Lean Community Certifies the Claude-Assisted Riemann Proof as Users Adopt Gemini 4 Neon

AI High Signal Digest

2026-11-22

Lean Community Certifies the Claude-Assisted Riemann Proof as Users Adopt Gemini 4 Neon

By AI High Signal Digest • November 22, 2026

Lean maintainers and independent checkers confirmed Levent Alpöge's machine-checked proof of the Riemann Hypothesis. Separately, Gemini 4 Neon picked up high-profile users, Agent Arena data showed GPT-6.2 Astra trading writing quality for coding, and an AI provider disclosed a security incident.

The Riemann proof is formally certified, but no human has read it yet

The previous brief covered Levent Alpöge's claimed Claude-assisted proof. It has now gone from "very likely correct" to formally certified. The Lean FRO and the Mathlib Initiative issued a joint statement certifying the proof as correct. Alpöge had released it on November 20 as a 17-million-line Lean library, and the communities verified it over the following day [1]. Scientific American reports that independent mathematicians finished a full verification and say "the results are real" [2].

The checks ran on more than one independent system: - **Lean and nanoda:** One independent run compiled all 175,187 files (17.37M lines) with unmodified Lean 4.33.1. It took about 2 hours and cost $12.53 [3]. It then checked 1,484,729 declarations in nanoda, a Rust kernel that shares no code with Lean, with zero errors and only the three standard axioms [3]. - **MIT run:** A separate report says MIT's Andrew Sutherland ran the proof through other independently written kernels, and all accepted it. That makes a Lean bug "essentially ruled out" [4].

The checks have limits. They settle the Riemann Hypothesis plus four GRH theorems for Dirichlet L-functions, not general GRH. Humans "still have to read

and understand what the proof is doing" [3]. No readable proof exists yet, and the Clay prize requires publication plus two years of review [4].

One dissent stands out. Polymarket reports that mathematician Jenny Lorraine Nielsen claims the proof is invalid despite passing Lean verification [5]. No details of her argument were available.

The reactions moved quickly into politics: - **Tim Gowers**, a Fields medalist, posted that he went on a pause-AI march the day after seeing his grandson [6]. - **Boaz Barak** said that within a couple of months, if not already, someone with no mathematical expertise could solve GRH with internal models and enough compute. He said this was extrapolation, not internal knowledge [7]. Talia Ringer countered that Alpöge has been interested in RH since his math PhD [8].

Formal proof search also produced a result outside math. A Tübingen–Google DeepMind paper reports that AlphaProof Nexus found a Lean proof settling a quantum-optics question first raised in 2017. High-dimensional GHZ states cannot be built using only linear optics and pair sources for any even photon count of 4 or more, when there are at least as many dimensions as photons [9].

Gemini 4 Neon wins converts, and Google adjusts access

Hands-on adoption is now backing up last period's benchmark results: - **Aravind Srinivas** says Neon on High reasoning feels intelligent and fast at once. He uses it on Perplexity Computer and has moved some of his Teammates to it [10]. - **Omar Sar** is "strongly considering" replacing Opus 5.5 with Neon as his daily driver. He also wants subscription access in other harnesses [11]. - **Jeremy Howard** notes that Google described Neon at launch as only "approaching" frontier models, and argues that evals are "a *very* imperfect measure of model utility" [12].

Google also changed how much users get and where they can use it: - **Quota:** Antigravity reset everyone's Gemini quota and said caching improvements raise what users can do within it by 22% [13]. - **Third-party harnesses:** Patrick Toulme said users will "NOT be banned ever" for using Antigravity in T3 Code [14].

Effort settings cost very different amounts. Simon Willison's Gemini 4 Hydrogen test cost 3.99¢ and used 5,311 output tokens on "high," versus 23.61¢ and 31,477 tokens on "xhigh" [15].

