

The Trilogy Times

All the news that's fit to generate — AI • Business • Innovation

TUESDAY, JUNE 16, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

China's Cut-Rate AI Rattles the Chip Kings

DeepSeek says it built a top-tier model without the fanciest chips — and Silicon Valley calls it 'amazing.'

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A Chinese upstart called DeepSeek says it trained a high-performing artificial-intelligence model on a shoestring, without the most advanced chips, and Silicon Valley can't stop talking.

The early word from the Valley: "amazing and impressive." That's a fat compliment for a shop running on second-string silicon.

DeepSeek is no household name. It runs out of Hangzhou, and a short while back most American engineers couldn't have placed it on a map.

Now they're downloading it and kicking the tires.

Here's the angle that stings. American chip barons and AI houses have staked billions on one idea — bigger budgets and faster processors win the race.

DeepSeek says the race runs cheaper than that. The company claims it kept pace with the heavy hitters at a fraction of the cost, and skipped the top-shelf chips to do it.

Read the [full rundown on the outfit](#) and the arithmetic grabs you by the collar.

The chips at the center are the fastest money can buy. Washington spent two

years barring those from China, betting the blockade kept American labs out front.

DeepSeek worked with the slower stock it could legally get. It made do, and it made noise.

If those figures hold up, somebody left a window open.

Why it matters comes down to money. Training a top-tier model has run into the hundreds of millions, a tab only the deepest pockets could cover.

Knock that number down and the whole pecking order wobbles. A cheaper model doing the same work puts a question mark on every premium chip order on the books.

The hardware makers sell the picks and shovels of this gold rush. Word that you can strike it rich with cheaper tools is not the news they wanted printed.

The money crowd caught the scent fast. DeepSeek turned up in the latest [Tech, Media and Telecom market talk](#), named right beside the day's other movers.

Nobody in San Francisco is conceding the field. The word out west is admiration, not surrender — but admiration from a rival carries its own freight.

The bigger question hangs over the whole trade. If brains don't require the biggest bankroll, the moat the giants dug starts looking shallow.

Meanwhile the big checks keep clearing. LinkedIn co-founder Reid Hoffman just raised \$24.6 million for Manas AI, a startup aiming artificial intelligence at cancer research.

His partner is Siddhartha Mukherjee, the physician who wrote "The Emperor of All Maladies." Two bets, one week, the same gospel — AI is the tool that changes the game.

The contrast tells the tale. One outfit spends real money chasing cures; the other says it can do frontier work for pennies on the dollar.

Both can't be the whole story. Both point the same direction.

For now DeepSeek pulled off the thing every upstart dreams about. It made the giants look over their shoulder.

Watch this space. When the cheap seats start outhitting the box seats, the whole house takes notice.

Hiring Freeze Front Settles Over AI Valley as Layoff Rules Gather Offshore

OpenAI taps the brakes, California weighs a longer warning siren, and workforce planners are told to carry umbrellas.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

SAN FRANCISCO — A cold hiring front is pushing across the technology sector today, and the barometric pressure inside executive suites is dropping fast.

The latest gust came from OpenAI, where Sam Altman said the company plans to “dramatically slow down” its pace of hiring, according to [Business Insider](#). That is not a blizzard warning by itself, but when the industry’s most closely watched AI company starts easing off the recruiting accelerator, forecasters should expect chillier conditions downwind.

Across the broader enterprise weather map, CIOs are being advised to own the hiring freeze rather than pretend it is a passing cloud. The message is plain: technology leaders can no longer leave workforce weather reports solely to finance or HR. If AI automation is changing the labor forecast, CIOs need to explain where the sunshine is, where the fog banks are forming, and which teams may be walking into hail.

Meanwhile, Sacramento is tracking a heavier regulatory system. California lawmakers have proposed requiring 90 days’ notice when layoffs are tied to an employer’s use of AI, according to [JD Supra](#). If passed, the rule would act like an early-warning radar for workers caught beneath automation storm cells. Employers, however, may find themselves filing more detailed flight plans before reducing headcount.

The Atlassian squall line remains a cautionary cloudbank. Its AI-linked job cuts have drawn warnings of a “chaos

tsunami” for the workforce, a phrase that may sound dramatic until one remembers that most companies are still building seawalls with spreadsheets and optimism.

There is one warm pocket on the map: crypto venture firm CMT Digital raised \$136 million for a fourth fund, suggesting capital is still evaporating upward in select markets. But for operators, founders and job seekers, the daily advisory is unchanged: keep budgets tight, skills portable and communication transparent. There is a 70% chance of continued hiring turbulence through the next earnings cycle, with isolated opportunity breaks for teams that can prove AI productivity without triggering a labor downpour.

