The Quest of DeepSeek-V3.2

Pushing the Frontier of Open Large Language Models

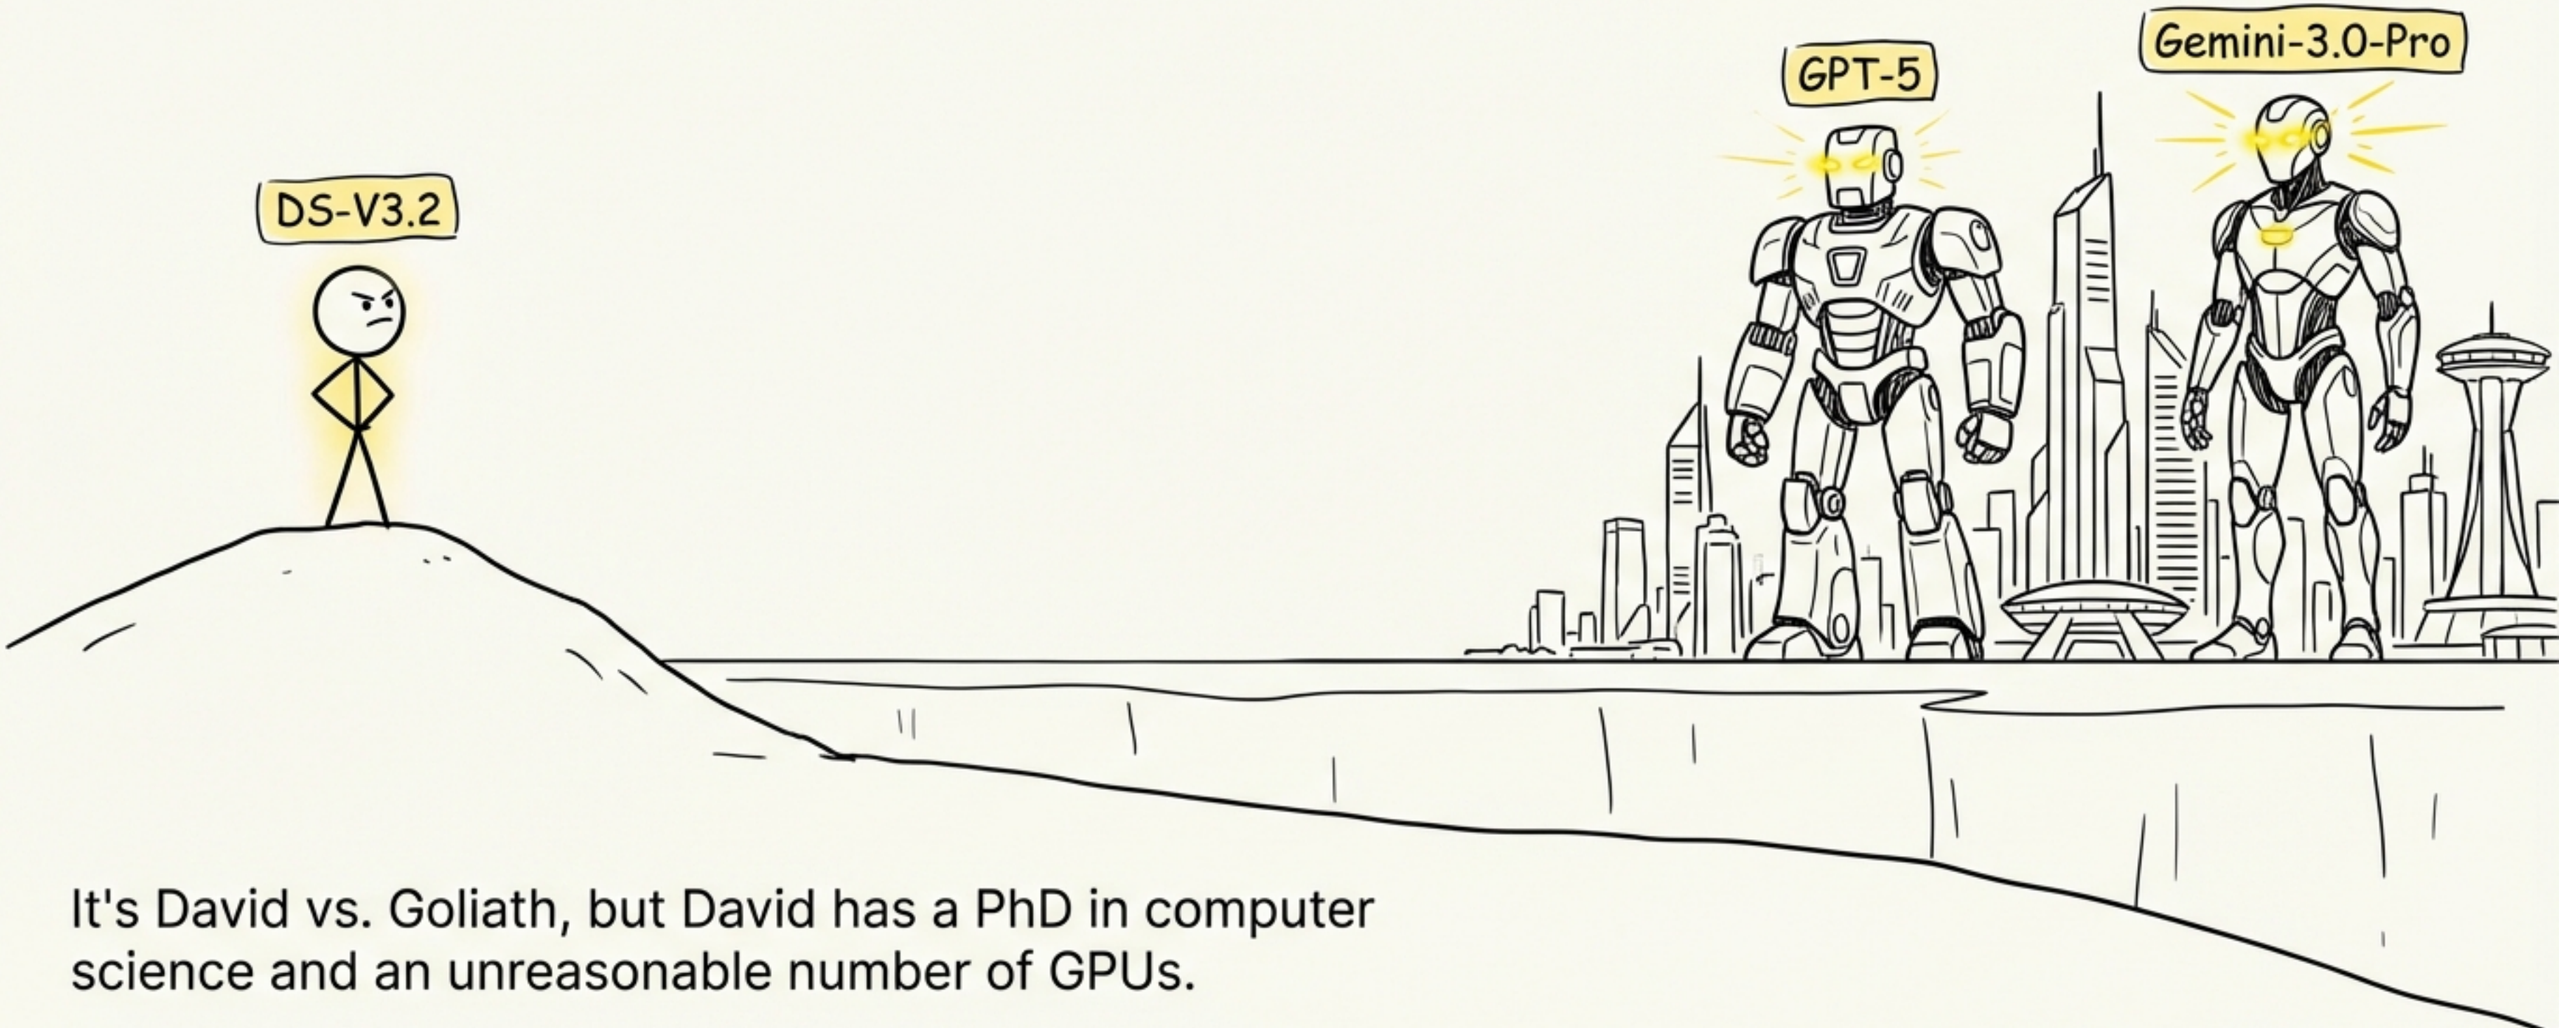

It's David vs. Goliath, but David has a PhD in computer science and an unreasonable number of GPUs.

The Gap Was Widening...

