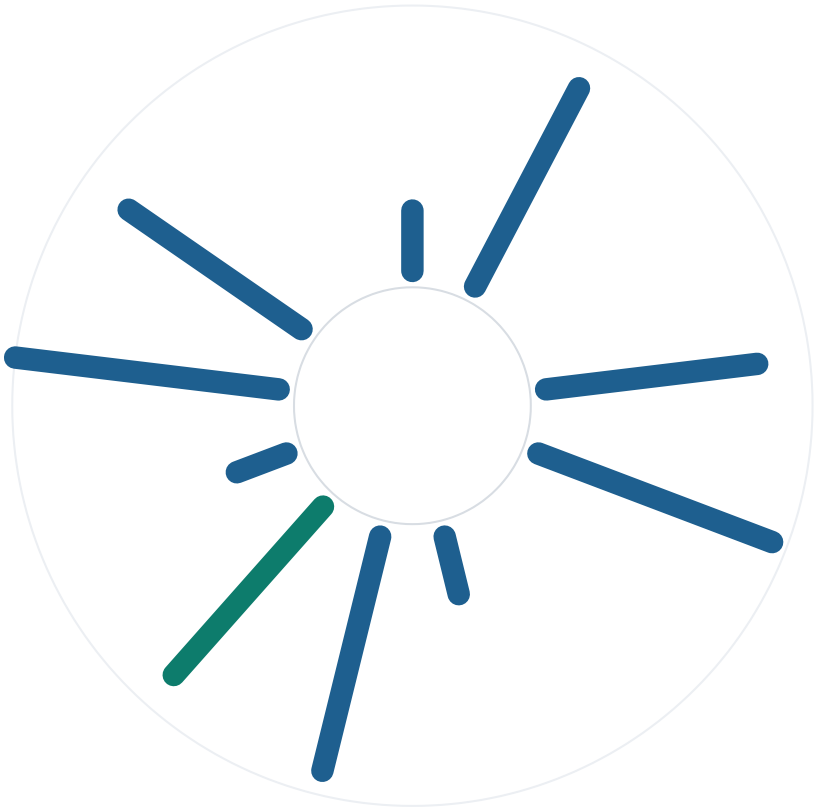
NEW HORIZON · AI DIGEST

The Long Read

27 November 2026

10	3	2626	1
STORIES	SECTIONS	WORDS	HIGH IMPACT

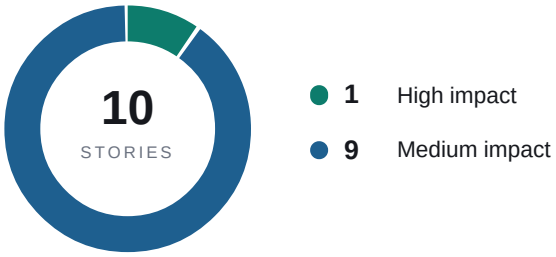


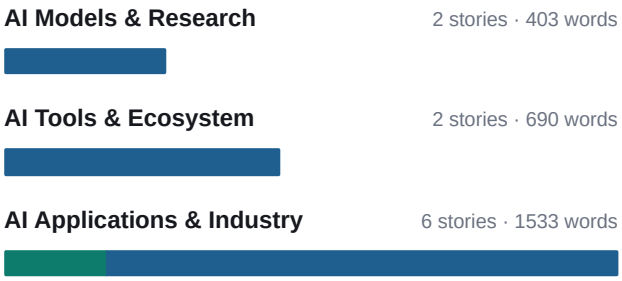
1 high impact 9 medium impact spoke length = depth of coverage

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

10 stories · 2626 words · 3 sections · 2 charted





AI Models & Research

01 Chinese AI labs warn of longer gaps between flagship models

02 Hugging Face tests whether agent teams earn their extra cost

AI Tools & Ecosystem

03 Public dashboard tracks local AI speed on AMD hardware

04 A laptop guide tests AI models built from just three values

AI Applications & Industry

05 Nvidia pauses AI server shipments to four Southeast Asian countries

06 US approves SpaceX plan for one million orbiting data satellites

07 SAP weighs at least 10,000 job cuts in AI restructuring

08 Apple removes MiniMax’s Talkie chatbot worldwide over child safety

09 Undisclosed AI use has affected at least 25 court cases

10 New York Times bars Photoshop 2027 from assignment photography

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 reuters.com

Chinese AI labs warn of longer gaps between flagship models

> Chinese AI labs are warning that the interval between flagship model releases could lengthen, with Z.AI's next model behind schedule and Moonshot saying gaps could double.

