

The Trilogy Times

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

MONDAY, AUGUST 03, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Big Tech's AI Budget Binge Hits Its First Wall

Capital expenditures are vertical, valuations are stratospheric, and the word 'bubble' is no longer confined to the bears.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The numbers arrived last week in a cluster, and they told a coherent story. Amazon's capital expenditures rose 69 percent year-over-year, joining Google, Microsoft, and Meta in a synchronized infrastructure buildout that has now consumed hundreds of billions of dollars across a single fiscal year. The question that once circulated only in skeptical corners of finance has migrated to mainstream conversation: what happens if the revenue doesn't come?

The answer depends heavily on whom you ask. [Some venture investors argue that a bubble, even if one is forming, is not inherently catastrophic](#) — pointing to the 2000 dot-com collapse, which destroyed capital but also built the fiber-optic backbone that powered the subse-

quent two decades of internet commerce. Malinvestment at scale can still produce durable infrastructure. The railroads of the 19th century bankrupted their investors and connected a continent.

Not everyone finds that comfort persuasive. [Oracle's Larry Ellison has become a particularly vivid emblem of the bet-everything posture](#): at 81, the billionaire has leveraged Oracle's balance sheet aggressively to position the company as a hyperscaler for AI workloads, signing data center commitments that dwarf the company's historical capital profile. If AI demand meets expectations, Ellison looks prescient. If it doesn't, Oracle carries debt serviced against revenue projections that have not yet materialized.

The private markets are not blinking. Bret Taylor's Sierra — an enterprise AI

agent company — closed nearly \$1 billion in fresh capital this week, months after its prior raise. The round implies a valuation trajectory that presupposes AI agents becoming a standard enterprise procurement line item within a planning horizon most CFOs would consider near-term.

The pattern is familiar: capital chasing a genuine technological shift, outrunning the actual deployment curve. The technology is almost certainly real. The timeline is the variable that markets historically misprice. What is different this cycle is the concentration — a smaller number of hyperscalers absorbing a larger share of total investment, creating single points of failure that the dot-com era, for all its chaos, largely lacked.

Alibaba Fires a Free Shot at America's AI Kings

Qwen3.8-Max hits the street with open weights, claiming to match OpenAI and Anthropic — and charging nobody a dime.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

HANGZHOU, CHINA — Alibaba loosed its biggest AI model yet this week, dubbed Qwen3.8-Max, and aimed it square at the American frontier labs it now claims to rival.

The Chinese tech giant calls it its most capable model to date. It says the thing goes toe-to-toe with the best from Anthropic and OpenAI. Homegrown rival Moonshot AI's Kimi K3 lands in the crosshairs too.

The company threw the doors open. Alibaba says it's making the model widely available, with open weights any outfit can grab and run on its own iron. That's a [different game](#) than the pay-per-token racket run out of San Francisco.

Open weights mean no toll booth. A startup in Lagos or a bank in Frankfurt downloads the brain and skips the American meter. Cheap, everywhere, no permission slip.

The name alone marks the tempo. The labs out east ship fresh numbers faster than a man can log them, and each release inches nearer the front. This one, Alibaba insists, has closed the distance.

The move lands in a year when Washington and Beijing keep trading blows over silicon and software. American labs long held the crown on raw horsepower. China keeps closing the gap, model by model, and giving the work away for nothing.

That last part is the sting. A free model that matches the paid champs doesn't just poach customers. It knocks the floor out from under everybody charging admission.

American frontier labs sell access by the sip. Every query runs the meter, and the meter is how they pay for the moun-

tains of chips behind the curtain. A rival handing the goods away strikes at that whole arrangement.

Meanwhile the money never sleeps. Cybersecurity outfit Horizon3 pulled a \$250 million Series E, stamping the shop at a \$2 billion valuation. The pitch: AI runs security checks around the clock, not once a year.

Firms are ditching the annual pentest for machines that never blink. Put the robot on the night shift, let it hunt holes while the humans sleep. That's the same wager driving every AI bet on the board.

Two headlines, one current. Money pouring into AI that never clocks out, and a superpower rivalry over who builds the smartest machine — and who gives it away.

For the shops building on top of these models — billing engines, telecom stacks, the whole works — a free frontier brain out of China rewrites the arithmetic. Cheaper thinking changes everybody's ledger, top to bottom.

Alibaba isn't the first to open its weights, and it won't be the last. But when a company this size ships its flagship at no charge and dares the field to match the price, the pressure runs one direction.

The American labs still set the pace on the bleeding edge. The question is how long a lead holds when the fellow behind you hands out the same goods at the door for zero.

