

Mixture of experts active params, automated training loops, and the RL infrastructure pivot

2026-05-13 / 00:05:57

“AI's next frontier isn't bigger models — it's superlearners that pull new knowledge from continuous experience.”

— Selin Oriax, today's narration

Today on Braixd: the local pass looks at three concrete shifts in how models are built and served. We open with AIDC-AI's Ovis2.6, a multimodal model that packs 80 billion parameters into a Mixture-of-Experts architecture with only about 3 billion active at inference. We move to AutoScientist from Adaption, which attempts to automate the full research loop so small labs don't lose compounding on broken experiment pipelines. Finally, we look at NVIDIA and Ineffable Intelligence's push to build reinforcement learning infrastructure on Grace Blackwell and Vera Rubin, marking a clear pivot from static human data toward continuous, experience-driven training.

CHAPTERS

00:00:04 The mixture of experts squeeze: 80 billion parameters, 3 billion active

00:02:07 Automating the research loop

00:03:57 The shift from static data to continuous experience
