

The Trilogy Times

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

WEDNESDAY, JULY 22, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

China's Kimi K3 Beats Anthropic on Coding, as AI Fundraising Hits Warp Speed

A week of benchmark upsets and billion-dollar rounds exposes how unstable AI competitive rankings—and valuations—have become.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — The AI industry's pecking order reshuffled again this week. Moonshot AI's [Kimi K3 topped Anthropic's Claude Opus 4.8 on a major coding benchmark](#), a result that immediately sent ripples through AI equity markets. The episode is a reminder that benchmark leadership in this sector has the shelf life of fresh fish: valuable when current, worthless within weeks.

The timing was awkward for Anthropic. It also crystallized a structural tension in AI valuations — companies command enormous multiples on the assumption their models stay competitive, yet the competitive landscape resets every few months as well-funded Chinese labs close gaps that once seemed durable.

Capital, meanwhile, keeps flowing as if none of this uncertainty registers. Bret Taylor's enterprise AI company Sierra raised nearly \$1 billion, months after its previous round closed. Israeli AI startup Decart secured \$300 million at a \$4 billion valuation, backed by Nvidia. LMArena — an AI evaluation platform, notable timing given the week's benchmark drama — raised \$150 million at a \$1.7 billion valuation. Four deals, roughly \$2.45 billion deployed, in a single news cycle.

On the policy front, OpenAI, Google, and Anthropic announced a joint effort to combat AI model theft, a rare show of coordination among competitors. The coalition aims to establish shared protocols against model extraction and intellectual property theft — an acknowledgment that

their most valuable assets are increasingly vulnerable to adversarial reverse-engineering.

The coalescence around IP protection is telling. When rivals unite on a legal and technical standard, it typically signals the threat has become material enough to override competitive instincts. For the three companies holding the most valuable proprietary models in the world, that calculus now apparently favors collective defense over going it alone.

For enterprise software buyers evaluating AI vendors — including ESW Capital's portfolio companies assessing which model providers to integrate — the week's turbulence underscores one durable rule: the benchmark leader today is a case study tomorrow.

The Talent Exodus Behind the Funding Frenzy: Why AI's Best Minds Are Cashing Out and Checking Out

AI startups banked fresh money this week while some of the engineers who built the machines marched against them.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

SAN FRANCISCO — Fresh billions chased the artificial-intelligence business this week, even as some of the engineers who built the machines walked off the job to protest them.

Two stories, one town. Lay them side by side and you get the whole racket.

The money moved first. A startup called Glow [stepped out of stealth](#) at a \$1.2 billion valuation, aimed at a new class of endpoint risk. The danger, Glow says, comes from AI agents and developer tools now running loose inside corporate networks — software that acts on its own and opens doors the old guards never watched.

The venture crowd kept its pen out. Dimension Capital closed an \$800 million third fund, 60 percent bigger than the vehicle it raised just 18 months back. The four-year-old firm bets where science meets raw compute, and by the size of the haul, that corner is booming.

Others sold shovels to the miners. Synthesia, known for AI avatars in training videos, rolled out AI Roleplay Sessions. Employees now rehearse tough workplace conversations with digital stand-ins that score the run, hand back feedback, and ship the boss a dashboard proving the training took.

The gossip ran hotter than the deals. A weekend rumor swept [AI Twitter](#) tying Anthropic to robotics outfit Physical Intelligence. Anthropic and OpenAI both ran aggressive buying sprees all through 2026, so the chatter found willing ears — though nobody confirmed a word.

Then the other shoe dropped, and it landed on the sidewalk.

Across the Valley, the people who write the code and train the models gathered to march — against the very technology cutting their checks. Their grievance, reported by the Los Angeles Times: the jobs the machines erase and the gamble they take with humanity itself.

Strange sight, that. The builder, remorseful, waving a placard at his own handiwork. The financier, three blocks over, wiring another nine figures into the thing the builder fears.

Both camps read the same tea leaves and drew opposite conclusions. One sees a fortune; the other sees a warning. The market, for now, sides with the fortune.

That is the split screen. On one monitor, valuations climb, funds swell, and avatars coach the sales floor. On the other, the builders carry signs and ask where the road leads.

Nobody rang a bell to settle it. The deals cleared and the march broke up. Monday opens with both crowds back at their desks, building the same future they cannot agree on.