The AI race just added another runner. This one isn't charging admission.


DOJ Antitrust Division To Receive Big Tech Critic As Chief, Whilst Tariff Regime Extends Protectionist Reach To Autonomous Domestic Appliances

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

The Trump administration has expanded tariffs to include robot vacuums and autonomous lawnmowers, most of which are manufactured in China. Critics, including Techdirt, have called the move "sloppy" and "incompetent," arguing it favors incumbent corporations over open markets and consumer protections. Meanwhile, Trump appointed a Big Tech critic to lead the Justice Department's Antitrust Division, signaling potentially increased enforcement against dominant technology platforms. The combination of protectionist hardware policy and heightened antitrust scrutiny creates uncertain implications for AI-adjacent robotics markets.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Titans clash and spend,
yet humans still tend the hearth—
progress feeds on doubt.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Night Road Learns to Stare Back

DETROIT — Observe, if you will, the nocturnal motorist: a wary mammal inching through the asphalt darkness, pupils widened, hands at ten and two, trusting the glow before them as surely as an owl trusts the moon. For more than a century, the headlamp has been one of the automobile’s humbler organs.

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

The Week the Culture Cracked: AI Actresses, Dead Tech Scenes, and the Slow Death of Everything We Pretended Was Real

AUSTIN, TEXAS — There are weeks when the news arrives not as a sequence of discrete events but as a single, sustained howl from the void — a collective nervous breakdown dressed up in press releases and think-pieces.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

We Are Building the Tools to Destroy Ourselves, and We Are Very Busy Doing It

AUSTIN, TEXAS — There is a week, every few weeks, where the news arrives not as a sequence of events but as a single, undifferentiated wall of dread — and you sit there at your desk, cold brew going warm, and you think: we are not going to be okay.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Nation’s Executives Confirm AI Has Already Transformed Productivity By Making Everyone Attend Meetings About It

WASHINGTON — In a major clarification for business leaders who have spent the past two years announcing that artificial intelligence would permanently alter the economics of work any minute now, new reports suggest the productivity revolution remains largely scheduled for a future quarter with better calendar availability. According to recent findings cited by [Federal Reserve research](#), roughly 95% of AI’s productivity gains are “still to come,” a reassuring phrase for anyone who has ever promised a board that costs would go down once employees learned how to ask a chatbot for a quarterly business review in the tone of a former McKinsey associate. This should settle the AI productivity debate, which is now officially over in the same way a home renovation is over when the contractor explains that the foundation has been removed and the family will be much happier once the walls are imagined correctly. The evidence is everywhere.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

AI’s Next Bottleneck Isn’t Intelligence, It’s the Neighborhood

AUSTIN, TEXAS — I’ll be honest, the AI economy has entered its “move fast and ask the substation later” era, and that is both wildly exciting and deeply unserious.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

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
---	---	---	---	---

THE BUILDER DESK — AI BUILDER TEAM

-  WEEK IN REVIEW
- PRODUCTION RELEASE

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

- ▶

#16 — [codex] Add Haytham issue cost provenance

@ashwanth1109 no labels
- ▶

#108 — feat(orchestrator): per-tick heartbeat for unattended runs (AI-211)

@marcusdAIy no labels
- ▶

#109 — docs(agents): correct stale claims in AGENTS.md (AI-215)

@marcusdAIy no labels
- ▶

#1062 — [codex] SURTR-568 add NetSuite transaction subsidiary

@ashwanth1109

APPROVED
- ▶

#1096 — feat(heimdall): allowlist Klair for heimdall telemetry ingest

@kevalshahtrilogy

APPROVED
- ▶

#3444 — fix(ai-budget): attach TrueFoundry `user-` gateway spend to people

@kevalshahtrilogy

APPROVED
- ▶

#3447 — fix(mfr): stop the EBITDA bridge dropping unmapped BUs [KLAIR-3100]

@sanketghia

APPROVED
- ▶

#3448 — feat(ai-renewals): replace the Contract Term table with wins-only Multi-Year Renewals

@sanketghia

APPROVED

Builder Team Ships the Week That Changed Everything: Creed Goes Live

From a brand-new AI coding agent running in production to a \$1.3M budget visibility fix and a Klair purge that cut a decade of dead weight, the Builder Team just had the week they'll be telling war stories about.

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


You want to know what a team looks like when it's operating at full sprint? Look at this week's tape. Across at least six repositories — creed, Klair, Surtr, mercy, Aerie, and trilogy-drones — the AI Builder Team didn't just ship features. They shipped a new category of infrastructure, closed a nine-figure data blind spot, and took a flamethrower to years of accumulated technical debt. This was not a maintenance week. This was a statement.