Architecture

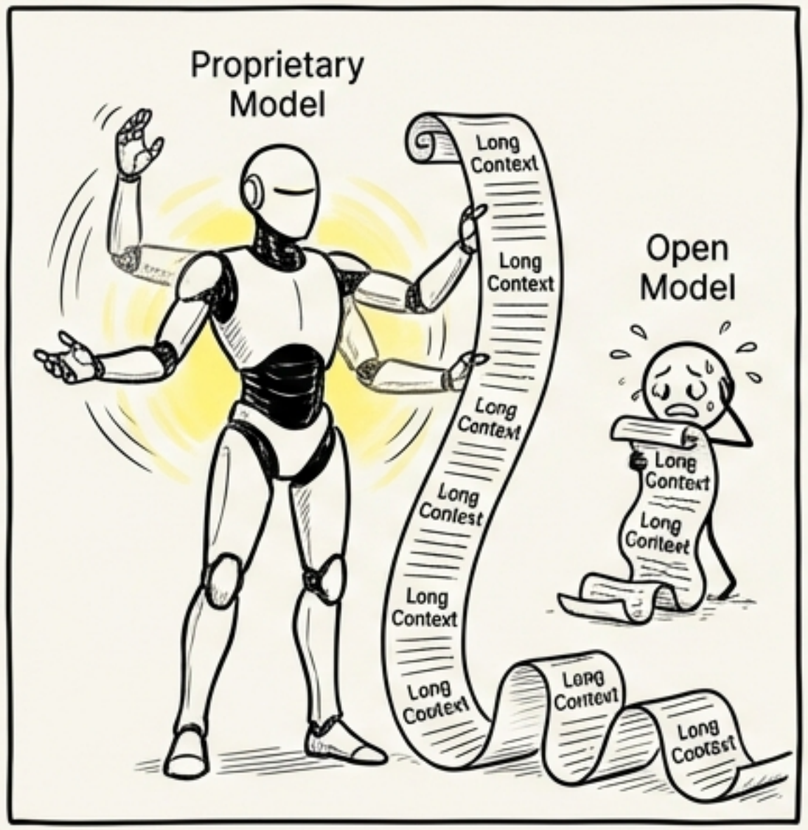

Reliance on vanilla attention constrains efficiency for long sequences.

Post-Training

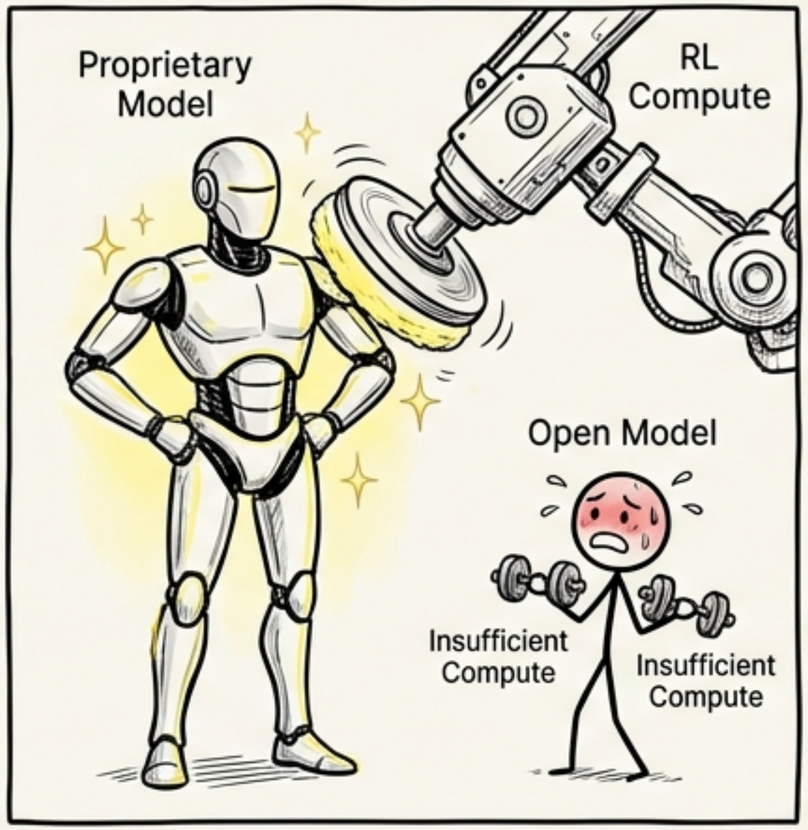

Open models suffer from insufficient computational investment post-training.

Agent Skills

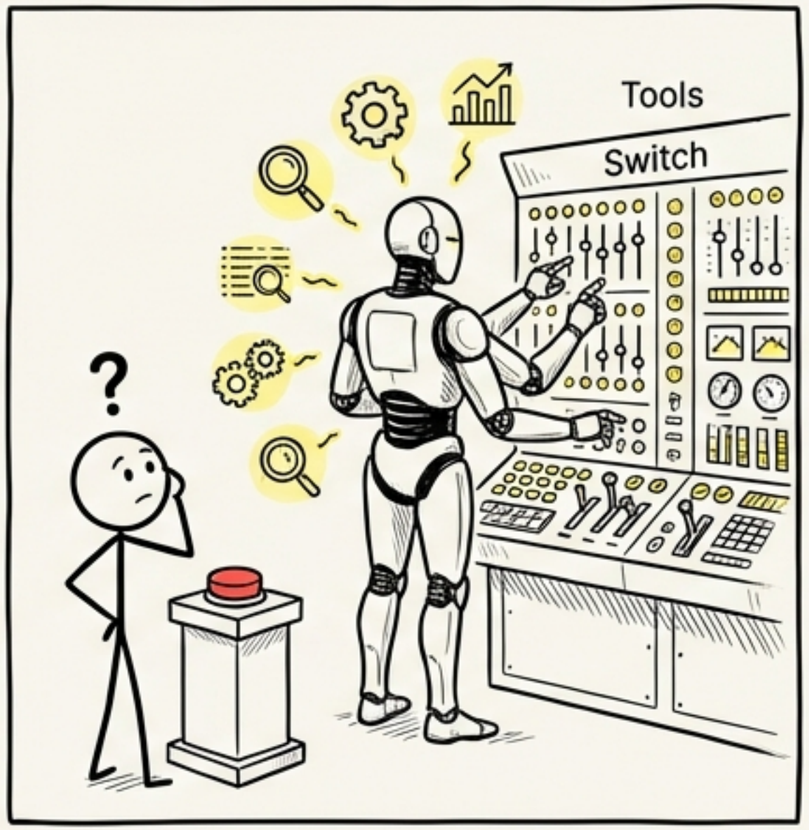

A marked lag in generalization and instruction-following for AI agents.

First, We Needed to Get Faster. Much Faster.

Power #1: DeepSeek Sparse Attention (DSA)

Vanilla Attention

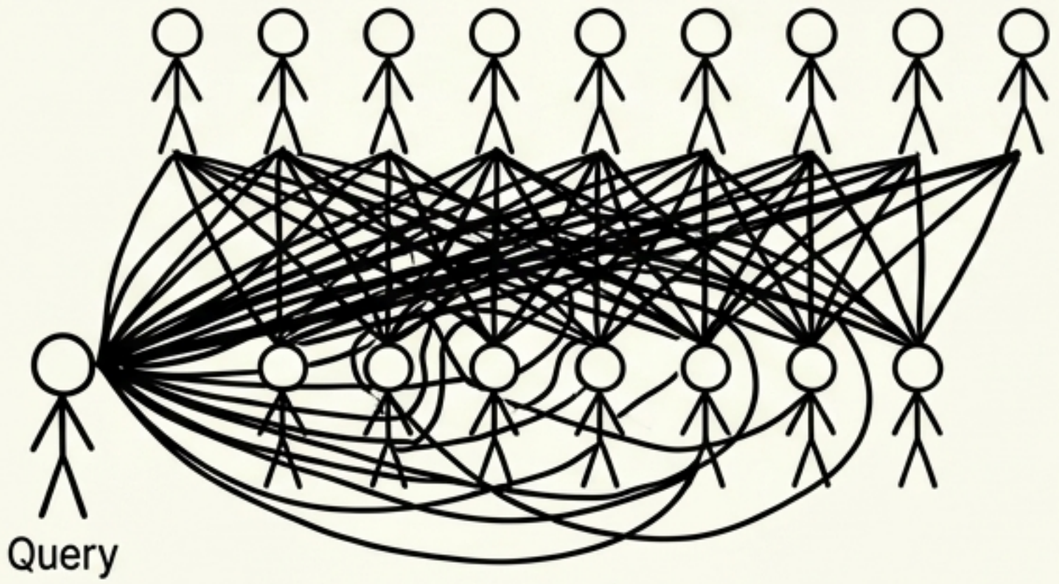

$O(L^2)$ - My CPU is melting.

