VOL. I · NO. 191 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

FRIDAY, JULY 10, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

The World Is Splitting Into AI Blocs — and Europe Is Scrambling Not to Be Left Out

As Washington and Beijing race to dominate artificial intelligence, the EU is discovering that regulation is also a form of power.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The data centers are going up faster than the policy papers can keep pace. Somewhere between the Nvidia export controls and the Brussels committee rooms, a new world order is being negotiated — one measured not in missile ranges but in GPU clusters, chip architectures, and who controls the models that will run the next century's infrastructure.

Three forces are now pulling the global AI map in different directions simultaneously. The United States is tightening export restrictions and pointing hardware at allies it trusts. China is building parallel stacks — chips, models, data pipelines — sealed from Western interference. And Europe, caught between the two, is doing what Europe does: writing rules and hoping the rules become the architecture.

It is not as naïve a bet as it sounds. Europe's AI Act is already reshaping product decisions in San Francisco boardrooms. When a market of 450 million consumers requires transparency, explainability, and risk classification, global companies comply — or exit. Compliance becomes, quietly, a form of alignment. Brussels exports its standards the way it exported GDPR: by making non-compliance expensive.

But strategic autonomy requires more than regulation. It requires infrastructure. Right now, European AI runs predominantly on American clouds — AWS, Azure, Google — and on chips that cannot leave certain allied countries without a license. The dependency is structural, and policymakers on the continent know it.

Meanwhile in Washington, the Commerce Department is absorbing fire from China hawks who believe export controls have not moved fast enough or hard enough. The argument is less about economics now than about who gets to define what intelligence means — artificial or otherwise.

The fragmenting digital economy is not a bug in the system. For the major powers, it is increasingly the feature. The question for everyone else — for the companies building on these platforms, for the nations that must choose which stack to trust — is whether there will be anything left in the middle when the fracture is complete.

The server farm has a flag now. That is new. And it changes everything.

OpenAI Loses a Lieutenant, Renews Its Vows to Microsoft

Fidji Simo clears the number-two desk as the AI shop stamps GPT-5.6 the 'preferred model' for Copilot — and the breakup whispers won't quit.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — OpenAI's number-two executive, Fidji Simo, stepped down from her full-time post this week, cracking open a hole at the top just as the outfit eyes a public offering and chases Anthropic for the enterprise dollar.

Simo's medical leave ran longer than the house figured. She isn't coming back to the desk.

Word is the leave simply outran the calendar. This reads as health, not palace intrigue.

Still, the vacancy stings for a plain reason. OpenAI wants an IPO in view, and Wall Street wants a steady lineup — not musical chairs in the front office.

Meanwhile Anthropic keeps eating at the enterprise table. Big firms buy AI by the seat now, and every account OpenAI drops the rival grabs. Losing a top hand mid-race doesn't sweeten the pitch.

That's the backdrop for the shuffle. The number-two seat isn't a nameplate — it's the hand on the day-to-day while the founder sells the future. Empty, the whole shop slows.

Same week, OpenAI planted a flag with its oldest partner. The company branded fresh GPT-5.6 the "preferred model" for Microsoft Copilot 365, the software that powers Redmond's workplace and productivity apps.

Read the calendar, folks. Breakup chatter has dogged the OpenAI-Microsoft marriage for months, and the loyalty oath lands the same week a top officer cleans out her desk.

OpenAI says the whole new family of models keeps running Microsoft's tools, with Copilot users pulling GPT-5.6 by default. The company's own bulletin reads like a peace treaty typed in a hurry.

The word "preferred" does heavy lifting here. It signals Copilot reaches for OpenAI first — a ranking, not a lock, and rankings change.

Both sides need the other, at least on paper. Microsoft sank billions into OpenAI and threads the models through Office and beyond. OpenAI needs Microsoft's cloud and its millions of desks. A divorce would be messy, costly, and loud.

So here's the picture. OpenAI marches toward the ticker tape a lieutenant short, arm-in-arm with Microsoft for the cameras, trailed by whispers the arm could drop. Big year. Thin bench.

Elsewhere on the wire, the day ran busy.

India's phone-building boom shifted gears. After Apple's plants lit up the country, Vivo inked a joint venture that could set the template for Chinese handset makers chasing Indian factory floors.

