NEW HORIZON · AI DIGEST

The Long Read

1 November 2026

12	3	2555	4
STORIES	SECTIONS	WORDS	HIGH IMPACT



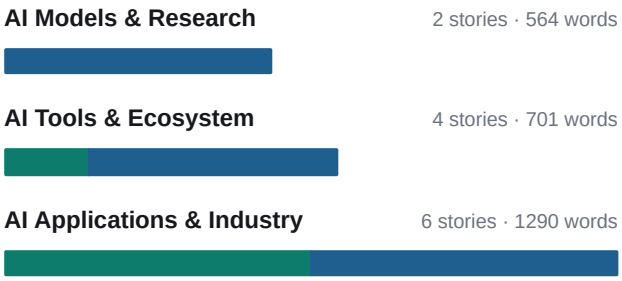
4 high impact 8 medium impact spoke length = depth of coverage

The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

12 stories · 2555 words · 3 sections





AI Models & Research

01 Meta Built a Humanoid Prototype in Three Months With a Chinese Partner

02 A New Test Rates Whether AI Can Name Its Sources

AI Tools & Ecosystem

03 A ChatGPT Share Link Could Run Code in Your Session

04 Altman Says Your PC's Idle RAM Should Feed His Data Centers

05 AI Inference Startup Baseten Doubles Its Valuation in Four Months

06 Publishers' AI Detectors Are Falsely Flagging Their Writers

AI Applications & Industry

07 OpenAI Cuts Its $30 Billion Round Target to $20 Billion as Investors Balk

08 S&P Flags Meta's AI Spending as a Credit Risk

09 Bond Markets Are Shutting the Door on AI Data-Center Lenders

10 Foxconn Lines Up a Bigger AI Deal With OpenAI

11 Apple's New CEO: No AI for AI's Sake

12 Court Clears NAACP's Pollution Suit Over xAI's Memphis Turbines

How to read this. Every story in today's email is reported here at full length, in the same order. Impact is the writer's judgement of whether a story changes what a reader — in business, at work, as an investor or as a builder — should do or believe this week. Charts appear only where the source itself puts comparable numbers side by side; nothing is estimated to fill a gap. Sources are listed in full on the final page.

SECTION 1 OF 3

AI Models & Research

2 stories · 2 medium

01 MEDIUM IMPACT Yahoo Finance

Meta Built a Humanoid Prototype in Three Months With a Chinese Partner

> Meta has built a humanoid robot prototype with Chinese startup Fourier in roughly three months and is already working on a second version.

The prototype, internally named Max, weighs about 106 pounds and stands around 1.5 metres tall. It runs on Nvidia's Jetson Thor chip, has a display that can mimic facial expressions, built-in speakers, and five-fingered hands designed to handle small objects, according to The Information, which cited people familiar with the project.

Meta's robotics team travelled to China to work with Fourier, and Li-Chen Miller, Meta's vice president of product for robotics, spent almost two months at the startup's Shanghai factory. Max was demonstrated to some Meta employees at an internal robotics summit in late August. Work on a second prototype, Max 2, began in September, with a redesigned torso, improved hands, a projector and additional sensors to help it navigate stairs.

The most striking detail is the speed: a working humanoid prototype assembled in about three months, with a second iteration already underway. That pace suggests Meta is treating this as a serious engineering effort rather than a research curiosity. The company sees humanoid robots as a potential physical extension of its Muse AI agent, with a long-term ambition of machines that can help with household tasks. Meta also strengthened its robotics capabilities through the May acquisition of Assured Robot Intelligence, a startup developing AI models that help robots understand and adapt to human behaviour.

Meta told The Information that its work with Fourier is focused on internal research and experimentation, and that it has no plans to sell the units. For now, the practical consequence is limited: no product is coming to market, and no revenue is attached. But the move signals that Meta is positioning itself in a field where competitors such as Tesla and a cluster of Chinese firms are already investing heavily, and where manufacturing partnerships may matter as much as software.

KEY FACTS

PROTOTYPE NAME
Max

WEIGHT
106 pounds

HEIGHT
1.5 metres

BUILD TIME
Roughly three months

PARTNER
Fourier, a Chinese robotics startup

SECOND PROTOTYPE
Max 2, development began September

WHY IT MATTERS

Meta is building robotics capability through a Chinese manufacturing partner at unusual speed, which could reshape competition in household robotics and shift how investors think about Meta's hardware ambitions.

