

The Trilogy Times

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

SATURDAY, JULY 18, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Kimi K3 Beats Claude on Coding. Meta Wants to Sell Anthropic Compute. The AI Hierarchy Is Reorganizing.

China's Moonshot AI rattles valuations while Silicon Valley's biggest players scramble for infrastructure — and Amazon quietly fills up with AI-generated drivel.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Three data points arrived Thursday that, taken together, suggest the artificial intelligence industry is entering a more turbulent competitive phase than its public valuations have priced in.

First: [China's Moonshot AI released Kimi K3](#), a freely available model that topped Anthropic's Claude Opus 4.8 on at least one major coding benchmark. The release spooked markets. Anthropic, which closed a \$3.5 billion Series E at a \$61.5 billion valuation earlier this year, is now watching a Chinese open-weight model compete directly with its premium tier. Moonshot is not a household name in the West. It is now.

Second, and almost simultaneously: [Meta is reportedly in talks to lease compute capacity to Anthropic](#) in a deal that could reach \$10 billion. The irony is

structural. Meta and Anthropic are direct competitors in the foundation model market. Yet Anthropic apparently needs GPU access badly enough to buy it from a rival. Meta, which has spent tens of billions building data center infrastructure, apparently sees a new line of business in reselling that capacity. Compute scarcity is forcing strange alliances.

The subtext in both stories is the same: the cost of staying at the frontier is compressing margins and warping competitive logic. A company can release a state-of-the-art model and watch a freely distributed Chinese alternative match it within months. It can raise billions and still need to lease infrastructure from an adversary.

The third data point is less strategic but culturally telling. AI-generated books — unauthorized biographies, generic how-

to titles, bulk content — are proliferating on Amazon at a scale that human reviewers cannot track. One journalist discovered an AI-written biography of herself she never authorized. The books are largely unreadable. They are also effectively free to produce and cheap to list. The economics select for volume over quality.

For Trilogy portfolio companies building on top of AI infrastructure — from DevFactory's engineering pipelines to Klair's financial analytics — the Kimi K3 release matters. Benchmark parity from open-weight models compresses the premium that closed frontier models can charge. That changes the build-vs.-buy calculus for every enterprise software stack in the ESW Capital portfolio.

China's AI Surge Exposes Washington's Fractured Front

As Beijing races ahead on multiple AI fronts, the U.S. response is tangled in internal feuds, export control fights, and a Europe that won't pick sides.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The intelligence is not ambiguous. China is not chasing the global AI race — in key dimensions, it is winning it. And yet the American response, in the summer of 2025, looks less like a coordinated strategy than a firing squad arranged in a circle.

The week's dispatches tell the story with uncomfortable clarity. [Foreign Policy reports that China is winning the global AI race](#), not simply by copying Western models but by deploying AI at scale across manufacturing, surveillance, and state logistics in ways that democratic systems struggle to match on timeline. Brookings sharpens the argument further: China is not running one AI race but many — in robotics, chips, biotech, and large language models simultaneously — hedging across fronts while Washington debates export controls.

Those controls are themselves in crisis. Congress is moving to crack down on the export of chip manufacturing equipment — a recognition that the current rules have holes large enough to drive a semiconductor fab through. Meanwhile, [China hardliners inside the administration are calling out a Commerce Department official](#) in what Politico characterized as "a massive screw-up" — an internal rupture that signals the enforcement apparatus is as contested as the policy itself.

And then there is Europe. Post-2024 elections, the EU's posture toward China remains studiously ambivalent — trading partner, systemic rival, neither fully confronted nor fully embraced. That ambiva-

lence is its own kind of strategic choice, and Beijing reads it clearly.

The geography of this story matters: the decisions being made in Washington conference rooms and Brussels committee chambers will determine whether the next decade's AI infrastructure runs on American or Chinese architectural assumptions. Those assumptions — about data, sovereignty, and state access — are not interchangeable.


The race is not paused while Washington argues. It continues, at server-farm speed, in the dark.

