

A Burbank startup just matched a datacenter with a MacBook in a benchmark test that could change the history of AI.

Punky Tiger Labs presents the results of more than 2 years of research in compact model optimization for local inference.



Palo Alto, California Apr 16, 2026 ([IssueWire.com](https://www.IssueWire.com)) - Punky Tiger Labs today publishes verified benchmark results for NYMPH AI, a 4.5 billion parameter model backed by 36+ AI patent applications filed with the USPTO. 11 benchmarks. 19,000+ items evaluated. Zero errors. Controlled laboratory conditions. April 14-15, 2026.

NYMPH AI matches GPT-4o — a model of over 200 billion parameters running on datacenter infrastructure — in code generation (HumanEval 90.24% vs 90.20%) and mathematical reasoning (GSM8K 92.65% vs ~93%). It surpasses GPT-4o in instruction following (IFEval 87.29% vs ~84%). A model 50 times smaller, running on a single Mac, producing the same results as a system that requires entire rooms of servers.

In its category there is no competition. NYMPH AI leads 7 of 11 benchmarks among all models under 20 billion parameters, outperforming Microsoft's Phi-3.5 (3.8B), Meta's Llama 3.1 (8B), Alibaba's Qwen 2.5 (7B), Google's Gemma 2 (9B), Zhipu's GLM-4 (9B), Mistral (7B), Anthropic's Claude 3 Haiku, and OpenAI's GPT-4o-mini. The closest model in its own size class — Microsoft's Phi-3.5 at 3.8 billion parameters — scores 62.8% in code generation where NYMPH AI scores 90.24%. A 27-point gap.

This is the first time a company has matched frontier AI models from a local device with a model of this size. The results open the door to more than 2 billion PCs and laptops worldwide that could run competitive artificial intelligence locally — devices that today depend entirely on the cloud for any AI capability.

Full benchmark results are presented in the attached dashboard.

The product

NYMPH AI is a private AI that does not share data to the cloud and works even without internet.

It runs entirely on the user's device. It uses the internet when available to stay current, search the web, and access real-time information — but it does not share your conversations with cloud servers. Your data stays on your device. Board a plane, lose your connection, or work in a classified facility — your AI keeps working.

It supports 119 languages, a 128K token context window, and 10 integrated research tools including Wolfram Alpha, PubMed, arXiv, CourtListener, and World Bank API. It generates PDF, DOCX, PPTX, and XLSX documents on-device with zero server persistence.

The engine is designed for domain customization. With customization work, NYMPH AI can be prepared for specific tasks — legal, finance, medicine, software engineering, academic research, and even gaming — turning a general assistant into a domain expert that runs entirely on the user's hardware. 154 out of 154 red-team security tests passed.

The company

Punky Tiger Labs, Inc. is a Delaware C-Corporation headquartered in Burbank, California. The company develops private AI infrastructure across software and hardware, with a single thesis: more than 95% of AI will be local.

According to the company's research, by 2030 there will be 2.4 billion AI users and more than 100 billion autonomous agents generating continuous inference demand. The energy required to serve that demand from centralized cloud infrastructure would require 712 times the world's current electricity production. The cloud will hit exponential demand against a physical limit. PTL's answer is the device you already own. 27.5 billion compute-capable devices sit in homes and offices worldwide, idle 85% of the time. Activating 10% of them with local AI would exceed any projected cloud demand. The datacenter the world needs is the one it already has installed.

App Store launch — May 2026

NYMPH Pulse arrives on the Mac App Store in May 2026. \$60/year or \$99 lifetime. The annual license includes all model sizes and all updates for current and future versions in the same category. Compatible with MacBook Air M1+, MacBook Pro, iMac, Mac Mini, Mac Studio, and Mac Pro. A professional tier (NYMPH Pro) is in development. Windows, iOS, and Android versions planned.

The market is demanding high-quality local AI and this is the first company in the world to deliver it. There are over 100 million active Mac users worldwide. If 1% adopt NYMPH Pulse at the lifetime price, that represents \$99 million in revenue from a single platform. At the annual price, 1% penetration generates \$60 million per year in recurring revenue. The broader addressable market — over 2 billion

PCs and laptops worldwide — is the next phase.

The professional era of local AI begins now.

"Every task, every device, every interaction will be connected to AI. That will create exponential demand that can only be met from local devices, with functional intelligences trained for specific tasks. That is our objective and our thesis."

— Alvaro Lopez Zuniga, Founder & CEO, Punky Tiger Labs, Inc.

Punky Tiger Labs, Inc. Delaware C-Corp. Burbank, California. 86 USPTO patent applications including 36+ in AI and AI hardware. The company is preparing the introduction of artificial intelligence hardware that would further improve these results and these interactions. "Cloud AI has been a very good first chapter, but this story is just beginning" — Rodney Waterman, Director of Channel Development, Punky Tiger Labs.

More information: rodney@punkytigerlabs.com | nymphtech.com

[Nymph Ai Benchmaks](#)

Media Contact

Punky Tiger Labs

*****@punkytigerlabs.com

+19013511534

<http://www.punkytigerlabs.com>

Source : Punky Tiger Labs

[See on IssueWire](#)