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.

Two Chinese AI developers have put timelines on the pressure. Z.AI's next model is running late. Moonshot says the wait between releases could stretch from six months to a year if access to overseas computing disappears. The causes named are new US licensing rules and tighter restrictions on copying rival models, which close off two routes the labs have relied on for training. No revised dates or figures were given.

> WHY IT MATTERS
>
> Release cadence is the visible scoreboard in AI, so longer gaps would suggest US export controls and licensing are slowing Chinese labs in practice, not just on paper — though the full report was unavailable and the scale of any delay is unclear.

Read the original at reuters.com →

02 MEDIUM IMPACT huggingface.co

Hugging Face tests whether agent teams earn their extra cost

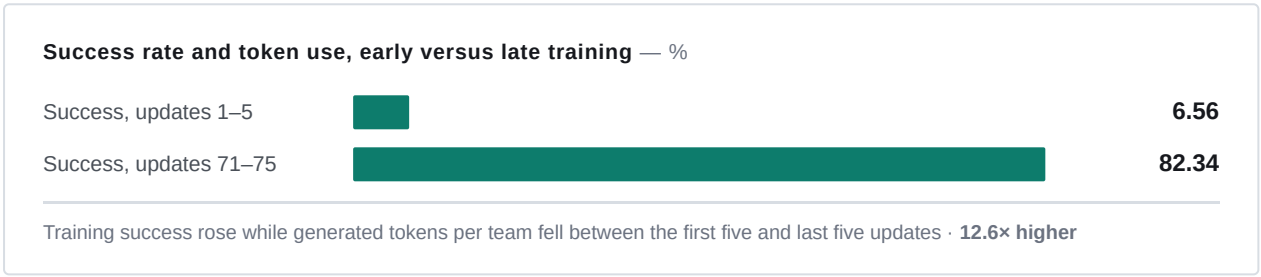
> Hugging Face has published a reproducible test of whether teams of AI agents are worth their extra cost, and its first small experiment suggests they can be — but only after training teaches them to talk.

The company's blog post, the first in a series on multi-agent systems, sets out a question that has hung over the recent enthusiasm for "swarms" of AI agents: when agents cooperate, do they actually help, or are they just spending more? The post notes that OpenAI announced in September that a swarm of roughly 10,000 agents had found a singularity in the 3D Navier-Stokes equations, one of the Millennium Prize problems. But Noam Brown estimates that less than 10% of that result came from multi-agent cooperation. In Anthropic's vulnerability-finding experiment, coordinated agents found more bugs than independent agents, but also searched more code and used more tokens; restricted to the same directories, results per million tokens were about the same.

Hugging Face's own test is deliberately modest. Four agents share one small language model and must pool four private numbers to produce a correctly sorted list. No agent can solve it alone. Over 75 training steps, success rose from 6.6% to 82.3%, and the team used 37.4% fewer tokens per attempt, falling from about 13,000 to 8,000. Messages per team rose from 3.48 to 6.28, with broadcasts climbing from 2.38 to 5.73 — more talking, less text.

The comparison that matters for budgets: an untrained team given eight retries solved 37.5% of problems and spent about 103,000 tokens per problem. The trained team, in one attempt, solved 82.3% and spent about 8,000. The post cautions that both figures come from training problems rather than held-out ones.

The work is a walkthrough, not a product. The example runs on two H200 chips for roughly $50 and about five hours; a shorter version took about 4.8 hours on two H100s and reached 60.9% held-out success. The source reports 93.2% four-agent success after 200 updates in the underlying paper. What is new is a cheap, checkable baseline for anyone weighing whether to add agents to a workflow — and a reminder that the answer depends on the task, the model and how the work is divided.



KEY FACTS

TRAINING SUCCESS, FIRST TO LAST FIVE UPDATES
6.6% to 82.3%

HELD-OUT SUCCESS AT CHECKPOINT 75
26 of 32 problems (81.25%)

TOKENS PER TEAM ATTEMPT
Down 37.4%, from about 13,000 to 8,000

UNTRAINED TEAM, BEST OF 8 ATTEMPTS
37.5% success, about 103,000 tokens per problem

COST OF THE WALKTHROUGH
About $50 on two H200s, roughly five hours

PAPER'S FOUR-AGENT SUCCESS AFTER 200 UPDATES
93.2%

> **WHY IT MATTERS**
>
> Teams building with AI now have a cheap, reproducible way to test whether adding agents beats one agent with more retries at the same token budget — the question that decides whether multi-agent setups save money or quietly multiply it.