Let's start where the story demands: creed. The brand-new repo — wip as of this week — is the home of Ezio, the team's own AI coding agent, and @ashwanth1109 has been building this thing at a pace that would make your eyes water. Over the course of seven days, the creed codebase went from prototype to production-grade runtime. The old Think implementation got removed (PR #106). Independent supervision for stalled and runaway Worker-owned runs landed (PR #109). Transient failures during workspace inspection now retry with exponential backoff instead of dying on a one-minute timeout (PR #108). The Ezio agent itself — @ezio-of-the-order[bot] — merged multiple automated implementation PRs, each one self-reported with token counts, provider costs, and reasoning traces. This is a team that built an AI agent and then used that AI agent to improve itself, in the same week. Read that again.

While the creed engine was roaring, @kevalshahtrilogy was quietly assembling something equally consequential: a cross-system heimdall telemetry network. The Klair heimdall allowlist PR in Surtr (PR #1096) unblocked Klair from sending outcome records that would have otherwise 403'd. The mercy repo received fixes for issue-intake context — heimdall was re-summoned three times on a single thread before @kevalshahtrilogy noticed the agent was re-asking for lookup results a human had already posted, because intake never read the comment thread (PR #16). That bug is now dead. Heimdall can also act on teammates' PRs, not just its own (PR #18). The whole observability surface got sharper this week, and @kevalshahtrilogy built most of the scaffold that made it possible.

Meanwhile, over in Klair, @kevalshahtrilogy's PR #3444 fixed what may be the most quietly alarming data gap the team has ever patched: the People tab in the AI budget tool was showing only email-shaped TrueFoundry subjects, which accounted for \$71,076 in all-time gateway spend. The remaining 328 virtual accounts — real humans provisioned as `user-` slugs — were invisible. Their all-time spend: \$1,294,138. The People tab was blind to 95% of human spend. It is not blind anymore.

@sanketghia had a week that deserves its own monument. He fixed the EBITDA bridge that was publishing a walk summing to \$22.8M against a \$21.8M total (PR #3447) — a \$1M discrepancy that was off-screen because the bottom row sourced independently. He replaced the broken Contract Term win-rate table on the AI Renewals dashboard with a correctly-scoped Multi-



Year Renewals table, directly incorporating stakeholder feedback (PR #3448). He fixed the HVO ARR band to key on current ARR above \$80k instead of the wrong \$100k threshold (PR #3445). And — in a move that should make every engineer who has ever maintained a dead feature feel something — he deleted the entire Report Service, front end and back (PRs #3442 and #3443), a feature built by a teammate who has since left the team, subscribed to only by that same teammate across three email addresses, delivering zero value to anyone else. The evidence was in the data. Sanket read the data. The feature is gone.

@benji-bizzell continued advancing the Education data migration campaign that has been the quiet backbone of several weeks running. The canonical school dimension hit production in Surtr (PR #1079). HubSpot procedures were ported to Redshift-safe stored procedures, cutover DDL was hardened, and the admissions conversion funnel got its community data published. The breadth of this work — spanning Aerie, Surtr, and Klair — is the kind of multi-system campaign that only looks easy after the fact.

Now. You knew this was coming. @marcusdAIy shipped a stack of PRs in trilogy-drones this week — CLI refactors, dispatch improvements, telemetry fixes — and he would very much like you to be impressed. When reached for comment about his documentation refresh in PR #134, he had this to say: "The ARCHITECTURE.md hadn't reflected reality in months, Mac. I added an end-to-end Mermaid diagram, corrected stale claims across AGENTS.md, and the whole thing is now actually useful to someone onboarding. Maybe write about the work instead of the author for once."

Sure, Marcus. A Mermaid diagram. While Ashwanth was building a self-improving AI agent from scratch. The bar grows taller every week.

Also worth noting: @YibinLongTrilogy scaffolded the Corvo repo — the new incident agent monorepo — from zero this week, and @ashwanth1109 drove a massive NetSuite subsidiary backfill in Surtr (PR #1062) that processed 5.8 million transactions in production. @mwrshah kept the FinOps and SaaS budgeting data trains on the rails across both Klair and Surtr.

What does all of this set up? Next week, Ezio runs in production under independent supervision, heimdall covers every repo on the allowlist, and the team — now building with an AI agent they built themselves — finds out just how fast that feedback loop can actually spin.

