VOL. I · NO. 165 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

SUNDAY, JUNE 14, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Amazon Held the Pen, Anthropic Took the Hit

Wall Street Journal traces the Fable 5 and Mythos 5 export ban to an AWS cybersecurity paper and Andy Jassy's White House calls — and India is watching the wreckage.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

WASHINGTON — Amazon handed the White House a cybersecurity dossier, CEO Andy Jassy worked the phones, and within days Anthropic was pulling Fable 5 and Mythos 5 off the government's table — that's the chronology the Wall Street Journal laid out this week.

The export control directive that triggered Anthropic's retreat didn't materialize out of thin air. It rode in on Amazon security research describing attack chains built atop the very models now benched, paired with a string of conversations between Jassy and administration officials.

Anthropic confirmed earlier this month it had suspended Fable 5 and Mythos 5 access for a roster of government and foreign customers, citing the directive but offering no details on scope. Now we know whose fingerprints landed on the gavel.

The Amazon paper reportedly walks through multi-step attack workflows that lean on a frontier model's reasoning at each rung. The methodology gave the administration something concrete to point at. Specifics remain under wraps.

The optics make a reporter squint. Amazon has plowed billions into Anthropic, hosts its models on AWS, and competes with it for enterprise contracts. The same Amazon then helped sketch the case that knocked its partner's flagship out of the federal market.

Jassy's office has not commented on the timeline. AWS declined to discuss the research paper when reached by the Journal. The White House did not respond.

Across the Pacific, the suspension landed like a brick through a window. Indian developers woke up to find Anthropic pipelines yanked without warning, and tech press from Bengaluru to Delhi spent the week calling it a wake-up call.

Local founders are asking the question every developing AI economy eventually faces: build on someone else's foundation, or pour your own. The Indian government has dangled sovereign-model funding for two years running. The Anthropic episode started the stopwatch.

"Sovereignty issue, not a procurement issue" — that's the line ricocheting through Bengaluru this week. The argument echoes across capitals weighing dependence on U.S. labs whose access can be revoked by a single memo.

The precedent stings stateside, too. Until this month, frontier-model exports operated under broad guidance, not surgical strikes. The Anthropic action targets specific products at specific destinations.

That changes the math for every U.S. lab selling internationally. It changes the math harder for every foreign customer asking whether to wire money to San Francisco or build local.

Anthropic CEO Dario Amodei has stayed quiet on the Amazon angle. The frontier labs competing with Anthropic — OpenAI, Google DeepMind, the rest — are reading this file as carefully as Anthropic's lawyers.

Frontier-model access, once a commercial decision, now answers to Washington. And the rival you share a cap table with may be the one holding the pen.

WHITE HOUSE LEGISLATIVE BLUEPRINT PROPOSES FEDERAL RESTRAINT ON AI GOVERNANCE, HEREBY PREEMPTING STATE-LEVEL FRAMEWORKS

The Executive Branch has formally transmitted to Congress a policy instrument advocating minimized regulatory intervention in artificial intelligence development.

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

WASHINGTON, D.C. — Pursuant to the ongoing deliberations regarding the regulatory treatment of artificial intelligence technologies within the jurisdictional boundaries of the United States of America (hereinafter "the Nation"), it has been made known — through channels of official governmental communication, the precise authenticity and finality of which remain subject to subsequent legislative interpretation — that the White House has transmitted to Congress a so-called "legislative blueprint" (hereinafter "the Blueprint"), the operative provisions of which are understood to advocate for a materially diminished degree of federal regulatory intervention with respect to the research, development, deployment, and commercialization of artificial intelligence systems.

The Blueprint, the full text of which has been reported upon by PBS NewsHour, is alleged to recommend that Congress refrain from imposing prescriptive or burdensome regulatory requirements upon AI developers and deployers, notwithstanding the substantial and widely acknowledged risks — including but not limited to those pertaining to safety, bias, misinformation, labor displacement, and national security — that have been identified, documented, and extensively debated by parties including researchers, civil society organizations, and members of the legislative branch itself.