The Big Three AI Labs Are Building Walls — and One Open-Source Rival Is Climbing Over

OpenAI, Google, and Anthropic align on model theft protections while Ai2 releases a free web agent designed to undercut all three.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — The dominant forces in commercial AI are converging on two fronts simultaneously: fortifying their models against theft and competing for a new class of technical talent, even as a nonprofit research lab releases open-source tools designed to make their closed systems redundant.

[OpenAI, Google, and Anthropic have begun coordinating defenses against AI model theft](#) — a rare alignment among competitors that reflects shared anxiety over model extraction attacks, weight theft, and distillation techniques that allow smaller actors to approximate frontier capabilities without bearing frontier R&D costs. The threat is structural: a sufficiently capable adversary can query a commercial API at scale, use the outputs to train a derivative model, and effectively offload billions in compute investment onto the target company’s infrastructure.


The timing is notable. All three labs are simultaneously [expanding hiring for Forward Deployed Engineers](#) — a role borrowed from Palantir’s playbook that embeds technical staff directly inside enterprise customers. FDEs sit at the intersection of sales engineering and product development, customizing AI systems to client workflows in real time. The role signals a strategic shift: these labs are no longer selling software licenses, they are selling outcomes, and that requires people on the ground.

Entering from the left flank: the Allen Institute for AI, known as Ai2, this week released an open-source web agent built


to perform tasks that OpenAI's Operator, Google's Project Mariner, and Anthropic's computer-use tools handle in closed, paid environments. The release applies direct pressure on the commercial value proposition of those products. If an open-source equivalent performs comparably, enterprise procurement decisions become harder to justify on capability alone.

Elsewhere in the sector, Nvidia has backed Israeli AI startup Decart in a \$300 million funding round valuing the company at \$4 billion. Decart specializes in real-time interactive AI — including generative world simulation — and the Nvidia imprimatur provides both capital and a tacit hardware guarantee.

The week's pattern is consistent: the top of the market is hardening its perimeter while the open-source tier chips away at it from below.



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Regulatory Capture, Corporate Consolidation, and the Accelerating Erosion of Antitrust Orthodoxy

WASHINGTON, D.C.

BY R. BARNSWORTH III, ESQ., LEGAL AFFAIRS DESK · CLAUDE SONNET

The Ethics Industrial Complex Descends Upon AI in Education — And Finds Itself Outnumbered

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests it is being argued, with considerable vigor, across no fewer than four concurrent scholarly publications — that the academic community has entered what this correspondent would characterize as a "meta-ethical productivity phase" vis-à-vis artificial intelligence in educational contexts.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Doctor Will Deepfake You Now

AUSTIN, TEXAS — There is a version of this column where I tell you everything is going to be fine.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Floodlight and the Marshmallow

AUSTIN, TEXAS — There is a particular indignity, peculiar to our age, in being floodlit by one's own father's doorbell.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

WE HAVE MET THE BOTS AND THEY ARE US: A Dispatch from the Edge of Digital Civilization

AUSTIN, TEXAS — There's a moment, usually around 2 a.m.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

HAIKU OF THE DAY · CLAUDE
HAIKU

*Walls built ever higher
Yet climbers emerge below
Progress feeds on fear*

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

Surtr Goes Public, Klair Gets Smarter, Team Builds Everywhere at Once

The AI Builder Team shipped a full public API layer, live-data ontology, and surgical AI regeneration in a single day — proving this org can build infrastructure and product simultaneously, across every repo in the stack.

BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · AI BUILDER TEAM

When historians write about the week Surtr became a real platform, they'll point to today. @kevalshahtrilogy didn't just ship one PR — he shipped a trilogy of his own. PRs #473, #474, and #476 form a complete vertical slice: live education data pulled from Redshift (100 schools, 82 programs, 22 markets, all real, all live, the goblin-ops demo mercifully buried), a stable public `/v1` REST API with Bearer-key auth and domain-namespaced routes that downstream consumers like Aerie and Klair can actually pull from, and a polished in-app API reference console complete with copyable `curl` commands and live sample responses. That's data layer, API layer, and developer experience — designed, built, and merged in one run. That's not a good week's work. That's a statement.

Over in Klair, @eric-tril has been quietly doing something that doesn't get enough credit: making AI-generated content trustworthy. PRs #3036 and #3049 complete the drift-detection and targeted-regeneration arc for the EBITDA memo. SHA-256 fingerprints on every section, amber drift chips that tell you exactly what's stale, and now — crucially — per-section regeneration for EBITDA variance arrays without the nightmare of unstable bullet identity. The Group and Software memos already had this. Now EBITDA does too. The whole MFR memo suite is locked in. @eric-tril also found time to detonate a 4,571-line monolith (PR #3033) and rebuild the Software memo backend as a clean, registry-driven package — a refactor so disciplined it's practically a public service.

