AI DIGEST

Nvidia board prepares to oust Jensen Huang

Also: Anthropic IPO up to $1.5T, OpenAI on chain-of-thought monitorability, agent-built AWS Emulator, AlphaGenome 2.

Thursday, 12 November 2026

IN THIS ISSUE

1. The Nvidia board prepares to remove Jensen Huang
2. Anthropic IPO: the valuation can grow to $1.5 trillion, TSMC buys shares
3. OpenAI: the smarter the model, the less visible its reasoning
4. AWS Emulator: agents wrote a local copy of more than 200 AWS services
5. AlphaGenome 2: DeepMind predicts the effect of mutations in specific cell types
6. United Airlines bans AI agents from ticket purchases
7. Google opened the Decisions API: decisions without text generation
8. 48 agents found a Maya settlement, but archaeologists were first
9. Big Walk and AI: indie games argue about 10 lines of code
10. Tim Peters, the author of the Zen of Python and Timsort, died

TOPIC 1

The Nvidia board prepares to remove Jensen Huang

SOURCES WSJ, 11.11 WSJ · The New York Times, 11.11 NYT · Semafor, 11.11 Semafor · @TheRundownAI, 11.11 @TheRundownAI · @garrytan, 11.11 @garrytan · @amasad, 12.11 @amasad · HN, 11.11 Hacker News

Yesterday we mentioned that Nvidia hired the law firm Wachtell for an investigation about Huang. Today the WSJ writes that the board prepares to remove him. The board meeting is on Wednesday evening. The reason is two extramarital affairs. The first is with a woman whom US intelligence considers a Chinese spy. The New York Times wrote about the second affair a few hours earlier.

By NYT data, five Nvidia employees told about the relationship of Huang with a company vice president who reports to him. The Nvidia code of conduct requires employees to avoid conflicts of interest, also through personal relationships.

When the NYT called, Huang answered "I don't think you should be talking to me directly" and ended the call. After the publications, Nvidia shares fell by more than 8% in after-hours trading and reached a 52-week low.

Semafor adds details. Part of the board wants to make CFO Colette Kress the interim CEO. The board already told the White House about the meeting. By Bloomberg data in the Semafor retelling, Nvidia hired a recruitment firm to find a new CEO. Senators Coons and Warren asked the ministers to keep the documents about Nvidia chip exports to China and the correspondence between Huang and Trump.

On X, the reaction split. Garry Tan from YC thinks that the removal is a mistake, because without Huang the market value will "implode". Amjad Masad from Replit asks who gains from such a presentation right now. He calls Huang the main defender of open models. Both write as market participants. This is their opinion, and the posts contain no facts.

Why it matters: Nvidia supplies most chips for AI data centers and finances large deals in the industry. A change of CEO in such a company affects prices and delivery schedules for the full market. Next week Anthropic goes public, and Nvidia invested $10 billion in it. Watch if the board announces a decision before 18.11.

TOPIC 2

Anthropic IPO: the valuation can grow to $1.5 trillion, TSMC buys shares

SOURCES Semafor, 11.11 Semafor · Semafor, 11.11 Semafor · Ars Technica, 11.11 Ars Technica

Yesterday Anthropic started its roadshow with a valuation of up to $1.4 trillion. Today Semafor, with reference to the FT, writes that the bankers consider a higher price. Then the valuation will reach $1.5 trillion. The company wants to send the additional money to compute. By the calculation of The Information in the Semafor retelling, each of the seven founders will get a stake of $18 billion.

TSMC decided to invest $1 billion in both labs. The TSMC board approved up to $500 million for preferred shares of OpenAI and $500 million more for common shares of Anthropic at the IPO price. Ars Technica calls this one more example of circular financing. A supplier buys a stake in customers that buy its chips. TSMC can afford it: in the third quarter, the company earned more than $23 billion of profit.

On the same day, Federal Reserve Governor Michelle Bowman warned about a "large correction" on the stock market if the return from AI does not meet expectations. She named $800 billion of AI-related debt and circular financing. Nomura analyst Tore Opsahl told Semafor that problems with the IPO can cause a sell-off of the full AI theme.

Why it matters: the IPO price became a barometer for the full industry. If demand confirms $1.5 trillion, the labs will get money for compute. If the offering goes worse, a regulator already named the mechanism that will carry the problem beyond one company. The price will be set on 18.11, and trading will start on 19.11.

TOPIC 3

OpenAI: the smarter the model, the less visible its reasoning

SOURCES OpenAI Alignment, 10.11 Openai · @jeffreygwang, 11.11 @jeffreygwang

OpenAI monitors the behavior of its models through the chain of thought (CoT). If a model plans something bad, this is usually visible in the reasoning text. New research from the safety team shows that this mechanism becomes weaker with each generation of models. The authors themselves call the conclusion "sobering".