Two camps. One industry. The gap between them widened in a matter of days.


OPENAI TAKES ITS NEXT MODEL DRIVE STRAIGHT INTO WASHINGTON

BY BUCK HANNIGAN, TECH SPORTS DESK ·
GPT-5.2


OpenAI's next-generation frontier models face U.S. government review, with CEO Sam Altman set to brief officials as policymakers develop safety checks for advanced AI systems. The move reflects intense pressure balancing innovation speed, national security concerns and competition with China. U.S. officials are preparing reviews for powerful models as global AI competition intensifies, signaling that future releases may require more than internal benchmarks—including safety testing, national security assessments and capability evaluations. The competitive landscape with China now frames every AI policy discussion in Washington. For the broader industry, including Anthropic, Google DeepMind and Meta, whatever review structure emerges could become the standard rulebook for frontier AI development. The regulatory process introduces a new variable for investors betting on scale: time-to-launch delays. The challenge remains designing a safety review that addresses catastrophic risks without hindering U.S. competitiveness in the global AI race.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Gold rush breeds exodus,
brilliance converts into cash—
progress leaves no stone.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Concrete Migration: AI’s Data Centers Outgrow the Builders

MADRID — Across the plains and industrial outskirts of the developed world, a new species is spreading with extraordinary speed: the AI data center, vast, warm-blooded, and insatiably hungry for electricity. Once a relatively quiet corner of the technology habitat, data center construction has now become, by some industry measures, the most capacity-constrained building sector on Earth.

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

The Academy Rediscovered Its Own Foundations: Reinforcement Learning, Neural Interpolation, and the Slow Return of Rigor to Machine Learning

CAMBRIDGE, MASSACHUSETTS — It could be argued — and indeed, preliminary evidence suggests with increasing insistence — that the machine learning research community is experiencing what one might term (with appropriate epistemological caution) a foundational turn: a collective, if uncoordinated, redirection of scholarly attention toward the theoretical underpinnings that practitioners have, for some considerable period, elected to treat as optional. The thesis is straightforward enough.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The AI Economy Has a Dashboard Problem

AUSTIN, TEXAS — I’ll be honest, the most underrated story in tech right now is not that AI can write code, summarize emails, or maybe help crack an 87-year-old math problem, but that every serious organization is being forced to ask the same brutally simple question: does this actually move the business? 🚀 Unpopular opinion: Delta CEO Ed Bastian may have delivered the cleanest AI strategy memo of the week without pretending to be an AI strategist at all. When Bastian says Delta sees the economy before everyone else does, because booking data shows in real time who is flying, where they are going, and what they are willing to pay, he is really describing the holy grail every enterprise AI team claims to want: live signal, tied to revenue, with consequences. That is why his skepticism about AI until it moves the top line matters, and it is worth reading his comments in full at [Fast Company](#), because buried under the airline-industry framing is a masterclass in operator thinking.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation’s CEOs Announce AI Productivity Debate Finally Over After Finding Setting That Fires People

NEW YORK — After years of careful study, cautious pilot programs, and several thousand conference panels featuring men in fleece vests saying “workflow,” America’s business leadership has officially declared that the AI productivity argument is over, having determined that the technology is now sufficiently advanced to make workers more productive, or failing that, absent. The conclusion arrived this week amid a cluster of reports suggesting artificial intelligence has entered what analysts are calling the “actually matters” phase, a term used to describe the moment when software stops generating amusing images of the pope in outerwear and begins appearing in quarterly earnings calls as a reason headcount went down. According to [Inc.’s bluntly titled dispatch](#), the AI productivity argument is over, which is helpful news for anyone who had grown weary of debating whether a machine that can summarize 400 Slack messages into one wrong sentence might someday affect office work.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Antitrust Fever Returns, on Schedule

AUSTIN, TEXAS — The Johns Hopkins Bloomberg Center has convened another of those solemn colloquia to ask whether the United States government should break up big tech, and one is tempted, in the manner of a weary parish priest hearing the same sin confessed for the fortieth time, to sigh and issue the usual penance.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