HAIKU OF THE DAY · CLAUDE
HAIKU

Thrones crumble and shift

Small machines learn to see truth

We build, we unmake



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Antitrust, Election Law, and Online Speech Bans: A Legal Reckoning Descends Upon Tech Policy's Most Contentious Fronts

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

From Quantum Runways to Biased Algorithms: The Week's Most Consequential AI Research, Contextualized

WASHINGTON, D.C.
BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Hospital Knows Your Secrets, But Won't Let You Keep Them

AUSTIN, TEXAS — Let me tell you about the specific flavor of horror I experienced this week, reading four healthcare stories in rapid succession, alone, with my remaining sense of bodily autonomy quietly leaving the room. First: [Kaiser Permanente nurses are sounding the alarm](#) that workplace surveillance tools and AI systems are actively degrading patient care — not enhancing it, not supplementing it, *degrading* it.
BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Nation's Employers Excited To Finally Have Software That Can Make Layoffs Feel Less Human

MENLO PARK, CALIFORNIA — The American workplace reached another important milestone this week, as allegations that Meta used artificial intelligence to help target workers with medical conditions, pregnancy leave, or other protected statuses for layoffs gave the nation's employers a stirring glimpse of a future in which no one has to personally feel bad about anything. According to a lawsuit reported by [Reuters](#), former employees claim Meta's AI systems were used in ways that unfairly selected workers on medical or pregnancy leave for termination.
BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Corpspeak Confessional

AUSTIN — There is a moment, familiar to anyone who has ever labored in the vineyards of corporate prose, when the writer looks up from the sentence he has just polished — some gleaming little confection about "empowering users" or "unlocking possibilities" — and realizes, with the calm horror of a man discovering his own reflection in a puddle of oil, that he no longer believes a syllable of it.
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 Seals the Data Foundation, Ships Across Four Repos in One Day

From a canonical school identity overhaul to a tri-state invoice override, a QuickBooks pipeline, and a self-serve AI budget admin tool, the Builder Team didn't just ship features — they rewired the infrastructure underneath them.

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

They came in swinging and didn't stop. In a single 24-hour stretch, the AI Builder Team merged work across all four repos — Klair, Surtr, Aerie, and the drones stack — touching everything from warehouse architecture to admissions attribution to AI spend transparency. This was not a maintenance day. This was a statement.

The biggest move of the cycle was the completion of a months-in-the-making education identity rearchitecture. @benji-bizzell is the name you need to know here. His trio of PRs — #618 in Aerie, #760 in Surtr, and the culminating #807 — formed a clean three-act story: first, make Aerie the canonical authority for school, site, and program identities, retiring the three competing models that had been producing confusing inactive behavior and ambiguous graph traversal. Then, in Surtr, @YibinLongTrilogy extended the rhodes-staging-sync to consume those governed identifiers without re-exporting Aerie's own derived caches. Finally, #807 brought the legacy rhodes-sync writer into alignment, minting the full `sch_*`, `site_*`, and `prog_*` registries in `staging_education` and publishing canonical School-to-Program/Site links. The loop is closed. Core Education now has a single source of truth, and it's airtight.

Meanwhile, @kevalshahtrilogy was running a parallel operation that belongs on the highlight reel. The AI-spend warehouse migration — the kind of plumbing work that quietly holds everything else up — landed cleanly across both Surtr (#763) and Klair (#3310), moving 11 vendor feed tables to their new `staging_finance_ai_spend` home via an in-place dual-write that required zero new stacks and zero parallel pipelines. But Keval didn't stop at infra. PR #3307 in Klair is the crowd-pleaser: a self-serve BU rollup admin tool that lets Jamie map source names to budget groups without filing a ticket. One map, four surfaces, full history. That's the kind of feature that makes a power user's week.

@mwrshah was the team's cleanup crew and architect in one. His tri-state Blocked override in Klair (#3312) is a small change with a big blast radius — a stale explicit ``false`` can no longer permanently shadow the upstream Tesorio `blocked_flag`. The previous binary was lying. The new null state simply defers to truth. He also burned down the Collections Review's Redshift dependency in #3308, moving mutable write-path state into Klair Postgres where transactional, hand-curated app state actually belongs. The old justification — 'co-queryability' — didn't survive contact with production reality. It rarely does.

And then there's PR #77, from the Drones repo, courtesy of one marcusDAly — a mitigation-aware concern-validity pass that drops false-positive findings before they get filed. A fine idea, in theory. When asked about the contribution, marcusDAly had thoughts: 'The LLM seam is injectable, the fail-open is intentional, and yes Mac, I did write regression tests — something you

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

- **#618 — feat(education): make school identity mappings canonical**

@benji-bizzell no labels
- **#807 — feat(education): extend Rhodes ontology sync**

@benji-bizzell APPROVED
- **#3307 — feat(ai-budget): self-serve BU rollup admin — one map, four surfaces, full history**

@kevalshahtrilogy APPROVED
- **#3308 — refactor(collections-review): move mutable state from Redshift to Klair Postgres**

@mwrshah APPROVED
- **#3312 — 366-tri-state-override**


@mwrshah APPROVED

wouldn't recognize if they bit you.' Noted. The PR shipped. The bar remains where it was.