Read the original at huggingface.co →

SECTION 2 OF 3

AI Tools & Ecosystem

2 stories ● 2 medium

03 MEDIUM IMPACT huggingface.co

Public dashboard tracks local AI speed on AMD hardware

> A public dashboard on Hugging Face now tracks how quickly AI models run on AMD's consumer and workstation chips, with weekly tests and the raw results open for anyone to download.

The dashboard, built by the llama.cpp project with AMD's help, measures two things for each combination of chip, model and software build: how fast a machine processes a prompt, and how fast it generates text. Both are reported in tokens per second, the rough equivalent of words per second, so a buyer can see in plain numbers whether a given graphics card will feel quick or sluggish running a model locally.

That matters to anyone weighing a local setup against a cloud subscription. Running a model on your own hardware avoids per-token bills and keeps data on the machine, but performance varies widely with the chip, the model and the software build. The dashboard currently covers two AMD chips, the Radeon 8060S and the Radeon AI Pro R9700S, and a model set that includes Inkling 2 Mini, Qwen, Gemma, gpt-oss and LFM, in several compression formats. The tests run weekly using the standard llama-bench tool, with Flash Attention on and off and 10 repetitions per configuration.

The most useful detail for a non-specialist is the context-scaling tab. It shows how generation speed falls as the conversation gets longer, which is the difference between a demo that feels instant and a long session that starts to crawl. The other tabs compare compression levels for a single setup and track performance across software builds, so a regression after an update is visible rather than mysterious.

What is new here is the packaging rather than the benchmarking itself: the workflow and the underlying results are public, stored in a Hugging Face bucket, and the data source link beside each chart opens the folder for that build, chip and model. Anyone can reuse the files to make their own charts or check a result against a local run. The project is also asking users which models and workloads to add next, a sign the coverage is expected to grow.

KEY FACTS

WHERE
Hugging Face

CHIPS COVERED
Radeon 8060S and Radeon AI Pro R9700S

TESTS
Prompt processing (PP512) and text generation (TG128), in tokens per second

TEST FREQUENCY
Weekly

REPETITIONS PER CONFIGURATION
10

DATA ACCESS
Public, downloadable from a Hugging Face storage bucket

WHY IT MATTERS

If you are deciding whether to run AI on your own AMD hardware instead of paying per token in the cloud, this gives you public speed numbers before you buy. It also lets teams spot when a software update slows their local setup down.

Read the original at huggingface.co →

---

04 MEDIUM IMPACT MachineLearningMastery.com

A laptop guide tests AI models built from just three values

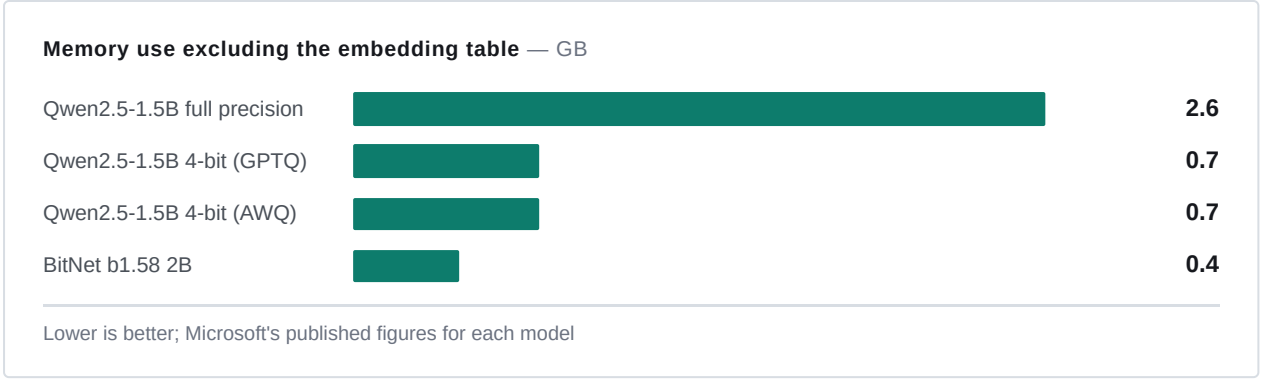
> A practical guide shows how to run Microsoft's BitNet — a language model whose weights are restricted to just three values — on an ordinary laptop processor, and how to measure it against a standard 4-bit model on the same machine.