Read the original at Yahoo Finance →

---

02 MEDIUM IMPACT arXiv.org

A New Test Rates Whether AI Can Name Its Sources

> Researchers have built a controlled test for whether language models can name the training data behind their answers, and found that a small change to how models score evidence sharply improves their ability to trace facts to the right source.

The work, posted to arXiv on 29 October, introduces the Controlled Attribution Benchmark, or CATT, a test bed in which language models are fine-tuned on entirely fictitious knowledge. Because every fact is invented for the experiment, researchers know exactly which training document a correct answer should point to — something impossible to establish with real-world data, where the same fact may appear in thousands of places.

The setup lets the team change one variable at a time, such as how often the same information is repeated across documents or how many similar-sounding distractors are present. Both consistently make attribution worse, the authors report. More striking is what fixes it. The true evidence inside a training document turns out to be highly concentrated in a few tokens, but standard methods dilute that signal by averaging scores across the whole document. Switching the loss function used to compute gradients — from cross-entropy to negative log-odds — surfaces the informative tokens, and aggregating only the top-scoring ones improved ground-truth recall by 35 to 72 percentage points across multiple baselines.

The correction also transferred to a separate benchmark, FTRACE-TREX, where it improved counterfactual retraining performance for every method tested, by up to 4.9 times.

For anyone building or buying systems that need to justify outputs — legal research tools, medical assistants, customer-support bots that cite policy documents — the finding is practical: attribution quality depends less on picking a fancier algorithm than on where you look for evidence inside a document and how you score it. The paper does not propose a new estimator, only a better way to use existing ones.

KEY FACTS

BENCHMARK INTRODUCED
Controlled Attribution Benchmark (CATT)

RECALL IMPROVEMENT
35 to 72 percentage points

RETRAINING PERFORMANCE GAIN
up to 4.9x

POSTED
29 October 2026

WHY IT MATTERS

Businesses relying on AI that must cite sources — in legal, medical or compliance settings — can get much more accurate attribution from existing tools by changing how evidence is scored, without new models or infrastructure.

Read the original at arXiv.org →

SECTION 2 OF 3

AI Tools & Ecosystem

4 stories · 1 high · 3 medium

03 HIGH IMPACT Simon Willison's Weblog

A ChatGPT Share Link Could Run Code in Your Session

> Security researchers found that a shared ChatGPT link could execute code inside another signed-in user's session, exposing private data and allowing actions on the victim's behalf.

Parameter Labs discovered and documented an attack chain against ChatGPT's HTML preview feature, which lets users share interactive custom apps inside a sandboxed frame. The researchers found that a flaw in how OpenAI integrated the MathLive JavaScript library allowed untrusted code from a sandboxed frame to inject HTML into the parent page. From there, the attacker could create a new frame without sandbox restrictions, then use YouTube's oEmbed endpoint to reflect JavaScript back in a way that bypassed OpenAI's content security policy and executed it.

Once the attacker had JavaScript running inside ChatGPT, fetching private user data from OpenAI's APIs and sending it out through an image reference was straightforward. The practical result: anyone could craft a ChatGPT share link that, when opened by another logged-in user, would exfiltrate that user's private data or take actions on their behalf. The victim only had to visit the link.

OpenAI has closed the hole and paid a $5,000 bug bounty. Simon Willison, who surfaced the write-up, called that amount "a bit cheap for a vulnerability this serious." The finding matters for any business where employees use ChatGPT with company or customer data: a single shared link could have become a data breach vector, and the attack required no action from the victim beyond opening it.

The episode is a reminder that the convenience features layered onto AI tools — shareable apps, previews, plugins — expand the attack surface faster than security review can keep up. For users and IT teams, the practical takeaway is to treat shared AI links with the same caution as email attachments, and to assume that any interactive preview feature may carry risk until the vendor proves otherwise.

KEY FACTS

BUG BOUNTY PAID
$5,000

RESEARCHER
Parameter Labs

STATUS
Patched by OpenAI

ATTACK VECTOR
Shared ChatGPT link

> **WHY IT MATTERS**
>
> Any business whose employees use ChatGPT with sensitive data should treat shared links as a potential breach vector and review what data staff expose through AI tools. The attack needed only a victim to open a link, making it a realistic phishing-style threat.

