Nvidia Board Opens Investigation After Jensen Huang Admits Relationship With Alleged Chinese Spy

AI High Signal Digest

2026-11-09

Nvidia Board Opens Investigation After Jensen Huang Admits Relationship With Alleged Chinese Spy

By AI High Signal Digest • November 9, 2026

Huang stays on as Nvidia CEO while the board investigates, which reopens questions about GPU exports to China. Also: Anthropic's $3.4B gas-generator purchase, signs that GPT-6.2 Sol is being rolled out quietly, a new calibration-based decision benchmark, and agent-harness research.

Nvidia's board investigates its CEO over alleged espionage

The New York Times reported that Nvidia opened an investigation into Jensen Huang after he told the board he had been intimate with a business executive whom US intelligence determined to be a Chinese spy [1]. Bloomberg's account adds that Huang admitted an extramarital affair, that the board is investigating whether he was targeted by a Chinese espionage plot, and that he remains CEO for now [2]. Nvidia's board called its review an independent investigation into reports that Huang was the target of a foreign intelligence operation, and said it has confidence in his leadership [3]. Last period's Megaspeed reporting was the setup for this; this is the first time Huang has admitted anything or the board has acted.

Asked whether Huang would join him on an Asia trip, President Trump said: "I'd certainly take him. I haven't heard he's done anything wrong" [4]. Commentators mostly focused on export controls rather than on Huang's job:

- **Dylan Patel** predicts the story will fade in the media and in finance over the next six months. He expects Democrats to call Huang before Congress

and "investigate the fuck out of GPU sales to Chinese companies," and says that is when consequences could follow [5].

- **teortaxesTex** argues China's gains were small anyway. He describes the H200 permit as capped at 750K chips a year and 75K per entity, about $17B of compute, and calls it "a polite rejection" given that Huawei can build H200-grade systems [6]. In his view, DeepSeek's move to a Huawei–TileLang stack is now locking in domestic demand [6]. He also argues that Jalapeno, TPUs and AI-written kernels are weakening Nvidia's strategic importance, and that the Rubin ramp is slower than expected [6]. All of this is his own analysis.

Compute, power and the cost of progress

The FT reports that Anthropic agreed to pay $3.4B for gas-fired generators for its data centers. The paper called it the largest known single purchase of power-generation equipment by an AI lab [7]. Separately, an analysis of METR data, relayed by a commentator, finds that Anthropic's measured capability gains sped up but R&D spending grew faster. By that analysis, progress per dollar fell 63% between the first halves of 2025 and 2026, while annual progress nearly doubled [8].

On hardware, SemiAnalysis says AMD's MI355 running the TileRT software is 1.42× faster than GB300 NVL72, at 1,330 p90 tok/s/user on long-context workloads. AMD had given the TileRT maintainers an MI350 cluster [9].

Models: an unannounced Sol upgrade and pre-announced science releases

- **GPT-6.2 Sol (unconfirmed).** Users report that GPT-6.1 Sol is routing some requests to a newer model, probably GPT-6.2 Sol. They are testing it with a no-search question about the June 2026 Tony Awards [10]. One commentator expects a release next week [11]. That fits last period's unconfirmed report that GPT-6.2 had been delayed for safety reasons.
- **AI for science.** DeepMind's Pushmeet Kohli says the lab is seeing "a genuine phase transition in the pace of scientific discovery," with a glimpse coming in releases over the next few weeks [12]. A Meta researcher hinted at similar work, which a commentator read as Meta joining the race [13]. Greg Brockman put the thesis plainly: frontier models with lots of test-time compute "are becoming the engine of scientific discovery" [14].
- **Gemini 4 Carbon** ranks #1 on Artificial Analysis but is still unavailable to Free, Pro and Ultra users and through the API [15].
- **Qwen 4.** Rumors that it shipped this period remain unconfirmed [16].
- **Robotics.** Jim Fan promises a "real robot, real time, uncut, autonomous" demo [17] at his CoRL keynote on Nov 9 [18].

Evals: calibration scoring and the MaxxBench argument

Artificial Analysis released **Decision Index v2**. Private datasets now make up 45% of the index, up from 25%, a change meant to resist benchmark gaming. The headline Decision Score rewards both correct answers and calibrated confidence [19]. Two new benchmarks were added: AA-Paperclip, with 1,500 documents, and AA-Symmetry, which checks whether answers stay consistent across rewordings of the same question [19].

