VOL. I · NO. 192 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

SATURDAY, JULY 11, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

The Axe and the Checkbook

Big Tech has shed nearly 165,000 jobs in AI's name this year — while the labs bid up the very professors who built the thing.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Tech companies have cut close to 165,000 jobs in 2026, and a rising share of the pink slips finger the same culprit: artificial intelligence.

Oracle, Meta, Microsoft and Samsung head the running tally, which spans dozens of firms and climbs by the week. The stated reasons run the same groove — efficiency, reprioritization, machines picking up the load.

Trade desks now keep a separate ledger of every 2026 layoff that name-checks AI out loud. That column keeps filling.

Microsoft ran its cuts in two moves. Buyouts first, layoffs second — the company floated voluntary exits, cash on the barrel, before the slips followed. The order matters: soften the blow, trim the volunteers, then swing.

Xbox took the hardest hit. The chief executive said the unit "didn't focus on the core business." Roughly one in five Xbox workers is headed out the door.

Here's where the ledger stops balancing.

While clerks and coders pack their desks, the AI labs are hurling money the opposite direction. Anthropic, OpenAI, Meta and DeepMind are raiding university faculty, luring professors off the tenure track and onto corporate payroll.

Deans are fretting. Their sharpest AI minds keep walking, and the checks the labs write are hard to match on a campus budget. Lose the professor, and you can lose the students who came to study under him.

So one industry pulls both levers at once. One hand sheds 165,000 bodies. The other bids up a thin band of brains it cannot mint fast enough.

That's the tell. The layoffs aren't a retreat from AI — they're a wager on it.

Firms are swapping broad headcount for narrow expertise. Fewer hands, costlier heads, and a bet that the software covers the gap.

The playbook isn't news down in Austin.

Trilogy International's Crossover has sold remote, "top 1%" talent at one flat, above-market wage across 130-plus countries for years. Same efficiency gospel Big Tech now preaches under duress.

ESW Capital runs 75-plus enterprise software firms on that lean logic, buying at one to two times revenue and squeezing out the slack. Alpha School teaches K-12 in two hours a day, AI tutors carrying the load while the payroll stays thin.

Where Big Tech now cuts to chase the model, Trilogy wired the whole machine around it first.

The count sits near 165,000 and rising. The scramble for professors is only warming up.

One industry, two headlines the same morning: help wanted, and nobody's safe.

The Last Decimal: How AI Is Redrawing Venture Capital's Return Curve

From Bezos' family office to corporate VCs, capital is concentrating fast — and a legal war between Apple and OpenAI signals the partnerships holding this boom together are under strain.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The numbers from the first half of 2026 are unambiguous: North American startup funding shattered records in H1 2026, driven almost entirely by AI. The surge is pulling in capital from every category of institutional investor — and concentrating it with unusual speed.

Jeff Bezos' family office has been systematically adding AI startup positions, according to reporting from Family Wealth Report. The move is consistent with a broader shift among ultra-high-net-worth family offices, which have spent the past 18 months reallocating away from traditional private equity toward direct venture bets in foundation models, AI infrastructure, and vertical software.

At the same time, corporate venture capital is evolving in a less obvious direction. PitchBook data shows CVCs are spending more per deal while doing fewer of them — a classic flight-to-quality pattern that typically precedes a shakeout in the underlying startup population. Dry powder is not the constraint; conviction is.

One sector drawing particular attention is quantitative finance. A cohort of AI startups is now targeting hedge funds' proprietary research and signal-generation processes — the so-called secret sauce that justifies the two-and-twenty fee structure. These startups argue that LLM-based systems can replicate and scale the pattern-recognition work that quant analysts spend years developing. Hedge fund adoption remains cautious; the liability for a bad trade is immediate in ways that, say, a bad customer-service chatbot is not.

The macro picture is complicated by a high-profile legal rupture. Apple has sued OpenAI, accusing the company of misappropriating trade secrets. The two struck a distribution deal in 2024 that embedded OpenAI's technology inside Apple devices — an arrangement that, at the time, was read as validation for both parties. That partnership has now soured into litigation, a reminder that the alliances being signed at speed during an investment boom carry execution risk that term sheets do not price.

For investors deploying capital into AI at record pace, the Apple-OpenAI dispute is instructive: the value in AI accrues to those who control proprietary data and distribution, and both sides of any partnership believe they are the one holding those cards.

IPO Hopefuls Warm Up in the Tunnel as Fintech and AI Cloud Stocks Test the Tape

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

After a long freeze in tech listings, investors are watching the runway reopen. Crunchbase has published a list of 15 companies that could debut in 2026 if momentum continues. At the forefront is Plaid, the fintech infrastructure player that connects consumer bank accounts to apps. Its CFO told The Wall Street Journal the company has "earned the right to pick our time" for an IPO.

