

Poolside Laguna M.1 and XS.2: A \$626M Startup's First Public Models Target Agentic Coding

Kabui, Charles

2026-05-04

[Read at ToKnow.ai](#)

**Poolside Ships Laguna:
225B MoE for
Agentic Coding**

First public models from the \$626M+ AI lab, XS.2 open-weight under Apache 2.0

72.5% M.1 on SWE-bench Verified	3B Active params in XS.2 (of 33B total)	30T Tokens used in pre-training
--	--	--

May 4, 2026

ToKnow.ai

Poolside, the AI lab [co-founded](#) by former GitHub CTO Jason Warner, released its first public models on April 28: [Laguna M.1](#) and [Laguna XS.2](#). M.1 is a 225-billion-parameter Mixture-of-Experts model with 23 billion active per token, trained from scratch on 30 trillion tokens

across 6,144 NVIDIA H200 GPUs. It scores 72.5% on SWE-bench Verified and 46.9% on SWE-bench Pro. XS.2 is far smaller: 33 billion total parameters, only 3 billion active, with 256 routed experts across 40 layers. It still hits 68.2% on SWE-bench Verified and 44.5% on SWE-bench Pro, and runs on a single GPU or a Mac with 36GB of RAM. XS.2 ships under Apache 2.0 on [Hugging Face](#). Both models are free to use for a limited time via [Poolside's API](#) and [OpenRouter](#).

Poolside trained both using the [Muon optimizer](#) instead of AdamW, reaching the same loss in 15% fewer steps. Their RL loop is fully asynchronous: actors run coding tasks in sandboxed containers while the trainer pulls trajectories continuously, with weight transfers under 5 seconds via RDMA. They also released [pool](#), a terminal coding agent implementing the Agent Client Protocol, and Shimmer, a cloud sandbox for building web apps. The company has raised over \$626 million.

XS.2 scoring 44.5% on SWE-bench Pro with only 3 billion active parameters shows how efficiently MoE routing compresses coding ability into a small inference budget. [GLM-5](#), another recent MoE release, made a similar case for sparse architectures. If coding skill keeps scaling with data and RL rather than raw parameter count, the cost of running capable software agents drops fast.

Sources:

- [Introducing Laguna XS.2 and Laguna M.1 \(Poolside Blog\)](#)
- [Laguna XS.2 and M.1: A Deeper Dive \(Poolside Blog\)](#)
- [Laguna XS.2 Model Card \(Hugging Face\)](#)
- [pool: Poolside's Coding Agent \(GitHub\)](#)
- [Poolside Founders Interview \(Upstarts Media\)](#)

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