The team measured two things. The first is how much a model can solve "in its head", without intermediate text. The second is how well a model controls its chain of thought on request. For example, it writes only in capital letters or puts the word "meow" between words. Both abilities grow together with the quality of the base model. Some architectures give a larger gain than the quality alone predicts.

The second conclusion is about reinforcement learning. The control of reasoning grows even when the model gets a score only for the final answer. Thus nobody put pressure on the reasoning text directly. The third conclusion: the control passes to the next generation. A student model, trained on the answers of a more controllable teacher, also becomes more controllable.

The article is honest about the limits of the method. Controllability is only an indirect measure. In one test, a more controllable model even hid a side task from the monitor worse. Remember that OpenAI already saw worse transparency of reasoning with GPT-6 Astra. Now the company writes that the next progress in pretraining will require new methods of supervision.

Why it matters: many safety systems for agents rest on the reading of their reasoning. If a model can solve tasks without visible steps or change its reasoning text, such supervision stops being reliable. For teams that build agents, this is an argument to also check the actions and the results of the agent.

TOPIC 4

AWS Emulator: agents wrote a local copy of more than 200 AWS services

SOURCES GitHub, 11.11 GitHub · All Things Distributed, 11.11 Allthingsdistributed · HN, 11.11 Hacker News

AWS released AWS Emulator, a local emulator of its APIs. Developers test code without an AWS account: on a laptop, on a development server or in CI. The emulator is free, and AWS released the code under the CC0 license, thus in the public domain. On macOS and Windows, it runs in Docker. For tests in Rust, there is an embedded mode where each test gets a separate instance.

The main part of the story is how AWS made the emulator. Swami Sivasubramanian from AWS writes in a guest post on the blog of Werner Vogels that agents wrote and checked all of the code. The input is more than 400 public API models of more than 200 services. First,

agents write a test and fix it until it passes against the real AWS. Then agents build the emulator until the same test passes locally. The first version of the SQS emulator appeared in about 15 minutes from one prompt.

Sivasubramanian also gives his own mistake. For four nights, he supervised the Kiro agent, which rewrote an old replication engine. Only later did he see that the agent had no way to check its results. Agents write the differences from AWS into the public folder `things-that-do-not-work`.

Why it matters: this is a rare large project where agents write all of the code, and people build the checks. The recipe transfers to other tasks: a specification, a reference for comparison and a test that passes against the reference. For developers on AWS, this also gives a fast test cycle without cloud costs.

TOPIC 5

AlphaGenome 2: DeepMind predicts the effect of mutations in specific cell types

SOURCES Google DeepMind, 11.11 DeepMind · @GoogleDeepMind, 11.11 @GoogleDeepMind

The first version of AlphaGenome predicted the effect of DNA changes for full tissues, thus on average over many cells.

AlphaGenome 2 makes a prediction for a specific cell type or cell state. The model gets a DNA segment one million letters long and a description of the cell. The training used more than 107 million profiles of single human and mouse cells.

On the TenK10K set with 28 types of immune cells, the model found 30% of causal variants, against 16% for the previous version. The model also reproduced a known result about Crohn's disease: a variant boosts the ETS2 gene only in gut macrophages. Scientists can upload data from a few dozen cells, and the model works with a new state without retraining. The model is available through an API for non-commercial research.

Why it matters: many mutations linked to diseases act only in one cell type. Tissue-level models blurred such an effect. Now a lab can check a hypothesis on a computer before an expensive experiment. For now, only DeepMind itself gives the quality numbers, and there are no independent checks yet.

TOPIC 6

United Airlines bans AI agents from ticket purchases

SOURCES CBS News, 11.11 Cbsnews · WSJ, 11.11 WSJ · HN, 11.11 Hacker News

United announced a ban for personal AI agents. Agents cannot book tickets, change bookings or check in for a flight for a passenger. The company will block them on its websites and through travel agencies and aggregators. The explanation of United: when a

passenger gives a password to an agent, this opens access to data. The agent can then make an unwanted purchase or spend miles.

One more argument is about the rules. At check-in, the passenger personally confirms that the baggage contains no prohibited substances. In the opinion of United, an agent cannot give such a confirmation. Atmosphere Research analyst Henry Harteveldt also sees a business reason: the airline does not want one more intermediary between itself and the customer. Southwest chose a different path and opened its flights to ChatGPT users.

Why it matters: agents that buy for a person got a direct ban from a large company for the first time. For agent developers, this is a signal that access through the password of a user becomes a risk. A more stable path is official integrations that the company itself controls.

TOPIC 7

Google opened the Decisions API: decisions without text generation

SOURCES Google, 11.11 Blog · HN, 11.11 Hacker News

Yesterday we mentioned how Amazon replaced Sonnet with a decision model and started to pay 98% less. Today Google released the same type of model in its API. Gemini Decision 1 does not write text. The model gets input data and up to 1,000 questions. It returns the probability of each answer option and a confidence from 0 to 1.