Read the original at Simon Willison’s Weblog →

04 MEDIUM IMPACT Tom's Hardware

Altman Says Your PC's Idle RAM Should Feed His Data Centers

> OpenAI chief Sam Altman says the RAM sitting idle in home computers is a wasteful use of the world's limited supply of memory chips — and that those chips would be better put to work in data centers like his own.

In an interview with the YouTube channel Cheeky Pint, Altman argued that running AI models locally on personal computers is "quite terrible" from a global allocation perspective. His reasoning: the world has a limited supply of DRAM, the memory chips that computers use for active tasks, and most home PCs use theirs only a small fraction of the time. Data centers, by contrast, can keep that memory busy around the clock.

"You want maximum utilization in the cloud if you're trying to deal with this very supply-constrained environment," Altman said. "And somehow many people's intuition is the opposite of that." The comment lands as OpenAI's primary business remains closed-source models that run in its own data centers, so the argument doubles as a defense of the subscription model that funds the company.

The stakes are straightforward. If local AI becomes good enough that people can run capable models on hardware they already own, they may pay less for cloud subscriptions. Altman's framing suggests OpenAI would rather see scarce memory chips flow to its data centers than to consumer desktops and laptops. That position has drawn sharp reactions from PC enthusiasts, who see it as an argument against owning your own hardware and in favor of renting computing power in the cloud indefinitely.

What the source does not address is whether data centers are actually more efficient users of that memory. Some commenters on the article pointed to reports that AI data centers have their own idle-capacity problems, buying hardware to keep it from competitors. The interview itself covers other ground — PHP vulnerabilities, sandboxing models during training, OpenAI's mathematics research — but the RAM remarks are the part likely to travel.

Altman's claim rests on an assumption that data center buildouts will continue at pace and that always-on cloud usage is inherently more worthwhile than memory used for spreadsheets, gaming, web browsing, or occasional local AI. Whether that assumption holds is the real question for anyone weighing a local AI setup against a cloud subscription.

KEY FACTS

ALTMAN'S DESCRIPTION OF HOME RAM USE
"Quite terrible"

ESTIMATED HOME PC RAM UTILIZATION
5% of the time

EXAMPLE HOME RAM CITED
32GB in a desktop

INTERVIEW PLATFORM
YouTube channel Cheeky Pint

WHY IT MATTERS

If Altman's view shapes policy or chip allocation, memory prices for consumer PCs could rise and local AI development could face headwinds — pushing more users toward cloud subscriptions. For businesses and individuals weighing local AI against cloud services, it signals where the industry's biggest player wants the market to go.

Read the original at Tom's Hardware →

---

05 MEDIUM IMPACT investing.com

AI Inference Startup Baseten Doubles Its Valuation in Four Months

> AI inference startup Baseten has doubled its valuation in four months.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

Baseten is raising about $2 billion across two tranches, according to Bloomberg. The round would value the company at $26 billion to $30 billion, double the $13 billion it reached in June. The company builds the infrastructure companies use to serve AI models.

WHY IT MATTERS

The reported jump suggests demand for AI-serving infrastructure remains strong enough to command sharply higher prices in a short period.

Read the original at investing.com →

---

06 MEDIUM IMPACT npr.org

Publishers' AI Detectors Are Falsely Flagging Their Writers

> Publishers' AI detectors are falsely flagging their own writers' work as machine-generated.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

A veteran ghostwriter told NPR that multiple AI-detection tools flagged her entirely human writing sample. Clients rarely mention using such screening. Some writers now deliberately mistype and drop their em dashes to seem more human. The flags are quietly reshaping how professionals write.

WHY IT MATTERS

If writers change their style to pass undisclosed AI checks, the quality of published work could suffer without clients knowing why.

Read the original at npr.org →

SECTION 3 OF 3

AI Applications & Industry

6 stories · 3 high · 3 medium

07 HIGH IMPACT TechCrunch

OpenAI Cuts Its $30 Billion Round Target to $20 Billion as Investors Balk

> OpenAI has lowered the target for its latest funding round from a reported $30 billion to $20 billion, and has secured only $11 billion so far, according to The Wall Street Journal.