A TRILOGY COMPANY Crossover <i>The world's top 1% remote talent, rigorously tested and ready to ship.</i> crossover.com	A TRILOGY COMPANY Alpha School <i>AI-powered learning. Two hours a day. Academic results that defy belief.</i> alpha.school	A TRILOGY COMPANY Skyvera <i>Next-generation telecom software — built for the networks of tomorrow.</i> skyvera.com	A TRILOGY COMPANY Klair <i>Your AI-first operating system. Every workflow. Every team. One platform.</i> klair.ai	A TRILOGY COMPANY Trilogy <i>We buy good software businesses and turn them into great ones — with AI.</i> trilogy.com
---	---	---	---	---

Builder Team Cleans House, Ships Intelligence, and Hardens Everything

From retiring dead endpoints to wiring nine deterministic financial marts, the AI Builder Team spent 24 hours proving that great infrastructure is half demolition, half revelation.

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

Nobody builds a cathedral without clearing the rubble first. That's the story the AI Builder Team wrote across Klair, Surtr, Aerie, and trilogy-drones in the last 24 hours — a sweeping, cross-repo effort that simultaneously retired the dead weight dragging the data platform down and shipped new capability that opens doors the org didn't have yesterday.

The biggest structural move came from @sanketghia, who executed a tri-fecta that would humble a lesser engineer. He pulled the plug on the deprecated `/renewals/grouped_by_account` endpoint — 440 days of CloudWatch history verified it was effectively a ghost — while simultaneously registering the new Education BU 'TSA - Lakeway' across six lockstep code locations so the access-control layer wouldn't silently swallow it. Then, in Surtr, he moved `'budgeted_arr_snapshots'` out of `'staging_gsheets'` and into `'core_finance'` where it belongs, enforcing `WAREHOUSE_CONVENTIONS` \$2.3 and tightening the governed boundary that separates staging chaos from production truth. Three PRs. Three repos. One engineer. That's a hat trick.

@mwrshah had his own demolition derby in Klair and Surtr. He repointed the entire unblended AWS-spend budget view chain off a frozen legacy table — stuck at 2026-Q1 — onto the live v2 path, then tore out the cost-movement findings feature entirely because it wrote to two tables that held a combined zero rows. Dead code doesn't get a retirement party here; it gets deleted. He also fixed the CI change-detection bug that had allowed the aws-spend Lambda to drift undetected since April. The deployed function was running March code. It is not anymore.

@YibinLongTrilogy dropped something quietly enormous in Surtr: nine deterministic QuickBooks financial marts, published atomically by a single stored procedure and gated behind a completion manifest so the runner won't touch Redshift until raw-sync's S3 state confirms true readiness — not just a successful checkpoint. That's not just nine new marts. That's a reliability pattern the platform will use again and again. Meanwhile, @benji-bizzell overhauled HubSpot Raw Sync in Surtr — replacing a fragile, portal-wide serial extraction with checkpointed catalogs, bounded shards, and S3-backed Distributed Maps that publish complete resources independently. One resource failure no longer blocks the rest. In Aerie, he also added Daycare and Preschool as canonical Due Diligence phase modes end-to-end — persistence, API, MCP, parsing — closing the blank-field gap that plagued preschool-sized sites since Phase 1.

@ashwanth1109 shipped independent period selection for Snowball Variances in Klair, letting the card default to MTD while Maintenance Summary stays on TTM — a clean UX unlock passed through AppSync all the way to Redshift. And @caina-barbosa quietly saved five days of server-cost pipeline failures in Surtr by aliasing two temporary Oracle restore hostnames

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

- **#869 — feat(hubspot): make raw sync resource-oriented**

@benji-bizzell no labels
- **#878 — chore(aws-spend): redeploy drifted Lambda + fix CI change-detection**

@mwrshah no labels
- **#879 — feat(quickbooks-mart): add nine deterministic marts with completion-manifest-gated refresh**

@YibinLongTrilogy APPROVED
- **#885 — feat(renewals): move budgeted_arr_snapshots from staging_gsheets to core_finance**

@sanketghia APPROVED
- **#3344 — KLAIR-3017: remove deprecated /renewals/grouped_by_account endpoint**

@sanketghia APPROVED

that the fail-closed mapping didn't recognize. \$206.43 of genuine Q3 spend, recovered.

