

Google Puts Gemini 4 Argon Back at the Frontier, With a Gated Launch to Cyber Defenders

AI High Signal Digest

2026-10-01

Google Puts Gemini 4 Argon Back at the Frontier, With a Gated Launch to Cyber Defenders

By AI High Signal Digest • October 1, 2026

Gemini 4 Argon matches GPT-6 Astra on independent indexes at launch-discount prices, though access is limited for now. OpenAI attributes a reasoning-distillation campaign to people linked to Moonshot, and agentic and open-model releases continued.

Gemini 4 Argon: Google rejoins the top tier, but on a short leash

Google DeepMind's Gemini 4 Argon is its first proprietary model above the Flash class in more than seven months. Independent benchmarks put it level with the best models available. On the Artificial Analysis Intelligence Index, Argon at high reasoning scores 53. That matches GPT-6 Astra (max), is one point ahead of GPT-6.1 Sol (max), and is 23 points above Gemini 3.1 Pro Preview [1]. Other results: - **Google's own comparisons:** first on 13 of 19 published benchmarks against Astra and Opus 5.5. That includes DeepSWE (77.9% vs. 74.2% for Opus 5.5) and the Vals Index (68.9% vs. 67.0%) [2]. - **Arena:** #1 in Text Arena with 1,525 points and #8 in WebDev [3]. - **Agentic and accuracy tests (Artificial Analysis):** #1 on AutomationBench-AA at 77.5%, but behind Sonnet 5.5, Opus 5.5 and Astra on Terminal Bench 4 at 57%. Its hallucination rate is 15%, compared with 51% for Astra. Its raw accuracy is 50%, 13 points below Astra [1].

Price and access. Standard pricing is \$4/\$20 per million input/output tokens. A launch promotion halves that to \$2/\$10, and cached input gets a 95% discount. Argon has a 1M-token context window, and a new Long Decode Continuation

feature lets reasoning run up to 1M output tokens [1]. At the promotional price, Artificial Analysis puts it at \$1.99 per Intelligence Index task, 60% of Astra’s cost but $2.7\times$ GPT-6.1 Sol’s. The cost would rise to \$3.98 once the discount ends, and Google has not said when that will be. The savings come from cheaper tokens, not from using fewer of them: Argon averages 62k output tokens per task against Astra’s 27k [1]. The model is not publicly available yet [1]. Sundar Pichai says it is with the US government and with trusted cyber defenders in the Fairwind Program, and that wider availability will come “as soon as we can and as safely as we can” [4].

Internal use. Google says Argon agents freed more than 300 TiB of memory across its data centers. They are also migrating more than 800,000 lines of C/C++ kernel code to Rust [5]. Vahab Mirrokni credits the model with helping complete the “CK conjecture” and says further results are queued for release [6].

Skepticism. One commenter flagged Argon’s 19.6% on Harvey’s Legal Agent Benchmark, below Muse Spark 1.2’s listed 25.42% [7]. @teortaxesTex expects Argon to disappoint in practice. He still thinks it matters strategically because it is “a completely new pretrain” that Google can iterate on [8].

The release fits a broader trend. Artificial Analysis says the gaps between successive record scores on its index have shrunk from 239 to 126 to 99 days [9].

OpenAI attributes a reasoning-extraction campaign to people linked to Moonshot

OpenAI published a report on a coordinated campaign to extract its models’ hidden reasoning. It attributes “a core cluster of the activity to individuals associated with Moonshot AI, the developer of Kimi.” It adds that it is unclear whether all the operators were one actor [10]. According to a summary of the report, OpenAI logged 16,000 extraction attempts from more than 4,000 users in two days, and found related activity across more than 15,000 users. Operators moved encrypted reasoning between conversations and prompted the model to reveal it [11].

Separately, academic researchers report that two months after they disclosed their own extraction attack, they could still pull reasoning from Astra and Sol 6.1 through third-party API providers. They have told OpenAI and Anthropic [12]. One author said patches are hard to roll out across every product version and third-party vendor [13].