The Financial Times reported Friday that OpenAI has cut its fundraising target to $20 billion, down from the $30 billion figure that had been reported earlier, when the company was said to be seeking a valuation of $1.4 trillion and then $1.25 trillion ahead of a planned IPO. A group of investors from the United Arab Emirates will lead the round, and a later retail offering for wealthy individuals could bring the total to the $20 billion target.

The Wall Street Journal added that OpenAI has raised $11 billion so far and that some large investors responded tepidly. The company has not commented publicly. The reported difficulty comes as Wall Street has grown wary of heavy spending on AI infrastructure, which has contributed to a sell-off in technology stocks. Anthropic's planned IPO may be facing similar headwinds, the reports suggest.

For investors and companies building on OpenAI's models, the smaller round is a signal about how the market now prices AI infrastructure bets. A $10 billion reduction in a single fundraising target is a material shift, and the gap between the $20 billion goal and the $11 billion secured so far suggests OpenAI may need to accept less favourable terms or rely more heavily on its UAE-based backers.

The most striking detail is the speed of the revision: reports of a $30 billion raise at a $1.4 trillion valuation have given way, within weeks, to a $20 billion target with a first close expected below that. That is not a company quietly adjusting its plans; it is a company responding to investors who are no longer willing to pay the prices they were discussing earlier this year.

KEY FACTS

PREVIOUS FUNDING TARGET
$30bn

NEW FUNDING TARGET
$20bn

AMOUNT SECURED SO FAR
$11bn

EARLIER REPORTED VALUATION
$1.4tn

LATER REPORTED VALUATION
$1.25tn

WHY IT MATTERS

A smaller round at a lower valuation would reset expectations for AI company valuations more broadly, and the tepid investor response suggests capital for AI infrastructure is becoming harder to raise.

Read the original at TechCrunch →

08 HIGH IMPACT CNBC

S&P Flags Meta's AI Spending as a Credit Risk

> S&P Global Ratings has revised Meta's credit outlook from stable to negative, citing uncertainty over whether the company's heavy AI spending will ever pay for itself.

The ratings agency said Friday that Meta's elevated capital expenditure and its first-ever negative free cash flow raise questions about the company's ability to monetize its AI investments relative to peers who are also pouring money into data centers. Meta's credit rating remains AA-, but the outlook downgrade signals that S&P sees a real chance of a full downgrade ahead. For context, S&P's scale runs from AAA at the top down to junk status below BBB, so Meta still sits comfortably in investment-grade territory — but the direction of travel has changed.

The numbers behind the concern are stark. Meta told investors last week that full-year capital spending will land between $135 billion and $145 billion. In the third quarter it posted negative free cash flow of $2.12 billion, the first time the company has gone into the red on that measure. Investors reacted by sending shares down 7% after the earnings report, and the stock finished October down 15%, its worst month in four years. On Monday Meta said it plans to raise $95 billion to fund AI initiatives, starting with $35 billion in public equity offerings and another $60 billion through an at-the-market program.

S&P's specific worry is timing. The agency noted that while Meta has launched its Muse AI agent app and related foundation models, "it is unclear when these revenue streams will generate material contributions that exceed their rising costs over the next few years." S&P said it expects capex growth to start flattening in 2028, but added it is "less certain that it will curtail its spending as previously anticipated." Meta pushed back, with a spokesperson saying the company "respectfully disagree[s]" and sees "significant product and business opportunities ahead."

The most striking detail is the scale of the financing: a company that generated enormous cash flow for years is now borrowing and issuing equity to fund AI, a reversal that would have been unthinkable before the current AI buildout. For investors, the question is no longer whether Meta can afford the spending, but whether the revenue from AI products will arrive before the market's patience runs out.

KEY FACTS

CREDIT OUTLOOK
Revised from AA- stable to AA- negative

Q3 FREE CASH FLOW
-$2.12 billion, first time negative

FULL-YEAR CAPEX GUIDANCE
$135 billion to $145 billion

PLANNED FUNDRAISING
$95 billion for AI initiatives

OCTOBER SHARE DECLINE
15%, worst month in four years

POST-EARNINGS SHARE DROP
7%

WHY IT MATTERS

Meta's outlook downgrade is an early warning that even the largest tech companies can strain their balance sheets on AI. Investors and business partners should watch whether AI revenue materializes fast enough to justify the spending, or whether more companies face similar credit scrutiny.