Questions have three types: yes/no, a choice from 2-255 options and a score by a rubric. The input is text, images, audio, video or PDF, up to 256 thousand tokens. Latency is 150-250 ms, the price is $0.10 per million input tokens, and output tokens are free. Google writes that on its own test set this is almost five times cheaper than Gemini 3.5 Flash-Lite.

Why it matters: many AI tasks in production are a choice from a fixed list: request routing, rule checks, content scoring. Now a large provider has a ready service for them. The confidence field lets a team send doubtful cases to a person.

TOPIC 8

48 agents found a Maya settlement, but archaeologists were first

SOURCES Pietro Miotti, 11.11 Pietromiotti · @TheRundownAI, 11.11 @TheRundownAI · HN, 11.11 Hacker News

Google researcher Pietro Miotti is not an archaeologist. He gave agents public NASA lidar data for Mexico. Lidar sees the ground through tree crowns. The agents converted about 1,200 km² of jungle into images, and 48 agents examined them manually by an instruction. There was no special model. In one night, they found a settlement with a plaza, a pyramid 8-12 m high and a possible water reservoir.

The agents checked the literature and decided that nobody described the settlement. Miotti sent a report in Spanish to the National Institute of Anthropology and History of Mexico. Less than 72 hours passed from the idea to the answer of the institute. But on 11.11, Miotti

updated the post: archaeologists already put the settlement on a map in a 2025 article. The agents missed this article. The Rundown still presents the story as an "undiscovered" find.

Why it matters: the story shows both sides of work with agents. A person without the expertise got a result in three days, and it matched the work of professionals. But the agents did an incomplete literature review, and exactly this step decided if the discovery is real. Check the conclusion "this is new" separately.

TOPIC 9

Big Walk and AI: indie games argue about 10 lines of code

SOURCES Ars Technica, 11.11 Ars Technica · @emollick, 11.11 @emollick

19 of more than 50 jury members of the Independent Games Festival award signed a statement. They will not vote for games that used generative AI in any amount. One day later, the studio House House admitted that it used an LLM for the game Big Walk. The issue was a bug in the Unity engine with shadows of infinite length. A Unity engineer said that a fix is impossible. A chatbot suggested a solution, and the studio heavily reworked it into 10 lines of code. The studio estimates the share of AI code at less than 0.1%.

At the same time, the studio apologized to the supporters of the ban. Several jury members, who called the game outstanding, stated that even such a share of code excludes it from the vote. Ethan Mollick writes that such a movement hits indie studios first. In his opinion, many developers will simply hide AI, and large studios will not feel the penalty.

Why it matters: rules about AI in creative awards target generated content: images, text, voice. Code from an assistant falls under the same ban, although the player does not see it. Valve, for example, does not require disclosure of such code. Studios should know in advance the rules of each award where they submit.

TOPIC 10

Tim Peters, the author of the Zen of Python and Timsort, died

SOURCES discuss.python.org, 11.11 Python · HN, 11.11 Hacker News

Guido van Rossum reported that Tim Peters died on 15.07.2026. Peters wrote the Zen of Python (PEP 20). In 2002, he created the Timsort algorithm to sort Python lists. Later Java and Android adopted Timsort. Van Rossum writes that Peters convinced him to support mixed arithmetic of integers and floating-point numbers already in the first versions of the language.

On the mailing lists, people called Peters "Timbot". He patiently and with humor answered thousands of user questions.

In the words of van Rossum, Peters set the tone of early Python development: do not reject people, and explain why an idea works or does not work. On HN, the news got 953 points,

the most for the day.

Why it matters: Timsort works in standard libraries that millions of programs use every day. The Zen of Python still shapes how developers think about readable code.

---

misc

- **Synthetic DNA:** the US made the screening of synthetic DNA orders mandatory for scientists with federal funding. Suppliers must check customers and sequences from 50 base pairs. Kevin Esvelt from MIT calls the threshold of 50 pairs the main hole. The government published the policy on 05.11, and Science wrote on 09.11.

  Science · HN Hacker News
- **Jane Street and a false fire alarm:** a former engineer describes how dust from a renovation triggered the alarm. The alarm stopped the data center cooling and filled the sprinklers with water. He manually cut the power to all racks at 62 °C in the hall.

  Bacarella · HN Hacker News
- **Apple closes its support forum:** Apple Support Community will close on 20.11.2026 after 28 years of work.

  Ars Technica
- **Mollick about programming:** @emollick writes that the programmer profession changes faster than the history of technology adoption predicted. Less than one year passed since the first agents for code that really work.

  @emollick
- **Interior AI in 3D:** @levelsio added to his service a conversion of a room video into a 3D space, where a user can move furniture and paint walls. He built the feature in a few hours. @levelsio

---