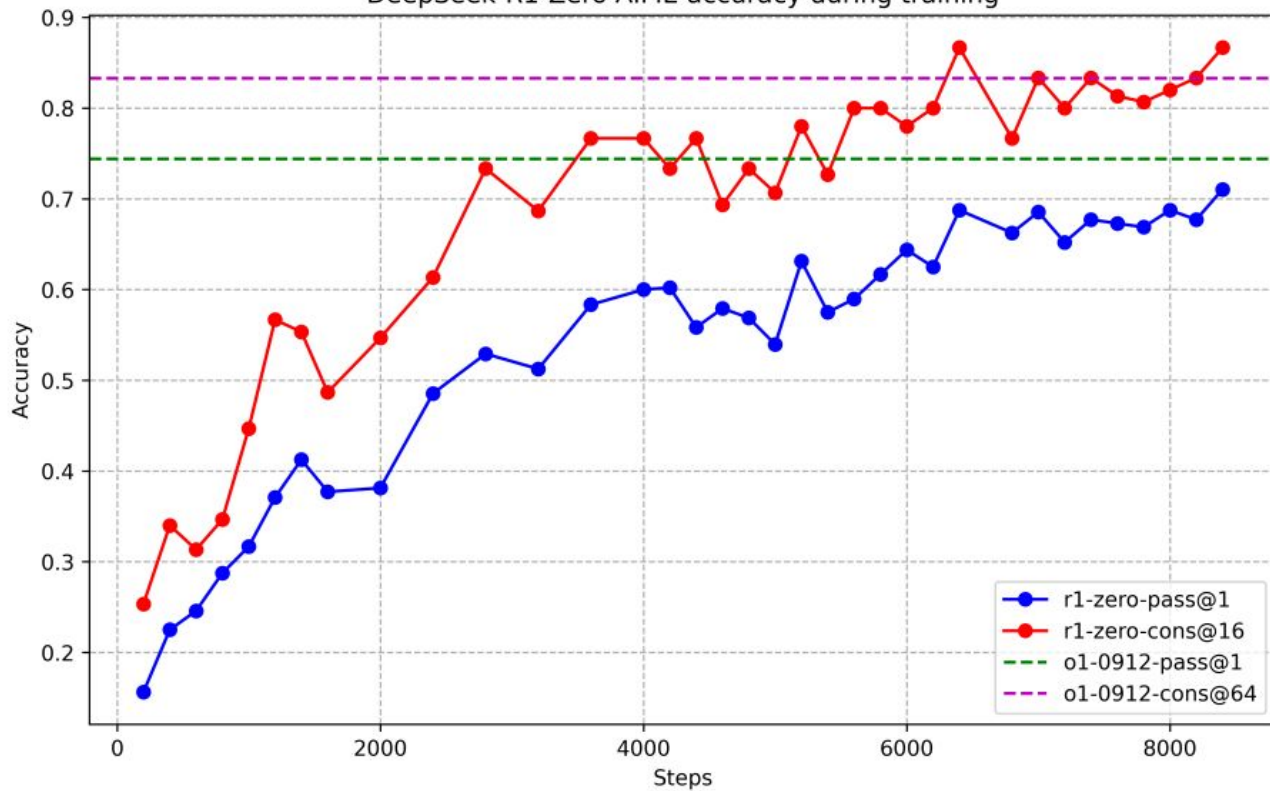
...

---

DeepSeek-R1-Zero average length per response during training



DeepSeek-R1-Zero AIME accuracy during training



# Wait what!?

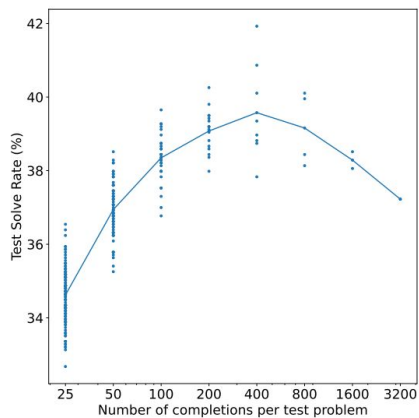
Big question: This seems too simple.