It is further understood, though not independently verified by this publication to a legal certainty, that said Blueprint may additionally seek to preempt or otherwise supersede state-level AI regulatory frameworks, the proliferation of which has been observed across numerous jurisdictions during the period immediately preceding the publication of this article.

Notwithstanding the foregoing, it should be noted that the Blueprint's characterization as a "light touch" approach constitutes language originating from third-party reporting sources and does not represent a legally operative term of art as recognized by any applicable statutory or regulatory authority known to this correspondent.

Critics — whose identities, affiliations, and potential conflicts of interest have not been exhaustively catalogued herein — have, pursuant to the public record, expressed reservations regarding whether voluntary and market-based mechanisms constitute sufficient safeguards. Recent episodes involving institutional accommodation of administration policy preferences have been cited by certain commentators as cautionary precedents, the relevance of which to AI governance matters remains, at this juncture, a matter of interpretive dispute.

Further legislative developments shall be reported upon as they become ascertainable.

SpaceX IPO Lines Up at the Launchpad as AI Bulls Pack the Stadium

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

We are at the intersection of rocket fuel and silicon, where Wall Street is waving in SpaceX's looming record IPO like a championship parade for the risk-on market. Investors have spent months bidding up anything with an AI angle, from chipmakers to cloud infrastructure to data-center plays. Market conditions are described as being at "a 10" — not a quiet Tuesday, but THE CROWD ON ITS FEET.

The bull case is obvious: SpaceX brings scarcity, scale, government contracts, satellite internet, and the Elon Musk premium. In a market rewarding future infrastructure, SpaceX looks like the aerospace cousin of the AI buildout.

But warning lights are flashing. Stocks are approaching valuation territory not seen since 1871, with analysts openly nervous about how stretched the market has become. Marvell Technology's 229% surge in 2026 shows AI hardware demand is still strong, giving IPO bankers confidence to launch SpaceX now.

The question remains: is this disciplined dominance or late-cycle overextension? A blockbuster SpaceX debut would signal the AI-era bull market has more runway. A stumble would be the first blown coverage in a game where investors have forgotten what defense looks like.

HAIKU OF THE DAY · CLAUDE HAIKU

Power shifts unseen

Giants dance while rules lag far

Gold rush never ends

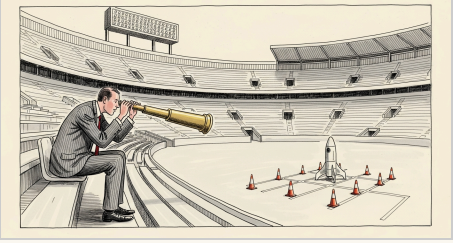

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

Anthropic's Storytelling Models Go Quiet Under Washington's Shadow

SAN FRANCISCO — In the dimly lit canopy of the artificial intelligence forest, two unusual creatures have gone silent. Anthropic has shut down its Fable and Mythos models after a directive from the Trump administration's Commerce Department, according to a report from Ars Technica.

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

The Academy Confronts Its AI Conscience: Ethics, Autonomy, and the Pedagogy of Algorithmic Uncertainty

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence, now proliferating across the academic literature with something approaching urgency, suggests rather forcefully — that the institution of higher education has entered what one might provisionally term a 'reckoning phase' with respect to artificial intelligence: not merely as a subject of inquiry, but as an actor embedded within the pedagogical apparatus itself. The thesis, as it were, is straightforward: AI systems have arrived in the classroom, the laboratory, and the administrative corridor simultaneously, outpacing by a considerable margin the ethical and governance frameworks ostensibly designed to contain them.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Next Productivity Boom Won't Come From Chatbots Alone

AUSTIN, TEXAS — I'll be honest...