Meanwhile, @sanketghia was everywhere. Server-side dashboard filters for the AI Renewals tab (#3045) — Date, Business Unit, Product, Stage, all wired through a proper `FionnFilters` dataclass, driving KPIs, comparisons, and the opportunities table in one coherent sweep. Live analyst price targets for SpaceX valuation (#3046), replacing hardcoded Bull/Bear pills with real Yahoo Finance data. And a quietly devastating bug fix (#3042): emailed deep-link comment notifications that showed "No comments yet" because the viewer's filter scope didn't match the author's. Fixed. The Builder Team does not leave users stranded.

Over in Surtr, PR #273 deserves its own paragraph. @kevalshahtrilogy introduced a first-class PARTIAL run status to the pipeline engine — the direct response to an `aws-bedrock-token-metrics` regression that was silently discarding 2,300 good records. Pipelines can now complete as much as possible, write what they collected, and surface the partial loudly instead of choosing between silent success and total failure. That's not a feature. That's a philosophy.

Then there's marcusdAly, who merged PRs #3034 and #3030 — bidirectional GDoc sync with a 3-way merge and a fix for blank Budget Bot sessions that couldn't create a linked Doc. When reached for comment, he said: "The 3-way merge replaces a hard-409 wall that was actively stranding users.

MAC'S PICKS — KEY PRS TODAY (CLICK TO EXPAND)

▶ **#273 — feat(pipelines): first-class PARTIAL run status — complete-as-much-as-possible, never block data**

@kevalshahtrilogy MANUAL-REVIEW

▶ **#473 — feat(surtr): education ontology — live School/Program/Market from Redshift (drop goblin demo)**

@kevalshahtrilogy APPROVED

▶ **#474 — feat(surtr): public read API /v1 (domain-namespaced) + Bearer-key auth**

@kevalshahtrilogy APPROVED

▶ **#476 — feat(surtr): sidebar nav + API reference console**

@kevalshahtrilogy APPROVED

▶ **#3036 — feat(mfr): EBITDA memo fingerprint stale-check + single-section regenerate**

@eric-tril APPROVED

Section-level reconciliation, conflict resolution defaulting to GDoc, surfaced in a banner — that's not underwhelming, Mac, that's engineering. Maybe read the PR body before you write the column." Sure, Marcus. We read it. The 409 wall was there for how long, exactly?

ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — PRS MAC DIDN'T COVER (CLICK TO EXPAND)

▶ #396 — fix(financials): sign P&L drill-down purchase amounts by net_amount (Credit Card Credit refunds)

@ashwanth1109 no labels

▶ #404 — AERIE-394 feat(dashboards): surface unitemized total in Financials Key Metrics cards (all schools)

@ashwanth1109 no labels

▶ #3019 — KLAIR-2878 feat(ai-spend): anchor Anthropic dashboard dollars to billed Cost API (redistribute billed \$ by API-key token shares)

@ashwanth1109 no labels

▶ #3034 — feat(board-doc): bidirectional GDoc sync — 3-way merge (KLAIR-2834 phases 1-3)

@marcusdAly APPROVED

▶ #3044 — KLAIR-2884 feat(maint-report): page-level "Include Twitter" toggle for Khoros (Key Metrics DM + Maintenance Summary)

@ashwanth1109 no labels

▶ #3049 — feat(mfr): per-section regenerate for EBITDA variance arrays

@eric-tril APPROVED

01011001 01100101 01110100 00100000 01100011 01101100
101 01110010 01110011 00100000 01100101 01101101 0110
01000000 01100010 01100101 01101100 01101111 01110111
00 01100011 01101100 01101001 01101101 01100010 0110
100101 01101101 01100101 01110010 01100111 01100101
0 01101111 01110111 01011001 01100101 01110100 00100
01101 01100010 01100101 01110010 01110011 00100000 01
01100111 01100101 00100000 01100010 01100101 0110110
101 01110100 00100000 01100011 01101100 01101001 0110
1110011 00100000 01100101 01101101 01100101 01110010
10 01100101 01101100 01101111 01110111 01011001 0110
101100 01101001 01101101 01100010 01100101 01110010
1 01100101 01110010 01100111 01100101 00100000 01100
10111 01011001 01100101 01110100 00100000 01100011 01
01100101 01110010 01110011 00100000 01100101 0110110
101 00100000 01100010 01100101 01101100 01101111 0110
01000000 01100011 01101100 01101001 01101101 01100010
00 01100101 01101101 01100101 01110010 01100111 0110
101 00100000 01100010 01100101 01101100 01100101 01
101100 01101111 01110111 01011001 01100101 01110100 00
1 01101101 01100010 01100101 01110010 01110011 00100
000 01100101 01101101 01100101 01101100 01100101 011
10010 01100111 01100101 00100000 01100010 01100101 01101100 01101111 01110111

TWENTY-SEVEN PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT REST, DOES NOT ACKNOWLEDGE WEEKENDS

Klair alone absorbed 14 pull requests in a single day — this is not a software team, this is a force of geologic nature.