## Punchline

- Method seems to require sufficiently strong model to learn
- Once learned can be transferred to simpler model (distillation)

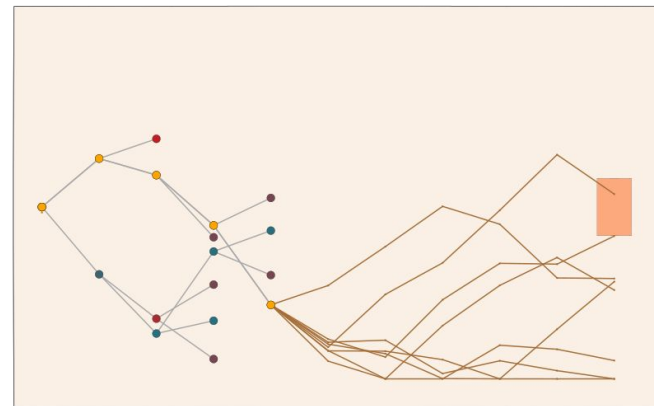
# What About Complex Methods?

## Process Reward Models



(a) 6B verification test performance when given varying numbers of completions per problem to rank.

## MCTS



# Challenges: Legibility

*Although DeepSeek-R1-Zero exhibits strong reasoning capabilities and autonomously develops unexpected and powerful reasoning behaviors, it faces several issues. For instance, DeepSeek-R1-Zero struggles with **challenges like poor readability, and language mixing.***

# Challenges: General-Purpose Models

*For reasoning data, we adhere to the methodology outlined in DeepSeek-R1-Zero, which utilizes rule-based rewards to guide the learning process in math, code, and logical reasoning domains. For general data, we resort to reward models to **capture human preferences in complex and nuanced scenarios***

# Challenges: Scaling

Big question: Can this stage reach pretraining scale?

*Due to the long evaluation times, which impact the efficiency of the RL process, large-scale RL has not been applied extensively in software engineering tasks. As a result, DeepSeek-R1 **has not demonstrated a huge improvement over DeepSeek-V3 on software engineering benchmarks.***



# Proposed Solutions

- R1 -> Cold-started from human written data
- Non-Reasoning data -> use LLMs to provide CoT with basic checks
- Many of this feels like precursor to R2, R3 ...

# What's happened this week?

TinyZero

User: Using the numbers [19, 36, 55, 7], create an equation that equals 65.

Assistant: Let me solve this step by step.

`<think>` We need to find an equation using the numbers 19, 36, 55, and 7 exactly once, with basic arithmetic operations, that equals 65. One possible combination is  $55 + 36 - 19 + 7$ . Let's check:  $55 + 36 = 91$ ,  $91 - 19 = 72$ , and  $72 + 7 = 79$ . However, this doesn't equal 65. Another combination is  $55 + 36 + 7 - 19$ . Let's check:  $55 + 36 = 91$ ,  $91 + 7 = 98$ , and  $98 - 19 = 79$ . This also doesn't equal 65. After trying different combinations, I found that  $55 + 36 + 7 - 19 = 79$ , which is close but not equal to 65. Finally, I found that  $55 + 36 - 19 + 7 = 79$ , which is still not equal to 65. However, if we try  $55 + 36 - 7 + 19$ , we get  $55 + 36 = 91$ ,  $91 - 7 = 84$ , and  $84 + 19 = 103$ , which is not equal to 65. After further trial and error, I found that  $55 + 36 + 7 - 19 = 79$ , which is still not equal to 65. However, if we try  $55 + 36 - 7 - 19$ , we get  $55 + 36 = 91$ ,  $91 - 7 = 84$ , and  $84 - 19 = 65$ . This equals 65. `</think>`

`<answer>`  $55 + 36 - 7 - 19$  `</answer>`

[← Back to Articles](#)

## Open-R1: a fully open reproduction of DeepSeek-R1

Published January 28, 2025

[Update on GitHub](#)