GPT-6.1 Sol: demand and cost data

Codex lead Thibault Sottiaux says GPT-6.1 Sol is OpenAI’s “most demanded model pretty much ever” across both the API and subscriptions. After heavy load, OpenAI has added capacity and expects serving speed to nearly double

[14]. ARC Prize verified scores of 94.2% on ARC-AGI-2 and 96.4% on ARC-AGI-3 with the provider’s adapter harness. On ARC-AGI-3 with the standard harness, it scored 52.7% [15]. Artificial Analysis says it costs about 30% less per task than GPT-6 Sol. The savings come from using fewer turns and a cheaper cache-read price, not from a price cut [16].

Biosecurity and policy

- **SynthID Bio.** Google DeepMind synthesized AI-designed proteins that are both functional and watermarked, published the work in *Nature*, and is open-sourcing the tools [17, 18].
- **FTC (unconfirmed).** An X account claims the FTC has confirmed an investigation into OpenAI, Anthropic and METR, and is drafting civil investigative demands about the dangers those labs have warned of [19]. No primary source was reviewed.
- **Brockman.** The NYT reports that Greg Brockman is no longer making a second \$25M donation he had promised to the super PAC Leading the Future [20].
- **SI czar (rumor).** Rumors say DNI Jay Clayton will be named White House superintelligence czar. He told CNBC that superintelligence “is a national security issue” [21].

Models, agents and infrastructure

- **Runway Praxis-1.** An open-weight “World Action Model” that uses video pretraining to produce robot policies for any embodiment. Runway is testing it with Noble Machines, Standard Bots and Ultra, and plans to release the weights in the coming months [22].
- **Ling-3.1-flash.** About 560B total parameters, about 25B active, with up to 1M-token context. Open-sourcing is planned [23].
- **Superhuman Stratego.** A *Nature* paper reports the first superhuman Stratego AI, built with general reinforcement-learning and test-time-compute methods for games with hidden information [24].
- **DeepSeek.** Released an open-source kernel toolkit for Huawei Ascend, with TileLang support optimized for the Ascend 950. It is framed as a CUDA alternative [25]. It also shipped a Harness v0.2 desktop preview for macOS and Windows. DeepSeek says Harness is the most-used coding agent among users of its official API [26].
- **Embeddings.** Perplexity’s open `pplx-embed-v2-context-9b-preview` claims state-of-the-art results on ConTEB and turbopuffer’s context-bench [27]. Cohere’s Embed 5 Pro leads the models Cohere itself measured on ViDoRe V3 [28].
- **Hardware.** Cognition is the first customer running NVIDIA Vera Rubin, on CoreWeave. It reports about 4.8× the token throughput of GB200 at the same decode speed [29]. Synopsys signed a multi-year deal worth more than \$1B to supply silicon IP for Amazon’s custom chips [30].

- **Video.** On Artificial Analysis’ new text-to-video benchmark, open-weights MiniMax H3 ranks #3, statistically tied with Seedance 2.5 at about one-seventh the price [31].
- **Factory vs. Cognition.** Factory removed Chris Degnan as board observer and advisor. It says he had been holding formal, recurring talks with Cognition executives while advising Factory on confidential matters. It also accuses Cognition engineers of faking interviews to extract product information [32].
- **Gemini app.** Skills will replace Gems, starting in November for personal accounts [33].

Sources

1. X post by @ArtificialAnlys
2. X post by @TheRunDownAI
3. X post by @arena
4. X post by @sundarpichai
5. X post by @kimmonismus
6. X post by @mirrokni
7. X post by @BlackHC
8. X post by @teortaxesTex
9. X post by @ArtificialAnlys
10. X post by @AndrewCurran__
11. X post by @kimmonismus
12. X post by @kotekjedi_ml
13. X post by @jonasgeiping
14. X post by @thsottiaux
15. X post by @arcprize
16. X post by @ArtificialAnlys
17. X post by @pushmeet
18. X post by @demishassabis
19. X post by @ns123abc
20. X post by @teddyschleifer
21. X post by @AndrewCurran__
22. X post by @runwayml
23. X post by @AntLingAGI
24. X post by @ssokota
25. X post by @kimmonismus
26. X article by @DeepSeekHarness
27. X post by @perplexity_ai
28. X post by @cohere
29. X post by @cognition
30. X post by @Synopsys
31. X post by @ArtificialAnlys

32. X post by @matanSF
33. X post by @GeminiApp