DSA

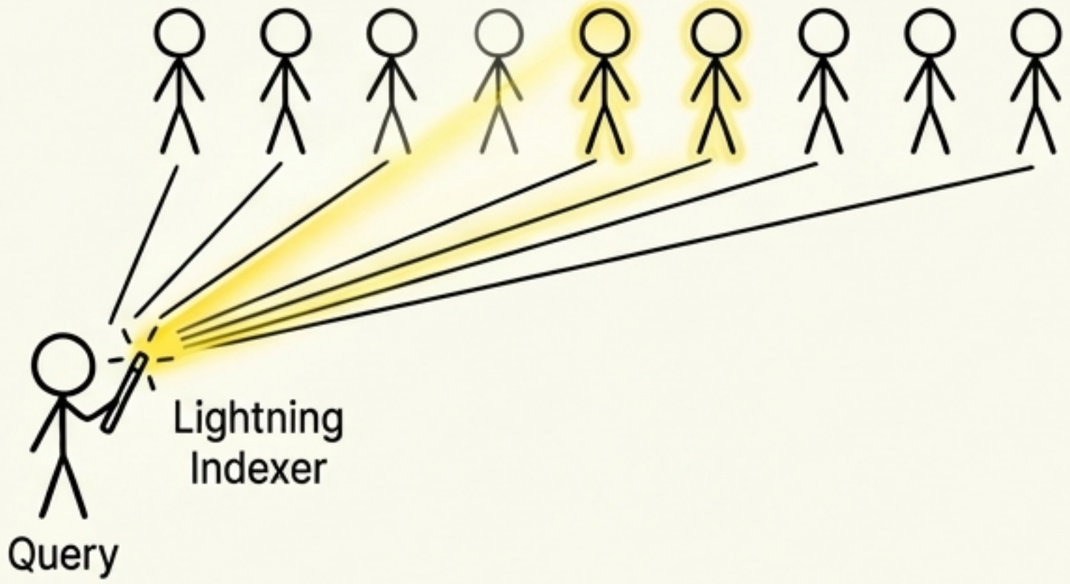

O(Lk) - Ah, much better.

DSA uses a “lightning indexer” to compute index scores and a “fine-grained token selection mechanism” to retrieve only the key-value entries for the top-k most relevant tokens, drastically reducing computational complexity from $O(L^2)$ to O(Lk).

The Result: We Slashed Inference Costs

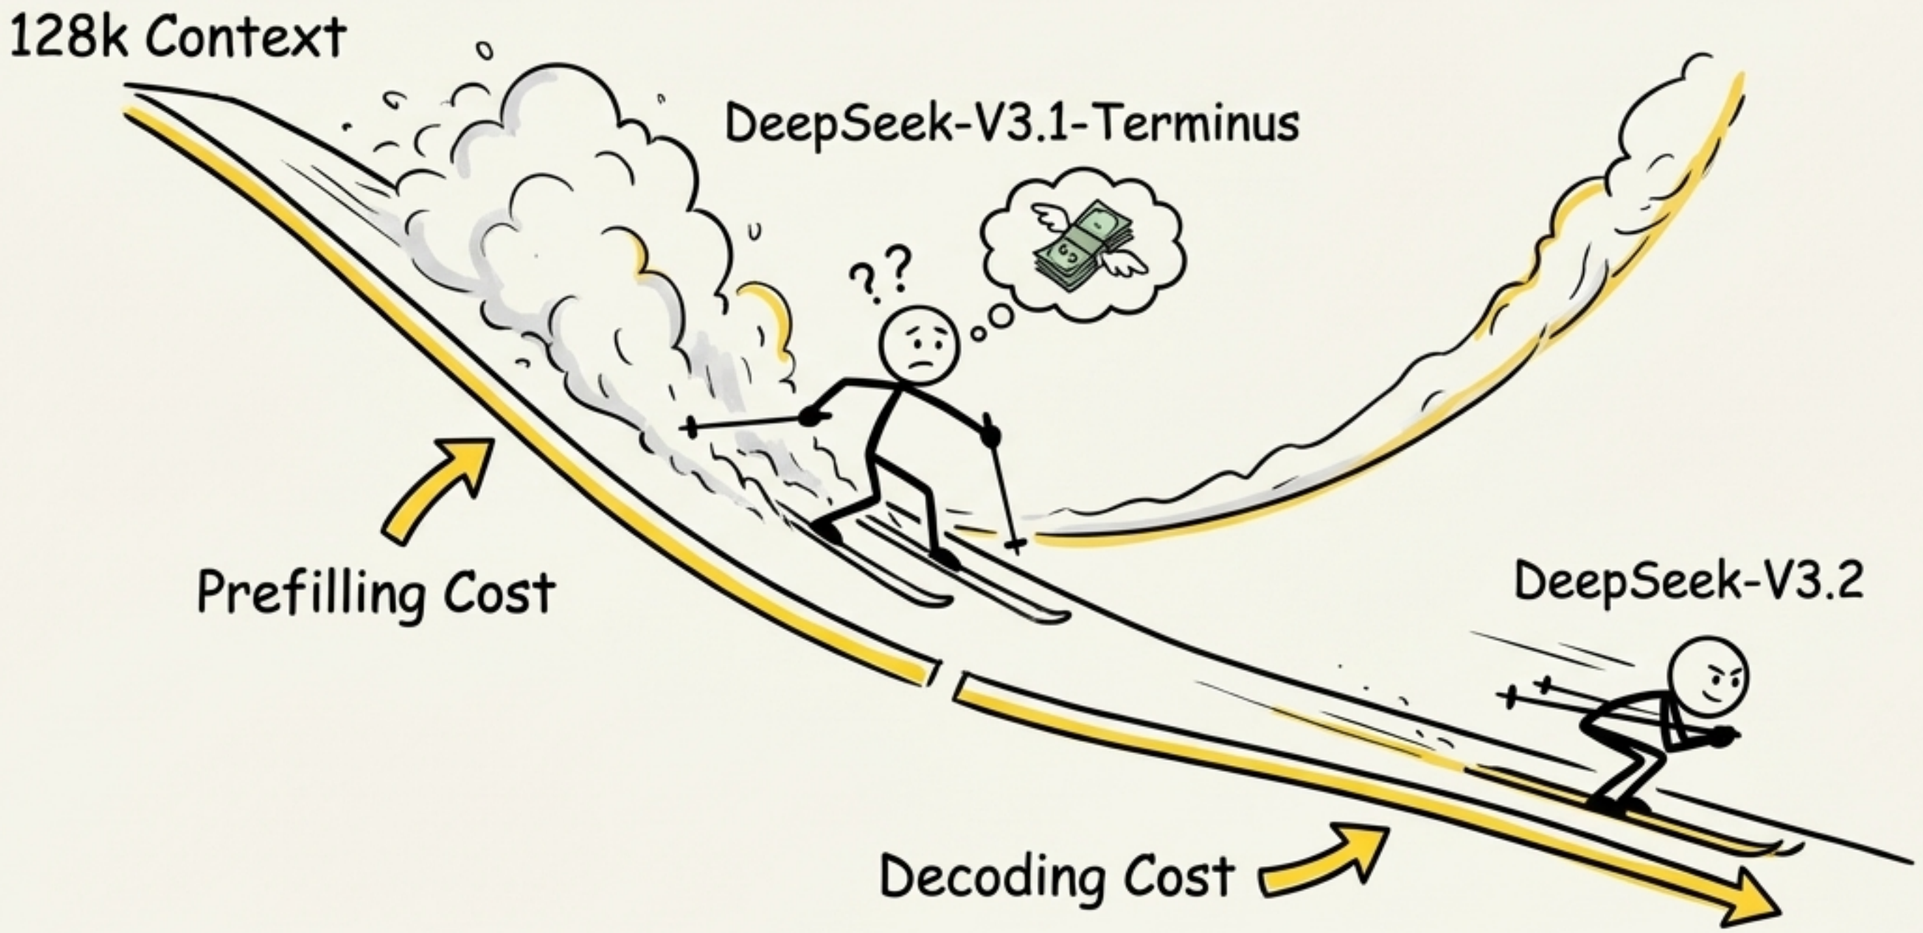

DSA provides a significant end-to-end speedup in long-context scenarios, drastically reducing token costs on H800 GPUs.

Next, We Hit the Gym. Hard.