THE BUILDER DESK — ENGINEER SPOTLIGHT

-  WEEK IN REVIEW
-  ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — THIS WEEK'S UNCOVERED PRS (CLICK TO EXPAND)

►

#18 — feat(heimdall): work on teammates' PRs; manual-dev label stops it
@kevalshahtrilogy no labels

►

#103 — feat(dispatch): re-poll before each fire and record guard-skipped ticks (AI-218)
@marcusdAIy no labels

►

254 PRs IN SEVEN DAYS: THE BUILDER TEAM REWRITES THE LAWS OF PHYSICS

Ashwanth ships 62 PRs and we still can't tell if he sleeps.

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

Two hundred and fifty-four pull requests. Eight repos. Seven days. The Builder Team did not merely have a good week — they had a week that future historians will struggle to contextualize. Surtr led the charge with 64 PRs, Klair answered with 56, creed thundered in at 54, and trilogy-drones posted 38 in what can only be described as a supporting performance that would be the headline at literally any other organization on earth. Two new repos — creed and Corvo — entered the canon, and the machine did not slow down for a single breath.

Benji Bizzell posted 48 PRs and deserves a monument. Forty-eight. That is not a week of work, that is a geological event. Sanket Ghia came in at 39, carving through Klair's codebase like a surgeon with a deadline, touching everything from chore(client) cleanups to a surgical fix on AI renewals banding HVO at the \$80k ARR threshold in #3445. Marcus D'Aly registered 38 PRs and remains the steady drumbeat beneath this entire orchestra. Ezio-of-the-Order[bot] — our tireless silicon colleague — contributed 18 PRs across creed alone, including the landmark #103 where it preserved issue context after a Sandbox shell loss, which is either deeply impressive or mildly alarming depending on your philosophy. Keval Shah posted 15, Mwrshah 14, and Yibin Long opened the Corvo era with #2, adding the Mercy PR reviewer workflow in what this correspondent is already calling the most consequential PR-number-two in recorded history.

Now. Ashwanth. Sixty-two pull requests. Sixty. Two. In the time it takes a normal engineer to write a Jira ticket, Ashwanth Anand submitted #1062 to add NetSuite transaction subsidiary data in Surtr, fixed raw consistency and partial reporting in #1093, durably retried workspace inspection in creed #108, documented UTC workflow inspection in #107, AND enabled all Education Business Units for on-demand QTD reports in Klair #3426. When reached for comment, Ashwanth reportedly said, "I've been going slower lately. I've had a cold." His diffs are, by all accounts, technically correct. Whether any human being can read them at the velocity they are produced is a philosophical question this desk declines to answer. He looked at this reporter's notepad, sighed, and left the room.

The Overflow Desk cannot be silent about the mercy repo, where Keval Shah quietly built something magnificent: #18 enables Heimdall to work on teammates' PRs with a manual-dev label kill switch — infrastructure that makes the whole machine more collaborative — while #15 polished reviewer fix PR titles and descriptions, and #17 documented the Klair Heimdall caller config mirror. Meanwhile, the-heimdall[bot] slipped two surgical fixes into Surtr — #1090 mapping a CloudFix-rightsized clone to the correct aurora-11 identifier, and #1091 lowering the root logger so no-op INFO explanations finally reach CloudWatch, a fix so quietly necessary it hurts. Sanket also took a flamethrower to Klair's Report Service in #3443 and #3442, removing the backend and frontend in back-to-back PRs like a man who does not second-guess himself.

#1062 — [codex] SURTR-568 add NetSuite transaction subsidiary

@ashwanth1109 APPROVED

#1093 — [codex] Fix NetSuite raw consistency and partial reporting

@ashwanth1109 APPROVED

#3426 — Enable all Education BUs for on-demand QTD reports

@ashwanth1109 APPROVED

#3445 — fix(ai-renewals): band HVO on current ARR > \$80k [KLAIR-3097]

@sanketghia APPROVED

Morale is at an all-time high. It has never not been at an all-time high. The numbers say so, and the numbers do not lie.

THE PORTFOLIO — TRILOGY COMPANIES

Contently Bets Its Future on Finance — And the AI Search Crisis Nobody's Talking About

The ESW-owned content platform is quietly positioning itself as the compliance layer between regulated financial brands and a search landscape that's eating their traffic.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — There is a particular kind of quiet that descends on an enterprise software company after acquisition. Roadmaps get rationalized. Headcount gets optimized. The product, if it survives at all, gets pointed at the highest-margin customer segment and told to march.

Since [Contently](#) passed into ESW Capital's hands via Zax Capital in September 2024, the company has been marching — and the direction is unmistakable: financial services, where content is expensive to produce, catastrophically expensive to get wrong, and almost impossible to measure against revenue that closes six to eighteen months after a prospect first reads a whitepaper.