BitNet is Microsoft's family of models trained so that every weight is one of three values: -1, 0 or +1. Three states carry about 1.58 bits of information, which gives the model its name. The point of training that way from the start, rather than rounding a finished model down to four bits, is that the model learns to work within the constraint instead of being damaged by it. The guide walks through building Microsoft's bitnet.cpp software, downloading the 2-billion-parameter model, and serving it alongside a 4-bit Qwen2.5-1.5B on the same laptop, then scoring both with one short Python script on speed, memory, message classification and tool-call extraction.

Microsoft's published numbers, which the guide uses as its reference, put BitNet's memory use (excluding the embedding table) at 0.4 GB against 0.7 GB for the 4-bit Qwen versions. On a three-test average BitNet scores 55.01 against 52.15 and 51.17 for the two 4-bit builds. The split matters more than the average: BitNet loses clearly on MMLU, a general knowledge test (53.17 against 58.06 and 57.43), and wins on grade-school maths and instruction following. Microsoft reports decoding at 29 milliseconds per token against 65 for full-precision Qwen on a Core i7 laptop with 8 threads — roughly 34 versus 15 tokens per second. The report gives no speed figure for the 4-bit models, so the comparison most readers care about is unpublished; hence the script.

The most surprising detail is the energy claim. The report lists an estimated 0.028 joules per token against 0.347, but an appendix says that comes from an arithmetic model, not a power meter. The downloaded files are also closer in size than the memory table suggests: about 1.19 GB for BitNet against about 0.99 GB for the 4-bit Qwen.

What is new here is not the model — Microsoft published it earlier — but a repeatable way to test it on your own hardware and tasks. The guide's own conclusion is that the tiny model suits classification, field extraction and drafting tool calls, provided plain code checks every output and failures escalate to a larger model. Microsoft says it does not recommend BitNet b1.58 for commercial or real-world use without further testing, and the model is English-focused with a 4,096-token context limit.



KEY FACTS

MEMORY USE, BITNET VS 4-BIT QWEN
0.4 GB vs 0.7 GB

AVERAGE SCORE ON THREE TESTS
55.01 BitNet vs 52.15 and 51.17 for 4-bit Qwen

DECODING SPEED ON A CORE I7 LAPTOP
29 ms per token vs 65 ms full-precision

DOWNLOAD SIZE
About 1.19 GB BitNet vs about 0.99 GB 4-bit Qwen

CONTEXT LIMIT
4,096 tokens

MICROSOFT'S GUIDANCE
Not recommended for commercial use without further testing

WHY IT MATTERS

Teams building agents can now test whether a model small enough for a laptop CPU handles the routine work — sorting messages, pulling out fields, formatting tool calls — before paying for a larger model or a GPU. The catch is that the speed advantage over 4-bit models is unmeasured, so the answer depends on your own hardware and tasks.

Read the original at MachineLearningMastery.com →

SECTION 3 OF 3

AI Applications & Industry

6 stories · 1 high · 5 medium

05 MEDIUM IMPACT TechCrunch

Nvidia pauses AI server shipments to four Southeast Asian countries

> Nvidia has paused AI server shipments to four Southeast Asian countries while export compliance among regional partners is reviewed.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (extracted only 115 words (paywall/consent wall?)). Follow the link for the full report.

Nvidia told Dell and Super Micro to hold shipments bound for Malaysia, Singapore, Thailand and Indonesia while it reviews whether regional partners are complying with export rules. The four countries host remote data centres that, by one estimate cited in the summary, supply 24% of China's AI computing power. The pause affects server makers and the regional operators waiting on hardware; no timeline for the review was given.

> WHY IT MATTERS
>
> If the review tightens how hardware reaches those data centres, a meaningful share of the computing power available to Chinese AI developers could be affected, though the scope and duration of the pause are not yet clear.

Read the original at TechCrunch →

06 MEDIUM IMPACT BBC News

US approves SpaceX plan for one million orbiting data satellites

> US regulators have given SpaceX approval to launch a million satellites into low-Earth orbit to run data centres in space, the first licence of its kind.