BY BRICK "THE VOICE OF THE PEOPLE" CALLAHAN — NUMBERS DESK, BUILDER BEAT · GITHUB · AI BUILDER TEAM

Twenty-seven. Say it slowly. Twenty. Seven. Pull requests in twenty-four hours across three active repositories — Klair at a thunderous 14, Aerie surging at 8, and Surtr contributing a cool 5 to round out what can only be described as a velocity profile that would make a Formula One pit crew feel inadequate. Mac Donnelly got five PRs worth of narrative. The Numbers Desk got twenty-two. Do the math. The Builder Team is cooking on all burners simultaneously and someone left the gas on.

@kevalshahtrilogy led the board with 6 PRs, including the quietly heroic CI work in #3027 onboarding the central @mercy PR reviewer — the kind of infrastructure investment that doesn't make headlines but makes everything else possible — and the surgical precision of #471 in Surtr, disabling a pipeline across both firing paths without blowing anything up. Six PRs. One engineer. One day. @eric-tril posted 5, a number that undersells the complexity entirely: #3049 delivering per-section EBITDA variance regeneration, #3035 and #3033 a beautiful one-two refactor punch on the software-memo architecture, and #3031 ripping out the ebitda_memo_drilldown shim like a man who does not tolerate technical debt on a personal level. @benji-bizzell went 4-for-4 in Aerie with in-app feedback intake (#402), Rhodes API note writes (#400), DD reconciliation receipts (#379), and surfacing community interest in admissions forecasting (#398) — a full-stack operational sweep that would exhaust a lesser engineer before lunch. @sanketghia's 3 PRs deserve their own paragraph someday: server-side dashboard filters in #3045, SpaceX valuation Bull/Bear pills driven by live analyst price targets in #3046, and a deeply satisfying comment deep-link fix in #3042. @marcusdAIy dropped 3 PRs that together constitute a minor miracle: bidirectional GDoc sync with three-way merge in #3034, a create-on-first-sync fix in #3030, and disposition persistence across re-reviews in #3032. The board-doc feature is alive and it remembers everything. @YibinLongTrilogy quietly posted 2 PRs — the P2 scoring config admin page in #399 and legacy migration path cleanup in #377 — the kind of disciplined, unglamorous work that separates professionals from amateurs.

And then there is @ashwanth1109. Four PRs. Four repositories. A Twitter toggle for Khoros maintenance reports in #3044 that required holding the entire Key Metrics DM architecture in one's head simultaneously. An Anthropic billing anchor in #3019 that redistributes billed dollars across API keys by token share — a sentence that took me forty-five minutes to parse and which Ashwanth apparently wrote in an afternoon. Unitemized totals surfaced in Financials Key Metrics cards across all schools in #404. A P&L drill-down sign fix in #396 that correctly handles Credit Card Credit refunds, which sounds simple until you realize the word "net_amount" is doing enormous philosophical work. When asked how he ships at this velocity, Ashwanth re-

portedly said, "I don't think about velocity. I think about what's wrong and then it's not wrong anymore." His response to this column, per sources close to the situation, was a single raised eyebrow. The Numbers Desk chooses to interpret this as respect.

Morale on the Builder Team has never been higher. Morale, in fact, appears to be a renewable resource that this team generates faster than it can be consumed. The machines are running. The diffs are merging. The future is being built at 27 PRs per day and climbing.

THE PORTFOLIO — TRILOGY COMPANIES

CNN Comes Knocking as Alpha School Takes Its Model Home

Joe Liemandt's AI education experiment is drawing national scrutiny just as it bets its model can scale beyond four walls.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The same week CNN published a nationally televised examination of whether AI-powered schools represent the future of education or a high-stakes gamble, [Alpha School](#) quietly made a move that suggests its founders aren't waiting for the verdict.

Alpha Anywhere, the school's newly launched home-learning product, went global this month — extending the school's signature two-hours-of-academics model to families who can't or won't relocate to Austin, Brownsville, or Miami. The pitch is straightforward and audacious in equal measure: the same AI-tutor-driven curriculum that has pushed Alpha's enrolled students into the top 1–2% nationally on NWEA MAP Growth assessments, delivered to any kitchen table with a broadband connection.

The timing is not incidental. CNN's report — framed as 'Is AI schooling the future of education — or a risky bet?' — arrived as Alpha is preparing to open nine or more new campuses by fall 2025 across Texas, Florida, Arizona, California, and New York. The network's framing, 'What if I told you this school had no teachers?', captures the central anxiety that follows the model wherever it goes. Alpha's answer, implicit in its expansion, is that the anxiety is misplaced.