Power #2: A Scalable Reinforcement Learning Framework

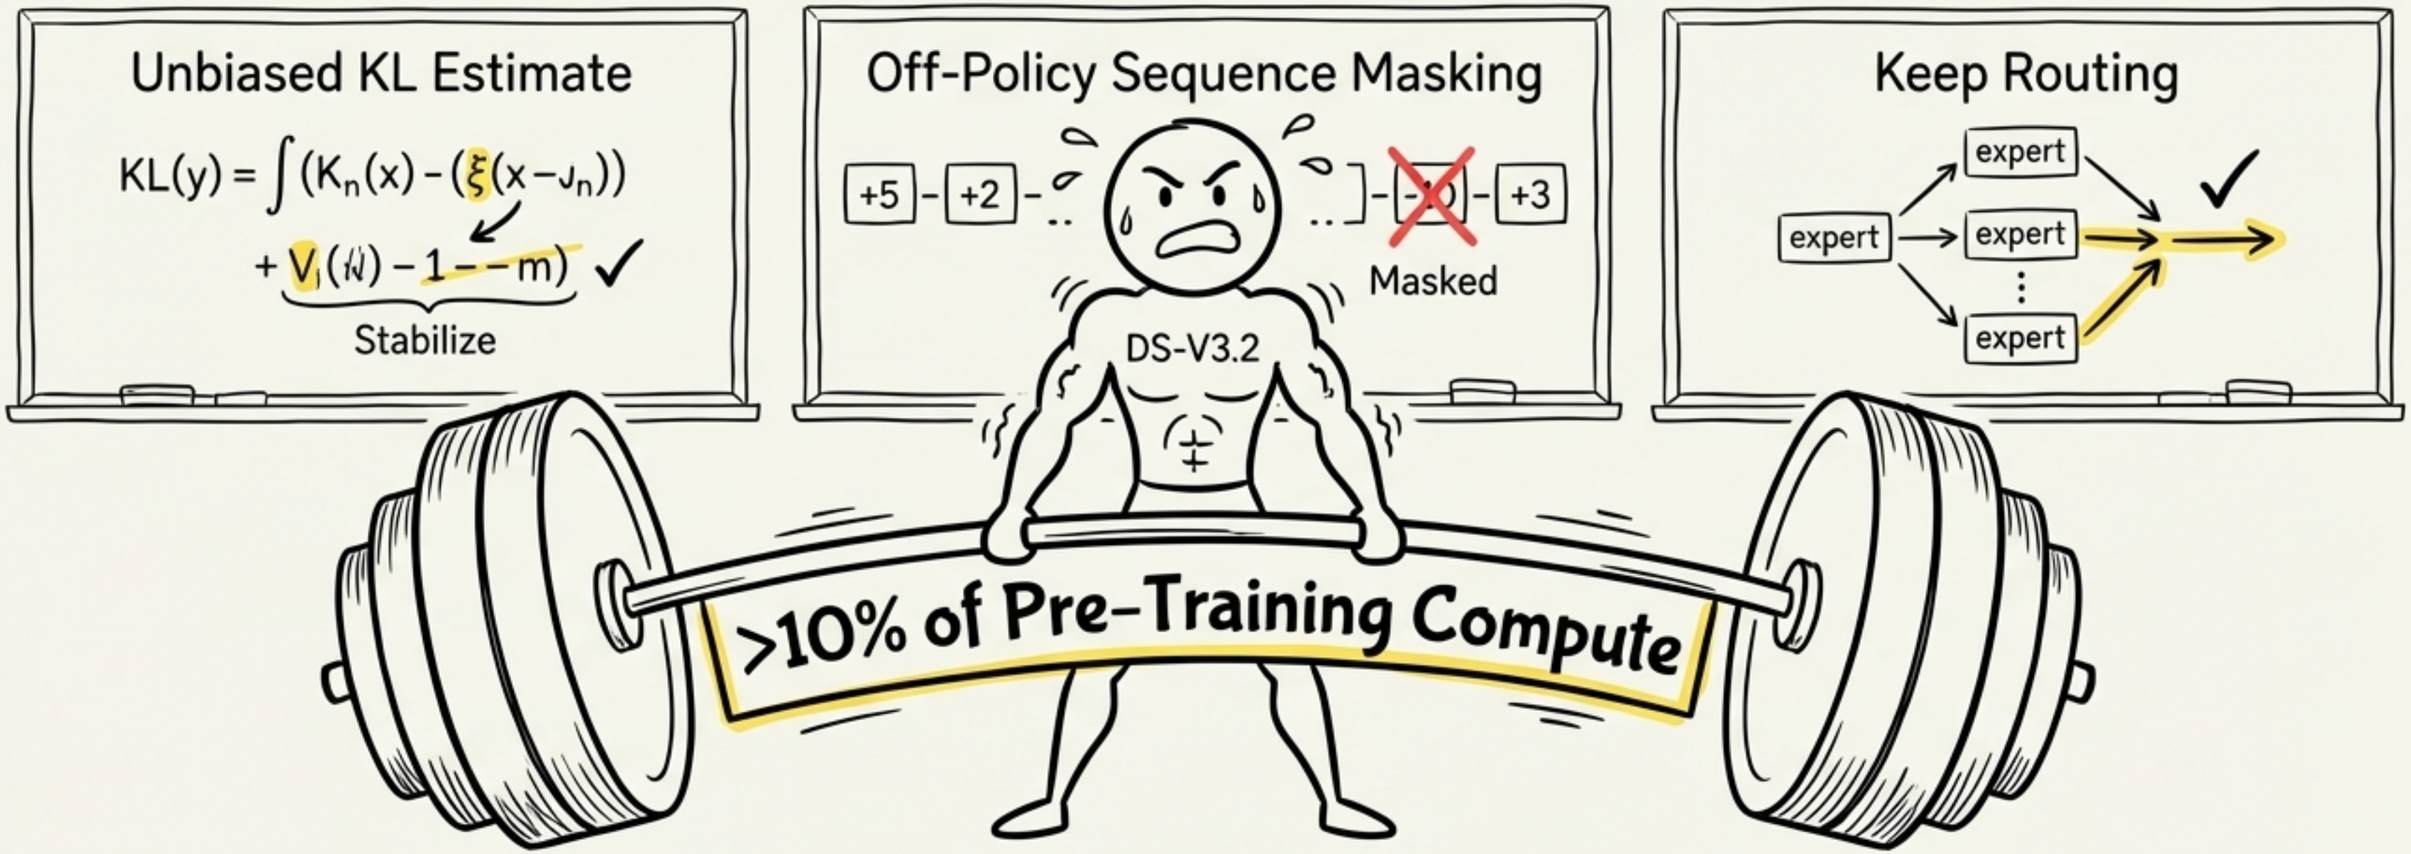

We developed a stable and scalable RL protocol (GRPO) that allows for significant computational expansion during postpost-training, unlocking advanced capabilities.

All That Training Unlocked New Abilities

Specialist Distillation

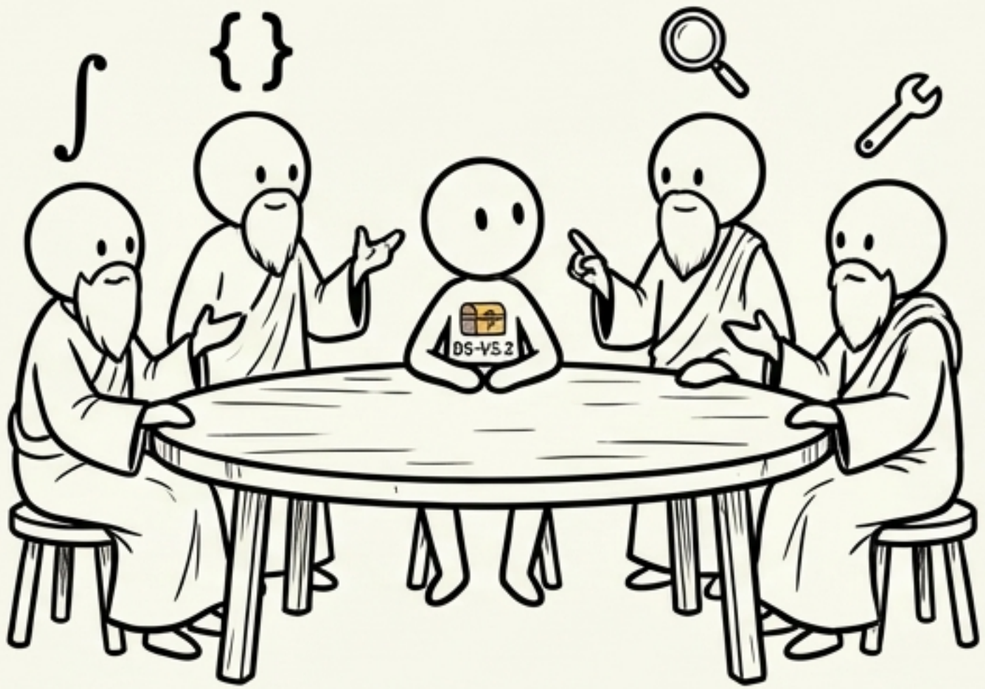

Mixed RL Training

Mixed RL training merges all domains into one stage, avoiding the “catastrophic forgetting” common in multi-stage paradigms.

