

Poolside Releases Laguna: Open Coding Models That Run on One GPU

Kabui, Charles

2026-06-12

[Read at ToKnow.ai](#)

**Poolside Releases Laguna:
Open Coding Models That
Run on One GPU**

A frontier lab's first public models, free under Apache 2.0

- 68.2%**
Coding tasks solved on SWE-bench Verified
- 1 GPU**
Runs the open-weight Laguna XS.2 locally
- Apache 2.0**
Open weights, free for commercial use

June 12, 2026

ToKnow.ai

Poolside, a foundation model lab that until now served only government and enterprise clients, released its first two public models on April 28. Both are built for agentic coding, where the model plans and carries out multi-step programming work on its own. The smaller [Laguna XS.2](#) has 33 billion total parameters but activates only 3 billion per token through a mixture-of-experts design, small enough to run on a single GPU. It ships with open weights under an

Apache 2.0 license, free to download, modify, and use commercially. On [SWE-bench Verified](#), a test of real GitHub bug fixes, it resolves 68.2% of tasks, ahead of Google's larger Gemma 4 at 52%. The bigger Laguna M.1, with 225 billion parameters and 23 billion active, reaches 72.5%.

A model that fits on one GPU lets a solo developer or small team run a capable coding assistant locally, for free, with no API bill and no code leaving their machine. XS.2 runs locally through Ollama and downloads from Hugging Face. Poolside also offers both models free for a limited time through its API and OpenRouter, and shares the heavier M.1 weights on request.

The strongest open-weight coding models have mostly come from Chinese labs like Alibaba's Qwen and DeepSeek. Poolside, which raised about \$2 billion last year at a \$12 billion valuation with backing from Nvidia, argues that Western labs have been too slow to open up. A well-funded US lab giving away capable coding weights is a genuine shift.

Read More: How Chinese open-weight models closed the gap with proprietary systems is covered in [Qwen3.6 and DeepSeek V4: China's Open-Weight Models Now Match Frontier Competitors](#).

Sources:

- [Poolside: Introducing Laguna XS.2 and Laguna M.1](#)
- [Laguna XS.2 Model Weights on Hugging Face](#)
- [Poolside: Laguna XS.2 and M.1, A Deeper Dive](#)
- [Reuters: Nvidia to Invest Up to \\$1 Billion in AI Startup Poolside](#)
- [Laguna XS.2 and M.1 on OpenRouter](#)

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