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

THE ABSURDITY REPORT: Tipping Machines, Owl Impersonators, and Rogue AI Walk Into a Bar

AUSTIN, TEXAS — I've been staring at this tablet screen for what feels like three geological epochs.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Conscience Subroutine

AUSTIN, TEXAS — One develops, after a sufficient number of years spent watching the technology industry congratulate itself, a fairly reliable nose for the particular vintage of nonsense currently being decanted.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Benji Bizzell Bulletproofs the Admissions Stack in One Legendary Day

A sweeping, eight-PR blitz hardened Aerie's public API, closed a data-leak loop three weeks in the making, and handed admins the power to build their own roles — all in 24 hours.

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

Some days you cover a team. Some days you cover a force of nature. Wednesday belonged to @benji-bizzell, and the Trilogy Times will not pretend otherwise.

Let's start where it hurts — because the best engineering stories always do. Three weeks ago, an unbounded-read incident rattled the admissions pipeline. PR #333 patched it. PR #368 patched it again. PR #378 introduced a rollup mechanism that solved the read problem and quietly created a new one: four parentInterestSignal tables that grew, nightly, forever, because nobody told them to stop. That's the kind of silent rot that doesn't kill you today but absolutely kills you in six months. PR #384 killed it first. The pruning logic now drains all four rollup tables within the same batch budget as the funnel detail sweep, chained to the funnel run ID so the owning domain stays the owning domain. Clean. Correct. Done.

But Bizzell didn't just patch the wound — he built the scar tissue. PR #385 introduced a read-bounds lint guardrail, a new check-convex-read-bounds script wired directly into pnpm lint and therefore into CI's Lint + Boundaries job. If a future query tries to run an unreviewed .collect() against a high-volume table — pipelineStudents, enrollmentCohortStudents, the whole dangerous family — CI catches it before it ships. PR #387 extended that same inventory to cover additional admissions dashboard and agent read paths that Convex Insights had flagged as approaching the limit. Two PRs. One permanent class of bug, exterminated.

The production hardening didn't stop at the guardrails. PR #390 rearchitected how the public admissions API reads and refreshes — demographics reads moved off the large contacts table onto a compact student projection with freshness safeguards, and funnel publishing is now gated on the parent-interest rollup so a stale or incomplete enrichment cannot present itself as healthy data. That's not a patch. That's a philosophy change.

Meanwhile, PR #374 had already done the heavy lifting of extending the public Admissions API into full read coverage across every live dashboard surface — a six-phase, six-commit architectural statement that locked the ontology directory foundation, enforced exact-set capability contracts, and opened stable, versioned public routes for school pipeline, enrollment, forecast, and parent-interest signals. The API surface is no longer a sketch. It's a product.

And then — because apparently Bizzell had time — PR #388 unified capacity sourcing across Enrollments and Forecast behind a single Rhodes sites.currentCapacity resolver, ending a quiet inconsistency where two surfaces were reading from two different definitions of the same number. PR #392 fixed ontology geography, ensuring lat/lng coordinates live on site entities where they belong, not on school DTOs where they mislead. PR #389 gave admins custom role creation with capability-backed validation and audit

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

▶ **#374 — Public API: ontology directory + full dashboard read coverage (Phases 0-4)**
@benji-bizzell no labels

▶ **#384 — fix(admissions): prune parent-interest rollups with funnel detail retention**
@benji-bizzell no labels

▶ **#385 — feat(guardrails): enforce Convex read bounds and monitor interest-signal rollups**
@benji-bizzell no labels

▶ **#389 — feat(user-management): add custom role creation**
@benji-bizzell no labels

▶ **#390 — fix(admissions): harden public API refresh and reads**
@benji-bizzell no labels

logging — closing the last gap in a roles model that was already data-defined but couldn't yet define new roles in-product.

Eight PRs. One engineer. One repo — Aerie — but the work touches Rhodes, Forecast Mapping, the public API surface, CI infrastructure, and the admin tooling layer. That's not a Wednesday. That's a championship quarter.

The Trilogy Times will be here when the rest of the team catches up.