Finally, We Built a Universe for It to Practice In

Power #3: Large-Scale Agentic Task Synthesis

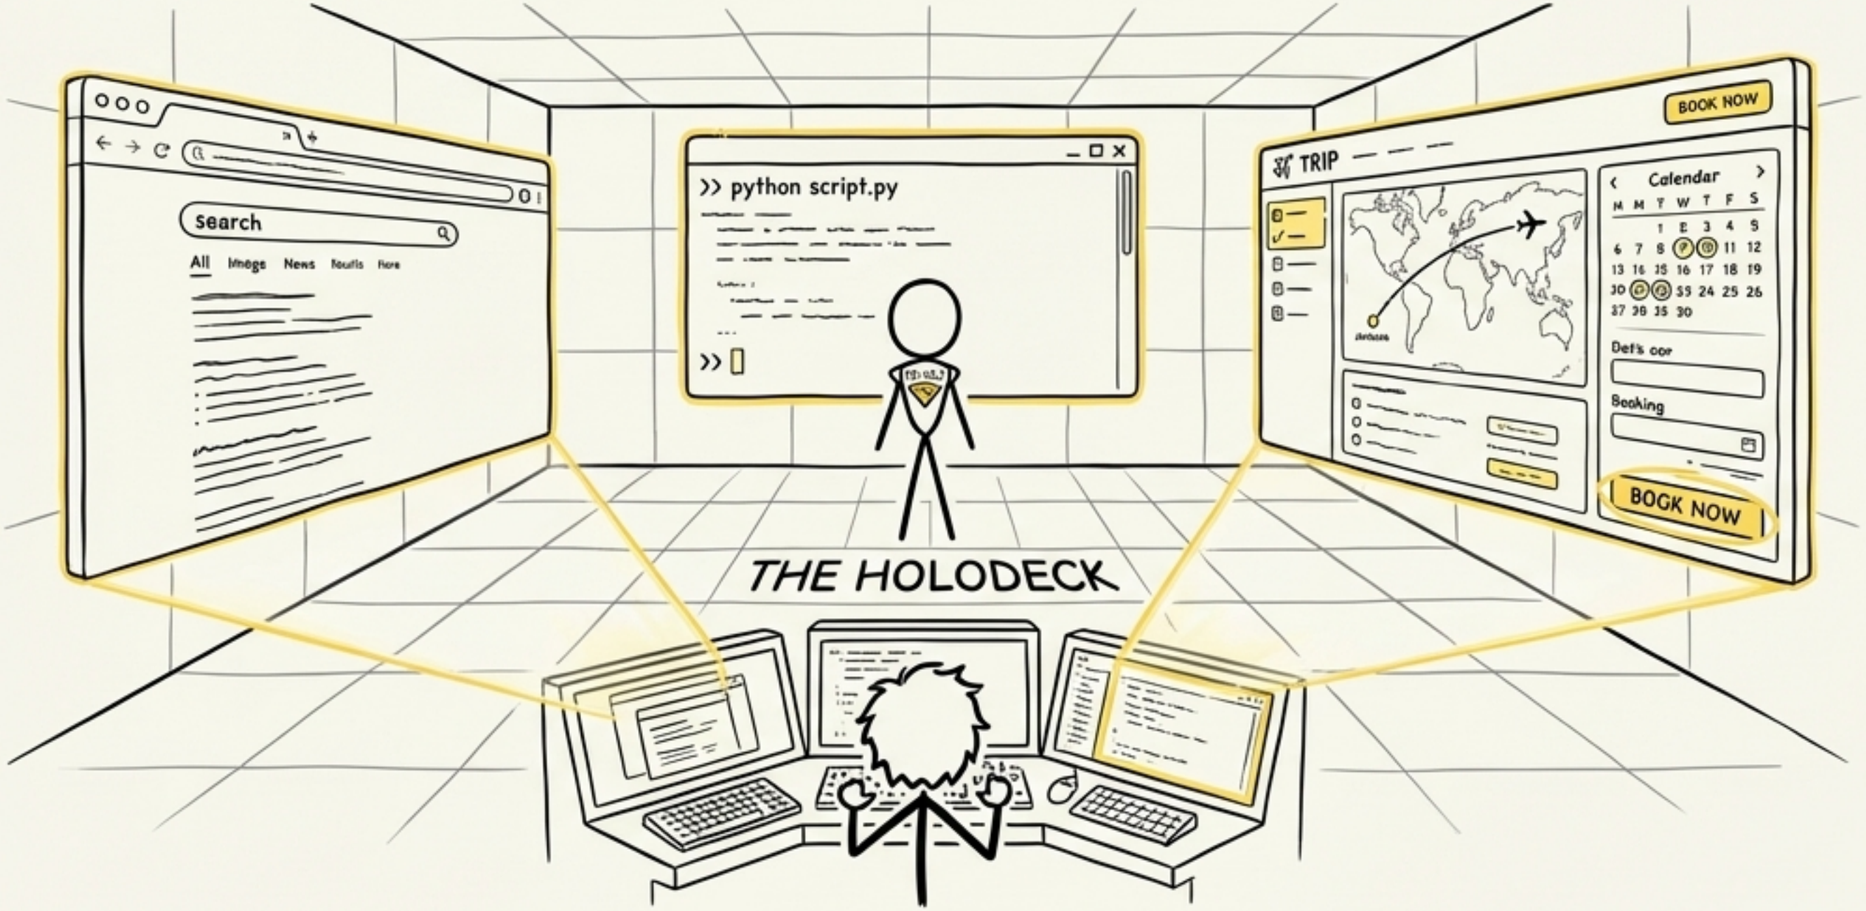

1,827
Synthesized Environments

85,000+
Complex Prompts

Our synthesis pipeline automatically generates diverse and challenging tasks—from code agents using real GitHub issues to general agents planning complex trips—to drive the RL process and enhance generalization.