Read the original at CNBC →

09 MEDIUM IMPACT wsj.com

Bond Markets Are Shutting the Door on AI Data-Center Lenders

> Bond markets are closing to AI data-center lenders, cutting off a key source of cheap financing.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

CoreWeave, Nebius and Nscale have effectively been shut out of the bond market in recent weeks, according to the Wall Street Journal. Lenders are demanding tougher terms. For an AI build-out increasingly financed with debt, the cost of money is becoming the binding constraint on expansion.

> WHY IT MATTERS
>
> If debt-funded AI infrastructure builders lose access to bond markets, the pace of data-center construction could slow sharply.

Read the original at wsj.com →

10 MEDIUM IMPACT reuters.com

Foxconn Lines Up a Bigger AI Deal With OpenAI

> Foxconn said it is working on a larger AI deal with OpenAI, expanding an existing partnership on data-center hardware.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

The Taiwanese manufacturer, best known for assembling iPhones, announced the planned tie-up at its annual technology event. The new agreement would build on a prior collaboration to co-design data-center racks for OpenAI. Chairman Young Liu also pushed back on concerns about overheated AI investment, saying there is no AI bubble. Foxconn's AI-server business already accounts for more than half of its revenue.

> WHY IT MATTERS
>
> A deeper Foxconn-OpenAI relationship would further anchor one of the world's largest electronics manufacturers to the AI infrastructure build-out, but the scope and value of the new deal remain unconfirmed.

Read the original at reuters.com →

11 MEDIUM IMPACT TechCrunch

Apple's New CEO: No AI for AI's Sake

> Apple CEO John Ternus defended the company's slower approach to AI in a Wired interview, saying Apple will use the technology in new products but has no interest in shipping it for its own sake.

Ternus, a longtime Apple executive who became CEO at the beginning of September, acknowledged that critics have spent the past couple of years describing Apple as behind on AI, particularly after the company delayed a more powerful version of Siri that finally launched this summer. "Everyone has an opinion," he told Wired's Steven Levy, adding, "I really like where we're at."

His argument is that Apple's job is not to showcase AI but to hide it inside products people already want. "We don't want to ship AI for AI's sake," Ternus said. "When we're doing it right, people won't have to think about whether it was AI that made this happen or more traditional software. What matters is, did we build something that they really love?"

The most revealing moment came when Levy asked whether AI agents might eventually make apps unnecessary. Ternus pushed back: "That could make life really easy, I suppose, but the creation of products is a deeply human and personal thing, something that we love to do." That answer matters for developers and investors betting on an agent-driven future where software interfaces dissolve. Apple's CEO is signalling that the company's ecosystem — and its App Store economics — are not about to be replaced by a chatbot that does everything.

For readers tracking the AI race, the interview is a clear statement of strategy rather than a product announcement. Apple is choosing to compete on integration and user experience, not on shipping the largest models or the most visible AI features. The risk is that rivals keep setting expectations Apple declines to match; the bet is that customers reward products that simply work.

KEY FACTS

CEO
John Ternus

BECAME CEO
Beginning of September

SIRI UPGRADE LAUNCHED
This summer

INTERVIEW
Wired's Steven Levy

> **WHY IT MATTERS**
>
> Apple is signalling that it will compete on polished, integrated products rather than visible AI features — a stance that affects developers building for its ecosystem, rivals setting the pace on AI, and investors weighing whether Apple's slower rollout is discipline or a liability.

Read the original at TechCrunch →

---

12 HIGH IMPACT TechCrunch

Court Clears NAACP's Pollution Suit Over xAI's Memphis Turbines

> A federal judge has cleared the NAACP's pollution lawsuit against xAI to move forward, ruling that the gas turbines powering its Memphis data center are not exempt from air-permit

requirements.

Judge Debra M. Brown denied xAI's motion to dismiss the case, which the Southern Environmental Law Center filed on behalf of the NAACP. The suit accuses xAI of installing and operating at least 35 combustion turbines and other sources of air pollution at its Colossus data center site outside Memphis without obtaining the required preconstruction or operating air permits, and of violating the Clean Air Act.

xAI had argued the turbines should be treated as mobile sources, which would exempt them from permitting, and that the plaintiffs lacked standing to sue. Judge Brown rejected both arguments, writing that the turbines are not mobile sources under Mississippi's State Implementation Plan and that the plaintiffs have statutory standing under the Clean Air Act. She also dismissed xAI's constitutional challenge to citizen lawsuits, ruling that such plaintiffs "sue on their own behalf, in their own name, and seek penalties that redress harm to themselves."