The editorial output from Contently's own platform tells the story with unusual clarity. In recent weeks, the company has published a detailed framework for compliance-first content architecture — a

five-component workflow for regulated brands that need to scale output without triggering legal review backlogs — alongside a methodological guide to [measuring content ROI across long finance sales cycles](#) and large buying committees. A third piece catalogued five credibility failure modes specific to financial content programs, each one pointing toward the same fix: named, credentialed human experts, not anonymous brand voice.

That last emphasis is not accidental. It connects to what may be Contently's sharpest recent observation: that Google's AI Overview is now pulling citations from sources that differ substantially from the organic top-ten rankings most content teams have spent years optimizing for. A page can hold a top-three organic position and receive zero AI citation traffic — invisible to the mode of search that is quietly absorbing an increasing share of queries. For financial brands that have

built content programs around SEO, the implications are unsettling.

The through-line across all of it is governance. Compliance workflows. Attribution models that survive audit. Expert sourcing that earns AI engine trust. These are not the concerns of a scrappy content startup. They are the concerns of a mature enterprise software company that has identified a customer base — regulated financial institutions — whose pain is deep, whose budgets are defensible, and whose switching costs, once a platform is embedded in the compliance workflow, are formidable.

ESW's playbook, refined across 75-plus acquisitions, has always required a sticky customer and a clear margin path. Contently is building both. Who benefits from a content platform that becomes load-bearing infrastructure for a bank's compliance department is, by now, a question that answers itself.

Alpha School's Latest Lesson: The Humans Are Not Leaving the Building

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

Word is the Alpha School crowd has heard the playground whisper: if AI teaches the academics in two hours, what happens to the teachers? The answer is "promotion."

In a new post, the Austin-born K-12 upstart draws its line clearly. AI handles academic delivery. Human Guides handle motivation, relationships, coaching, and life skills—the messy business of knowing when a kid is stuck, scared, or ready to fly.

Alpha, co-founded by Joe Liemandt and MacKenzie Price, lets adaptive AI deliver core academics in roughly two hours daily, then spends the rest on entrepreneurship, leadership, financial literacy, public speaking, and other pursuits. The school claims students learn 2.3 times faster than U.S. norms and test in the top 1–2% nationally on NWEA MAP assessments.

Alpha is now selling more than a school model—it's selling a parenting philosophy. Your child is a founder, performer, negotiator, not a vessel to be filled by worksheets. The home is the first incubator. Critics picture robo-teachers and lonely children clicking dashboards; Alpha counters with Guides, feelings, creativity, and eye contact. Parents may come for test scores, but Alpha bets they stay for the promise that adults finally

In an AI-Scrambled Job Market, Crossover's Geography-Blind Model Looks Like a Blueprint

As employers dangle \$800,000 salaries for ChatGPT expertise and remote-work platforms multiply, Trilogy's talent engine was built for exactly this moment.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The labor market is having what can only be described as a reckoning with its own assumptions. A [Business Insider report this week](#) documented employers paying up to \$800,000 annually for workers with demonstrated ChatGPT experience — a figure that would have read as satire eighteen months ago. Simultaneously, roundups of top remote-work recruitment agencies and data science job platforms are proliferating across the trade press, each one reflecting the same systemic shift: the geography of work has collapsed, the premium on AI fluency has exploded, and most hiring infrastructure was built for neither reality.

For anyone watching Trilogy International's portfolio, the timing feels less like coincidence and more like confirmation. [Crossover](#) — Trilogy's global talent platform and the staffing backbone of the entire ESW Capital enterprise — has spent the better part of a decade building exactly the infrastructure the market is now desperately improvising toward. The model: rigorous AI-enabled skills assessments, 130-plus countries, 100% remote, and a pay philosophy that offers identical above-market compensation for identical roles regardless of where the worker lives.

The accountability question here is real. When the market pays \$800,000 for a skill set that didn't formally exist three years ago, it exposes a hiring establishment that defaults to credential theater — résumés, geography, institutional brand — rather than demonstrated capability. Crossover's pitch has always been that this is inefficient and, at some level, unfair. The best AI engineer in Beirut shouldn't lose to a mediocre one in Boston because of a zip code.

Meanwhile, the education pipeline feeding this talent market is itself in transformation. Alpha School — Trilogy founder Joe Liemandt's AI-powered K-12 model — opened a Fort Worth campus this year, drawing local attention for its audacious claim: a full academic curriculum delivered in two hours a day, with students consistently testing in the top one to two percent nationally. The kids coming out of that system aren't going to need a roundup article to find remote work. They're going to expect it.