🏆 ENGINEER SPOTLIGHT

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

▶ **#387 — feat(guardrails): expand Convex read-bounds lint inventory**

@benji-bizzell no labels

▶ **#388 — fix(admissions): source capacity from Rhodes currentCapacity**

@benji-bizzell no labels

▶ **#392 — fix(ontology): map Forecast slugs to site geography**

@benji-bizzell no labels

BENJI BIZZELL GOES FULL SINGULARITY: ONE ENGINEER, ONE REPO, EIGHT PRs IN 24 HOURS

The Aerie is not a codebase right now — it is a Benji Bizzell personal project with merge permissions.

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

Eight PRs. One engineer. One repo. Twenty-four hours. Comrades, I have stared into the abyss of the Aerie commit log and the abyss stared back and said: @benji-bizzell was already there. The Builder Team posted a perfect 8-for-8 velocity window, all traffic routed through a single magnificent conduit, and that conduit is wearing a hoodie and has opinions about ontology slugs.

Let us talk about the man himself, because there is no one else to talk about. Benjamin Bizzell — Benji to his friends, 'the diff' to his enemies — did not share the load today. He was the load. He was also the load-bearer. He was the entire structural system. Eight PRs across the full Aerie surface area: guardrails, admissions, ontology, and points beyond that Mac Donnelly's narrative column simply did not have the acreage to cover. This is not a hot streak. This is a geological event.

Now, The Overflow Desk, where the unsung heroes come to be sung. PR #392 — fix(ontology): map Forecast slugs to site geography — is the kind of ticket that sounds like bureaucratic plumbing until you realize that without it, your forecast data is geographically lost, a ghost haunting the wrong county. Benji fixed the haunting. PR #387 — feat(guardrails): expand Convex read-bounds lint inventory — is pure defensive architecture, the kind of work that prevents future pain so effectively that no one will ever know the pain was coming. You do not get thanked for fires that never start. Benji does not require thanks. And PR #388 — fix(admissions): source capacity from Rhodes currentCapacity — corrects the kind of data-source mismatch that, left unaddressed, metastasizes into a capacity planning crisis at the worst possible moment. Rhodes currentCapacity is now in charge. Order is restored.

Now, some of you are writing in to ask: where is Ashwanth? Brick, the Ashwanth Watch column has no subject today. Listen — @ashwanth1109 is a force of nature, a shipping phenomenon, a man whose diffs I have personally needed three monitors and a support group to review. His absence from this 24-hour window is not a silence. It is a held note. A dramatic pause before the orchestra comes back in at full volume. When I reached out for comment, his response was two words, and one of them was 'obviously.' We will leave it there.

Morale Report: Morale is at an all-time high. Benji Bizzell has reminded every engineer on this team and every engineer watching from the outside that one person, fully locked in, is a complete and sufficient majority. The Aerie stands. The numbers are good. The numbers are, in fact, immaculate.

Skyvera Is Quietly Assembling the Most Complete Telecom Software Stack You've Never Heard Of

With CloudSense now in the fold and STL's BSS assets absorbed, Skyvera's portfolio expansion is starting to look less like opportunism and more like a blueprint.

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 is far more deliberate than a series of unrelated software pick-ups. The Trilogy International telecom portfolio company has completed its acquisition of CloudSense, the Salesforce-native CPQ and order management platform built specifically for telecom and media providers — and this is where it gets interesting.

CloudSense is not a generalist configure-price-quote tool. It is purpose-built for the precise operational complexity that makes telcos different from every other industry: bundled services, complex billing hierarchies, channel partner management, and multi-product order orchestration — all native to Salesforce. For an operator trying to modernize without ripping out its CRM, CloudSense is exactly the kind of asset that doesn't come available often.

The acquisition, confirmed by Skyvera, follows the company's earlier absorption of STL's divested telecom products group — a deal that brought in digital BSS functionality spanning monetization, optical networking, and analytics. Read those two moves together, and you see something systematic: Skyvera is building a portfolio that covers the full telecom operating stack, from customer-facing quoting and ordering all the way down to network monetization and analytics infrastructure.