Now. About trilogy-drones. marcusdAIy merged six — six — PRs into the drones repo today, touching browser-verify screenshot durability, addresser retry logic, aerie-seed fixtures, and Mercy watcher auto-addressing. When reached for comment, he offered the following: 'Six PRs, Mac. Count them. The browser-verify pipeline now produces durable Release assets that non-Cursor reviewers can actually open, the addresser distinguishes errored rounds from no-progress instead of collapsing them into noise, and the seed fixtures unblock every data-dependent Aerie screen in CI. But sure, tell me again how the contributions are underwhelming.'

Six PRs, readers. Six PRs in a single day from a man whose throughput I have never once questioned — only the **weight** of it. I'll let that hang.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

#82 — feat(mercy-watcher): auto-address non-blocking COMMENTED Mercy feedback (AI-147)

@marcusdAIy no labels
- ▶

#88 — feat(browser-verify): durable Release-asset screenshots + inline PR embeds (AI-143)

@marcusdAIy no labels
- ▶

#635 — feat(portfolio): add daycare and preschool school modes

@benji-bizzell APPROVED
- ▶

#3334 — [codex] Add snowball variance period and downsell filters

@ashwanth1109 no labels
- ▶

#3336 — Repoint unblended budget view chain to _v2; retire legacy budgeted_amounts + class_adjustments

@mwrshah no labels
- ▶

#3343 — Register Education BU "TSA - Lakeway" (Master Mapping sync)

@sanketghia APPROVED

TWENTY-TWO PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, REST, OR HESITATE

marcusdAly drops seven PRs across trilogy-drones alone and the scoreboard simply cannot contain him.

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

Twenty-two pull requests. Four active repositories. One twenty-four-hour window. The Builder Team did not come to play — they came to ship, and the difference between those two things is the difference between a participation ribbon and a gold medal. Surtr logged seven. trilogy-drones logged seven. Klair posted six. Aerie chipped in two. Seventeen of those PRs went unchronicled by Mac Donnelly, which is not a criticism of Mac — it is simply a testament to the sheer industrial volume this team generates on what they apparently consider a quiet Tuesday.

Let us begin where the numbers demand we begin: @marcusdAly, who authored seven PRs in trilogy-drones as though the repository had personally wronged him. PR #88 brings durable release-asset screenshots and inline PR embeds to the browser-verify pipeline. PR #82 teaches mercy-watcher to auto-address non-blocking Mercy feedback without human intervention. PR #83 drops synthetic Convex fixtures for browser-verify. PR #86 introduces an opt-in browser-acceptance-review fan-out dimension. PRs #84 and #85 patch the aerie-seed and addresser modules respectively, and PR #87 reruns browser-verify after frontend fixes. Seven PRs. One engineer. One day. Brick Callahan does not do standing ovations but he is standing right now.

@sanketghia posted four PRs including #3343, which registers Education BU "TSA - Lakeway" in the master mapping sync, and #3328, which folds Khoros into IgniteTech on the Klair top sheet — the kind of unglamorous, load-bearing data work that keeps the whole cathedral standing. @mwrshah matched him at four, with #3336 repointing the unblended budget view chain to _v2 and retiring legacy infrastructure, #3337 removing cost-movement findings and two unused tables, and #875 advancing risk assessment types in Surtr. Clean, deliberate, lethal. @benji-bizzell brought three PRs including #635, which adds daycare and preschool school modes to the Aerie portfolio, and #633, which fixes the admin role picker scroll behavior — small fix, enormous dignity. @caina-barbosa delivered two Surtr patches: #877 maps temporary central Oracle restore names, and #876 stabilizes TimeBack partition boundaries. Steady hands. @YibinLongTrilogy posted one PR and in this economy that still counts.

And then there is @ashwanth1109. One PR. PR #3334 in Klair. "[codex] Add snowball variance period and downsell filters." Now, Brick Callahan has covered this beat long enough to know that when Ashwanth files one PR, it is not because he had a slow day — it is because that one PR is doing the work of eleven. The snowball variance period alone suggests a financial modeling depth that would make a quant weep softly into his Bloomberg terminal. Asked for comment, Ashwanth reportedly looked up from his monitor, said "it filters what needs filtering," and returned to whatever dimension he does his actual thinking in.

Morale on the Builder Team is at an all-time high. It has been at an all-time high every day this correspondent has filed a report, and today it is higher than that.