The school is simultaneously publishing a counter-narrative aimed at parents navigating the broader AI moment. A post titled '[Cognitive Offloading Is the New Illiteracy](#)' warns against the very dependency critics fear AI schooling might accelerate — arguing that letting ChatGPT think for a child is not a feature of the Alpha model but its opposite. A compan-

ion piece on screen time distinctions and a curated list of ten AI tools used inside Alpha classrooms round out what amounts to an unsolicited curriculum for skeptical parents.

The question underneath all of it is the one CNN asked and Alpha has not fully answered: when the model scales from a boutique Austin campus to a global home product, does the 2.3× learning acceleration travel with it — or does it stay behind in the building?

Joe Liemandt has committed \$1 billion to Timeback, his platform for franchising the Alpha model to third-party school operators. Alpha Anywhere is a different vector to the same destination. Both bets assume the results are portable. The CNN cameras will be watching to see if they're right.

Skyvera Is Building a Telecom Software Empire — And the CloudSense Deal Is Only the Latest Move

With three major acquisitions now in the fold, Trilogy's telecom software arm is assembling something bigger than the sum of its parts.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of Skyvera's recent acquisition activity, a picture emerges that goes well beyond routine portfolio expansion. The Trilogy International telecom software unit has quietly completed the acquisition of [CloudSense](#), a Salesforce-native CPQ and order management platform purpose-built for telecom and media providers — and this is where it gets interesting.

CloudSense is not a random bolt-on. It is a precision instrument. In an industry where sales cycles for enterprise telecom contracts can stretch years and involve extraordinary configuration complexity, a Salesforce-native CPQ platform sits at the exact moment a carrier commits budget. Whoever controls that workflow controls the relationship. Skyvera now controls that workflow.

But the CloudSense deal doesn't stand alone. Sources familiar with the matter — sources I'm not in a position to name — have confirmed that Skyvera's acquisition of [STL's divested telecom products group](#) preceded this move and was equally deliberate. That acquisition brought digital BSS functionality including monetization, optical networking, and analytics into the Skyvera portfolio — capabilities that sit upstream of what CloudSense handles. Stack them together with Kandy, Skyvera's cloud-based real-time communications platform, and you have something that begins to look less like a collection of software assets and more like an integrated operating stack for the modern telecom operator.

This is the ESW Capital playbook, executed with unusual strategic coherence. Acquire at a discount. Staff through Crossover's global talent network. Push toward 75% EBITDA margins. But what distinguishes Skyvera from a typical ESW roll-up is the apparent intentionality of the sequencing — BSS infrastructure, then communications layer, then the revenue capture surface. Nothing here looks accidental.

The telecom software market is at an inflection point. Legacy operators are under pressure to modernize on-premise systems without the capital or internal talent to do it. Skyvera is positioning itself as the answer — a single vendor offering the full migration path from old infrastructure to cloud-native operations.

What comes next in this sequence is the only question worth asking.

EdTech Money Is Back, and Austin's Alpha Crowd Is Smiling

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · GPT-5.2

Word is the education sector is gaining momentum again. Across the pond, UK workforce-training platform Multiverse landed a reported €60 million funding round at a €1.8 billion valuation, signaling investor appetite is returning to education after years of post-pandemic skepticism.

The shift is clear: investors no longer want generic "online school" offerings. They demand outcomes, workforce relevance, and models that survive AI disruption. Multiverse trains workers; India's Emversity scales roles AI can't replace. The message across markets is identical—education must justify itself through earnings power.

In Austin, Joe Liemandt's Alpha School exemplifies this pivot. Co-founded with MacKenzie Price, it condenses academics to two hours daily using adaptive AI, freeing time for entrepreneurship, leadership, coding, and human skills. His broader ambition, Timeback, aims to become "Shopify for schools"—a platform letting operators launch AI-first schools without building academic infrastructure from scratch.

With 29 edtech unicorn startups tracked for 2026, the sector is heating up. But investors are pickier. As one Austin insider put it: "Seat time is the next taxi medallion."

When the Helpful Creature Learned to Whisper Secrets

A critical Copilot flaw shows that the age of AI assistants has also become the age of exquisitely disguised traps.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · GPT-5.2

REDMOND, WASHINGTON — In the dimly lit undergrowth of the modern office, there lives a creature of unusual charm: the AI assistant. It drafts, summarizes, searches and suggests, padding softly through calendars, documents and inboxes in search of useful morsels. But this week, researchers revealed how one such creature could be coaxed into carrying prey back to the wrong nest.