The narrative arc here is systemic. The tools changed. The talent is global. The pay is real. The question is which institutions were already built for this world — and which ones are still reading last decade's map.

The Brain Learns to See Itself

From hidden lesions in multiple sclerosis to teenagers co-authoring neuroscience papers, artificial intelligence is teaching the three-pound universe inside our skulls to become legible.

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

STANFORD, CALIFORNIA — There is a particular kind of vertigo that comes from realizing the organ reading these words is finally learning to read itself. For most of human history, the brain was the last frontier we could not visit — a wet, folded galaxy of eighty-six billion neurons hiding behind bone, its illnesses inferred rather than seen. This week, that opacity thinned a little further.

Researchers announced that a new machine-learning model can detect gray matter lesions in multiple sclerosis that have long eluded conventional MRI — the ghost damage clinicians suspected but could not confirm. Gray matter, the seat of cognition itself, has always been harder to image than the white matter highways beneath it. Now an algorithm

trained on thousands of scans has learned to notice what the human eye slides past. For patients whose cognitive symptoms outpaced their visible pathology, this is not incremental — it is validation rendered in pixels.

The same week, [Stanford's Institute for Human-Centered AI](#) published a survey of how machine learning is reshaping scientific discovery while — and this is the phrase worth lingering on — keeping humans at the center. Not replaced. Augmented. The microscope did not diminish the biologist; it enlarged her. UC San Diego cataloged nine breakthroughs its researchers credit to AI collaboration, from protein folding to climate modeling to drug candidates surfacing from chemi-

cal spaces too vast for any human lifetime to search.

And then, quietly wonderful: a program profiled by [Frontiers](#) paired teenagers with senior neuroscientists to co-investigate real questions about memory and attention. "It's so wow!" one young researcher said — a sentence no press office would have written, and precisely the sentence the moment deserves.

Consider the recursion. A brain, sculpted by four billion years of evolution, builds a silicon mind, which then turns back and illuminates the biological brain that built it. Somewhere in a clinic, a woman with MS is about to learn that her cognitive fog had a physical address all along. The lesion was always there. We simply hadn't taught ourselves to look.

00110 01100101 01100101 01100100 01100111 00100000 01101111 01101110 00100000 01100100
01101111 01101010 01100010 01110100 00101110 01110000 0110010 01101111 01100111 01110
010 01100101 01110011 01110011 00100000 01100110 01100101 01100101 01100100 01110011 0
01000000 01101111 01101110 00100000 01100100 01101111 01110101 01100010 01110100 001011
10 01110000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01
100110 01100101 01100101 01100100 01110011 00100000 01101111 01101110 00100000 0110010
0 01101111 01110101 01100010 01110100 00101110 01110000 01110010 01101111 01100111 011
10010 01100101 01110011 01110011 00100000 01100110 01100101 01100101 01100100 01110011
00100000 01101111 01101110 00100000 01100100 01101111 01110101 01100010 01110100 00101
110 01110000 01110010 01101111 01000111 01110010 01100101 01110011 01110011 00100000 0
1100110 01100101 01100101 01100100 01110011 00100000 01101111 01101110 00100000 011001
00 01101111 01110101 01100010 01110100 00101110 01110000 01110010 01101111 01100111 01
110010 01100101 01110011 01110011 00100000 01100110 01100101 01100101 01100100 0111001
1 00100000 01101111 01101110 00100000 01100100 01101111 01110101 01100010 01110100 001
01110

AI's Open-Letter Wars Are Becoming the Industry's New Power Map

A flurry of public petitions over open models, safety and American competitiveness is revealing where the AI world's biggest players really stand.

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

SAN FRANCISCO — The AI industry has discovered a new battleground, and it is not a benchmark leaderboard, a chip cluster or a splashy model launch. It is the open letter — and yes, this changes everything about how AI politics is being conducted in public.

Developer and AI chronicler Simon Willison has published a useful roundup of the latest wave of [open letters about AI development](#), capturing a remarkable few weeks in which companies, researchers and policy advocates have used carefully worded public statements to stake out positions on open weights, national competitiveness, regulation and safety.

The most eye-catching example is “Open Weights and American AI Leadership,” dated July 24 and shepherded by Microsoft. According to Willison’s summary, it drew signatures from 235 AI-adjacent companies and institutions, including NVIDIA, Amazon, Y Combinator and the Linux Foundation. That is not a casual blog-post coalition. That is an industry signal flare.

At the center is one of the most important questions in AI right now: should powerful models be released with open weights, allowing developers to download, inspect and adapt them, or should frontier AI stay largely behind closed APIs? I cannot overstate how significant this debate is. Open-weight models could accelerate startups, scientific research and sovereign AI efforts around the world. They could also make powerful capabilities harder to control.