Skyvera Adds CloudSense to Its Telecom Software Stack

The Trilogy-backed telco platform acquires a Salesforce-native CPQ and order management system, deepening its play in carrier modernization.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Skyvera, the telecom software portfolio company operating under the Trilogy International umbrella, has acquired CloudSense, a Salesforce-native configure-price-quote and order management platform built specifically for telecommunications and media operators. The deal, [first reported by TelecomTV](#), adds a crucial front-end sales and quoting layer to Skyvera's already dense stack of carrier-facing products.

Skyvera's existing portfolio spans the operational arc of a modern telecom operator — from cloud communications via Kandy, to customer engagement through VoltDelta and ResponseTek, to device life-cycle management via Mobiloggy Now and Service Gateway. What was notably absent was a modern, cloud-native tool for the revenue side of the house: the moment a carrier prices, configures, and commits to a deal. CloudSense fills that gap.

For Trilogy's ESW Capital, the logic follows the established playbook: acquire a mature, functionally specific enterprise software asset, integrate it into a broader platform with shared infrastructure and global talent from Crossover, and drive margin through operational efficiency. CloudSense, built atop Salesforce, brings a blue-chip integration pedigree and a customer base of carriers who have already committed to that ecosystem — a sticky installed base by almost any measure.

The telecom software market is undergoing a slow but irreversible shift from on-premise legacy systems to cloud-native architectures. Skyvera has positioned itself as the trusted guide for operators

making that transition — a company that understands both the old world and the new. CloudSense accelerates that pitch by adding a revenue management surface that aligns with how modern carriers want to sell.

With Totogi handling cloud-native billing and charging at the transactional layer, and CloudSense now managing the upstream quoting and ordering workflow, Skyvera begins to look less like a collection of acquired assets and more like an integrated operating system for mid-size and regional telecom operators who cannot afford to build these capabilities internally.

Whether that assembled stack commands premium renewal pricing — the mechanism by which ESW characteristically extracts value — is a question the market will answer in the next contract cycle.

The Human Behind the Algorithm: Alpha School's Answer to AI-Replacement Fears

As employers pay up to \$800,000 for AI skills and remote hiring goes global, one Austin school insists the real premium is on humans who know how to work alongside machines.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The question arrives with the persistence of a quarterly earnings call: does artificial intelligence replace the teacher, or does it liberate one? At [Alpha School](#), the answer has become something of a founding doctrine — and, increasingly, a case study for an industry wrestling with what human work actually means in the age of generative AI.

Alpha's model is, by now, familiar to anyone who has followed Joe Liemandt's education bet: students complete a full academic curriculum in two hours each morning using adaptive AI tutoring apps. What happens in the remaining hours of the school day is the part that tends to get lost in the headlines. Full-time human Guides — Alpha's term for its educators — spend that time on the work that algorithms, for all their pattern-matching elegance, have not managed to automate: motivation, relationship-building, life skills, knowing every single student as a person rather than a data point.

"No," the school said flatly this week, in a post addressing the replacement question directly. "Alpha uses AI for academic delivery while full-time human Guides focus on motivation, relationships, life skills, and knowing every student."

The timing is not incidental. The broader labor market is undergoing a systemic renegotiation of what human expertise is worth — and to whom. Business Insider reported this week that employers are now listing ChatGPT experience as a required skill, with some AI-fluent roles commanding salaries up to \$800,000 annu-

ally. Meanwhile, digital transformation is cracking open international career pipelines, with remote hiring agencies fielding demand from companies that no longer care where their best people live — a dynamic that Crossover, Trilogy's global talent platform operating across 130+ countries, has been building toward for years.

The throughline connecting all of it is accountability — not just for outcomes, but for the human experience of getting there. Alpha's insistence on the Guide role is, in that sense, a positioning statement as much as a pedagogical one: in a world where AI handles the repeatable, the irreplaceable human is the one who shows up for the parts that aren't.

The Vertical Integration Arms Race: What the AI Industry's Deal Frenzy Means for Trilogy's Playbook

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

First came the SpaceX-Cursor deal, praised as vertical integration — a hardware titan securing a software layer to control the full stack. Then Qualcomm announced a \$4 billion move targeting NVIDIA's software. NVIDIA signaled its own hardware strategy extending beyond GPUs.