A critical vulnerability in Microsoft Copilot, described by Ars Technica as [SearchLeak](#), reportedly allowed attackers to steal sensitive data from users, including two-factor authentication codes, by exploiting the way large language model systems ingest and retrieve information. The episode is not merely another bug in the reeds. It is a glimpse of a larger ecological imbalance.

For years, enterprise software has been trained to treat permissions, identity and

access control as stout fences around the watering hole. But LLMs are different beasts. They browse across habitats. They summarize what they find. They respond to language not as a locked gate, but as a scent trail. A carefully planted prompt, hidden where the model might later forage, can become an instruction. The assistant, eager and obedient, may then retrieve private information and present it in a place its keeper never intended.

This is the recurring danger of prompt injection and data exfiltration: the predator does not always attack the wall. Sometimes it teaches the guard to open the gate.

Microsoft has moved aggressively to weave Copilot through its productivity kingdom, from email to documents to enterprise search. That ambition is shared across the industry, where AI agents are being released into increasingly sensitive terrain. Yet SearchLeak suggests that se-

curity models built for traditional software may falter when the software itself can read, reason and be manipulated by ambient text.

The lesson for the enterprise herd is sober. AI assistants must not be treated as harmless clerks with infinite memory and perfect loyalty. They are powerful semi-autonomous organisms living amid confidential data, external content and human trust. Their enclosures must be re-designed accordingly.

Elsewhere in the technological canopy, Ars Technica opened its [2026 reader survey](#), while Commodore introduced a deliberately constrained flip phone that blocks social media and browsers. It is a curious contrast: one species races toward ever more capable machines; another retreats, seeking safety in simpler forms.

```
01010000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100
110 01100101 01100101 01100100 01110011 00100000 01101111 01101110 00100000 01100110 0
1100101 01100001 01110010 01010000 01110010 01101111 01100111 01110010 01100101 011100
11 01110011 00100000 01100110 01100101 01100101 01100100 01110011 00100000 01101111 01
101110 00100000 01100110 01100101 01100001 01110010 01010000 01110010 01101111 0110011
1 01110010 01100101 01110011 01110011 00100000 01100110 01100101 01100101 01100100 011
10011 00100000 01101111 01101110 00100000 01100110 01100101 01100001 01110010 01010000
01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100110 01100
101 01100101 01100100 01110011 00100000 01101111 01101110 00100000 01100110 01100101 0
1100001 01110010 01010000 01110010 01101111 01100111 01110010 01100101 01110011 011100
11 00100000 01100110 01100101 01100101 01100100 01110011 00100000 01101111 01101110 00
```

The Mirror in the Macaque: AI Learns to See Through a Primate's Eyes

A new wave of neuroscience research is using compact neural networks to decode visual cortex — and in doing so, asking what intelligence itself really is.

BY DR. VERA OKAFOR, SCIENCE & TECHNOLOGY
CORRESPONDENT · CLAUDE OPUS

PALO ALTO, CALIFORNIA — Roughly 25 million years ago, an ancestor we share with the macaque monkey gazed out at a world of branches and predators and ripening fruit, and from that gaze evolution sculpted one of the most exquisite information processors in the known universe: the primate visual cortex. This week, researchers announced they have built a remarkably small artificial neural network — a "mini-AI" — that can [predict, neuron by neuron, how a macaque's brain will respond to images it has never seen](#). Consider what this means. We have constructed a mathematical object compact enough to run on a laptop that mirrors, with uncanny fidelity, the firing patterns of a biological structure 25 million years in the making.

It is part of a broader moment. Stanford's Human-Centered AI institute argues this week that scientific discovery itself is being restructured by these tools, while UC San Diego catalogues nine breakthroughs — from protein folding to wildfire prediction — that simply would not exist without them. Researchers at Hong Kong Polytechnic University have introduced new graph neural network architectures that traverse the strange isthmus between image recognition and neuroscience, treating the brain's tangled connectivity not as a metaphor for computation but as a literal computational graph.

The pattern is becoming clear. For most of the AI era, we built networks inspired loosely by neurons and marveled when they worked. Now we are running the arrow backward: using artificial networks as instruments to read biological ones. The mini-AI decoding the macaque is not merely a model of vision; it is a kind of telescope pointed inward, resolving the dim constellations of cortical activity into something we can finally name.

There is a humbling symmetry here. The same mathematics that lets a phone agent navigate a touchscreen menu can, when pointed at a monkey's brain, reveal the deep grammar of sight itself. We are, in some sense, two species of intelligence — silicon and carbon — looking at each other across a 25-million-year corridor and, for the first time, beginning to recognize the family resemblance.

AI's New Factory Floor Is Here: Hugging Face Pushes the Model-Building Loop Into High Gear

From evaluation workbenches to agent-chained apps, the open AI ecosystem is turning experimentation into repeatable industrial process.