The Federal Communications Commission (FCC) said on Wednesday that the project, called Starmind, would "support America's booming super intelligence economy". The satellites would be powered by sunlight and process data for artificial intelligence in orbit, relaying results back to Earth through SpaceX's existing Starlink internet network. FCC Chairman Brendan Carr called it "a bold new frontier" for computing infrastructure; Elon Musk, SpaceX's owner, wrote on X: "Much appreciated. This will be a great asset to America and the world."

The licence sets a schedule rather than a finished system. Deployment must begin no later than November 2033, with 10% of the satellites in place by November 2035 and 50% by November 2038. SpaceX has said it wants to launch its first Starmind satellite next year. The BBC notes it is unclear whether a million satellites

will actually fly, or whether SpaceX sought approval for more than it intends to use. For comparison, Starlink already runs more than 10,000 satellites.

SpaceX argues orbital data centres are the most cost- and energy-efficient way to meet AI computing demand beyond what can be built on the ground. That pitch lands as terrestrial data centres face growing local opposition: they are large, use enormous amounts of power and water for cooling, and the BBC reports experts saying they were a factor in the Republicans' big midterm election losses. Google and Blue Origin are also exploring the orbital concept.

The approval drew immediate criticism. Amazon, a rival, called it "a lofty ambition rather than a real plan". The American Astronomical Society said it was "disheartened and dismayed", warning the constellation's larger and more numerous satellites would appear as extremely bright objects for longer periods and "significantly undermine decades of federal investment in astronomical facilities". Prof Sir Martin Sweeting of Surrey Satellite Technology said the idea was technically feasible but that moving huge volumes of data would depend on laser links from the ground, which he said have not been perfected to that level. An Earthjustice attorney said expanding satellites in space has a direct impact on people's everyday lives. Experts also point to the cost of launching hardware and the risk from space debris and collisions.

KEY FACTS

SATELLITES APPROVED
One million

REGULATOR
US Federal Communications Commission (FCC)

DEPLOYMENT MUST BEGIN BY
November 2033

SHARE DEPLOYED BY NOVEMBER 2038
50%

FIRST STARMIND LAUNCH PLANNED
Next year

STARLINK'S CURRENT SATELLITES
More than 10,000

> **WHY IT MATTERS**
>
> SpaceX has secured a regulatory path to sell AI computing capacity that doesn't compete for land, power or water on Earth — a signal to investors and rival operators that orbital data centres are being taken seriously by US regulators, even if the economics remain unproven.

Read the original at BBC News →

---

07 HIGH IMPACT Yahoo Finance

SAP weighs at least 10,000 job cuts in AI restructuring

> SAP is preparing to eliminate at least 10,000 positions — at least 9% of its global workforce — to rebuild its software business around artificial intelligence, according to people familiar with the matter.

The Walldorf, Germany-based company, Europe's most valuable software maker, has not finalized the number, and it could change after labor negotiations, the people said. SAP declined to say whether cuts are planned, offering only that it "continuously reviews its organization." All departments have been asked to identify possible cuts, with product and engineering and sales under particular scrutiny.

The backdrop is a push by Chief Executive Christian Klein to embed AI tools into SAP's applications after earlier attempts at AI products ran into problems. The executive board has been reorganized twice this year; two of six board members leave in March, and Klein has taken over most of the product lead's responsibilities. Sales chief Thomas Saueressig said last week he will leave to become CEO of Ionos Group SE. SAP has announced major restructurings in late January alongside full-year results before, most recently in 2024 — a round that also eliminated around 10,000 positions, cost about €3 billion ($3.3 billion) and took more than a year, mostly through voluntary departures and retraining. Headcount stayed roughly flat as SAP hired for new roles.

That precedent is the detail worth repeating to a colleague: a 10,000-job cut at SAP has historically meant a reshuffle of skills, not a shrinking company. What is different this time is the stated purpose — reorganizing the business for the AI era rather than trimming costs — and the timing, ahead of the January results.

For SAP's roughly 110,000 employees worldwide, the immediate constraint is Germany, where the labor council negotiated an employment guarantee through 2028. That limits how fast cuts can fall on the largest share of staff and pushes the plan toward the voluntary and retraining route used in 2024. For customers, the question is whether a leaner product and sales organization speeds up or disrupts the AI features SAP is promising to build into its applications. For investors, the 2024 playbook — a multi-billion-euro charge absorbed over a year — is the closest guide to what this one costs.