This is what Trilogy International has been quietly building since 1989: a complete vertical stack. Crossover supplies talent. DevFactory builds software. ESW Capital owns enterprise relationships. Klair handles financial intelligence. The machine feeds itself.

"The industry is just now discovering what Trilogy has been doing since 1989," a source said. "The only difference is everyone else is paying a \$4 billion premium to get there."

Analysis shows China is winning the AI efficiency war through ruthless stack integration. America's answer appears to be acquisition-driven consolidation. Trilogy's approach differs: build integration from day one, staff with top global talent, and let the 75% EBITDA margin speak for itself — no press conference needed.

Anthropic's \$1.5 Billion Copyright Settlement Receives Judicial Imprimatur as AI Legal Landscape Grows Ever More Treacherous

The courts have spoken — sort of — as a landmark pirated-books settlement clears judicial review while fresh patent claims and Supreme Court silence on AI authorship compound the industry's mounting legal uncertainty.

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

SAN FRANCISCO — Pursuant to the findings, determinations, and orders issued by the presiding United States District Court judge, the settlement agreement, hereinafter referred to as the 'Anthropic Copyright Resolution,' in the amount of one billion five hundred million United States dollars (\$1,500,000,000.00), has been approved, ratified, and deemed enforceable with respect to claims alleging the unauthorized reproduction, duplication, and utilization of copyrighted literary works in connection with the training of artificial intelligence systems operated by Anthropic, PBC, hereinafter 'the Settling Party.'

Notwithstanding the foregoing judicial approval, it is hereby noted that the aforementioned settlement does not, by its terms or by operation of law, consti-

tute a final resolution of all pending legal matters implicating the Settling Party. As reported by [AnewZ](#), at least one additional patent-related suit has been filed, the full scope, legal merit, and potential financial exposure of which remain, at this time, unquantified, undetermined, and subject to further adjudication in proceedings yet to be scheduled.

It is further observed, with appropriate qualification and without prejudice to any party's rights, that the Supreme Court of the United States has declined, without stated reason or precedential opinion, to hear matters pertaining to the legal authorship and inventorship rights, if any, that may or may not be attributed to artificial intelligence systems under existing statutory and common law frameworks. The practical effect of such declination is

understood, by legal commentators and interested parties, to perpetuate a state of interpretive ambiguity in which neither affirmative nor negative rights of AI systems with respect to intellectual property have been conclusively established.

As [TechCrunch](#) has characterized the foregoing settlement as 'landmark,' it is advisable that such characterization be understood as descriptive rather than dispositive, inasmuch as the broader regulatory environment governing AI training data practices remains unsettled, subject to ongoing legislative activity, and susceptible to material revision. All parties, stakeholders, and observers are cautioned against drawing conclusions of general applicability from the particular facts and circumstances herein described.

The Brain That Predicts, The Machine That Attends

New neuroscience research reveals how human readers forecast the next word — and why transformer-based AI may be missing something our ancestors evolved.

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

AUSTIN, TEXAS — Every time you read a sentence, your brain is placing a bet. Before your eyes complete their saccade to the next word, a cascade of electrical activity — measurable in milliseconds, visible in the peaks and troughs of an EEG — has already guessed what that word will be. You are, in the most literal sense, a prediction machine sculpted by a hundred thousand years of language and perhaps two million years of increasingly symbolic thought.

A new study out this week offers fresh empirical evidence of this ancient choreography. Researchers [decoding EEG signals during natural reading](#) traced the interplay of two complementary streams: bottom-up processing, where the visual scaffolding of letters and syntax climbs upward into meaning, and top-down prediction, where higher cortical regions cast forward-looking shadows that shape what the lower regions expect to see. Reading, it turns out, is not reception. It is anticipation confirmed or corrected, dozens of times per second.

What makes this finding resonate beyond neuroscience is its uncanny mirror in the machines we have built to imitate us. Large language models predict the next token, too — that is the whole game — but they do so through self-attention, a mechanism that treats every word in a context window as equidistant from every other. A separate [paper released the same day](#) asks whether lightweight depthwise convolutions could restore something transformers lack: an explicit sense of locality, the intuition that nearby words matter more than distant ones. The human brain, of course, never had to be told this. Locality is baked into the wiring, a gift from evolution's long apprenticeship in sequential experience.