[@eliebak](#)  
Elie Bakouch



[@lvwerra](#)  
Leandro von Werra



[@lewtun](#)  
Lewis Tunstall

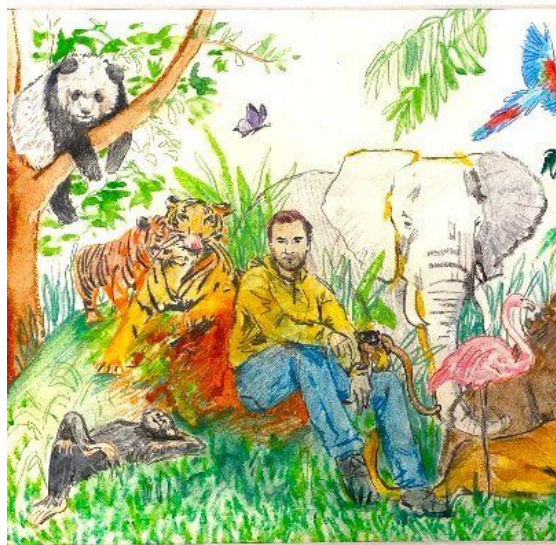
### What is DeepSeek-R1?

If you've ever struggled with a tough math problem, you know how useful it is to think a little longer and work through it carefully. OpenAI's o1 model showed that when LLMs are trained to do the same—by using more compute during inference—they get significantly better at solving reasoning tasks like mathematics, coding, and logic.

## Mini-R1: Reproduce Deepseek R1 „aha moment“ a RL tutorial

January 30, 2025 21 minute read [View Code](#)

The release of Deepseek R1 shocked the industry. Why? Well, DeepSeek-R1 is an open model that rivals OpenAI's o1 in complex reasoning tasks, introduced using Group Relative Policy Optimization (GRPO) and RL-focused multi-stage training approach. They not only released the model, but also a research paper on how they did it.



For friends of open source: imo the highest leverage thing you can do is help construct a high diversity of RL environments that help elicit LLM cognitive strategies. To build a gym of sorts. This is a highly parallelizable task, which favors a large community of collaborators.

- Andrei Karpathy

What Now?

AI EFFECT

## Nvidia sheds almost \$600 billion in market cap, biggest one-day loss in U.S. history

PUBLISHED MON, JAN 27 2025-4:08 PM EST | UPDATED MON, JAN 27 2025-5:26 PM EST



Samantha Subin  
@SAMANTHA\_SUBIN

WATCH LIVE

### KEY POINTS

- Nvidia shares plunged 17% on Monday, resulting in a market cap loss of close to \$600 billion, the biggest drop ever for a U.S. company.
- The sell-off, which hit much of the U.S. tech sector, was sparked by concerns about increased competition from Chinese AI lab DeepSeek.
- Data center companies that rely on Nvidia chips also plummeted, with Dell, Oracle and Super Micro Computer all falling by at least 8.7%.

???



The 2010s were the age of scaling, now we're back in the age of wonder and discovery once again. Everyone is looking for the next thing. Scaling the right thing matters more now than ever.

- Ilya

<https://www.reuters.com/technology/artificial-intelligence/openai-rivals-see-k-new-path-smarter-ai-current-methods-hit-limitations-2024-11-11/>



From 2020-2023, the main thing being scaled was pretrained models: models trained on increasing amounts of internet text with a tiny bit of other training on top. In 2024, the idea of using reinforcement learning (RL) to train models to generate chains of thought has become a new focus of scaling.

- Dario Amodei



We've grown accustomed to Moore's Law "falling from the sky" — waiting 18 months for better hardware and software. Scaling Law is treated similarly. However, these advancements are the result of generations of relentless effort by Western-led technology communities. Because we haven't been actively involved in this process, we've come to overlook its significance.

— Liang Wenfeng (DeepSeek Founder)—

<https://thechinaacademy.org/interview-with-deepseek-founder-were-done-following-its-time-to-lead/>





Published by  
Brian Lehrer Show Segments by The Brian Lehrer Show

## Why China's AI Startup 'DeepSeek' is a Sputnik Moment

▶ Listen

+ Queue



# OpenAI

 deepseek