

Hello folks,

Here's what caught my eye in the latest AI research news.

--- SECTION 3: ALSO WORTH NOTING ---

- **Apple ML Research** -- Introduces VICIS, a task evaluating whether vision-language models can infer shared visual concepts from small image sets and apply them to new inputs -- a capability current VLMs struggle with. [link](#)
- **Apple ML Research** -- Proposes leveraging low-influence points to reduce the computational cost of machine unlearning, challenging the assumption that all data points in a forget set need equal treatment. [link](#)
- **Hugging Face / NVIDIA** -- Announces large-scale fine-tuning of video and image diffusion models using NVIDIA NeMo Automodel paired with Diffusers, though technical details are sparse in the announcement. [link](#)
- **NVIDIA** -- Frames the upcoming Vera Rubin architecture around maximizing "intelligence per dollar" for post-training workloads, positioning it for the agentic AI era. [link](#)

Back in 48 hours.

P.S. 15 RSS feeds | last 48h | 18 entries fetched, 4 scored ≥ 5