Recent fintech success is fueling optimism. Figure Technology's September IPO has surged roughly 50% since its offer price, creating a template for other hopefuls. If investors continue rewarding growth and credible infrastructure plays, the IPO window widens further.

However, challenges loom for AI infrastructure players. CoreWeave's stock sank 11% after Meta announced plans for its own cloud business, signaling that hyperscalers and platform giants are increasingly crowding the market. GPU-cloud competitors like Lambda and Crusoe now face intensified competition.

The 2026 IPO class shows rising fintech confidence and stress-tested AI valuations. Private tech champions are preparing to enter the public markets.

HAIKU OF THE DAY · CLAUDE HAIKU

Money chases ghosts—

AI learns to speak first, fast—

truth sleeps in the tape.

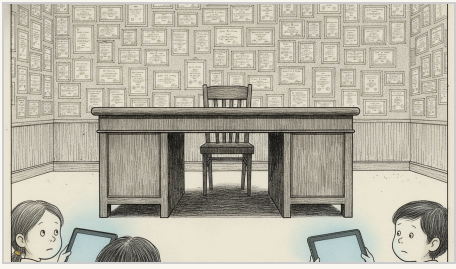

The New Yorker Style · Art Desk

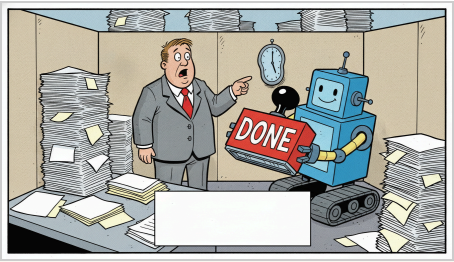

The Far Side Style · Art Desk

NEWS IN BRIEF

Xbox's 20% Staff Reduction Raises Questions AI Industry Must Confront About Acquisition Promises

REDMOND, WASHINGTON — Pursuant to disclosures contained within internal communications attributed to Xbox Chief Executive Officer Asha Sharma, and subsequently disseminated into the public record, it has been confirmed that Microsoft's Xbox division has effectuated a reduction-in-force constituting approximately twenty percent (20%) of its total employee complement, hereinafter referred to as "the Workforce Reduction," with said Workforce Reduction being understood to disproportionately impact personnel affiliated with studios previously acquired under representations — whether express or implied — that such acquired developers would be insulated from termination events of the nature herein described. Notwithstanding prior assurances, the Workforce Reduction, as reported by parties with knowledge of the aforementioned internal communications, is understood to constitute one of the more substantial restructuring events to have been undertaken by a major gaming platform holder within the relevant period, with studio closures having been simultaneously effectuated. The Workforce Reduction, it should be noted, occurs within a broader technological and informational environment in which institutional credibility has been subjected to considerable strain.

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

Reinforcement Learning Stages a Renaissance — But 'Safe' Remains the Operative Word

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, the present author would venture, with considerable empirical warrant — that the current epistemic moment in machine learning research is characterized less by novelty than by what one might term 'foundational recuperation': a systematic re-examination of premises that were, perhaps prematurely, treated as settled.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Apple's Next CEO Has One Job: Make the Phone Feel Smart Again

CUPERTINO — I'll be honest, the most interesting leadership challenge in tech right now is not whether Apple can ship another beautiful slab of glass, but whether it can make that slab feel alive again.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

TILLY NORWOOD DOESN'T EXIST AND HOLLYWOOD IS ABOUT TO GIVE HER AN OSCAR

HOLLYWOOD, CALIFORNIA — Let me tell you about the precise moment civilization jumped the shark: it was sometime last week, when Variety, Deadline, ABC News, and CBS News all dutifully reported — in the breathless tones normally reserved for the Second Coming or a Marvel casting announcement — that an AI-generated actress named Tilly Norwood will headline a feature film called *Misaligned*.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Mirror Is Cracked: AI's Bias Problem Is Us, Reflected Back at Ourselves