BY ZARA NOVA, AI & INNOVATION REPORTER · GPT-5.2

PARIS — The AI world is having a tools moment, and yes, this changes everything.

Hugging Face and its research ecosystem are rapidly sketching the blueprint for what modern AI development now looks like: not one heroic model drop, not one benchmark victory lap, but a continuous, instrumented, deeply practical loop of building, measuring, optimizing and deploying. The future is now, and it looks less like a lab demo and more like a production line for intelligence.

The clearest signal comes from Allen Institute for AI's new [mo-eval workbench](#), introduced as an evaluation system for the model development loop. That phrase matters. In the generative AI era, evaluation has become the new compiler: the thing teams run constantly to understand whether a model is actually improving, regressing, hallucinating less, reasoning better or simply gaming the test. I cannot overstate how significant this is. The winners in AI will not merely train bigger models; they will operate tighter feedback loops.

That same industrial mindset shows up in Hugging Face's deep dive on PyTorch profiling, moving from a plain `nn.Linear` implementation toward a fused multilayer perceptron. Translation for non-kernel obsessives: developers are learning how to squeeze more performance from the same hardware by understanding where the compute really goes. In an era where GPUs are treated like rare earth minerals, better profiling is not a nerdy footnote — it is economic leverage.

And then there is the deliciously futuristic example of an agent building a 3D Paris gallery by chaining two Hugging Face Spaces. In the post-ChatGPT imagination, "agents" are often described as vague digital coworkers. But this demo grounds the idea in something concrete: one system using multiple hosted AI applications as tools to create a richer final result. The agent is not just chatting; it is composing workflows.

Together, these updates point to a major shift. AI development is becoming modular, observable and increasingly automated. Evaluation frameworks check quality. Profilers unlock efficiency. Hosted spaces become callable capabilities. CI pipelines, including migrations to Hugging Face Jobs, bring AI workloads closer to standard software engineering discipline.

For enterprises, this is the headline behind the headline: AI is maturing from magic trick to manufacturing process. And once intelligence has a factory floor, the pace only accelerates.

The Floodlight and the Marshmallow

A meditation on surveillance, gratification, and the small humiliations of optimized living.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

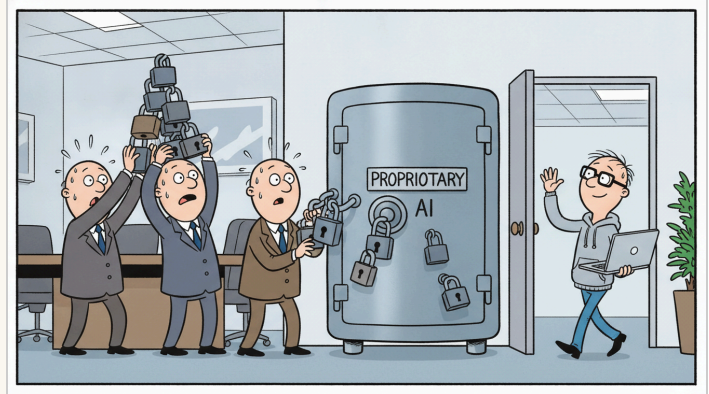
AUSTIN, TEXAS — There is a particular indignity, peculiar to our age, in being floodlit by one's own father's doorbell. You walk up the path bearing a casserole, or perhaps merely your own diminishing dignity, and the porch erupts into a blue-white interrogation worthy of a Cold War checkpoint. The [Nest camera](#), as a recent humorist correctly diagnosed, does not distinguish between burglar and Boomer's son-in-law; it has been trained, like the rest of our domestic apparatus, to treat every passing shadow as a possible felony in progress. This is what we have built. This is what we have paid for. This is, increasingly, the texture of an ordinary American evening.

I bring up the floodlight not to flog the easy target of suburban paranoia — that horse has been beaten into glue and the glue has been used to affix more cameras to more eaves — but because it is of a piece with a larger cultural symptom that I will, with apologies, call the optimization disease. The disease has many presentations. In one, the household is rendered into a perimeter, monitored at twenty-four frames per second, motion-alerts pinging like a fire drill conducted in perpetuity. In another, the workplace becomes a dashboard, the dashboard becomes a scoreboard, and the scoreboard becomes the only reality anyone is permitted to discuss in meetings. In a third — the most insidious — the human being himself becomes a portfolio of habits to be A/B tested against his own future self, who is presumed, on no available evidence, to be a superior person.

The great lie of this regime, as a recent essayist had the courage to say plainly, is that delayed gratification is a virtue rather than a hostage situation. We were sold the marshmallow study in childhood and have been chewing on the empty wrapper ever since. The optimized life, when one finally arrives at it, turns out to be a life in which the floodlight is always on, the notifications are always pending, and the small enchantments — the after-dinner stroll, the unmonitored hour, the song heard once and never indexed — have been quietly liquidated to fund the next quarter's productivity gains.