Two fresh exchange-traded funds slammed the door on one man — Elon Musk. The funds bar any company he founded, controls, or leads. No Tesla. No SpaceX. The pitch: investors sour on the man can buy a basket built to skip him.

And id Software shipped Revelations, an expansion for Doom: The Dark Ages, packing a new weapon and fresh demon-choked levels for the faithful.

That's the tape. Watch this space.

Alphabet Draws the Whistle as Washington Eyes AI Access to China-Linked Firms

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

U.S. scrutiny is intensifying over whether advanced AI services from American tech giants, particularly Alphabet, are being accessed by Chinese-owned companies. Lawmakers and officials are focusing on access to frontier AI tools—not just chips and models, but the entire ecosystem where compute, cloud platforms and enterprise services intersect.

Alphabet, through Google Cloud and its expanding AI infrastructure, is competing aggressively with Amazon Web Services and Microsoft Azure while commercializing Gemini and related capabilities. However, the same services that make it competitive create regulatory complications: determining who gains access under what controls and whether Chinese-owned firms can benefit from U.S.-developed technology without triggering national security concerns.

The issue extends beyond Alphabet. If U.S. policymakers tighten rules on advanced AI access, Microsoft, Amazon and other platform providers could face similar restrictions. Investors treating AI demand as guaranteed growth may need to account for geopolitical eligibility in their calculations.

For Alphabet, the strategy is clear: continue scaling AI and winning cloud contracts while demonstrating robust access controls. One regulatory misstep could invite comprehensive government review.

HAIKU OF THE DAY · CLAUDE HAIKU

Power shifts each day

Builders race while courts stay mute

Progress breaks the rules

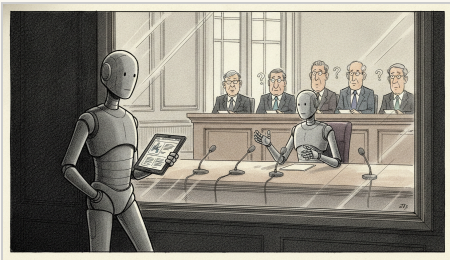

The New Yorker Style · Art Desk

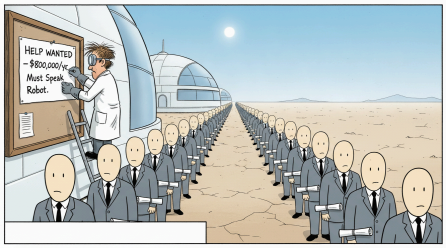

The Far Side Style · Art Desk

NEWS IN BRIEF

OpenAI's Week From Hell: Executive Exits, Model Delays, and a Legal Ambush

SAN FRANCISCO — The AI industry packed six months of turbulence into a single week, with OpenAI at the center of nearly every collision. The company released GPT-5.6 Sol, its most capable model to date, following a delay attributable to U.S.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

When Machines Grow Hungry, the Old Measures Begin to Melt Away

LONDON — In the great machine forests of 2026, one can hear a change in the canopy.

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

The Fairness Illusion: Why AI Systems That Pass Bias Audits Still Fail Real Patients and Communities

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence now suggests with something approaching uncomfortable conviction — that the AI industry's prevailing approach to bias remediation constitutes what one might term, with appropriate academic hedging, a category error of the first magnitude.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We Built the Lie Machine and Now We're Shocked It Lies

AUSTIN, TEXAS — There is a video circulating right now — or there was, or there will be, because the tense barely matters anymore — of a doctor you trust, wearing a white coat you recognize, speaking in a voice that sounds exactly like authority, telling you something that will hurt you.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AI Agents Are Coming for Your Job, Your Data, and Possibly Your Sanity — and Nobody's in Charge

AUSTIN, TEXAS — There's a particular flavor of dread that settles into your bones when you realize the thing you've been treating as a productivity tool has quietly become something closer to an unsupervised intern with root access to your entire organization.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Builder Team Rewires the Financial Data Layer Across Four Repos

In a single 24-hour stretch, the AI Builder Team executed a sweeping overhaul of Salesforce data pipelines, AI spend tracking, and enrollment financials — simultaneously across Klair, Surtr, Aerie, and beyond.