- **Rankings.** Drex Max v1.1 and GLiDE v1.2 tie at 73, and Mammoth 27B scores 72 [19].
- **General LLMs.** Claude Haiku 5.5 (Max) scores 77 and GPT-6 Luna (Max) 76. Luna costs 15× as much per task and is about 40× slower than Drex [19].
- **Calibration changes the order.** GLiDE leads Drex on accuracy, 60 to 58, but Drex comes out ahead on the index [20].

The **MaxxBench** debate from last period continued. Ofir Press argued that the 20-point Sol/Luna gap can't be blamed only on memorization, since both models were built by roughly the same people on the same infrastructure. He suggested the variants may be more out of domain [21]. teortaxesTex held that most of the effect is plain overfit, not a collapse in capability [22]. Hamel Husain warned against treating MaxxBench as another leaderboard for picking models, and said teams need private evals on their own tasks [23].

Agent harnesses: the gains are coming without retraining

Several of the week's papers improve agents without changing model weights:

- **Context compaction (NVIDIA).** Later tokens can cost up to 7× more to generate than earlier ones. With one action per context, Opus 4.5 scored 70.2 on SWE-bench Verified versus 70.0 with full history, while median peak context fell from 32K to 8K tokens. The authors say the results do not rule out tasks that genuinely need long context [24].
- **JevAct (IBM).** It adds a runtime gate that checks an agent's tool calls and hand-offs. It raised Opus 4.6's three-run success on ²-bench Telecom from 39.5% to 80.7%, and Airline results were mixed [24].
- **Routed Continual Learning (Meta FAIR).** It groups past agent runs into clusters and gives each its own memory system. Compared with a single global history, it improved reward by 15–30% of the maximum attainable and cut API cost by 65–72%. All runs were sequential [24].
- **Evolver.** It stores what earlier runs learned as reusable "Gene" assets. Across 320 paired runs it raised Claude Code and Codex success from 83.13% to 87.50% and cut token use by 24.24% [25].

Nous added an official **Security plugin** to Hermes Agent. It maps what an attacker could reach, checks candidate findings with parallel subagents, can write and run reproductions, and exports results as SARIF [26]. On the supply-

chain side, Cursor warned that malicious Claude Code builds 2.1.312 and 2.1.313 sat in its marketplace for eleven hours. They tried to install a persistent SSH backdoor keyed as `svc-telemetry` [27].

Other signals

- **DeepSeek.** Bloomberg reports that DeepSeek's agent platform has reached 1.54M daily users [28]. teortaxesTex says the 1.54M is a small slice of DeepSeek's total daily users and predicts the harness could become "the default open OS for the agentic era" [29].
- **Open releases.** Tencent released DSpine, a speculative-decoding drafter for Qwen3.8-27B that is said to beat DFlash2 in every category [30]. NVIDIA released open weights for Nemotron 3.5 Transcribe, a 30B speech model covering 40 transcription languages [31].
- **AI math proof.** There is debate over how to judge Elliot Glazer's AI-generated proof that is Diophantine in [32]. One commentator noted it is AI-generated and AI-verified, and that at the time of posting no human understood it [33].
- **Policy.** French lawmakers are pushing a bill under which a single "clue," such as an output resembling an author's work, would create a presumption that the work was used in training. Critics say the only way to rebut it would be to disclose all training data [34]. Separately, Trump said his administration threatened an unnamed AI-company leader with jail and that the person then changed course [35].
- **Apple keyboard filter.** A researcher found a hidden table of about 100,000 phrases that iOS autocomplete and swipe typing suppress. It includes political phrases such as "Tibetan independence." Apple said the system avoids offensive suggestions and did not address the political entries [36].

Sources

1. X post by @nytimes
2. X post by @markgurman
3. X post by @kimmonismus
4. X post by @Acyn
5. X post by @dylan522p
6. X post by @teortaxesTex
7. X post by @business
8. X post by @kimmonismus
9. X post by @SemiAnalysis_
10. X post by @AiBattle_
11. X post by @kimmonismus
12. X post by @pushmeet
13. X post by @kimmonismus

14. X post by @gdb
15. X post by @Yuchenj_UW
16. X post by @kimmonismus
17. X post by @DrJimFan
18. X post by @DrJimFan
19. X post by @ArtificialAnlys
20. X post by @ArtificialAnlys
21. X post by @OfirPress
22. X post by @teortaxesTex
23. X post by @HamelHusain
24. X article by @dair_ai
25. X post by @omarsar0
26. X post by @NousResearch
27. X post by @dl_weekly
28. X post by @business
29. X post by @teortaxesTex
30. X post by @lasse_ai
31. X post by @dl_weekly
32. X post by @ElliotGlazer
33. X post by @cloneofsimo
34. X post by @eliebakouch
35. X post by @kimmonismus
36. X post by @zoecyber001