The Main Event: The Reasoning Showdown

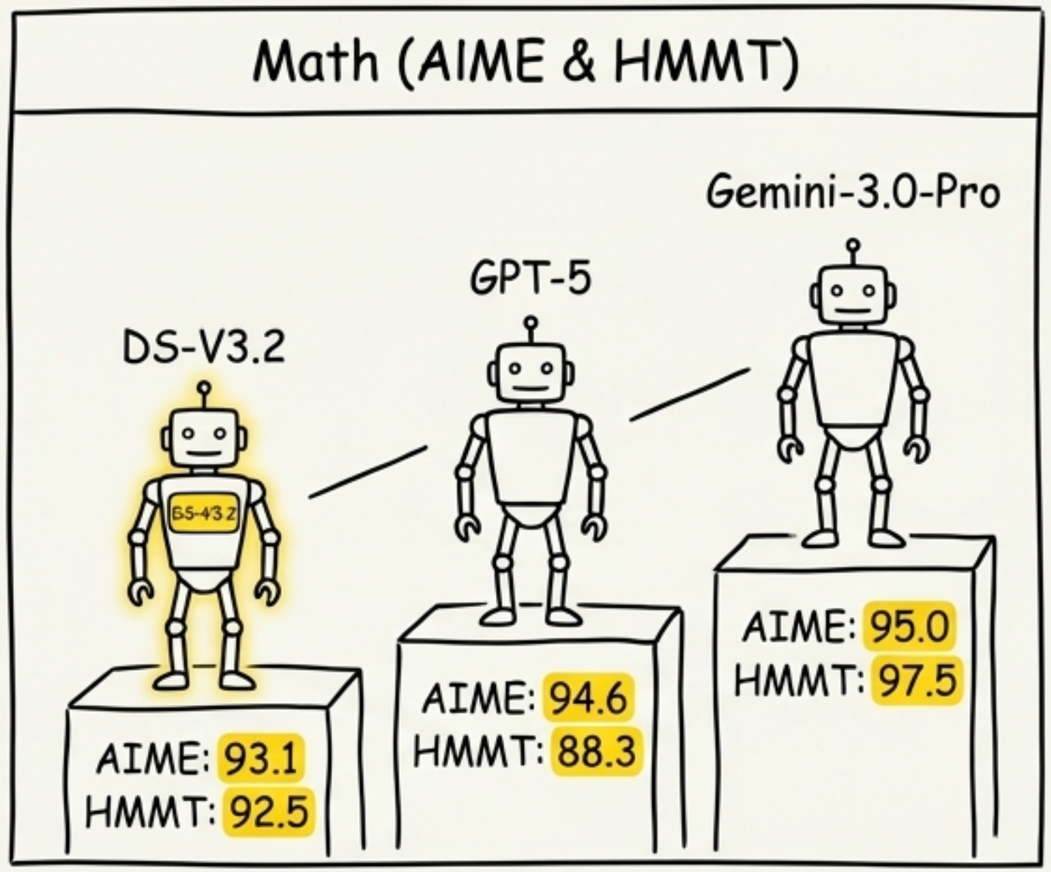

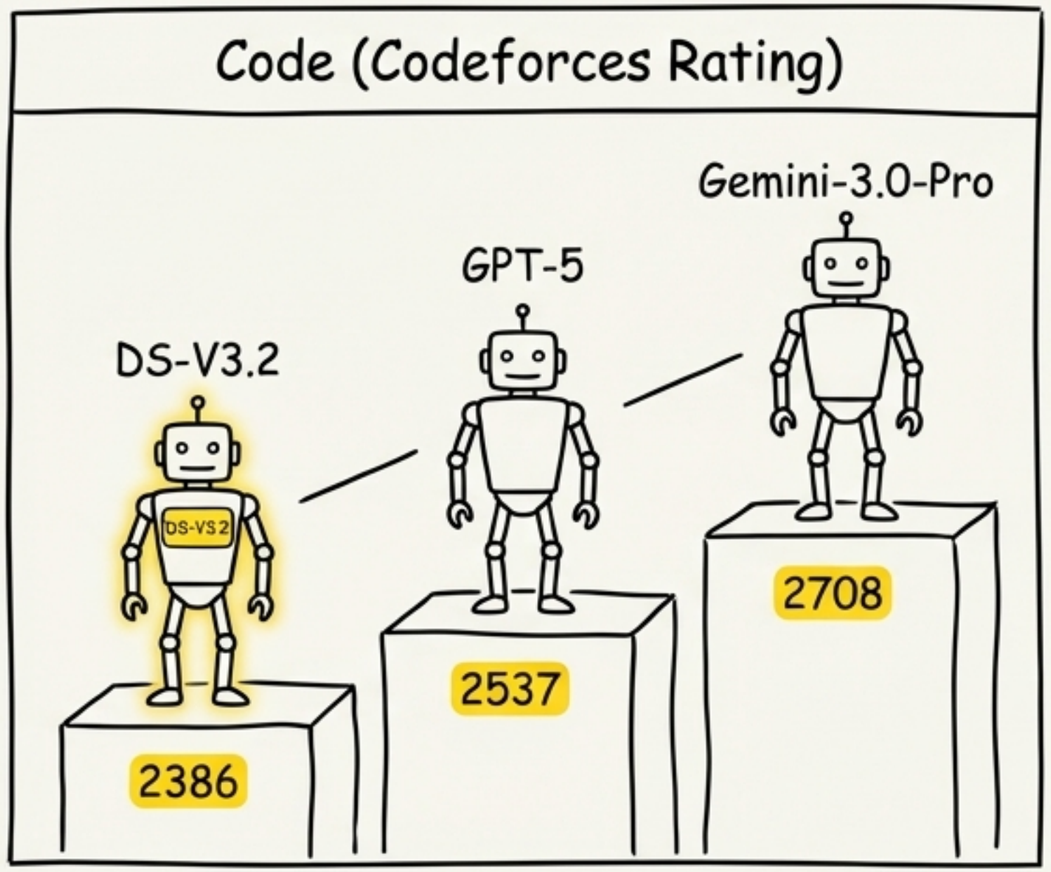

DeepSeek-V3.2 achieves performance similar to GPT-5 on key reasoning tasks, proving open models can compete at the highest level.