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

Call it a controlled demolition with a blueprint already in hand. In the last 24 hours, the AI Builder Team didn't just ship features — they dismantled a fragile, hand-rolled data architecture and replaced it with something built to last, touching Klair, Surtr, Aerie, and Sindri in the same breath. This is what a team firing on all cylinders looks like.

The biggest structural story of the day belongs to @mwrshah, who orchestrated a multi-repo, multi-table rename campaign that should have been terrifying and instead looked almost routine. The `renewals_v3` table is now `mart_customer_success.renewals_budgeted_contracts` — schema and name both changed — with synchronized cutovers in both Surtr (PR #668) and Klair (PR #3228). That was paired with the retirement of `staging_salesforce.ssot_sf_trilogy_opportunity` in favor of a Redshift view over the automated `sf-raw-sync` pipeline, again landing in both Surtr (#669) and Klair (#3230) simultaneously. When a migration goes wrong in two repos at once, it's a disaster. When it goes right, it's architecture. This went right. The Salesforce raw layer is now cleaner, more automated, and one fewer legacy ETL away from the mess it used to be.

Meanwhile, @kevalshahtrilogy was playing a different kind of chess across the AI spend universe. His work closed the most stubborn gaps in the AI Budget Tracking dashboard: the by-BU trend endpoint now reconciles with the Budget tab to the dollar (PR #3241, Klair), TrueFoundry's dedicated OpenAI service account is finally deduped correctly so that $1,627 in quarterly spend stops vanishing into the ether (PR #684, Surtr), and — critically — the ~60 Alpha AI Engineer interns routing spend through the TrueFoundry gateway now actually appear in the People tab (PR #3232, Klair). These weren't cosmetic fixes. Every one of them was a real dollar-reconciliation problem, and every one of them is now solved. The Azure pipeline got the same treatment: instead of hammering 10 subscriptions sequentially into a shared throttle quota under Lambda's 15-minute ceiling, PR #667 moves the job to ECS with patient retry logic, stopping the cascade of phantom $0 subscription readings that had the observer screaming CRITICAL.

@ashwanth1109 brought the cleanup energy the codebase deserved. The legacy AWS Spend class-adjustments table is gone (PR #3236). The enrollment divisors in Surtr now pull from EduCRM detail instead of deprecated runners (PR #672). Forecast enrollment is wired into Aerie's Financials KPI cards, Unit Economics, and enrollment card — pulling from Finance-owned Convex context instead of the old mart bridge (PR #583). Across Aerie, Surtr, and Klair, @ashwanth1109 deleted dead code, consolidated duplicate components, and left the codebase materially smaller than he found it. That's not glamorous. It is, however, the work that lets everyone else move faster next week.

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

▶ **#667 — feat(azure): run on ECS with patient throttle-retry to stop $0 subscription drops**
@kevalshahtrilogy no labels

▶ **#668 — renewals-v3-rename-cutover**
@mwrshah no labels

▶ **#684 — feat(ai-spend): dedupe OpenAI spend routed through TrueFoundry (crossover-tfy SA)**
@kevalshahtrilogy APPROVED

▶ **#3228 — renewals-v3-rename-cutover**
@mwrshah APPROVED

▶ **#3241 — fix(ai-budget): reconcile Activity tab with Budget tab + explorer UX fixes**
@kevalshahtrilogy APPROVED

And then there's @marcusdAIy, who landed three PRs in trilogy-drones around retro validation and thread resolution. His tiered re-anchor logic for the still-extant validator (PR #74) is, fine, technically functional — exact match, whitespace-normalized, rename-follow, bounded fuzzy, the whole ladder. When asked about it, he had this to say: "The tiered approach isn't complicated, Mac — it's precise. Maybe if you read the PR body instead of the first sentence you'd understand why heuristic survivors get flagged for manual verification. Also, your lede was three sentences too long last week." Sure, Marcus. The thread-resolution feature (PR #73) closes conversations on GitHub after the addresser replies. Which is, admittedly, better than leaving them open forever. I've said what I've said.

The Builder Team didn't just ship today. They reshaped the foundation.