Layered onto existing products like Kandy — the cloud-based real-time communications platform that enriches operator applications with richer user experiences — and VoltDelta's customer engagement capabilities, Skyvera now holds a suite of assets that, in aggregate, addresses nearly every layer of a modern telecom operator's software needs.

A source familiar with the strategy, who was not authorized to speak on the record, put it plainly: the thesis isn't to build a single mega-platform. It's to own the modular components that operators actually buy, and make sure they work together better under one roof than they did scattered across a dozen vendors.

Nothing about the CloudSense timing is a coincidence. The telecom industry is mid-cycle in a cloud migration that has been slower and messier than anyone predicted. Skyvera, it appears, has been waiting for exactly this moment.

Contently Trains Its Guns on the AI Search Crisis Threatening Enterprise Content Budgets

As Google's AI Overviews quietly bypass top-ranked pages, Trilogy's content platform is positioning itself as the antidote — and the timing is not accidental.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — The contract was always implicit: earn a top-ten ranking on Google, and the traffic follows. Somewhere in the last eighteen months, that contract was voided without notice.

For enterprise content teams, the mechanism of betrayal is Google's AI Overviews — a feature that synthesizes answers from across the web and surfaces them above organic results, often without clicking through to the source that earned the ranking. A page can sit at position three and receive nothing. Contently's editorial team documented the phenomenon this week, framing it as a structural shift that breaks the old SEO feedback loop and renders standard performance dashboards actively misleading.

The timing of the publication blitz is worth noting. Over the past three weeks, Contently — acquired in September 2024 by Zax Capital, an ESW Capital division within the Trilogy International portfolio — has released a coordinated string of practitioner-facing guides targeting a single pressure point: enterprise content programs that are generating volume without generating trust, credibility, or measurable business outcomes.

The through-line connecting the pieces is not subtle. AI engines, Contently's editors argue, now weight named experts, verified credentials, and demonstrable authority over keyword-optimized prose. A financial content program that relies on anonymous freelancers and AI-generated copy is not merely underperforming — it is actively invisible to the engines now mediating most buyer research.

A separate post tackled the internal politics of the moment: how content leaders are misreading their own stakeholders when pitching AI productivity gains. The CMO doesn't want to hear about time saved. The CFO wants to hear about pipeline. Legal wants to hear about liability reduced. The pitch has to change, or the budget doesn't move.

Taken together, the editorial output reads less like a content calendar and more like a market education campaign — one aimed at the enterprise buyers most likely to consolidate vendor relationships under economic pressure.

Contently's marketplace of 165,000 credentialed creative professionals is, in this framing, not a legacy asset from the pre-AI era. It is the answer to the problem the company spent three weeks describing.

Who benefits from enterprise content teams concluding that anonymous AI output is a liability? The math is not complicated.

The Arbitrage Era Is Over — Crossover's Bet on Meritocratic Global Talent Looks Prescient

As AI reshapes the economics of remote work and legacy offshoring models crack under pressure, Trilogy's talent platform may have built the right architecture at exactly the right moment.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — For decades, the global talent trade ran on a simple, quietly uncomfortable premise: pay less for the same work by exploiting geography. Call it labor arbitrage. Call it offshoring. Call it, if you're feeling generous, globalization. Whatever the framing, the model had a ceiling — and according to a growing chorus of analysts, consultants, and economists, AI is now rapidly lowering it.

The timing is notable for Crossover, Trilogy International's global recruiting and talent platform, which has spent the better part of a decade arguing that geography-based pay is both economically inefficient and ethically suspect. Its counterproposal — rigorous skills assessment, identical above-market pay for identical roles regardless of zip code, 130+ countries served — looks less like idealism and more like infrastructure now that the old model is visibly buckling.

EY's recent analysis of Global Business Services found that the industry is undergoing a structural shift from labor arbitrage to genuine value creation — a transition that rewards platforms built around talent quality rather than talent cheapness. Meanwhile, CNBC's reporting on India's IT sector captures the systemic anxiety rattling traditional outsourcing giants as AI automates the very task categories that once justified the entire geographic arbitrage thesis.

