

# DeepSeek V4 Lands on an Unsteady Floor

👉 UNSTEADY FLOOR

2026-04-24 · 00:15:49

*“DeepSeek's V4 loss curve kept catching fire, and the team kept putting it out with bandages.”*

— LENAR KESS, TODAY'S NARRATION

DeepSeek V4 ships hours after GPT-5.5, and the technical report tells a more interesting story than the benchmark bars. Susan Zhang reads the paper out loud: anticipatory routing, logit clamps, and a training run that kept catching fire at 33 trillion tokens. I walk through what the fragility actually means for anyone planning to finetune on top of it.

On the OpenAI side, GPT-5.5 lands with a quiet thud on Victor Taelin's LamBench. Codex picks up a proper reviewer agent. A plugin called endless-toil makes your editor groan at bad code. Sapiens2 admits it trained on half of Flickr's humans. And Fireship spends a week automating his mom's IT support with a voice-cloned agent called OpenClaw.

— Lenar Kess

## CHAPTERS

- 00:00:04 Cold open — the floor under V4
- 00:01:27 DeepSeek V4, by the numbers
- 00:02:58 Susan Zhang reads the paper
- 00:05:22 GPT-5.5 lands, quietly
- 00:07:18 The V4-over-GPT-5.5 timing is probably not a coincidence
- 00:08:40 Codex grows a reviewer
- 00:10:23 Qwen 3.6 is not the release the benchmarks suggest

00:12:22 Sapiens2 and the quiet dataset admission

---

00:14:03 Endless-toil, the groaning plugin

---

CANONICAL

<https://braid.opentangle.com/episodes/2026-04-24.html>