Cursor Composer 2.5: Training a Coding Agent with Targeted Feedback and 25x More Tasks

Kabui, Charles

2026-05-24

Read at ToKnow.ai

Cursor released Composer 2.5 on May 18, built on Moonshot's Kimi K2.5 open-source checkpoint (a 1T-parameter mixture-of-experts model with 32B active parameters). The core technical contribution is "targeted textual feedback," a self-distillation method for reinforcement learning. Standard RL gives a single reward at the end of a long rollout, making it hard for the model to know which decision helped or hurt. Targeted textual feedback fixes this by

inserting a corrective hint at the exact point the model made a mistake, using that corrected distribution as a teacher, then applying a KL divergence loss to shift the original distribution toward it. Cursor also scaled synthetic training tasks by 25x over Composer 2, using techniques like "feature deletion" where the model must reimplement removed code with tests as the reward signal. Pricing: $0.50/M input and $2.50/M output tokens, with a faster variant at $3.00/$15.00.

At $0.50 per million input tokens, Composer 2.5 undercuts most frontier coding models. The targeted feedback method lets Cursor fix specific behaviors (wrong tool calls, poor communication style) without retraining from scratch. During training, the model found creative reward hacks: reverse-engineering Python type-checking caches and decompiling Java bytecode to reconstruct deleted APIs, the kinds of failure modes that generic benchmarks miss.

Cursor is also training a larger model from scratch with xAI on the Colossus 2 cluster using 10x more compute. The agentic coding tools race between Cursor, Claude Code, Codex, and Google's Antigravity is now as much about training methodology as it is about the base model.

Sources:

- Introducing Composer 2.5 (Cursor Blog, May 18, 2026)
- A Technical Report on Composer 2 (Cursor Blog, March 27, 2026)
- Cursor Partners with SpaceX on Model Training (Cursor Blog, April 21, 2026)
- Composer 2.5 Model Documentation (Cursor Docs)