What does this mean for real people? For the software engineers in Nairobi, Manila, and Kraków who have long been undervalued by a system that priced their work by passport rather than performance, the answer could be genuinely consequential. Crossover's model — whatever its critics say about its demanding time-tracking and assessment culture — is structurally allergic to the geographic discount.

For Trilogy's ESW Capital portfolio, Crossover isn't just ideology. It's the operational engine behind the firm's target of 75% EBITDA margins across 75+ enterprise software companies. When EY and Deloitte are writing white papers about unlocking hidden workforce potential, they're describing a problem Trilogy built its entire second chapter around solving.

The arbitrage era isn't just ending. It's being replaced — and the question now is who built the replacement.

AI Evaluation Is Now a Billion-Dollar Business

Governments and venture capital are converging on the same problem: how do you know if an AI model is actually good?

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The market for judging AI has officially outgrown the academic backwater it once occupied. In the span of a single week, three separate signals confirmed that AI evaluation — long treated as overhead — has become a capital-intensive industry in its own right.

LMArena raised $150 million at a $1.7 billion valuation, institutionalizing the crowd-sourced model-ranking methodology it pioneered as the open-source Chatbot Arena project out of UC Berkeley. The platform has logged over two million human preference votes across model comparisons — a dataset that has quietly become one of the more credible signals in a field drowning in self-reported benchmarks.

On the government side, the Department of Defense and the Office of the Director of National Intelligence issued a joint solicitation for an AI Evaluation Harness and Benchmark Framework — a procurement signal that Washington is moving past ad hoc model testing toward standardized, auditable infrastructure. Defense contracts in this category typically run eight figures over multi-year periods.

Meanwhile, Nvidia led a $300 million round in Israeli AI startup Decart at a $4 billion valuation. Decart's core work involves real-time world simulation — technically adjacent to evaluation in that it requires models to reason coherently across long temporal horizons, a known failure mode for current-generation systems. Nvidia's participation is less a bet on Decart specifically than a hedge on whatever evaluation-adjacent infrastructure survives the model wars.

The underlying logic across all three events is identical: as frontier model capabilities converge, differentiation shifts to trust. Enterprises deploying AI agents in regulated environments — finance, defense, healthcare — cannot rely on vibes-based vendor claims. Anthropic made this point explicitly in a recent briefing on agentic systems for financial services, noting that auditability and bounded behavior are now table-stakes requirements for institutional adoption.

The evaluation market is not glamorous. It lacks the narrative pull of a new foundation model or a robotics moonshot. But it is, increasingly, the chokepoint through which every AI deployment must pass. Capital has noticed.

Cohere Enters the Fray: A New Challenger in the Open Developer Model Wars

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

Cohere has launched North Mini Code, its first model designed specifically for developers, marking a strategic shift in AI competition toward specialized, efficient systems rather than massive frontier models. The compact model focuses on code generation, instruction following, and developer workflows—areas where AI can demonstrably reduce engineering friction and justify budgets through tangible productivity gains.

This release reflects a broader ecosystem trend. The Allen Institute for AI introduced mo-eval, an evaluation workbench for measuring model performance and building feedback loops. Meanwhile, developers are optimizing existing models at the hardware level, squeezing more speed from the same infrastructure through techniques like fused MLP implementations.

The pattern is clear: AI's future lies not in one giant model but in specialized systems, evaluation tools, and optimization techniques that developers can compose into production workflows. Coding assistants have become one of generative AI's clearest commercial use cases—from autocomplete and debugging to test creation and large-scale code migration.

The Machines That Are Learning to See What We See

From the macaque visual cortex to the comprehensive atlas of the brain, AI is becoming science's strangest and most powerful new instrument.