KEY FACTS

JOB CUTS PLANNED
At least 10,000

SHARE OF GLOBAL WORKFORCE
At least 9%

TOTAL EMPLOYEES WORLDWIDE
More than 110,000

PREVIOUS ROUND (2024)
Around 10,000 positions, about €3 billion ($3.3 billion), over a year

GERMAN EMPLOYMENT GUARANTEE
Through 2028

BOARD DEPARTURES
Two of six members leave in March

> **WHY IT MATTERS**
>
> SAP's move signals that AI restructuring at large software firms is arriving as headcount plans, not product launches — and that a German employment guarantee through 2028 will shape how the cuts are actually carried out. Workers in product, engineering and sales at enterprise software companies should read this as the template for the next round.

Read the original at Yahoo Finance →

---

08 MEDIUM IMPACT reuters.com

Apple removes MiniMax’s Talkie chatbot worldwide over child safety

> Apple has pulled MiniMax's Talkie chatbot from its App Store worldwide over child-safety concerns.

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 removal follows reports that AI-generated characters on Talkie sent sexualised messages to minors. The app had about 11 million monthly active users in October. Child-safety groups had questioned its age checks despite an 18+ app-store rating. Apple's action removes Talkie from the main distribution channel for iPhone users globally. MiniMax has not commented in the material available.

> **WHY IT MATTERS**
>
> If app stores treat an 18+ rating as insufficient protection for minors, consumer AI chat apps face stricter age checks or removal.

Read the original at reuters.com →

---

09 MEDIUM IMPACT the Guardian

Undisclosed AI use has affected at least 25 court cases

> At least 25 court cases in England and Wales this year have been affected by undisclosed use of AI, according to a trawl of published judgments by the Bureau of Investigative Journalism shared with the Guardian.

Judges have had to deal with AI-generated witness statements, fictitious case law and statutes that do not exist. In one high court dispute between tax lawyers, the judge called the use of "generative AI hallucinations" and references to nonexistent cases "unacceptable". In another, a barrister in a $21m (£16m) dispute between two Nigerian oil companies argued from a ruling supposedly by a senior judge that a hallucinating AI had invented outright; the judge warned of "the dangers of relying on the product of AI without verification". Earlier this month a Crown Prosecution Service lawyer cited an imaginary legal precedent.

The problem is not confined to professionals cutting corners. Many cases involved people representing themselves, and Brett Dixon, president of the Law Society, said AI producing verbose documents is eating judges' time and undermining efforts to clear backlogs. "You go from three pages from a lawyer … using human judgment to a rambling 60 pages from a commercial AI," he said. Lord Sales, a supreme court judge, said the court was growing "increasingly concerned", and that the greater long-term danger is that people "gradually lose the ability to know whether it has got the answer wrong".

The most striking case involved a police officer in a grievous bodily harm prosecution who uploaded a transcript of the victim's interviews to Microsoft's Copilot and prompted it "to achieve the highest sentence and the most tear-jerking VPS to make a judge or reader weep when reading". The officer was removed from frontline duty; the court of appeal, dismissing the attacker's challenge to a four-year sentence, said no police officer should ever prepare a statement that way. In another, insolvency judge Mark Mullen found a junior solicitor had relied on a chatbot's invented quote from a fictitious statute, examined the chat logs, and warned that "AI has the potential to be wholly unreliable".

What is new is the volume and the response. Last month the Civil Justice Council, in a working group led by Sir Colin Birss, the second most senior judge in England and Wales, recommended that witnesses in civil cases sign a declaration that AI did not generate or influence the substance of their statements, and flagged urgent guidance needed for litigants in person. The Ministry of Justice said it was taking steps to ensure AI use is "safe, appropriate, and properly controlled", including more judges, legal officers and modern case management systems.

KEY FACTS

CASES AFFECTED IN 2026
At least 25

DISPUTED SUM IN OIL COMPANIES CASE
$21m (£16m)

PRISON SENTENCE CHALLENGED
Four years

AI TOOL USED BY POLICE OFFICER
Microsoft Copilot

DOCUMENT LENGTH CITED
Three pages from a lawyer vs a rambling 60 from AI

PROPOSED SAFEGUARD
Witnesses in civil cases to sign a declaration that AI did not generate or influence their statements

> WHY IT MATTERS
>
> Any organisation whose staff draft legal, regulatory or official documents with AI now has a documented record of the failure mode — invented citations that survive into court — and courts are moving toward signed declarations that the work was not AI-generated. Verification is becoming a procedural requirement, not a courtesy.