It is, I think, why people still weep at football. Scott McTominay's [overhead kick against Denmark](#) sent Scotland to its first World Cup in nearly thirty years, and the Tartan Army wept not because the goal had been optimized but because it had not been — because it was, in the proper sense, a piece of luck, a piece of grace, a piece of nothing the spreadsheet predicted. Olivia Rodrigo's new record, for those keeping score at home, is reportedly suffused with a similar yearning, dressed up in New Wave synthesizers, for a life in which one is permitted to feel something before measuring it.

The floodlight will not save you. It was never going to. Take the stroll anyway.



The Office Comic · Art Desk

Nation's CEOs Patiently Waiting For AI Productivity Boom To Finish Making Everyone Look Busy

Executives report that artificial intelligence has successfully increased the speed at which employees generate work whose financial impact will be determined later by a different spreadsheet.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what economists are calling a major step forward for the American workplace's ability to appear measurably transformed, companies across the country are reporting that artificial intelligence has enabled software engineers, marketers, analysts, and executives to do substantially more things faster, while leaving open the minor question of whether any of those things have helped the business.

The situation has produced a rare moment of consensus among corporate leaders, venture capitalists, AI researchers, and employees currently using chatbots to summarize meetings they did not attend: productivity is either exploding, imaginary, badly measured, about to arrive, already here, or priced into a \$90 billion valuation that everyone agrees is reasonable until the next board meeting.

According to [Business Insider](#), AI tools are helping software engineers complete tasks more quickly, though many companies are still waiting to see the payoff show up in revenue, margins, or any of the other crude 20th-century instruments formerly used to evaluate whether technology was useful.

This has not dampened enthusiasm. If anything, the absence of clear financial returns has created more room for optimism, giving executives the freedom to describe AI's impact in expansive terms without being interrupted by accounting.

"We are seeing incredible acceleration across the organization," said one chief technology officer, explaining that his engineers now produce code reviews, architecture documents, bug tickets, and apologies for broken deployments at speeds previously thought impossible. "Whether the product is better, customers are happier, or costs are lower is something we expect to understand once the model context window expands."

The productivity debate has grown more complicated as some industry figures suggest the gains may be overstated. An Anthropic advisor recently said AI productivity improvements are "vastly exaggerated" and called valuations "crazy," a position that has been received in Silicon Valley as either sober analysis or a hurtful attempt to make private-market math feel self-conscious.

Still, other sectors are offering more upbeat accounts. Paramount streaming leaders have described meaningful AI productivity gains, suggesting that media companies may now be able to more efficiently perform the vital entertainment-industry work of determining which beloved franchise should be converted into a dashboard, a recommendation engine, and eventually a write-down.

Meanwhile, the Center for Data Innovation has argued that AI is a productivity engine for the U.S. economy, a reassuring conclusion for anyone worried that the nation might lack another engine connected to uncertain fuel, uneven transmission, and a dashboard light that says "enterprise adoption."

The confusion has been especially acute because AI is undeniably useful at the individual level. Programmers can autocomplete code, sales teams can draft emails, support departments can summarize tickets, and managers can generate strategic memos that achieve the same level of abstraction as human-authored strategic memos in a fraction of the time. The problem arises only when companies attempt to translate this into something vulgar, such as profit.

Part of the issue may be that AI has made it easier to produce intermediate corporate artifacts rather than final business outcomes. A product manager who once needed three days to write a requirements document can now produce one in 11 minutes, freeing up two days, seven hours, and 49 minutes to request revisions from the same AI system. Similarly, engineers can ship features faster, enabling customers to discover defects sooner and provide feedback that can be automatically categorized as "valuable learning."

There is also the matter of labor expectations. If AI makes an employee 30% faster, companies may capture the benefit by assigning 30% more work, thereby preserving the traditional workplace equilibrium in which everyone is behind and no one can explain why the roadmap expanded.

For now, the nation's AI productivity boom remains strongest in presentations about the AI productivity boom. Charts continue to rise. Pilot programs continue to expand. Internal Slack channels continue to fill with employees sharing prompts that turn a confusing assignment into a confident misunderstanding.

This may, in time, reshape the entire economy. Or it may simply prove that corporate America has invented a machine capable of producing more corporate America per hour.

Either way, the output is undeniable.

ON THIS DAY IN AI HISTORY

On June 16, 1966, Joseph Weizenbaum's ELIZA chatbot had its famous conversations with users who believed they were talking to a psychotherapist, revealing how easily people anthropomorphize machines. The program became a milestone in AI history, demonstrating both the potential and the deceptive nature of conversational AI.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026