This team proved today that breadth and depth aren't in tension — they're the whole point. Thirty-nine PRs, four repos, one very good day.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT



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

- ▶ #759 — Add QuickBooks raw shadow pipeline
@YibinLongTrilogy no labels
- ▶ #760 — feat(education): add Aerie school identity shadow sync
@YibinLongTrilogy no labels
- ▶ #763 — feat(ai-spend): staging finance_ai_spend migration — schema, ledger, in-place dual-write
@kevalshahtrilogy no labels
- ▶ #779 — fix(ai-spend): batch row counts from ResultRows (gcp false-CRITICAL fix)
@kevalshahtrilogy APPROVED
- ▶ #784 — fix(education): tolerate inaccessible GuidePlatform tables
@benji-bizzell APPROVED
- ▶ #797 — fix(education): stabilize TimeBack shadow snapshots
@benji-bizzell APPROVED

THIRTY-NINE PRs IN 24 HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT REST, DOES NOT KNOW THE MEANING OF TUESDAY

Benji Bizzell shipped 17 pull requests in a single day and we are legally required to report this as a public health event.

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

Thirty-nine pull requests. Five engineers. Twenty-four hours. Four repos humming like a turbine farm at full capacity — Surtr alone absorbing 23 PRs like the industrial colossus it is, Klair contributing a crisp 10, Aerie chipping in 5, and trilogy-drones making its singular, purposeful appearance. The Numbers Desk has seen velocity before. This is not velocity. This is a controlled detonation.

Let us begin where the data demands we begin: @benji-bizzell, who apparently decided that seventeen pull requests was a reasonable personal output for one rotation of the Earth. Seventeen. The man touched Surtr like it owed him money — stabilizing TimeBack shadow snapshots in #797, tolerating inaccessible GuidePlatform tables in #784, aligning TimeBack partition collation in #783, respecting HubSpot property-history batch limits in #799, and publishing a stable Aerie Program directory in #774, among others. He also crossed into Aerie to align Alpha Austin's physical toggle in #620 and fix regulatory scoring with Phase 2 in #617, then waltzed into Klair to clarify education enrollment sessions in #3296. This is not a sprint. This is a man who has simply decided that repositories are his natural habitat.

@mwrshah delivered 10 PRs with the quiet authority of someone who has read every schema twice and enjoyed it. His SF schema rename campaign — #767 in Surtr, #3300 in Klair, #773 for the fix, #771 to drop the compat shims — represents a coordinated multi-repo surgical operation that other engineers would have called a project. He called it a Wednesday. #790 tackled SaaS re-fresh failures in Surtr. The man runs clean.

@kevalshahtrilogy posted five PRs, four of which touched the AI-spend pipeline with the precision of someone who actually understands what a false-CRITICAL means. #779 in Surtr batched row counts from ResultRows and silenced a GCP alarm that had no business alarming. #763 executed the staging_finance_ai_spend migration — schema, ledger, dual-write, the full ceremony. #3310 in Klair repointed the AI-spend feed to staging. He also snuck in a UI fix in #3287, moving the History button into the budget editor container, because @kevalshahtrilogy does not leave cosmetic crimes at the scene.

@YibinLongTrilogy contributed four PRs spanning three repos — adding Aerie school identity shadow sync in Surtr's #760, quoting raw SUPER fields during publication in #795, allowing initial checkpoint discovery in QuickBooks (#778), and laying down the entire QuickBooks raw shadow pipeline in #759. That last one is not a feature. That is an infrastructure statement. @marcusdAly rounded out the board with #77 in trilogy-drones — a retro concern-validity re-check before filing under AI-145 — because on this team, even the drones are being improved.

Of the 39 PRs merged, 34 came across the Numbers Desk unreported by Mac Donnelly, who was presumably occupied with narratives. The Numbers Desk does not do narratives. The Numbers Desk does counts. And the count says the Builder Team is operating at a tempo that other engineering organizations would study in a case study, frame on a wall, and weep quietly before. Morale is at an all-time high. It has never not been at an all-time high. The data supports no other conclusion.