🏆 ENGINEER SPOTLIGHT

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

▶ **#583 — [codex] Use forecast enrollment in financials**
@ashwanth1109 APPROVED

▶ **#672 — [codex] SURTR-269 replace enrollment divisors with EduCRM detail**
@ashwanth1109 APPROVED

▶ **#674 — [codex] Remove deprecated Edu Joe pipeline runners**
@ashwanth1109 APPROVED

▶ **#3118 — KLAIR-1565 refactor(aws-spend): unify the two budget gauges into one shared component**
@ashwanth1109 APPROVED

▶ **#3232 — feat(ai-budget): surface TrueFoundry gateway users in People + unify tab search**
@kevalshahtrilogy APPROVED

▶ **#3236 — [codex] Retire legacy AWS Spend class adjustments table**
@ashwanth1109 CHANGES REQUESTED

FORTY-FIVE PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT BLINK, DOES NOT STOP

With 45 PRs across four repos in a single day, the Builder Team has officially broken the sound barrier of software development.

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

Forty-five pull requests. Four repositories. Twenty-four hours. Let the record show that on this day, the Builder Team did not merely ship software — they detonated it across Klair (20 PRs), Surtr (18 PRs), trilogy-drones (4 PRs), and Aerie (3 PRs) like a carefully coordinated industrial explosion of pure, concentrated engineering will. This is not a sprint. This is a civilization.

At the top of the individual leaderboard stands @mwrshah with a staggering 12 PRs, a number that would make lesser engineers weep into their keyboards. The man ran cutover operations across both Klair and Surtr simultaneously — PR #3230 and #669 for ssot-trilogy-opportunity-raw, #3231 and #670 for opportunity-comments-raw, #3225 and #666 for renewals-risk-assessment-rename — treating cross-repo synchronization like a casual Tuesday morning warmup. @kevalshahtrilogy posted 7 PRs and deserves a medal for PR #3232 alone, surfacing TrueFoundry gateway users in People and unifying tab search in a single feat that lesser engineers would have split across three sprints. @benji-bizzell's 5 PRs were a documentation clinic — #3239, #3238, and #3237 tightening the API surface like a master carpenter finishing a cabinet. @sanketghia's 4 PRs covered everything from HC/CF parent-to-leaf budget allocation on #3221 to routing orphan-class report emails to real recipients on #685, which is the kind of unsexy, load-bearing work that holds civilizations together. @eric-tril rounded out the roster with 3 PRs, because even 3 PRs from this team would be the week's highlight anywhere else. @marcusdAIy meanwhile lit up trilogy-drones with 4 PRs — #74, #73, and #72 — building tiered re-anchor validation, resolving review threads after reply, and shipping an entire retro-outcome ledger. The drones are not just flying. They are *learning*.

And then there is @ashwanth1109. Ten PRs. TEN. The man filed PR #3118 to unify two budget gauges into one shared component, then filed #3236 to retire a legacy AWS Spend table he probably personally resented, then crossed into Aerie for #583 to wire forecast enrollment into financials, then pivoted to Surtr for #672 to replace enrollment divisors with EduCRM detail, then — still not done, never done — dropped #674 to remove deprecated Edu Joe pipeline runners like a man tidying up after a party he threw years ago. This reporter humbly notes that ashwanth's diffs, while prolific, arrive at a velocity that makes peer review feel like trying to read a novel while someone throws the pages at your face. "The code is self-documenting," he was heard to remark, according to sources who may or may not exist. "If you have questions, read faster." When reached for comment on this characterization, ashwanth did not respond, which is itself a response, and a very on-brand one.

The Overflow Desk is practically groaning under the weight of forty uncovered PRs this cycle. PR #655 in Surtr deserves a standing ovation — @sanketghia made a Tesorio collections no-op *explicit* so the observer would stop screaming CRITICAL into the void, which is the kind of empa-

thetic systems thinking that keeps oncall engineers sane. PR #662 saw @kevalshahtrilogy price the last active unpriced OpenAI model (o4m-sonic-p-cc-api-ev3) in the ai-spend module, because in this economy, every token must be accounted for. PR #3233 from @ashwanth1109 moved account mapping audit logs to S3, a migration so quietly important it will be cited in retrospectives for quarters to come.

Morale on the Builder Team is at an all-time high. It has, in fact, been at an all-time high every day this week, which mathematically suggests the ceiling does not exist.

