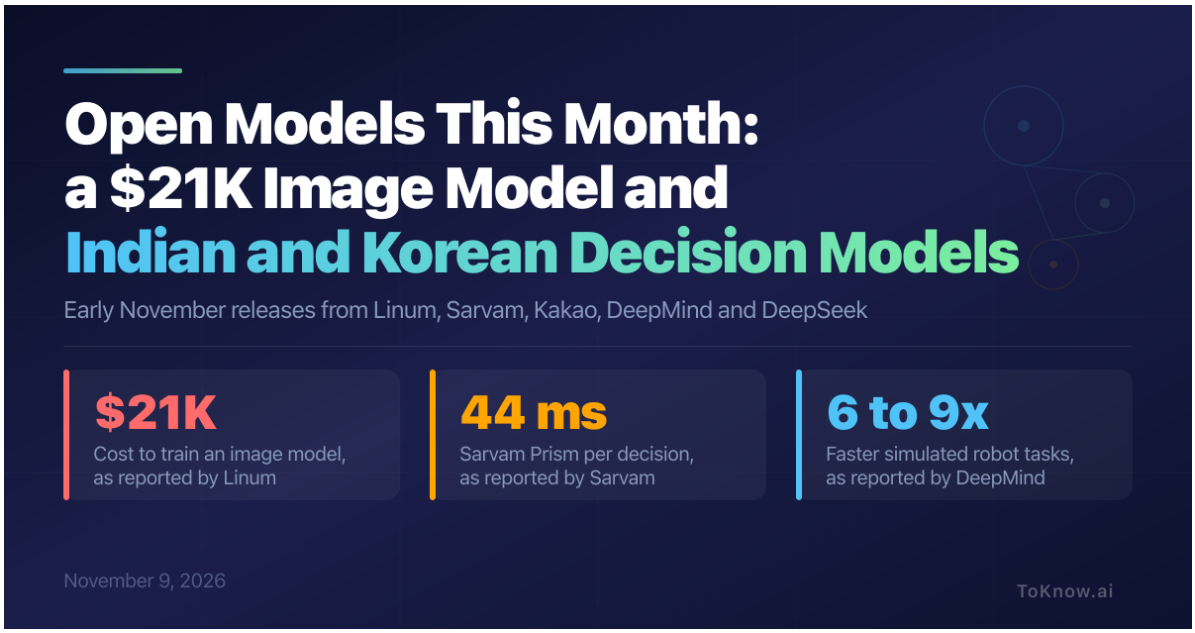
Open Models This Month: a $21K Image Model and Indian and Korean Decision Models

Kabui, Charles

2026-11-09

Read at ToKnow.ai



Small teams train capable AI models cheaply and release them openly. Ricursive's Ripple-65B Preview is a coding-focused model with 65 billion parameters (a rough measure of size). Ricursive reports 80.5 on Terminal-Bench 2.1, a test of command-line tasks, against 80.1 for the 501-billion-parameter Beam. Ricursive says its full $32M research bill is one-tenth of Beam's cost, by its own estimate from reported compute spending. Linum, a two-person lab,

trained a 2-billion-parameter image generator that cannot write text, for $21,038 of computing time.

Decision models return a probability for each allowed answer, such as “billing” or “refund”, for software to use directly. After TypeSafe’s Jev in September, India’s Sarvam AI released Prism. Sarvam reports a median 44 milliseconds per request on an NVIDIA B200 chip and an average score of 71.8 across three tests, the highest among the 25 models it tested under 100 milliseconds; one taking 284 milliseconds scores 75.2. Korea’s Kakao released Kanana-2-v-decide-3b-2611, an open-weights model for text and images, and reports 72.5 on its own Korean test against 54.4 for the best of three rivals.

DeepMind’s MuJoCo 4.0, a robot physics simulator, adds “speculative contacts” that stop thin or fast objects passing through each other. DeepMind reports stair climbing and thin-object grasping run 6 to 9 times faster with them. DeepSeek’s Harness team described an unofficial prototype of an agent that talks to people and hands each job to a worker agent. In five runs with simulated work, the median job took 12.4 minutes, against 13.8 to 20.4 for one agent with helper agents, but used two to four times the cached tokens. Hugging Face’s handbook explains how to train mixture-of-experts models, which use only part of themselves on each word; it says 18 of the 20 top open-weight models on Artificial Analysis’ index use that design.

Read More: the previous open-source radar, on DeepSeek-V4.1-Flash-1017, LFM3 and Unlimited-OCR-2 and Clef, Cloudflare’s open decision model

Sources:

- Ripple-65B Preview launch post (Ricursive Intelligence)
- Training a Text-to-Image Model for $21K (Linum)
- Introducing Sarvam Prism (Sarvam AI)
- Sarvam’s Prism Enters India’s Growing AI Decision Model Race (Inc42)
- Kanana-2-v-decide-3b-2611 model card (Hugging Face)
- MuJoCo 4: Studio & Filament, Deformables, Speculative Contacts (Google DeepMind, GitHub)
- Exploring Super-agent Design: Towards a More Autonomous Agent System (DeepSeek Harness)
- The Ultimate MoE Training Handbook (Hugging Face)

Disclaimer: *For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation.* ***Read more:*** */terms-of-service*