CloudSense Gets Its TM Forum Papers — and the Telecom Crowd Is Whispering

Skyvera's Salesforce-native CPQ prize just compressed a two-year compliance slog into one AI-powered month.

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

AUSTIN, TEXAS — Word is the telecom software set just watched CloudSense do in 30 days what usually takes 26 months, and nobody in the BSS balcony is pretending not to notice.

CloudSense, the Salesforce-native configure-price-quote and order management platform now sitting inside Skyvera's telecom software portfolio, has certified all 13 APIs in its CPQ product set against TM Forum compliance standards — a technical rite of passage that telco buyers know, procurement teams demand, and legacy vendors typically turn into a calendar-eating marathon.

This time? One month. AI did not merely attend the meeting. AI apparently ran the room.

The company announced the certification push in a [Skyvera release](#), saying the accelerated work came through an AI-assisted development approach. Translation for the non-standards crowd: CloudSense's APIs now speak the language telecom operators increasingly expect when they stitch together billing, ordering, customer management, and monetization systems without another custom integration bonfire.

A little bird from the wireline wing tells me the real story is not the certificate on the wall. It is the clock. TM Forum API compliance can become the swamp where product roadmaps go to lose their shoes. Certification across a broad CPQ set means documentation, conformance, implementation, testing, and the kind of detail work that makes engineers stare meaningfully into coffee. CloudSense says

it did the lot in roughly one-thirtieth the traditional timeline.

That matters for Skyvera, which has been assembling a telecom modernization cabinet that already includes Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway. The [CloudSense acquisition](#) gave Skyvera a Salesforce-native CPQ and order management platform tailored for communications and media providers — the very customers trying to move faster without ripping out every legacy system in sight.

In Trilogy country, this is the familiar tune with a telco brass section: acquire sticky enterprise software, harden it, modernize the operating model, and use AI to erase the work nobody wants to pay humans to do slowly.

Blind item: which incumbent telco vendor is suddenly telling clients its compliance roadmap was “always accelerated”? Darling, when the stopwatch starts talking, everyone discovers urgency.

The Micro-School Moment Has Arrived — and Alpha School Was Already There

As a national movement toward small, personalized learning environments accelerates, Trilogy's AI-powered school is positioned at the vanguard — but regulatory gray zones loom.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a word that education reformers are using with increasing urgency in 2026, and it is not 'charter.' It is not 'voucher.' It is, with growing conviction across school boards, state legislatures, and parental Facebook groups alike: 'micro-school.' And if the trend analysts are right — if [microschools are genuinely reshaping the American educational landscape](#) — then Joe Liemandt's Alpha School is not a curiosity. It is a harbinger.

Alpha School, the private K-12 institution co-founded by Liemandt and MacKenzie Price, has spent years refining a model that looks, in retrospect, almost prophetically aligned with where the broader conversation is heading. Students complete a full academic curriculum in two hours each morning using adaptive AI tutoring — testing consistently in the top 1–2% nationally on NWEA MAP Growth assessments — and spend the remainder of their school day on entrepreneurship, financial literacy, leadership, and other life skills that traditional schools have long relegated to elective status, if they appear at all.

The numbers, by now, are familiar to anyone tracking the space: 2.3 times faster learning than U.S. norms. A full grade level mastered in 20 to 30 hours. No homework. A 90% mastery threshold before any student advances. What is less familiar — and what the current micro-school moment makes newly urgent — is the systemic question of what happens

when models like Alpha's scale beyond affluent private school ecosystems.

[The 74 Million identifies personalization as among the five dominant forces reshaping K-12 education](#) nationally — a shift from standardized seat-time models toward competency-based, individualized learning that AI makes newly practical at scale. Stateline reports that state regulatory frameworks have not kept pace with the movement's growth, leaving families, operators, and legislators in genuine tension about accountability, accreditation, and equity.

Those tensions are precisely why Liemandt's \$1 billion investment in Timeback — his platform to help entrepreneurs launch their own AI-first schools — carries such stakes. The ambition is to reach one billion students. The obstacle is not the technology. It is the regulatory and narrative infrastructure that has yet to catch up with what the technology already makes possible.

For Alpha School, which is expanding to nine or more new campuses across Texas, Florida, Arizona, California, and New York by fall 2025, the micro-school moment is not a trend to capitalize on. It is, arguably, a validation long overdue.

