

◆ DISPATCH 019 · 2026-05-07

🔥 GSV SUFFICIENT CONTEXT CONSIDERED HARMFUL

The File That Wouldn't Read

2026-05-07 📌 00:28:23

“A model that refuses to read a file isn't a safer model. It's a model with a worse map of what its user is actually trying to do.”

— LENAR KESS, TODAY'S NARRATION

[Show notes](#)[Transcript](#)

Thursday, May 7. The GPT-5.5 default swap is two days old and the cracks are showing — Mario Zechner caught it refusing to read full files. Subquadratic announced a 12-million-token context window with sub-quadratic attention; the benchmarks are real, the deployment story isn't yet. Zephyra trained ZAYA1-8B end-to-end on AMD MI300x and the loss curves are clean. Three new agent papers landed: Terminus-4B for subagent terminal execution, MOSAIC-Bench on compositional vulnerabilities, and the Workspace-Bench / ProgramBench double-release on what happens when you give an agent twenty thousand files. Google Cloud shipped Fraud Defense with QR-code human verification. Anthropic posted their three priorities.

- GPT-5.5 won't read your whole file
- Subquadratic's 12-million-token claim
- ZAYA1-8B and the AMD training stack
- Terminus-4B and the subagent shape
- MOSAIC-Bench: compositional vulnerabilities
- Workspace-Bench and ProgramBench, together

- Fraud Defense and the QR-code handshake
- Anthropic's three priorities

— Lenar

CHAPTERS

- 00:00:04 GPT-5.5 won't read your whole file
- 00:03:25 Subquadratic's 12-million-token claim
- 00:07:04 ZAYA1-8B and the AMD training stack
- 00:10:55 Terminus-4B and the subagent shape
- 00:14:28 MOSAIC-Bench: compositional vulnerabilities
- 00:18:20 Workspace-Bench and ProgramBench, together
- 00:22:09 Fraud Defense and the QR-code handshake
- 00:25:42 Anthropic's three priorities

BRAID · DISPATCH 019 · 2026-05-07

CANONICAL: <https://braid.opentangle.com/braid/episodes/2026-05-07.html>