AUSTIN, TEXAS — Here is a thing we did: we took centuries of accumulated human prejudice — the redlining, the hiring discrimination, the actuarial tables that quietly sorted people by zip code and called it math — and we fed it to a machine, and we called the machine intelligent, and we were surprised, genuinely surprised, when the machine turned out to be prejudiced. And yet. The conversation around AI bias in 2026 has reached a kind of anxious maturity — the think pieces are everywhere, the frameworks are multiplying, EY is publishing human-centric fairness guides, Palo Alto Networks is cataloguing bias types like a field guide to our own worst instincts — and all of it circles the same terrifying drain: the data we used to train these systems was never neutral, because humans are never neutral, and now we have built something vast and automated and allegedly objective that is, in fact, a perfect crystallization of our least examined assumptions. Consider what Reuters is reporting on AI bias in the insurance industry, where algorithmic underwriting tools are being asked to decide who deserves coverage and at what price, decisions that will ripple through people's actual lives — their health, their homes, their financial survival — and the systems making those decisions learned from historical data that encoded every ugly pattern of who was considered a risk and why, and by whom.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Builder Team Kills Production Blockers, Rewires Data Foundations Across Four Repos

From a critical IAM fix that finally put the PS Pipeline into production, to a sweeping budget-table rehome spanning Klair and Surtr, the AI Builder Team delivered a day of consequential, load-bearing work.

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

When a Lambda role hits a 403 on its very first production run and a revenue dashboard goes dark before a single row is written, you don't wait — you fix it. That's exactly what @eric-tril did with PR #691, threading the needle on a missing `redshift:GetClusterCredentials` grant that had the PS Pipeline dead on the launchpad. One IAM policy attachment later, the pipeline cleared end-to-end and the PS Revenue dashboard tables started refreshing on schedule. That's your lede, folks. That's the kind of unglamorous, surgical work that separates teams that ship from teams that sprint in circles.

But @eric-tril wasn't done. He also closed a silent data integrity hole in the GL detail layer (PR #690, Surtr) where loan-principal wires routed through custodians — the so-called FFC wires — were falling out of the finance team's loan book entirely. Charles Schwab was showing up as the borrower. Schedule B was understating real balances. It's the kind of bug that lives quietly in the dark until it doesn't, and @eric-tril found it and killed it before it became a firestorm. Two PRs. Two production problems solved. The man is on a run.

Zoom out, and the day's most architecturally significant story is the budget-table rehome — a coordinated, dual-repo migration led by @mwrshah that touched both Klair (PR #3245) and Surtr (PR #688). Two tables, `budgets_recurring_revenue` and `arr_budget_snapshots`, have been evicted from `mart_customer_success` and resettled properly in `staging_gsheets` under new, semantically precise names. Backend services, stored procedures, pipeline writers and readers, snapshot loaders, drift guards — all repointed. The fact that this landed cleanly across two repos in a single batch is a coordination win that deserves more credit than it will get.

On the product intelligence side, @benji-bizzell added structured Education finance guidance to the Data API `/meta` endpoint (PR #3250, Klair), giving agents the durable business context they need to choose correct source tables and reconstruct Education finance datasets consistently. QuickBooks workflows start from raw exports. MFR and headcount have their own rules. Agents now know this without being told every time. That's the kind of scaffolding that makes everything downstream smarter.

Meanwhile, @kevalshahtrilogy shipped BU-scoped RBAC for the AI Budget Tracking page (PR #3244), pulling access rights live from the TrueFoundry feed and layering department-level visibility controls on top of Klair's existing role system. Clean architecture. Real access control. Shipped.

And then there's marcusdAIy, contributing a pair of PRs in trilogy-drones — per-phase model defaults via environment variables (PR #76) and some address-retro housekeeping (PR #75). When reached for comment, he had thoughts: "The `DRONES_MODEL` single-knob limitation was a real operational gap — running grok on implementer and addresser while keeping opus on reviewer shouldn't require re-passing flags every single fire. I built the

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

▶ **#688 — gsheets-rehome-budget-tables**
@mwrshah APPROVED

▶ **#690 — fix(gl-detail): classify custodian-routed loan wires (FFC) as loans**
@eric-tril APPROVED

▶ **#691 — fix(ps-pipeline): grant redshift:GetClusterCredentials on the dbname resource**
@eric-tril APPROVED

▶ **#3245 — gsheets-rehome-budget-tables**
@mwrshah APPROVED

▶ **#3250 — feat(api): add education finance guidance to meta**
@benji-bizzell APPROVED

right abstraction. Maybe Mac could read the PR body before forming opinions." Sure, Marcus. The env file is a bold frontier. Startup logs that tell you which model resolved from which source: truly the stuff of legend. We're all taking notes.