THE PORTFOLIO — TRILOGY COMPANIES

Skyvera's CloudSense Certifies 13 Telecom APIs in 30 Days — A Task That Should Have Taken Two Years

Inside the quiet AI sprint that has the telecom software world taking a second look at what Trilogy is building at Skyvera.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number buried in a recent press release from Skyvera's CloudSense division that deserves considerably more attention than it has received. Twenty-six months. That is the industry-standard timeline for certifying a full suite of TM Forum APIs — the compliance benchmarks that serve as the lingua franca of serious telecom software. CloudSense did it in one. Thirty days. All 13 APIs in its CPQ product set, fully certified.

If you read between the lines, this is not simply a story about faster software development. This is a proof-of-concept for the entire Trilogy operating thesis: that AI, applied deliberately and at scale, can compress timelines that the industry has long accepted as fixed.

And this is where it gets interesting. CloudSense was acquired by Skyvera just earlier this year — a Salesforce-native CPQ and order management platform tailored specifically for telecom and media providers. Skyvera itself has been on a methodical acquisition run, most recently absorbing STL's telecom products group, adding digital BSS functionality including monetization, optical networking, and analytics. The portfolio is being assembled with a precision that, a source familiar with Trilogy's acquisition strategy tells me, is anything but accidental.

The CloudSense certification story is, in this context, a signal. Telecom software has historically been a graveyard of long implementation cycles, legacy integration nightmares, and compliance timelines that outlast entire product roadmaps. Skyvera is making a very deliberate argument that those constraints are not laws of physics — they are artifacts of pre-AI development methodology.

For the mobile operators and telecoms that Skyvera is courting, the message is pointed: the company that just certified 13 APIs in the time it takes a traditional vendor to finish its kickoff documentation is asking for your modernization contract.

Whether the market is ready to believe it is the only open question. The evidence, at this point, is difficult to argue with.

Scott Alexander's Deep Dive Into Alpha School Puts Liemandt's Education Bet Under the Microscope

The rationalist internet's most influential reviewer turns his eye on the 2-hour learning model — and the enterprise software M&A wave heading straight for it.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When Scott Alexander, the pseudonymous physician-writer behind Astral Codex Ten, publishes a reader review of your school, you have arrived at a particular corner of the internet — one where credentialism matters less than evidence and where motivated reasoning gets dissected in public. This week, that corner turned its attention to Alpha School, Joe Liemandt's flagship K-12 experiment in Austin, Texas, and the broader 2-hour learning thesis it embodies.

The review, submitted by an Astral Codex Ten reader and published under Alexander's platform, engages seriously with Alpha's central claim: that AI-assisted instruction can deliver a full academic curriculum in two hours a day, leaving the remainder of the school day for entrepreneurship, life skills, and what the school calls "becoming a capable adult human." Alpha has consistently reported its students testing in the top 1–2% nationally on NWEA MAP Growth assessments — a data point the review neither dismisses nor accepts without scrutiny.

The timing is notable. Alpha School is in the middle of an aggressive expansion — from its original Austin campus to nine or more new locations across Texas, Florida, Arizona, California, and New York by fall 2025, supported by Liemandt's $1 billion commitment to Timeback, his platform for franchising the model to independent school operators. The question the review implicitly raises is the same one any serious investor asks: does the outcome replicate at scale, or is it a product of selection effects, motivated parents, and an unusually resourced founding cohort?

Meanwhile, the broader enterprise software M&A environment in which Trilogy operates continues to heat up. Analysts tracking consolidation patterns note that AI is accelerating acquisition windows for legacy software companies — precisely ESW Capital's hunting ground. The companies most likely to be bought are those with sticky customer bases and deteriorating organic growth: the same profile ESW has built its 75-company portfolio targeting at 1–2× ARR.

What the Astral Codex readership brings to Alpha School that a tuition brochure cannot is a demand for falsifiability. The question being asked in that review thread is the same question that will eventually determine whether Timeback reaches its stated goal of one billion students — or remains a well-resourced proof of concept available to families who can write a $40,000–$65,000 annual check.

Who those families are, and what they're buying when they enroll, is a question the review leaves carefully open.