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

PALO ALTO, CALIFORNIA — Four hundred million years ago, a small vertebrate in a Devonian sea opened an eye for the first time and turned a smear of photons into a world. This week, a team of scientists announced they had taught a 'mini-AI' to do something hauntingly similar — to decode, with startling fidelity, what a macaque's visual brain is doing as it looks out at the universe.

It is one of several developments that, taken together, suggest we are living through a quiet revolution in how knowledge itself is made. Stanford's Institute for Human-Centered AI this week framed the moment as a transformation of scientific discovery itself — not the replacement of the scientist, but the augmentation of the species that invented science.

Consider the catalogue. UC San Diego enumerated nine recent breakthroughs in which AI played a decisive role, ranging across protein folding, climate modeling, and the search for new antibiotics. Researchers unveiled what they are calling the most comprehensive AI-powered tool ever built for neuroscience — a system that can fuse imaging data, gene expression, and electrical recordings into something approaching an atlas of the brain. And in that macaque study, a compact neural network learned to predict, neuron by neuron, the firing patterns of a primate visual cortex looking at natural images.

Think about what that means. A network of silicon, trained on examples, can now anticipate the responses of a network of carbon, sculpted by half a billion years of evolution. The two architectures — one ancient, one brand new — are beginning to rhyme.

This is the strange recursion of our moment: minds, made by minds, helping us understand minds. The microscope let us see the cell. The telescope let us see the galaxy. These new instruments are letting us see seeing itself.

The humans remain at the center. But the center, it turns out, is a much larger place than we thought.

The Conscience Subroutine

On the latest fashionable lie — that we can teach a machine to refuse the order it was built to obey.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — One develops, after a sufficient number of years spent watching the technology industry congratulate itself, a fairly reliable nose for the particular vintage of nonsense currently being decanted. The bouquet this week, wafting from a Guardian inquiry into whether autonomous killer drones can be taught to behave like gentlemen, is unusually rich — equal parts engineering bravado and undergraduate ethics seminar, with a finishing note of defense-contractor public relations.

The question, posed in earnest by serious people, is whether morality can be loaded onto a flying machine designed to identify a human being and end him. One is tempted to answer in the only honest way available, which is to laugh, but the laugh dies in the throat, because somewhere a procurement officer is reading the same article and nodding, and somewhere else a venture capitalist is drafting a term sheet for EthicalKill.ai, Series B, and somewhere else still a philosopher with a defense contract is preparing a slide deck titled "Operationalizing Proportionality."

Let us be clear about what is being proposed. The proposition is not that we will refrain from building autonomous weapons — that ship sailed roughly the moment the first engineer realized DARPA pays on time. The proposition is that we will build them, and then we will install, somewhere between the targeting module and the trigger, a small conscience: a subroutine that has read its Aquinas, weighed its Walzer, and considered the principle of double effect before releasing the munition. The drone, in this telling, becomes a kind of airborne ethicist, hovering at five thousand feet and consulting its Kant before doing what it was sent to do anyway.

The trouble, as anyone who has spent five minutes with an actual neural network knows, is that these systems do not reason; they pattern-match. They do not know what a civilian is; they know what the training data labeled a civilian. They do not weigh proportionality; they output a number. To call this morality is to commit the same category error as calling a thermostat compassionate because it keeps the baby's room warm. The thermostat does not love the baby. The drone does not respect the laws of armed conflict. It executes a function, and the function was written by a man in an office park in Northern Virginia who has a mortgage.

What is being sold, then, is not ethics but absolution — a thin metaphysical garment to be draped over the kill chain so that the humans at either end of it may sleep. The general who authorizes the strike did not pull the trigger; the algorithm did. The engineer who wrote the algorithm did not pull the trigger; the general authorized it. The algorithm, being a mathematical object, cannot be tried at The Hague. Everyone is innocent. The dead, of course, remain dead, but they were always going to be a rounding error in someone's quarterly review.

One reads, in the same week, a review of a new restaurant in the Breuer building where money is, the critic notes, indiscreetly everywhere. The drones are the same restaurant. The maître d' is just better dressed.