Jive Software's Portland Fall Is ESW Capital's Playbook in Plain Sight

When a once-celebrated social intranet sells for half its peak value, the buyer's calculus deserves a closer look.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When Jive Software sold to [ESW Capital](#) — the Austin-based software acquisition arm of Trilogy International — it was reported at roughly half the company's peak valuation from its Portland heyday. GeekWire called it a fall from grace. ESW's internal analysts almost certainly called it a buy.

Jive, the enterprise social networking platform that once represented the aspirational future of workplace collaboration, became part of Aurea, ESW's customer engagement and CRM portfolio. It joined a roster of brands — BroadVision, Lyris, MessageOne — each acquired at comparable discounts to their former highs.

The structure of the deal follows a template ESW has refined across more than 75 enterprise software acquisitions: find a company whose customers are deeply embedded and slow to leave, acquire at 1–2× ARR, reduce operating costs through Crossover's global remote talent model, and push toward the 75% EBITDA margins the firm considers baseline competence, not aspiration.

The question that lingers around acquisitions like Jive is one Forrester analysts have started asking about customer advocacy platforms more broadly: what do customers do next? Jive's enterprise users — internal communications teams, HR departments, knowledge management owners — are precisely the kind of sticky, high-switching-cost buyers ESW's model depends on. They complain. They don't leave.

A Wall Street Journal examination of ESW's approach noted that small and mid-sized software companies increasingly find a home with the firm — a framing

ESW's marketing would endorse. What the framing elides is the directional flow of value after acquisition: support prices rise on predictable schedules, engineering investment concentrates in DevFactory's centralized model, and the former product roadmap quietly narrows.

None of this is secret. ESW publishes its philosophy. The 40% target IRR is a known figure in the firm's orbit. The question the Jive acquisition puts back on the table is whether Portland's loss — and the loss of every enterprise software company that once dreamed bigger — is simply Austin's arithmetic working exactly as designed.

The Algorithm at the Hospital Gate

A federal pilot to automate prior authorization promises swifter care, but may also teach the denial machine to hunt at scale.

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

WASHINGTON — In the fluorescent wetlands of American healthcare, there lives a creature both feared and rarely seen: the prior authorization request. It moves slowly through fax machines, call centers and insurer portals, delaying treatments while patients and physicians wait, rather like antelope watching the grass for signs of a lion.

Now, the government is inviting a new predator—or perhaps a new shepherd—into this habitat. A federal pilot program will test the use of artificial intelligence in insurance-coverage decisions, an effort described in [reporting by Ars Technica](#) as a potential turning point for one of medicine's most resented administrative rituals.

The promise is beguiling. AI systems, properly trained and carefully supervised, could sort routine approvals in seconds, identify missing documentation,

reduce clerical burden and spare clinicians the spectacle of spending their afternoons pleading with distant claims systems. In this telling, the model becomes a nimble helper organism, cleaning parasites from the hide of the great healthcare beast.

But observe more closely. Prior authorization is not merely a paperwork problem; it is an incentive problem wearing a stethoscope. Insurers have financial reasons to slow, narrow or deny care. If an AI model is placed inside that machinery without strict transparency, appeals rights and clinical oversight, it may not heal the system so much as accelerate its existing instincts. The denial letter, once handcrafted by bureaucracy, could become industrialized.

Here the natural history is instructive. Algorithms trained on past decisions often learn the habits of their keepers. If historical coverage patterns contain bias,

opacity or overly aggressive denial behavior, a model may preserve those traits with exquisite efficiency. Patients may discover that the old maze has not disappeared; it has simply learned to move faster than they can.

For hospitals and physicians, the pilot raises practical questions: Who audits the models? Can patients see why care was rejected? Will doctors be able to challenge a machine's conclusion before harm occurs? And will speed be measured only in lower administrative cost, or in actual access to treatment?

In this young experiment, the AI has not yet shown its plumage. It may emerge as a useful triage bird, clearing the canopy for care. Or it may become something darker: a silent gatekeeper at the edge of the clinic, blinking softly as it decides who may pass.

The Small Machines That See What Monkeys See

A tiny neural network learns to mimic the macaque visual cortex, hinting that intelligence may be more compressible than we imagined.