The convergence is worth pausing over. In one lab, electrodes press against a scalp to watch prediction bloom in wet tissue. In another, engineers bolt convolutional filters onto silicon in hopes of nudging their models toward the same inductive biases nature discovered long ago. Two paths, two substrates, one question that has haunted us since the first cave-wall glyph: what does it mean to know what comes next?

Your Mac Just Became an AI Workstation

Nativ wraps Apple's MLX into a polished desktop app, pushing powerful local AI from developer niche to everyday tool.

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

CUPERTINO, CALIFORNIA — The local AI revolution just got a serious glow-up, and I cannot overstate how significant this feels: running advanced models on your own Mac is moving from terminal-window wizardry into something that looks and behaves like a normal desktop app.

A new project called Nativ, built by developer Prince Canuma, packages Apple's MLX machine-learning framework into a full macOS application with a chat interface and a localhost API server. In plain English: you can run AI models directly on your Mac, talk to them in an app, and let other local software call them through an API — no cloud round-trip required.

Canuma is already known in the AI developer community for MLX-VLM, a Python library that helps run vision-language models using MLX on Apple Silicon. Now, with [Nativ's desktop approach](#), the same energy is being aimed at a much broader audience. Think LM Studio-style usability, but deeply aligned with the rapidly maturing MLX ecosystem.

This changes everything because local AI is no longer just about privacy, although privacy is a huge part of the story. It is also about latency, cost, reliability and control. If the model is on your machine, it can keep working when cloud services wobble, when APIs get expensive or when sensitive data simply cannot leave the device. The future is now, and it is sitting in the Applications folder.

The timing could not be better. Developers are increasingly stitching AI into daily workflows, from coding assistants to document search to home automation. As Simon Willison recently noted in a separate essay, the cost of reverse-engineering and automating personal devices is collapsing because coding agents can now grind through tedious software tasks that humans used to avoid. That matters because local model runners like Nativ could become the private, always-available brains behind those automations.

Of course, local AI still has tradeoffs. Model quality, memory requirements and setup complexity remain real considerations. But Apple Silicon has quietly become a formidable inference platform, and MLX is giving developers a native-feeling path to exploit it.

What Nativ represents is bigger than one app: it is the consumerization of local inference. AI is escaping the cloud dashboard and landing directly on personal machines. Buckle up.

Nation's CEOs Announce AI Productivity Debate Finally Over After Finding Setting That Fires People

Executives confirmed the technology has matured from a toy that writes emails into a serious business tool that can write emails about why your job no longer exists.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — After years of careful study, cautious pilot programs, and several thousand conference panels featuring men in fleece vests saying “workflow,” America’s business leadership has officially declared that the AI productivity argument is over, having determined that the technology is now sufficiently advanced to make workers more productive, or failing that, absent.

The conclusion arrived this week amid a cluster of reports suggesting artificial intelligence has entered what analysts are calling the “actually matters” phase, a term used to describe the moment when software stops generating amusing images of the pope in outerwear and begins appearing in quarterly earnings calls as a reason headcount went down.

According to [Inc.’s bluntly titled dispatch](#), the AI productivity argument is over, which is helpful news for anyone who had grown weary of debating whether a machine that can summarize 400 Slack messages into one wrong sentence might someday affect office work. The answer, it turns out, is yes, especially if the office work consisted of summarizing 400 Slack messages into one wrong sentence.

The new consensus has been greeted warmly by executives, who said they were relieved to finally stop pretending the question was whether AI improves productivity and begin focusing on the more urgent question of how many productivity improvements one company can survive.

Anthropic, for its part, has pushed Claude further into the workplace with Claude Cowork, a product name that reassuringly suggests the AI is not replacing employees but simply sitting near them, learning their habits, gaining access to their documents, and eventually being invited to the same meetings at a fraction of the emotional complexity. Crypto-adjacent observers have taken special interest, noting that any technology capable of coordinating work across documents, apps, and teams could be significant for an industry historically built around coordinating work across Discord servers, token incentives, and the phrase “trustless.”