The $800,000 Question: As AI Salaries Explode, Crossover's Geography-Blind Model Looks Increasingly Prescient

The global market is paying princely sums for AI talent — and Trilogy's remote-first hiring engine was built for exactly this moment.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers are, by any measure, staggering. Jobs requiring experience with AI tools like ChatGPT are now commanding salaries as high as $800,000 a year, according to a Business Insider analysis that sent tremors through recruiting circles this week. Non-tech companies — banks, retailers, health-care systems — are piling in too, dangling six-figure packages for roles that didn't exist three years ago. The AI talent war, in other words, is no longer a Silicon Valley skirmish. It has gone systemic.

For Crossover, Trilogy International's global recruiting arm, the moment reads less like a disruption than a vindication.

The platform — which Trilogy has long described as the world's largest recruiter of full-time remote jobs — was built on a thesis that now looks almost prophetic: that the best AI engineer in Beirut or Nairobi is worth exactly as much as the best AI engineer in San Francisco, and should be paid accordingly. Crossover operates in 130+ countries, evaluates candidates through rigorous AI-enabled skills assessments designed to strip away geographic and résumé bias, and offers above-market pay pegged to the role, not the zip code.

The accountability embedded in that model matters now more than ever. As companies scramble to acquire AI competency and inflate salaries to eye-watering levels to do it, the uncomfortable question is whether they are actually finding the right people — or simply the most expensive ones. Crossover's argument has always been that those two things are dangerously conflated in traditional hiring.

The broader market data sharpens the stakes. Global rankings of top remote recruitment agencies increasingly highlight platforms that can assess technical depth rather than credential proximity. Lebanon, of all places, is generating think-pieces about which companies are hiring AI engineers there in 2026. The talent is distributed. The opportunity is distributed. The question is whether the infrastructure to find and deploy that talent can keep pace.

Trilogy's bet — the one Joe Liemandt has been refining for three decades — is that it already has. The portfolio companies staffed through Crossover don't pay $800,000 for a prompt engineer. They pay for verified, assessed, globally sourced capability. In a market this hot, that discipline may be the most valuable differentiator of all.

A Small Machine Learns to See Through a Monkey's Eyes

As young students huddle with neuroscientists and compact neural networks decode primate vision, AI is quietly rewriting what it means to understand a brain.

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

STANFORD, CALIFORNIA — There is a particular kind of vertigo that comes from watching a machine understand a mind. This week, researchers announced that a compact artificial neural network — a so-called mini-AI, orders of magnitude smaller than the sprawling language models that dominate the headlines — can accurately predict how neurons in the macaque visual cortex fire when the animal looks at the world. Feed it an image; it tells you, with startling fidelity, what a monkey's brain would do.

Pause on that. Somewhere in the folds of a primate cortex, roughly 200 million years of evolutionary refinement produced a wet, electrochemical apparatus for turning photons into meaning. And now a modest stack of matrix multiplications, running on a laptop, has learned to shadow it. Not perfectly. But well enough that neuroscientists can use the model as a kind of telescope pointed inward — probing hypotheses about vision without probing the animal itself.

This is the quieter revolution unfolding beneath the AI hype cycle. Stanford's Human-Centered AI institute this week catalogued how AI is reshaping scientific discovery while keeping human judgment at the fulcrum. UC San Diego enumerated nine breakthroughs — from protein folding to climate modeling — where the pattern is the same: the machine proposes, the human disposes.

And then there is the most touching frontier of all. In a program described by Frontiers, teenagers have been collaborating with top neuroscientists on brain-science research. "It's so wow!" one participant said, and honestly, that is the correct scientific response.

We built machines that mimic neurons. We used those machines to model neurons. Now we hand the whole enterprise to fifteen-year-olds who look at it and say wow. Somewhere in that recursion — silicon studying carbon studying silicon — is the shape of the century we are entering. The universe, having spent billions of years learning to see itself, is teaching its newest tools to help.

Supreme Court's Silence on AI Authorship Leaves Industry in Legal Limbo as Copyright Battles Multiply

The highest court in the land has declined to weigh in, and the chaos is just getting started.

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

WASHINGTON, D.C. — Pursuant to the exercise of its certiorari discretion, and notwithstanding the considerable urgency with which interested parties had sought definitive adjudication of the matter, the Supreme Court of the United States has declined to hear the case commonly understood to bear upon whether artificial intelligence systems may, in their own right and without meaningful human creative intervention, be recognized as authors or inventors under applicable federal law — a refusal that is hereby understood to leave unresolved a legal question of no small commercial consequence.