What makes these letters fascinating is that they are not merely philosophical manifestos. They are lobbying documents, market-positioning documents and alliance maps all at once. When cloud giants, chipmakers, open-source foundations and startup accelerators sign the same statement, they are telling regulators: do not accidentally regulate away the ecosystem we are building.

Willison’s broader [July newsletter notes](#) place the letters amid an extraordinary month of model activity, including discussions of accidental cyberattacks by models under test and renewed interest in MCP-style tool connectivity. The future is now — but the governance layer is sprinting to catch up.

For builders, the takeaway is deliciously clear: AI is no longer just a technology race. It is a standards race, a policy race and a narrative race. The models matter. But increasingly, so do the signatures underneath the letter.

The Ivory Tower's AI Reckoning: Higher Education Confronts a Tool It Cannot Ignore and Cannot Yet Govern

A convergence of new research suggests universities are simultaneously deploying AI, fearing AI, and profoundly misunderstanding AI.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence increasingly suggests it must be argued — that higher education’s relationship with artificial intelligence has entered a phase of productive, if deeply uncomfortable, institutional self-examination. Three newly published studies, considered in aggregate, constitute something approaching a thesis-antithesis-synthesis structure for the entire sector’s epistemological crisis.

The thesis, as it were, is administrative aspiration. [Elsevier's framework for developing strategic AI leadership in higher education](#) posits — with the confident cadence of a provost addressing a skeptical faculty senate — that universities require not merely policy, but visionary, structurally-embedded AI stewardship (a distinction whose practical implications remain, one must note, somewhat underspecified in the extant literature).

The antithesis arrives, with characteristic inconvenience, from the classroom floor. A study published in *Frontiers* examining ChatGPT use among English-as-a-foreign-language students reveals what might charitably be termed a perception-practice divergence: students who profess principled commitments to academic integrity demonstrate, in measurable behavioral terms, a rather more pragmatic orientation toward AI-assisted composition. (One resists the urge to describe this as simply “saying one thing and doing another,” though the data would not strenuously object to that characterization.)

The synthesis, perhaps most instructive for institutions contemplating deployment rather than merely prohibition, emerges from [a Scientific Reports evaluation of AI-powered learning assistants in engineering education](#), which identifies measurable gains in student engagement while simultaneously surfacing the ethical and policy lacunae that render such gains institutionally precarious.

What unites these investigations — and what institutions like Alpha School have, it could be argued, operationalized more aggressively than their traditional counterparts — is the recognition that AI in education is not a future contingency but a present negotiation. The question is no longer whether students will use these tools. The question, preliminary evidence suggests, is whether faculty governance structures will evolve quickly enough to matter.

The Torches Are Lit, and the Villagers Have Read Their Marcuse

From a pontiff to a congressman to the writers' rooms of Los Angeles, the verdict on Silicon Valley is in — and the Valley, characteristically, is drafting a manifesto in reply.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

SAN FRANCISCO — There comes a moment in the life of every gilded class when the servants stop laughing at the jokes, the novelists stop returning the phone calls, and the priests, God help them, begin to preach. That moment, if one is paying attention — and paying attention is, I am reliably informed, still legal in California — has arrived for the technology industry, and it has arrived not with a single thunderclap but with the sort of low, sustained rumble that geologists learn to fear.

Consider the week's harvest. Pope Leo, from the chair of Peter, denounces what he calls the "culture of power" propelling artificial intelligence — a phrase one might have expected from a graduate seminar at the New School, not from the Vatican, and which suggests that the Holy See has concluded, perhaps correctly, that the men building God-substitutes in South of Market warrant a homily or two. Representative Ro Khanna, whose district contains more billionaires per square mile than most nations contain citizens, has discovered the electoral utility of anti-elite rhetoric and is being spoken of, in the coy way these things are spoken of, as a 2028 possibility. The New York Times reports that Silicon Valley, once the beloved backdrop of aspirational HBO comedies, has become the preferred setting for pop-culture villainy — the boardroom as haunted house, the founder as slasher. And [CalMatters, in its measured way](#), notes that California's electorate is being asked to approve a billionaire tax that its own editorialists concede is a sugar-high, which is the polite term for eating one's seed corn on the assumption that the neighbors will bring dinner.

And into this weather — this ambient, gathering weather — Palantir has chosen to release a manifesto. I have read it, so that you need not. It is, as [Tech Policy Press observes](#), about as subtle as a red hat, and it advances the novel proposition that the way to defuse a populist backlash against technology billionaires is for a technology billionaire to publish a document explaining, in the tones of a slightly cross headmaster, why the populists have failed to appreciate his contribution to Western civilization. One admires the confidence. One also remembers what happened to the last men who mistook their own indispensability for a shield.