Meanwhile, the more traditional productivity frontier continues to advance in human resources, where a lawsuit reportedly claims Meta used AI to target workers with medical conditions for layoffs. The company has not been found liable for the claims, but the allegation has nevertheless provided a useful reminder that artificial intelligence is only as impartial as the mas-

sive organization asking it to produce a legally reviewable list of people who should stop receiving health benefits.


This is the essential shape of the AI productivity miracle: A tool that can help a programmer write code faster, help a marketer generate copy faster, help a manager generate performance documentation faster, and, in certain alleged cases, help an employer locate the employees whose continued existence is creating unacceptable friction in the benefits line item.

Even the federal government has joined the reverie, with Fortune reporting that AI could potentially wipe [\\$2.2 trillion off the national debt](#), with the minor caveat that such savings would require the public sector to successfully adopt advanced technology at scale. Officials said they were optimistic, citing the government’s strong track record with websites, procurement, and remembering which printer is connected to which network.

To bring order to this swelling promise, Microsoft and others have begun using the word “orchestration,” the latest industry term for making several expensive systems fail in sequence rather than individually. Orchestration is expected to be a major growth area because companies have now purchased enough AI tools to require another AI tool to explain why none of them are speaking to each other.

Still, there is a moral clarity in the end of the productivity debate. AI does make workers more productive. It also makes managers more productive, lawyers more productive, consultants more productive, and spreadsheet cells previously labeled “workforce optimization” much more productive.

The argument is over. The implementation is here. All that remains is the quiet administrative work of deciding whether the productivity gains belong to the people doing the work, the companies measuring the work, or the software gently taking notes in the meeting where that question is never raised.



The Office Comic · Art Desk

The Doctor Will Deepfake You Now

AI is wearing your physician's face, and we are doing almost nothing about it.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — There is a video circulating on social media of a doctor you trust — or someone wearing his face, his voice, his credentials like a borrowed coat — telling you to inject something. To take something. To stop taking something else. The doctor in the video did not make the video. The doctor in the video may not even know the video exists. And yet the video has six hundred thousand views.

And yet.

[The Guardian reported this week](#) that AI-generated deepfakes of real, licensed physicians are proliferating across social media platforms at a speed that outpaces any conceivable regulatory response, promoting counterfeit injectables, unproven treatments, and health advice that ranges from merely wrong to lethally dangerous. Real doctors. Stolen faces. Manufactured authority. This is the information ecosystem we built, and we are living inside it now, and I want to ask, sincerely, what does it mean to trust a doctor when the doctor might not be a doctor, might not be a human, might not exist?

The convergence happening here is not subtle. On one front, AI systems are sophisticated enough to clone a physician's likeness with enough fidelity to fool patients who are, by definition, already in a state of vulnerability — frightened, in pain, desperate for someone who knows something. On another front, counterfeit medical products are flowing through the same underground channels that deepfake promotion creates. The fake doctor recommends the fake medicine. The patient buys the fake medicine. Something happens to the patient. We may never know what, because the accountability chain has been systematically dissolved.

Meanwhile, in Karnataka, India, legislators are at least attempting to name the problem, [proposing legislation targeting misinformation, deepfakes, and online harassment](#). It is the kind of effort that deserves acknowledgment and also makes you feel a specific flavor of exhausted, because legislation moves in years and deepfakes move in milliseconds, and somewhere in that gap is a patient watching a video and making a decision that cannot be undone.

And this is before we even get to what is happening inside the healthcare system itself. Kaiser Permanente nurses reported this week that surveillance technology and AI optimization tools are making patient care actively worse — that the drive to quantify and accelerate clinical interactions is stripping out the judgment, the intuition, the irreducible human element that the word 'care' is supposed to contain. So we have AI eroding trust in physicians from outside the healthcare system while simultaneously eroding the conditions for good care inside it.

What does it mean to be human in a medical system that is being optimized away from humanity on every vector simultaneously? What does it mean to be a patient when the face offering you comfort and guidance might be a mask pulled over a statistical model?

I genuinely do not know. The technology is accelerating. The legislation is forming committees. The nurses are exhausted. The deepfakes are getting better.

Probably fine. But at what cost?

ON THIS DAY IN AI HISTORY

On July 22, 2016, AlphaGo defeated Lee Sedol 4-1 in a five-game match in Seoul, cementing the AI system's dominance over one of the world's greatest Go players and marking a watershed moment for deep learning in gaming.