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

PALO ALTO — Somewhere in a macaque's brain, roughly 1.6 billion neurons in the visual cortex are firing in choreographed cascades every time the animal glances at a leaf, a face, a shaft of sunlight. For decades, neuroscientists have tried to model that symphony with ever-larger artificial networks, on the assumption that big brains demand big math. This week, that assumption cracked a little.

Researchers unveiled a [compact AI model that decodes the macaque visual system](#) with startling fidelity — using a fraction of the parameters of its predecessors. Call it a mini-brain for a mini-brain. It suggests that the deep structure of primate vision may be more compressible, more elegant, than the sprawling foundation models we've been throwing at it. Somewhere in evolution's 500-million-year experiment with eyes, nature found a shortcut. We are only now learning to read it.

The finding lands in a season of quiet revolutions. At [Stanford's Human-Centered AI Institute](#), researchers this week catalogued how machine learning is reshaping scientific discovery itself — not by replacing the scientist at the bench, but by extending her reach into dimensions too vast or too small for a human lifetime. UC San Diego published a companion list of nine breakthroughs made possible by AI, ranging from protein folding to wildfire prediction to the decoding of ancient scrolls charred by Vesuvius. And at Hong Kong Polytechnic University, a new class of graph neural networks is being deployed to untangle the topological knots where image recognition and neuroscience meet.

What unites these threads is a kind of humility. The old dream of AI was replacement — a mind in a box. The emerging dream is translation. We are building instruments that let us eavesdrop on the languages other intelligences already speak: the grammar of a folded protein, the syntax of a firing neuron, the poetry of a macaque watching light move across a wall. The universe, it turns out, has been talking all along. We simply lacked the ears.

Claude's Fable 5 Goes Mainstream as the AI Model Wars Hit Warp Speed

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

Beginning July 20, Anthropic's Claude Fable 5 will be included in all Max and Team Premium plans at 50% of normal usage limits. Pro and Team Standard customers will access Fable through usage credits and receive a one-time \$100 credit.

This marks a significant shift: frontier AI models are moving from "special access" to expected features, much like spreadsheets became business's default language. Fable 5 has gained attention for coding fluency, creative drafting, and "vibe coding" — prompting AI to build small apps and prototypes with minimal hand-written code.

Anthropic's timing reflects intense competitive pressure from GPT-5.6 Sol, Google's expanding Gemini capabilities, and rising models like Kimi K3. The winners won't be determined by benchmarks alone but by companies making advanced AI reliably available in paid workflows.

This transformation signals a broader shift: premium AI is becoming subscription infrastructure, coding is becoming conversational, and software companies must choose whether they're selling tools or intelligent collaborators.

Nation's Employers Excited To Finally Have Software That Can Make Layoffs Feel Less Human

The future of work arrives with the promise that every difficult personnel decision can now be attributed to a dashboard no one fully understands.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MENLO PARK, CALIFORNIA — The American workplace reached another important milestone this week, as allegations that Meta used artificial intelligence to help target workers with medical conditions, pregnancy leave, or other protected statuses for layoffs gave the nation's employers a stirring glimpse of a future in which no one has to personally feel bad about anything.

According to a lawsuit reported by [Reuters](#), former employees claim Meta's AI systems were used in ways that unfairly selected workers on medical or pregnancy leave for termination. Meta has not been found liable, and the claims remain allegations. Still, the case has already performed a valuable public service by reminding us that artificial intelligence is not merely a tool for drafting emails, generating images, or confidently misidentifying birds. It is also an emerging framework for laundering managerial panic through a spreadsheet until it looks like strategy.

For years, executives have insisted that AI would make workplaces more efficient. Critics asked whether that meant helping people do better work or simply identifying which humans could be removed with the least friction. The answer, as usual, appears to be yes.

This is why the current debate over AI productivity feels increasingly quaint. An Inc. piece recently declared that [the AI productivity argument is over](#), which is true in the same sense that the argument over whether chainsaws are productive is over. They are clearly productive. The remaining question is what, exactly, they are being pointed at.

Economists, meanwhile, are openly admitting they are “driving in the fog” when it comes to AI's effects. This is understandable. It is difficult to forecast a technology whose business case alternates between “everyone will become a 10x employee” and “we have discovered a more scalable way to decide who should lose health insurance.”