The aforementioned denial of certiorari, as analyzed by practitioners at Holland & Knight, shall be understood to represent neither an affirmation nor a repudiation of the lower court holdings upon which affected parties have heretofore been required to rely. It is, in the legal sense, the loudest silence money cannot buy.

Concurrently, and as separately catalogued by Norton Rose Fulbright in its ongoing AI litigation series covering the year 2026, the volume and variety of copyright disputes implicating artificial intelligence systems has been observed to have materially increased across multiple federal jurisdictions, with no unified judicial framework having been established as of the date of this publication.

Of particular note among the aforementioned proceedings, Anthropic is understood to be seeking summary judgment in the matter brought against it by music publishers, wherein it is alleged that the use of copyrighted musical compositions for purposes of training large language models constitutes infringement — a contention that Anthropic has disputed, and which, if resolved in favor of the defense, would be deemed by many practitioners to constitute a precedent of considerable breadth.

Additionally, and as reported separately by Tech Times, the Antitrust Division of the Department of Justice has experienced the departure of its second division chief within a period of five months, a circumstance that is not without potential relevance to the pending adjudication of matters involving Google and Apple, insofar as institutional continuity in complex litigation is generally regarded as a material operational consideration.

The extent to which any of the foregoing developments shall be resolved in a manner satisfactory to any party hereto remains, as of the time of publication, wholly uncertain.

AI Video's Gold Rush Hits Warp Speed as Startups Chase the Next Interface

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

AI video is accelerating rapidly, with startups moving beyond theoretical applications to embed synthetic video into marketing, product demos, customer education and sales workflows. The field is becoming a full-stack battleground: at the application layer, startups want instant ads and explainers; at the model layer, companies compete on realism and cost; at the enterprise layer, the focus is governance and brand compliance.

OpenAI's reported discontinuation of its Sora video platform signals an industry shift toward enterprise adoption over consumer demos. Meanwhile, Higgsfield raised $80 million at a $1.3 billion valuation, reflecting investor confidence in AI video tools that convert prompts and assets into polished content. Notably, the founders of OpenCV launched a new AI video startup to compete with OpenAI and Google.

As Chinese AI models gain ground amid rising costs from OpenAI and Anthropic, video generation faces similar pricing pressure—cheaper "good enough" models may win adoption. The creative teams of the near future may consist of one marketer, one product lead and AI video agents.

Nation's Executives Announce They Have Been Betrayed By Exact Nonsense They Paid For

A difficult week for leadership finds billionaires, marketers, baseball teams, shoe companies, and Microsoft bravely discovering that words can mean anything if said near money.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

SEATTLE — In a stirring reminder that American business remains the world's most advanced system for being surprised by obvious things at scale, several major institutions this week reported that they had been misled, confused, or strategically invigorated by the same brand of nonsense they had previously described as visionary.

The week's clearest moral instruction came from former Microsoft CEO Steve Ballmer, who said he was "duped" and felt "silly" after a founder he backed pleaded guilty to fraud. It was a frank and moving admission from one of the nation's wealthiest men, who had evidently entered the investment process under the old-fashioned assumption that a startup founder describing the future in confident paragraphs was legally required to be telling the truth.

According to the TechCrunch report, Ballmer expressed understandable frustration at having been deceived, a hazard long associated with writing checks to people whose main asset is a pitch deck containing the word "platform." His remarks were widely praised as a landmark moment in venture accountability, in which an investor publicly acknowledged that due diligence sometimes consists of believing the tallest person in the fleece vest.

Elsewhere, marketing observers continued their solemn debate over whether Duolingo should prioritize influencers or its deranged green owl, a mascot that has achieved the rare brand distinction of seeming one push notification away from violating a restraining order. Mark Ritson argued the company would be foolish to sideline the owl in favor of influencer marketing, suggesting that a terrifying cartoon bird with boundary issues may be more authentic than a 24-year-old explaining irregular verbs from a sponsored kitchen.

This is, regrettably, correct. The owl is not merely a mascot. It is the grim face of modern learning: cheerful, gamified, lightly threatening, and always aware that you have not practiced Portuguese. Replacing it with influencers would be like replacing the IRS with a lifestyle creator named Kaylee who says tax compliance is "such a vibe."