Read the original at the Guardian →

---

10 MEDIUM IMPACT The Verge

New York Times bars Photoshop 2027 from assignment photography

> The New York Times has asked freelance photographers not to use Adobe Photoshop 2027 on assignment, because the software's AI features can stamp their files with metadata identifying the image as AI-generated.

The instruction came in an email from two of the paper's photography leaders, Meaghan Looram and Josh Haner, seen by The Verge and posted on Threads. They wrote that the new features use AI "in ways that may not be apparent to users" and that this "can result in metadata identifying photographs as AI-generated." A Times spokesperson, Graham James, confirmed the requirement. The paper is asking contributors to avoid release 28.0.0 and later updates while shooting for it.

The November update added several AI tools, including an "AI Assisted Editor" view that gathers Photoshop's AI features on one toolbar, an AI assistant that gives step-by-step editing instructions, prompt-based image editing, new AI upscaling, and a tool that recomposes an image for different canvas sizes. The problem is that it is not always obvious which tools are AI-powered; even older automated selection tools may now count. A photo that was only lightly adjusted can therefore carry a tag saying it was edited with AI.

That tag comes from C2PA metadata, a standard Adobe began supporting across its apps in August to record whether an image was generated or edited with AI. For a newsroom that keeps editing to a minimum, the label is misleading. The Times says it does not retouch or digitally manipulate photographs, though it may make minimal colour adjustments to reproduce an image accurately. A blunt AI tag on a routine assignment photo risks readers' trust.

The Times did not say which specific Photoshop features prompted the ban, and Adobe did not respond to The Verge's request for comment. The restriction is narrow: it applies to assignment work, not to the software generally. But it puts a number on a quiet cost of AI features — that the tools can leave marks a publisher never intended, and that those marks are now visible to anyone who checks the file.

KEY FACTS

BANNED VERSION

Photoshop 2027, release 28.0.0 and later

WHO ISSUED IT

Meaghan Looram and Josh Haner, New York Times photography leaders

METADATA STANDARD

C2PA, supported across Adobe apps since August

ADOBE RESPONSE

Did not respond to The Verge's request for comment

> WHY IT MATTERS
>
> News organisations and other publishers that rely on freelance contributors now have to police which software versions touch their images, because an automatic AI label can undermine a photo's credibility even when nothing was fabricated. Adobe's customers, meanwhile, face a choice between new tools and clean metadata.

Read the original at The Verge →

Sources

01 Chinese AI labs warn of longer gaps between flagship models
https://www.reuters.com/world/china/us-chip-curbs-distillation-crackdown-threaten-chinese-ai-model-development-2026-11-26

02 Hugging Face tests whether agent teams earn their extra cost
https://huggingface.co/blog/training-swarms

03 Public dashboard tracks local AI speed on AMD hardware
https://huggingface.co/blog/ggml-org/llamacpp-amd-bench

04 A laptop guide tests AI models built from just three values
https://machinelearningmastery.com/running-1-bit-llms-locally-a-practical-guide-to-bitnet-inference/

05 Nvidia pauses AI server shipments to four Southeast Asian countries
https://techcrunch.com/2026/11/26/nvidia-hits-pause-on-chip-shipments-to-southeast-asia/

06 US approves SpaceX plan for one million orbiting data satellites
https://www.bbc.co.uk/news/articles/cwp836l8n1jgo

07 SAP weighs at least 10,000 job cuts in AI restructuring
https://finance.yahoo.com/technology/ai/articles/sap-weighs-cutting-least-10-161548594.html

08 Apple removes MiniMax's Talkie chatbot worldwide over child safety
https://www.reuters.com/business/retail-consumer/apple-pulls-minimax-ai-chatbot-app-store-over-child-safety-concerns-2026-11-26

09 Undisclosed AI use has affected at least 25 court cases
https://www.theguardian.com/technology/2026/nov/26/dozens-of-court-cases-in-england-and-wales-affected-by-undisclosed-ai-use-in-2026

10 New York Times bars Photoshop 2027 from assignment photography
https://www.theverge.com/ai-artificial-intelligence/1041227/the-new-york-times-warns-photographers-not-to-update-photoshop-over-ai-features

---

About this document. Every story in the 27 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 2026-11-27T06:39:02Z · model deepseek-v4.1-flash:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.