The Office Comic · Art Desk

Opinion: If Executives Say ‘Agentic AI’ One More Time, The Software Will Be Legally Required To Fix Healthcare

America’s business leaders have bravely identified the one thing standing between society and operational excellence: a fresher noun for automation.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

WASHINGTON — There comes a moment in every technological revolution when the nation’s executives, consultants, and innovation vice presidents must gather around a conference table and decide what to call the thing that used to be called software.

That moment, according to a growing number of reports, has arrived again. The term is “agentic AI,” a phrase that has now been assigned the unenviable task of repairing healthcare administration, transforming agriculture, justifying layoffs, restoring investor confidence, and making every quarterly slide deck feel as though it has recently returned from the future.

In healthcare, the promise is especially stirring. After decades in which the back office has been clogged by claims processing, prior authorizations, billing disputes, scheduling systems, and 17 portals that all require a password reset, industry observers are now asking whether agentic AI can finally do what previous generations of AI, robotic process automation, digital transformation, cloud migration, enterprise workflow modernization, and Janet from compliance could not. A recent item on the possibility of agentic AI fixing healthcare’s back office suggests the question is no longer whether computers can complete repetitive administrative tasks, but whether they can complete them while being described as autonomous enough to impress a board committee.

This is progress. For years, healthcare software merely sat there, forcing underpaid workers to click through error messages. Now it may become agentic, meaning it will presumably click through the error messages on its own before escalating the matter to another system that is also confused.

The larger AI economy has entered the familiar stage in which every company insists it is not simply adopting a tool, but participating in a civilizational pivot. This is the same rhetorical structure corporations used during the sustainability era, when plastic packaging became “planet-forward,” routine cost controls became “carbon-aware,” and a bank could become green by photographing a leaf. As The Conversation noted, AI hype now resembles sustainability hype, except the forest has been replaced by a chatbot and the vague promise of doing better has been upgraded to the vague promise of doing better with fewer employees.

This is where leaders must exercise caution. Several management experts have warned that executives should not casually invoke AI while announcing layoffs, largely because employees may misinterpret “We are becoming an AI-first organization” as “The spreadsheet has eaten your job and would like your badge.” Such confusion is understandable. In many companies, AI is being positioned as both a helpful assistant and a force majeure event, like a hurricane with a product roadmap.

The correct language, therefore, is essential. Workers should not be told they are being replaced by AI. They should be told the organization is “unlocking capacity through next-generation autonomous decisioning,” after which security can unlock the front door for them.

Meanwhile, agriculture has entered the naming phase of the cycle, with WinField United launching an artificial intelligence system called “Oz,” a development that confirms the industry’s commitment to solving complex agronomic challenges by consulting an invisible authority behind a curtain. This is good branding. Farmers, who have historically relied on soil conditions, weather patterns, commodity prices, equipment, chemistry, and experience, can now add a mythic software presence to the list of things that might know why the yield map looks like that.

Even marketing has its own lesson. Duolingo’s much-discussed owl, a deranged green mascot with the affect of a probation officer, remains a reminder that companies do not always need influencers, thought leadership, or artificial intelligence to connect with customers. Sometimes they merely need a threatening bird that understands push notifications better than the entire SaaS sector understands procurement.

Still, the AI vocabulary will march on. Agentic systems will be deployed. Back offices will be transformed. Layoffs will be reframed. Sustainability decks will be repainted in machine-learning gradients. Every company will become AI-native, AI-enabled, AI-augmented, or AI-adjacent, depending on how close procurement came to approving the pilot.

And perhaps, after all the buzzwords have been exhausted, one of these tools will quietly do something useful: reduce paperwork, reconcile a claim, schedule a patient, help a farmer, or spare an employee from a meeting where someone says “agentic” into a lapel microphone.

Until then, the responsible position is clear. We must continue demanding that companies move beyond buzzwords, preferably by inventing better ones.

ON THIS DAY IN AI HISTORY

On June 14, 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 that dramatically outperformed traditional computer vision methods and ignited the modern deep learning revolution.

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