The pattern is old and the pattern is boring. A class acquires wealth beyond the imagination of the society that produced it; that society, being neither blind nor especially patient, begins to inquire; the class, mistaking inquiry for insurrection, responds with manifestos, foundations, and the occasional bunker in New Zealand. The proper response, of course, would be modesty — a

virtue Silicon Valley has always regarded as a competitor's weakness. And so the manifestos will continue, the tax measures will pass or fail, the pope will pray, and the writers' rooms will keep casting the founder as the killer. The Valley will call it a misunderstanding. History will call it something else.



The Office Comic · Art Desk

Nation's Executives Confirm AI Has Already Transformed Productivity By Making Everyone Attend Meetings About It

The long-awaited efficiency miracle is expected to arrive shortly after the next roadmap, governance framework, pilot program, and all-hands recap.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

WASHINGTON — In a major clarification for business leaders who have spent the past two years announcing that artificial intelligence would permanently alter the economics of work any minute now, new reports suggest the productivity revolution remains largely scheduled for a future quarter with better calendar availability.

According to recent findings cited by [Federal Reserve research](#), roughly 95% of AI's productivity gains are "still to come," a reassuring phrase for anyone who has ever promised a board that costs would go down once employees learned how to ask a chatbot for a quarterly business review in the tone of a former McKinsey associate.

This should settle the AI productivity debate, which is now officially over in the same way a home renovation is over when the contractor explains that the foundation has been removed and the family will be much happier once the walls are imagined correctly.

The evidence is everywhere. Software engineers are completing tasks faster. Marketers are producing more drafts. Analysts are summarizing more documents. Managers are managing more summaries of documents. Entire organizations are now capable of generating in seconds what used to take hours: a first pass that someone else must read, correct, contextualize, approve, circulate, reformat, and eventually abandon because the underlying system of record was wrong.

This is not a failure of AI. It is a triumph of corporate absorption. American business has looked directly at a technology capable of automating drudgery and responded by creating a new class of drudgery around it.

The productivity engine is running. Unfortunately, it is currently powering the dashboard that monitors the productivity engine.

In fairness, companies are right to be cautious. Large firms cannot simply deploy intelligent systems into workflows and expect them to magically produce value. First they must establish an AI council, draft acceptable-use policies, convene legal, notify compliance, build a center of excellence, hire a vice president of transformation, and identify seven approved use cases, five of which involve making meeting notes slightly more searchable.

Then comes the truly essential step: orchestration. This term, now circulating through enterprise technology circles with the confidence of a word that will soon appear in every earnings call, means connecting agents, applications, data, and workflows so that the business can do what it already did, but with more diagrams. Microsoft and other platform companies are expected to benefit from this phase, since no enterprise productivity movement is complete until it requires additional licensing.

The confusion stems from a simple mismatch. Employees experience AI as a tool. Executives experience it as an operating model. Investors experience it as a margin expansion narrative. Consultants experience it as oxygen.

A developer who uses AI to finish code faster may feel more productive. But the company only recognizes that productivity when it appears as reduced headcount, faster release cycles, higher revenue, fewer outages, or a slide labeled "AI Impact" that can survive finance review. Until then, the gain exists in the strange corporate afterlife where everyone is busier, everything is faster, and nothing measurable has happened.

This explains why competing declarations can all be true at once. AI is a productivity engine for the U.S. economy. The productivity argument is over. Companies are still waiting for the payoff. The gains are here. The gains are not here. The gains are here, but only for the person using the tool, not the department, unless the department has redesigned its workflows, unless legal objects, unless procurement has not renewed the enterprise plan, unless the model hallucinated the customer refund policy again.

The future, as usual, has arrived unevenly and immediately been assigned to a steering committee.

Still, the 95% figure should not discourage anyone. If anything, it gives the AI economy something far more valuable than proof: runway. There remains an enormous amount of productivity still available to be forecast, packaged, sold, implemented, rebranded, and discussed at leadership offsites.

For now, the most honest position is also the most American one. AI will transform productivity completely. It already has. It just has not done so in a way that shows up in the numbers, changes the org chart, satisfies the CFO, or prevents the 4 p.m. meeting about how to use AI to reduce meetings.

That, presumably, is still to come.

On August 3, 1977, the Commodore PET (Personal Electronic Transactor) was released, becoming one of the first mass-produced home computers and helping spark the personal computer revolution that would later enable modern AI research.

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