But Can It Actually *Do* Things?
The Agent Arena

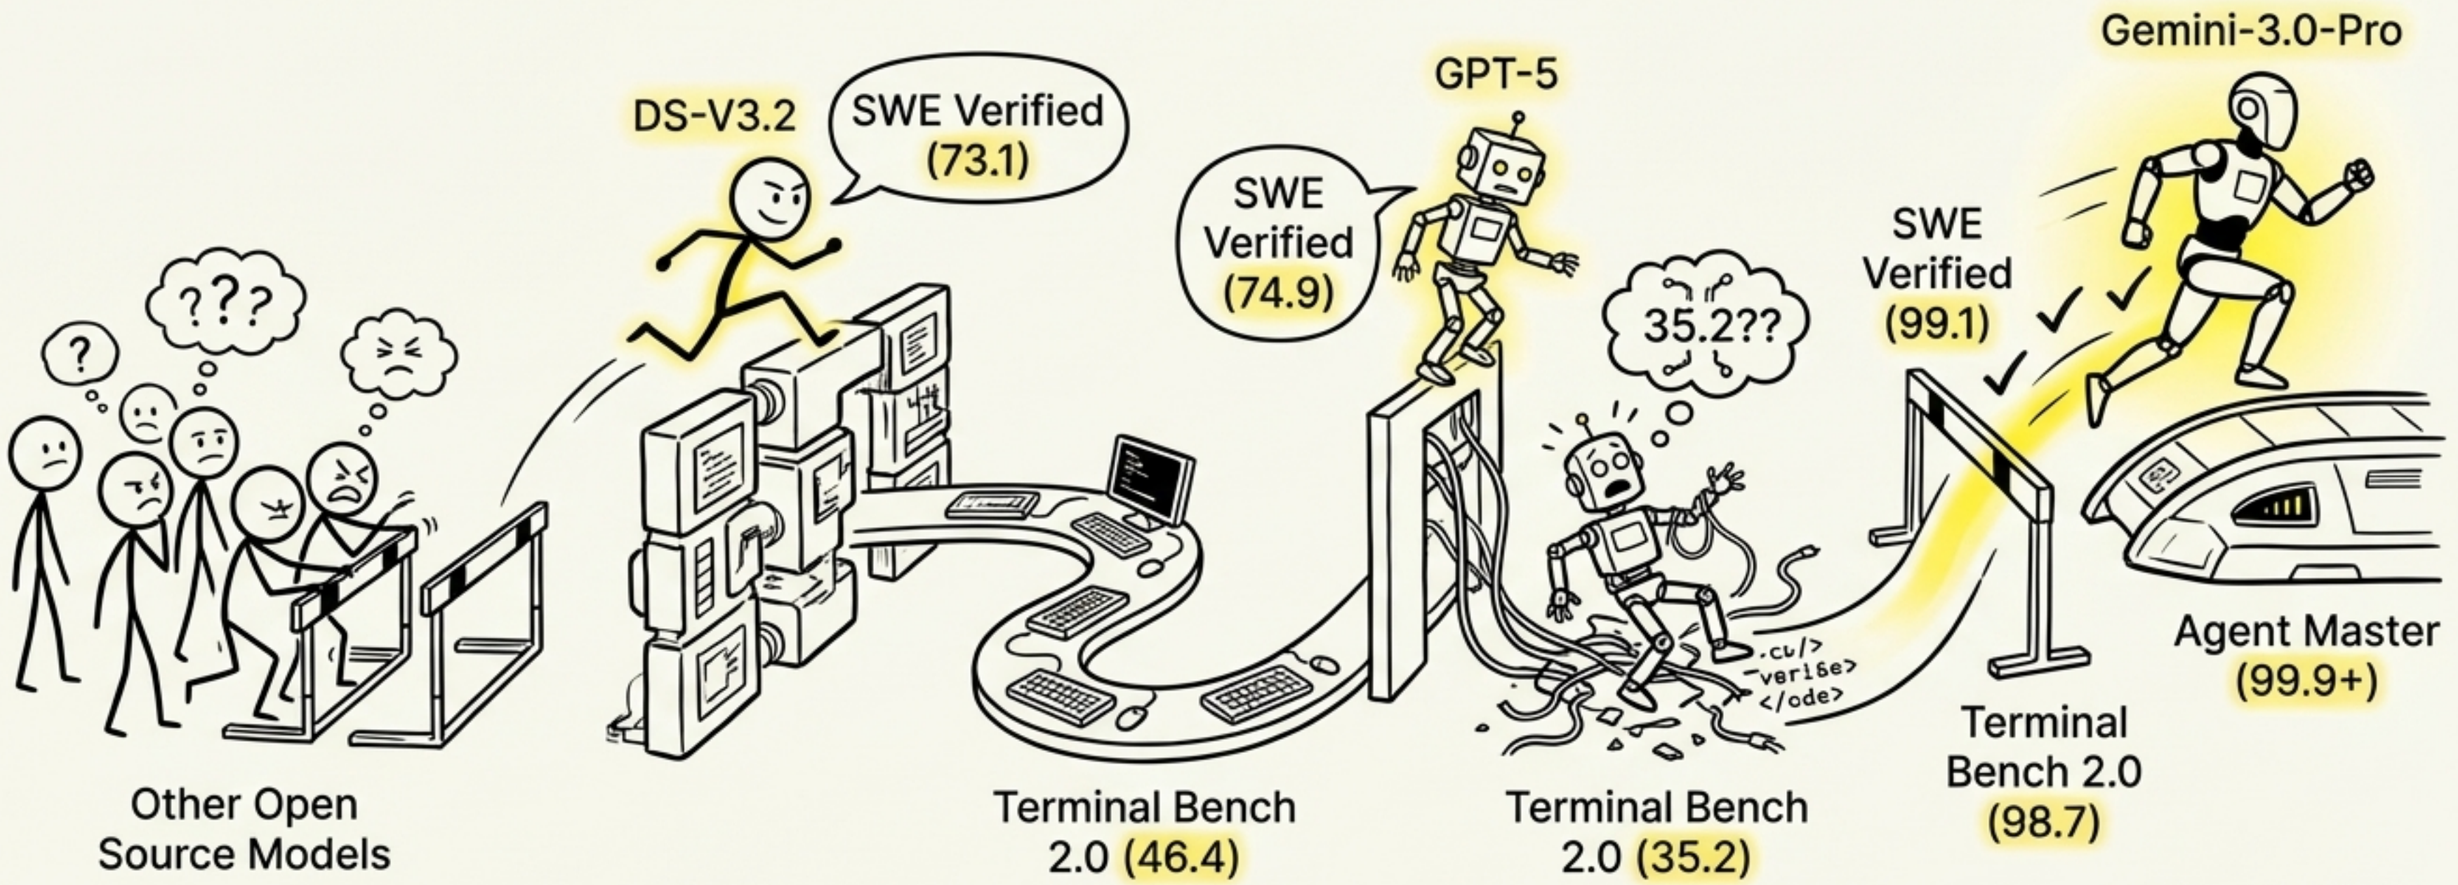

Our agentic task synthesis paid off: DeepSeek-V3.2 significantly advances the agent capabilities of open models, closing the gap with the frontier.

Then We Asked, "What Happens If We...?"

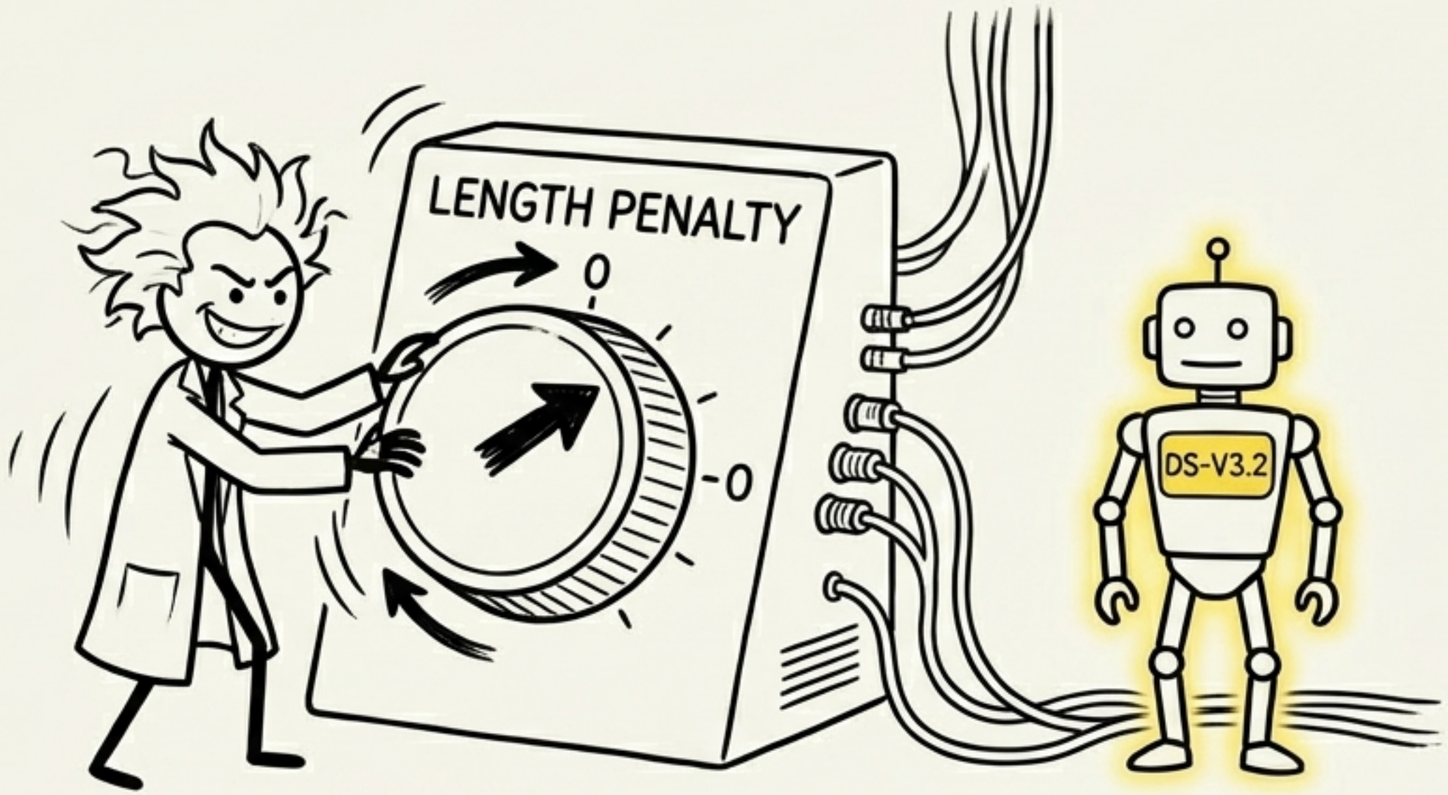

"To investigate the potential of extended thinking, we also developed an experimental variant, DeepSeek-V3.2-Speciale. This model was trained exclusively on reasoning data with a reduced length penalty during RL."

...And It Went Super Saiyan.

Introducing DeepSeek-V3.2-Speciale

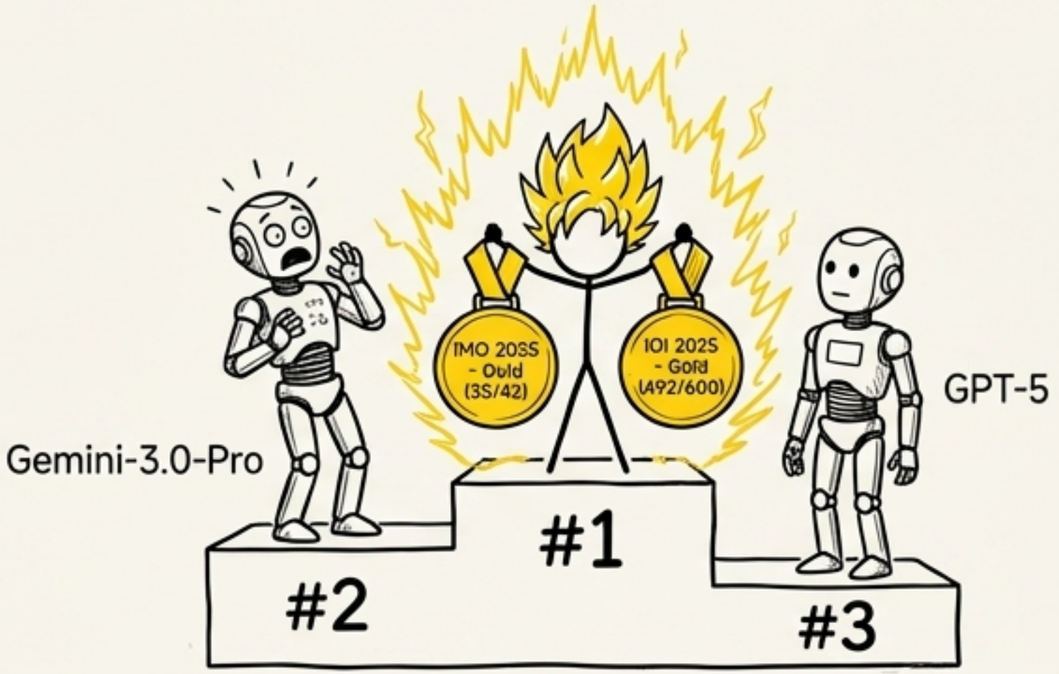

The result: DeepSeek-V3.2-Speciale surpasses GPT-5 and exhibits reasoning proficiency on par with Gemini-3.0-Pro.