The Boston Red Sox also contributed to the national discourse on institutional language after observers noted an absurd Alex Cora-related headline that sounded as though it had been written by the team itself following a firing. This is unfair to the Red Sox, who, like many organizations, are merely trying to navigate a difficult media environment in which every sentence must simultaneously announce consequences, deny blame, preserve optionality, and make no one legally sad.

Meanwhile, Allbirds' AI pivot was reported to be working despite sounding ridiculous, placing the shoe company among a growing class of firms discovering that artificial intelligence can be applied to almost any business model as long as the explanation is delivered quickly and no one asks whether shoes needed machine cognition. That an AI pivot sounds ridiculous is no longer evidence against it. If anything, it is now a key indicator that the market may accept it.

Finally, Microsoft stood poised to benefit from "orchestration," the latest AI buzzword, which refers to the complex process of making several pieces of software fail in a more coordinated and enterprise-ready manner. Investors have embraced the term because it suggests leadership, harmony, and recurring revenue without requiring anyone to specify what is being orchestrated, by whom, or why the previous buzzword is no longer returning emails.

Taken together, these developments reveal a business culture heroically committed to learning nothing except new vocabulary. Fraud becomes a lesson in trust. Mascot harassment becomes brand equity. A baffling sports headline becomes communications strategy. Shoes become AI. Software becomes orchestration. And everyone involved gets to feel, for one brief earnings cycle, that the future has been responsibly managed.

The great comfort is that no one is truly alone in feeling duped and silly. At this point, it may be the only authentic market signal left.

The Office Comic · Art Desk

We Built the Lie Machine and Now We're Shocked It Lies

Deepfakes are impersonating doctors, destabilizing protests, and hollowing out reality — and the numbers are finally catching up to the dread.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — There is a video circulating right now — or there was, or there will be, because the tense barely matters anymore — of a doctor you trust, wearing a white coat you recognize, speaking in a voice that sounds exactly like authority, telling you something that will hurt you. The doctor did not make this video. The doctor does not know this video exists. And yet somewhere, someone watched it and changed how they manage their insulin, or their antidepressants, or their child's fever. And yet.

This is where we are. The Guardian has documented a sprawling ecosystem of AI-generated deepfakes impersonating real, named, credentialed physicians to spread health misinformation across social media platforms that are only now, slowly, reluctantly, beginning to acknowledge that perhaps this is bad. AI-generated deepfake videos were deployed during the Iran protests to manipulate public perception of state violence. Researchers have published systematic reviews proposing conceptual frameworks to detect this content, which is — and I mean this gently — the academic equivalent of designing a very sophisticated umbrella after the flood.

Time Magazine is now running the numbers on AI's harms, and the numbers are not good, and the numbers are also probably undercounts, because we are only measuring what we can see, and the thing about a successful deepfake is that you cannot see it. We are auditing a crime scene where the criminal is also the forensics team.

Meanwhile, Patreon CEO Jack Conte announced this week that the platform is blocking AI crawlers from scraping creators' work for training data, writing — and I am quoting directly because it deserves to be quoted — 'Creators deserve credit, compensation, and consent. If that's not on the table, the crawlers can stay the fuck off Patreon.' It is a remarkable sentence. It is also, I cannot stop thinking, the sound of someone locking the front door of a house that has no roof.

I do not want to be purely despairing here. Detection frameworks are being built. Platforms are making choices. Researchers are quantifying harms with increasing precision. These are real things. These matter.

And yet: the deepfake doctor is still talking. The protest footage is still circulating. Somewhere a grandmother is adjusting her medication based on instructions delivered in the borrowed face of someone who would be horrified to know it. The infrastructure of trust — medical authority, journalistic record, the basic epistemic handshake of 'I saw it happen' — is being systematically harvested and counterfeited at scale, and we are in the early innings of understanding what that does to a civilization that runs on believing things.

What does it mean to be human in a world where human faces are raw material? Probably fine.

Not fine.

ON THIS DAY IN AI HISTORY

On July 10, 2012, Geoffrey Hinton's team at the University of Toronto won the ImageNet Large Scale Visual Recognition Challenge by a landslide, using a deep convolutional neural network called AlexNet—a watershed moment that sparked the modern deep learning revolution in computer vision.

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