

Open algorithms, closed weights, and the arithmetic of AI tooling

2026-05-15 / 00:09:30

“Open-sourcing the algorithm changes the conversation from trust us to inspect it yourself. But the part that actually decides your feed — the ranking model — is still proprietary.”

— Selin Oriax, today's narration

[X open-sourced its recommendation algorithm](#) — but the model it calls isn't public. [Cloudflare ran a benchmark](#) showing SDK-based agent coding costs 8.4× less than MCP dispatch. [arXiv drew a hard line](#) on unchecked LLM output. And [Anthropic's Mythos](#) raises the cost-vs-safety tension we keep running into.

Also: [Osaurus](#), the local-plus-cloud Mac harness, and [Figure AI's 30-hour robot run](#). A Friday of infrastructure stories.

CHAPTERS

- 00:00:04 X's open-sourced recommendation algorithm
- 00:01:47 Cloudflare's Code Mode vs MCP benchmark
- 00:03:52 arXiv's LLM policy enforcement
- 00:05:36 The Mythos cost-and-safety question
- 00:07:44 Smaller stories: Osaurus, hardware, and robot endurance