The ruling is the second recent setback for xAI on this front: a similar motion to dismiss from the Department of Justice was denied earlier in October. The NAACP's director of environmental and climate justice, Abre' Connor, called the decision "another important victory in our effort to hold xAI accountable" and said the courts have repeatedly rejected attempts by xAI and the DOJ to dismiss the case.

The lawsuit is proceeding as public opposition to data centers hardens. New York recently passed a moratorium on new data centers, other state and local governments are weighing similar measures, and Amazon and Microsoft have promised to stop using non-disclosure agreements in their dealings with government agencies. TechCrunch has reached out to SpaceX, xAI's parent company, for comment.

KEY FACTS

TURBINES AT ISSUE
At least 35 combustion turbines

LOCATION
Colossus data center, outside Memphis

RULING
Motion to dismiss denied

JUDGE
Debra M. Brown, federal court

RELATED RULING
DOJ motion to dismiss denied earlier in October

NEW YORK ACTION
Moratorium on new data centers

WHY IT MATTERS

For companies building or operating large AI data centers, this ruling signals that gas-turbine permitting exemptions will face serious legal scrutiny, and that local environmental groups now have a viable path to challenge expansion plans in court.

Read the original at TechCrunch →

Sources

01 Meta Built a Humanoid Prototype in Three Months With a Chinese Partner
https://uk.finance.yahoo.com/news/meta-builds-humanoid-robot-china-180010281.html

02 A New Test Rates Whether AI Can Name Its Sources
https://arxiv.org/abs/2610.40655

03 A ChatGPT Share Link Could Run Code in Your Session
https://simonwillison.net/2026/Oct/31/how-we-took-over-chatgpt-accounts/

04 Altman Says Your PC's Idle RAM Should Feed His Data Centers
https://www.tomshardware.com/tech-industry/data-centers/openai-ceo-says-its-quite-terrible-for-your-pc-to-have-ram-sam-altman-says-data-centers-offer-better-utilization

05 AI Inference Startup Baseten Doubles Its Valuation in Four Months
https://www.investing.com/news/company-news/baseten-seeks-30bn-valuation-in-new-2-billion-funding-round--bloomberg-4988882

06 Publishers' AI Detectors Are Falsely Flagging Their Writers
https://www.npr.org/2026/10/31/nx-s1-5963520/ai-writing-detectors-pangram

07 OpenAI Cuts Its $30 Billion Round Target to $20 Billion as Investors Balk
https://techcrunch.com/2026/10/31/openai-reportedly-lowers-funding-target-for-new-round/

08 S&P Flags Meta's AI Spending as a Credit Risk
https://www.cnbc.com/2026/10/30/meta-outlook-to-revised-to-negative-by-sp-after-cash-flow-negative.html

09 Bond Markets Are Shutting the Door on AI Data-Center Lenders
https://www.wsj.com/finance/the-new-bottleneck-in-americas-ai-build-out-a-shortage-of-cheap-money-0c14c84d

10 Foxconn Lines Up a Bigger AI Deal With OpenAI
https://www.reuters.com/world/china/foxconn-working-new-tie-up-with-openai-says-there-is-no-ai-bubble-2026-10-30

11 Apple's New CEO: No AI for AI's Sake
https://techcrunch.com/2026/10/31/apple-ceo-says-the-company-doesnt-want-to-ship-ai-for-ais-sake/

12 Court Clears NAACP's Pollution Suit Over xAI's Memphis Turbines
https://techcrunch.com/2026/10/31/judge-allows-naacp-lawsuit-over-xai-data-center-pollution-to-move-forward/

About this document. Every story in the 1 November 2026 New Horizon AI Digest, reported at length. Each entry is written from the publisher's own article text; where a source could not be retrieved the entry is explicitly marked and kept short rather than padded.

Images and licensing. Figures are used only where the source licence permits redistribution, and are credited in the caption. Publisher artwork is not reproduced. All titles link to the original publication.

Vocabulary. Common AI terms are defined at new-horizon.tech/glossary/.

Generated 2026-11-01T06:36:47Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.