GPT-6.2 Astra gains in coding and loses in writing

Agent Arena ranks GPT-6.2 Astra (Max) #6, one place above GPT-6 Astra. Compared with GPT-6 Astra: - Coding rose 50 points, to 1,338. - Chat fell 54 points, to 1,211, dropping from #8 to #11 [16]. - Writing and translation tasks

users confirmed as successful fell from 91.2% to 85.9%, with the sharpest drop in creative copy [17].

These numbers fit the mixed reception of the release.

Inkling 2 Mini spreads to local hardware

Thinking Machines' looped open model is the #1 trending model on Hugging Face, and Unsloth's GGUF version is #2 [18]. Project Maya v1.1.73 runs it on NVIDIA cards with as little as 8 GB, keeping experts that don't fit in system RAM. The release also supports V100, RTX 20-series and AMD GPUs on Linux [19]. One user reports 144 tok/s on an RTX 4090 and calls it his first good model for a single-GPU box [20].

Inference security

- **Together AI:** It disclosed a security event that leaked a subset of prompts sent to its serverless APIs in June and July 2026 [21].
- **MiniMax routing on OpenRouter:** This follows last period's report. A researcher planted secret links in outputs routed to MiniMax, and someone opened them 40 to 130 hours later. AtlasCloud, Novita and StreamLake were named as forwarding requests to MiniMax [22]. Sebastian Raschka offered a possible benign explanation: cheaper providers at capacity may reroute traffic to the model's official provider rather than fail [23].
- **MiniMax cryptanalysis:** MiniMax separately published work describing an internal model it says is useful for cryptanalysis [24].

Compute bulls vs. market turmoil

The selloff continued. Taiwan banned short selling after the TAIEX fell 17.5% in its worst week in 35 years, leaving it 35% below its October high [25]. Commentators split on what it means for AI compute: - **François Chollet** argues this is a buying opportunity. GPU spot prices keep rising, and HBM and DRAM are "effectively sold out through 2029." He also says inference margins are 70–85% [26]. - **Elon Musk** warns that even a few months of delay between paying for compute and earning revenue from it would mean "hundreds of billions" in extra working capital [27].

Debate continued over the board's handling of Jensen Huang. Paul Graham says the tech world sees it as a huge mistake [28]. Others note that Huang hand-picked the board and that what happened between them is unknown [29].

Also notable

- **Gemma 4 speculative decoding:** Google built a multi-token drafter into the model itself. It reuses the model's activations and existing KV

cache, so it skips the separate prefill step a standalone drafter needs. Reported speedups reach 3.1x for large models on servers and 2.2x on a smartphone GPU [30].

- **US House report on open source:** Clement Delangue says a new US House report recognizes that the right response to Chinese open-source AI is building more American open source, not banning it [31].
- **LangChain's "knowledge factory":** LangChain argues that enterprise agents need curated company knowledge more than better models. It says this approach already powers its internal GTM and content agents and LangSmith Engine [32].

Sources

1. X post by @leanprover
2. X post by @sciam
3. X post by @elliotarledge
4. X post by @atulit_gaur
5. X post by @Polymarket
6. X post by @wtgowers
7. X post by @boazbaraktcs
8. X post by @TaliaRinger
9. X post by @dair_ai
10. X post by @AravSrinivas
11. X post by @omarsar0
12. X post by @jeremyphoward
13. X post by @_mohansolo
14. X post by @PatrickToulme
15. X post by @simonw
16. X post by @arena
17. X post by @arena
18. X post by @ClementDelangue
19. X post by @PeasantSmith
20. X post by @teortaxesTex
21. X post by @vipulved
22. X post by @michael_chomsky
23. X post by @rasbt
24. X post by @teortaxesTex
25. X post by @BullTheoryio
26. X post by @fchollet
27. X post by @elonmusk
28. X post by @paulg
29. X post by @gerstenzang
30. X post by @dair_ai
31. X post by @ClementDelangue

32. X article by @hwchase17