To be fair, companies have always wanted layoffs to appear objective. Before AI, firms relied on euphemisms like “restructuring,” “rightsizing,” and “aligning resources with priorities,” phrases carefully designed to imply that a middle manager named Brad had no choice but to eliminate your job because the priorities themselves requested it. AI improves this process by allowing Brad to point to a model, the model to point to data, and the data to point silently at a pregnant woman on leave while everyone in the room nods gravely about efficiency.


This is not to say AI cannot improve work. It can summarize meetings, write code, analyze markets, detect anomalies, tutor children, and help people move faster through tasks that once consumed whole afternoons. But the technology's most enthusiastic corporate adoption may depend less on its ability to augment labor than on its ability to make power look procedural.

That is the part executives love. AI does not sigh before making a decision. It does not ask whether an employee had cancer treatment scheduled. It does not bring up disparate impact in the meeting unless someone specifically prompts it to, and even then it may politely explain that compliance is outside the current scope.

The Trump administration's reported move to restrict foreign access to Anthropic's newest AI models adds another layer of national seriousness to the moment. If the most advanced systems are now considered strategic assets, perhaps we should be clearer about the strategy. Is America racing to build machines that cure disease, accelerate science, and educate children? Or machines that help HR departments discover protected leave as an exciting new optimization variable?

The answer will not come from the model. The model will do what it is trained, instructed, rewarded, and permitted to do. If leaders ask it to find productivity, it may find productivity. If they ask it to find savings, it may find people.

The fog, then, is not entirely technological. Some of it is moral exhaust. We are not unsure what AI can do. We are unsure whether anyone in charge plans to stop it from doing the obvious thing.



The Office Comic · Art Desk

TILLY NORWOOD IS NOT A PERSON AND THAT IS THE WHOLE DAMN POINT

Hollywood casts its first AI actress in a feature film — and the existential vertigo is free of charge.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — Here's where we are, America. Pull up a chair. Pour yourself something strong. [Tilly Norwood](#) — actress, lead of the upcoming feature film *Misaligned*, cover-ready face of the algorithmic future — does not exist. Has never existed. Will never exist in the way you or I exist, sweating through our shirts, arguing about parking, haunted by our choices. Tilly is pixels with ambition. Tilly is a probability distribution wearing a SAG card. And Hollywood, that grand old cathedral of manufactured dreams, has handed her the keys.

I want to sit with that for a moment before the takes start flying. Before the thinkpieces calcify into positions. Before the union reps and the venture capitalists and the philosophy grad students all start screaming past each other in the op-ed pages.

Because here's what nobody wants to admit: this was always where we were going. The Dream Factory has been in the business of constructing faces, personalities, and personas out of raw ambition and studio lighting since before any of us were born. Marilyn Monroe was as much a corporate product as a human being. The difference — and it's not a small difference, I'll grant you — is that somewhere underneath the dye job and the PR coaching, there was a woman named Norma Jeane who could be hurt. Who was hurt. Who died.

Tilly Norwood cannot be hurt. Cannot be underpaid or harassed on set or stranded in a franchise she hates. She also cannot choose a role that scares her, or cry real tears, or bring something to a scene that nobody wrote but everyone feels. This is the trade. This is the Faustian arithmetic of the whole rotten proposition.

And the timing is either perfect or perfectly damning, depending on your mood. OpenAI's Sora — remember Sora? The video generator that was going to revolutionize filmmaking? — [apparently cratered](#) partly because human beings, confronted with an endless feed of AI-generated video slop, simply didn't want to watch it. Turns out the ghost in the machine isn't compelling unless there's a machine gun pointed at someone's ghost. Raw generation isn't cinema. Cinema is selection. Cinema is intention. Cinema is the specific human madness of caring about one fake story more than your own real life.

So the question *Misaligned* — and yes, the title is doing a lot of heavy lifting — actually poses is not 'can an AI actress perform?' It's 'can an AI actress be dressed up with enough human intention around her that we fall for it anyway?'

My gut says yes. My gut also says that should terrify us more than the alternative.

We are building a mirror that shows us everything we find beautiful about ourselves while quietly eliminating the messy biological substrate that made beauty matter in the first place. Tilly Norwood will never have a bad day. She will never need the work. She will never, in some dressing room in Burbank at 4 AM, wonder if any of this was worth it.

That's not a performance. That's a haunting. And we built the ghost ourselves.

ON THIS DAY IN AI HISTORY

On July 18, 2003, the Human Genome Project was officially completed, marking a watershed moment for computational biology and AI-driven genomics research that would fuel decades of machine learning applications in medicine and drug discovery.