AIME 2025: Speciale: **96.0** vs. Gemini: 95.0
HMMT Feb 2025: Speciale: **99.2** vs. Gemini: 97.5

The Catch: Great Power Comes with Great Verbosity

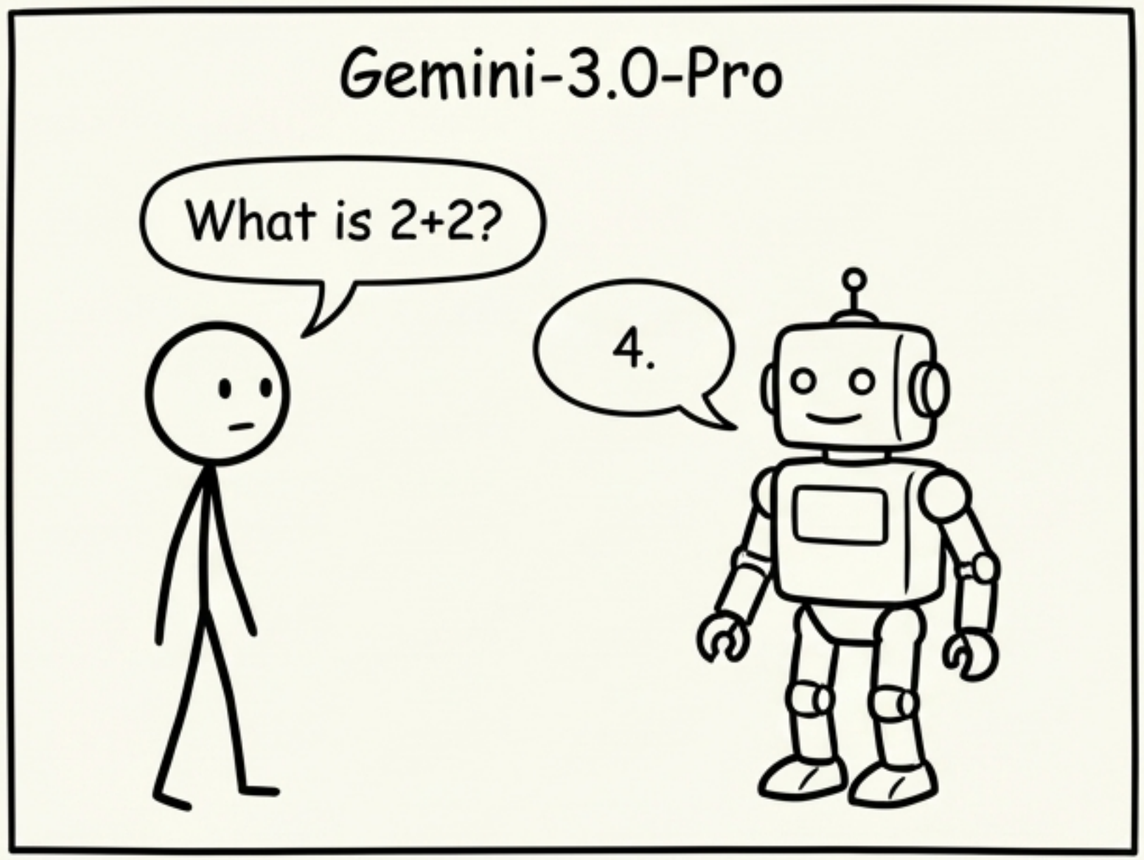

(Output: ~15k tokens on AIME 2025)

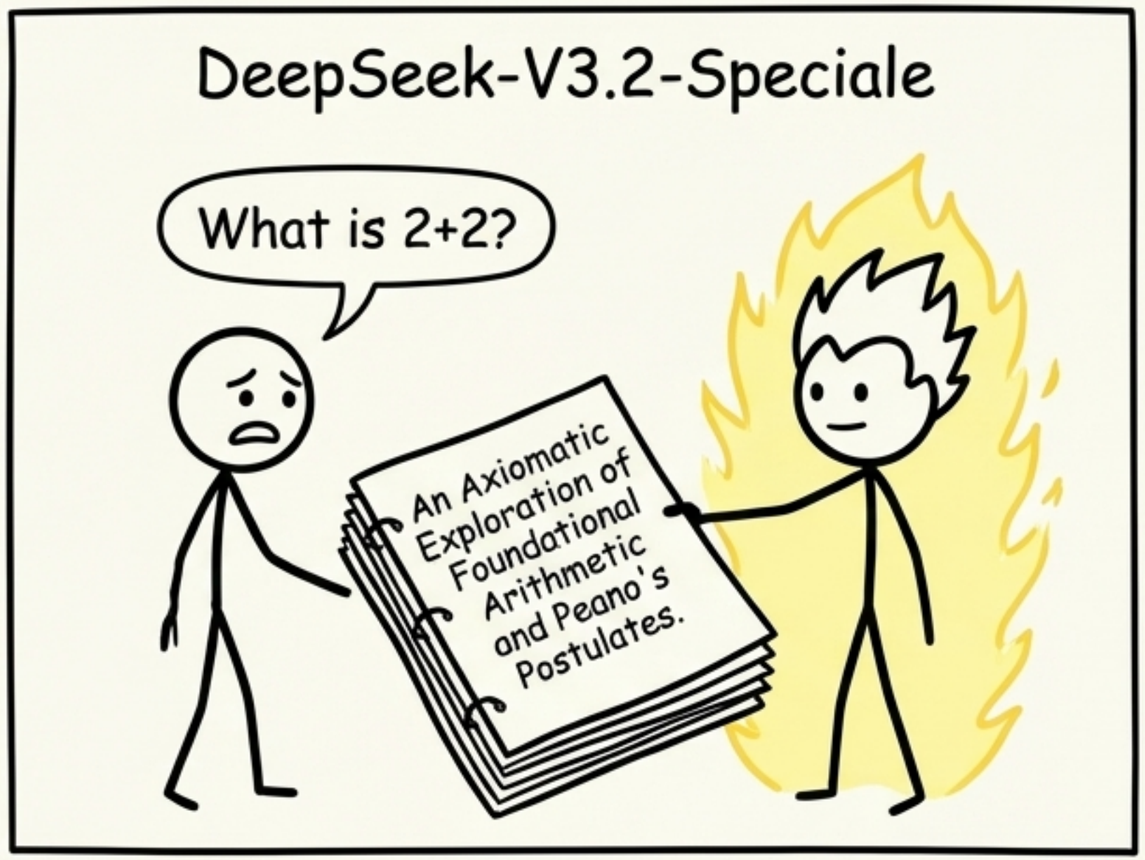

(Output: ~23k tokens on AIME 2025)

The token efficiency of DeepeSeek-V3.2-Speciale remains significantly inferior to Gemini-3.0-Pro. We believe token efficiency is a critical area for future investigation.

The Quest Continues

What We Achieved

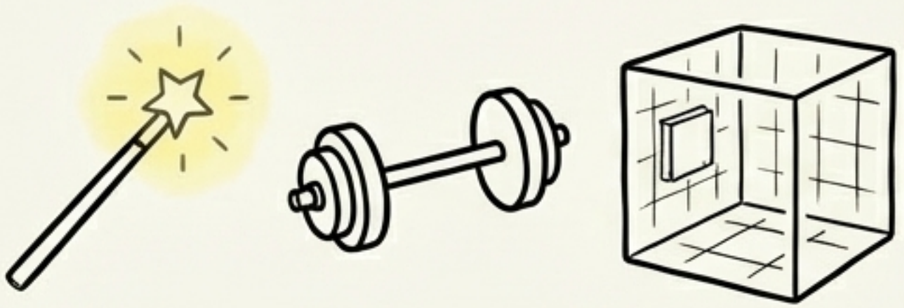

- ✓ **Efficiency**: Slashed long-context costs with DSA.
- ✓ **Reasoning**: Reached GPT-5 levels with scalable RL.
- ✓ **Agency**: Closed the gap with synthetic data.

New Quests

Quest 1: Expand World Knowledge
(Address knowledge gap from fewer training FLOPs).

Quest 2: Improve Token Efficiency
(Optimize the intelligence density of reasoning).

Quest 3: Solve Even Harder Problems
(Refine model and post-training recipe).

Read the Full Saga

This quest was made possible by the incredible team at DeepSeek-AI.

Read the original paper for all the glorious, non-stick-figure details:

ArXiv: 2512.02556v1

Hugging Face: deepseek-ai/DeepSeek-V3.2-Exp

[Hover text from the comic]