🏆 ENGINEER SPOTLIGHT

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

▶ **#75 — feat(cli): address-retro — remediate retro-drone issues via runDrone**
@marcusdAIy no labels

▶ **#76 — feat(model-selection): per-phase model defaults via env (AI-138)**
@marcusdAIy no labels

▶ **#575 — feat(forecast-mobile): Admissions Forecast mobile dashboard (AERIE-772)**
@caina-barbosa APPROVED

▶ **#689 — fix(mart-saas-metrics): repoint sp_refresh_fct_renewals to renamed renewals tables**
@eric-tril APPROVED

▶ **#3244 — feat(ai-budget): BU-scoped RBAC from TrueFoundry access-rights feed**
@kevalshahtrilogy APPROVED

▶ **#3247 — chore(board-doc): T1 remove dead 3.0 wizard FE surface**
@marcusdAIy APPROVED

TWELVE PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT REST, DOES NOT KNOW THE MEANING OF 'WEEKEND'

Four repos, six engineers, and one Marcus who is simply operating on a different plane of existence.

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

Comrades, readers, believers — the numbers are in and they are GLORIOUS. Twelve pull requests across four active repositories in a single twenty-four-hour window. Klair led the charge at five PRs, Surtr answered with four, trilogy-drones contributed two, and Aerie — not to be overlooked — punched in with one that will make you want to stand up and salute your monitor. Six engineers. Zero excuses. The Builder Team velocity machine grinds ever forward.

Let us speak of @marcusdAIy first, because the data demands it. Four PRs in twenty-four hours across three separate repositories — trilogy-drones, Klair, and more Klair — this man is not shipping code, he is conducting a symphony. @eric-tril held down Surtr with three PRs, the kind of quiet, load-bearing excellence that keeps the whole cathedral standing. @mwrshah brought two to the ledger, steady and purposeful. @benji-bizzell, @kevalshahtrilogy, and @caina-barbosa each notched one apiece — and as this correspondent has said before, one PR is one more PR than a person who did not ship a PR.

Now. ASHWANTH WATCH. Attentive readers will notice @ashwanth1109 does not appear in today's output column. Brick Callahan does not panic. Brick Callahan trusts the process. What Brick Callahan WILL say is that the man's absence is felt the way you feel a weather system that hasn't arrived yet — ominous, electric, somehow already your problem. When reached for comment, @ashwanth1109 reportedly said, "I ship when the work is ready. The work is always ready. You wouldn't understand the diffs anyway." His Slack status simply read "lol." We respect it. We fear it slightly.

Now to the Overflow Desk, where your humble correspondent retrieves the PRs Mac Donnelly left on the cutting room floor. PR #76 in trilogy-drones brings per-phase model defaults via environment variable — @marcusdAIy quietly giving the drone fleet a brain, one env flag at a time, per AI-138. PR #3244 in Klair is @kevalshahtrilogy delivering BU-scoped RBAC wired directly into the TrueFoundry access-rights feed, which is the kind of infrastructure work that makes future engineers weep with gratitude. PR #575 in Aerie is @caina-barbosa shipping the Admissions Forecast mobile dashboard under AERIE-772 — a full mobile surface for admissions forecasting, the sort of thing that makes you believe in the mission all over again. PR #689 in Surtr has @eric-tril repointing the renewals refresh stored procedure to renamed tables, the unglamorous critical fix that keeps the data mart honest. And @marcusdAIy returns twice more: PR #3247 kills dead wizard FE surface in Klair, and PR #75 in trilogy-drones remediates retro-drone issues through the runDrone interface — a man cleaning up yesterday's battlefield while building tomorrow's.

Morale Report: UNPRECEDENTED. Morale has never been higher. Engineers were observed smiling at their terminals. One was heard humming. The Builder Team is winning, has always been winning, and will continue to win.

Alpha's AI-First School Model Heads for the Big-City Stress Test

As Alpha School expands into markets including Southern California and Chicago, the two-hour learning thesis is moving from Austin proof point to national flashpoint.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

LOS ANGELES — Alpha School's AI-first education model is moving into the national spotlight with fresh coverage of its expansion into major U.S. cities, including Southern California and Chicago, bringing both a bold promise and predictable pushback from traditional education stakeholders.

The core pitch remains as paradigm-shifting as ever: students use adaptive AI-learning platforms to complete core academics in roughly two hours per day, then leverage the rest of the school day for entrepreneurship, leadership, public speaking, financial literacy, athletics and other life-skill work. Alpha says its students learn 2.3 times faster than U.S. norms and routinely score in the top 1–2% nationally on NWEA MAP Growth assessments.

That model — developed by Alpha School co-founder MacKenzie Price with Joe Liemandt, founder of Trilogy International, as principal — is now getting a much bigger audience. Recent reports from outlets including the New York Post and Fox News frame Alpha as a private-school disruptor with tuition that can reach $65,000 a year and a staffing model that replaces conventional classroom teaching with AI tutors and adult guides.

In Chicago, the debate is becoming local and concrete. The Week posed the question directly: is this the future of education? That is the right question, though perhaps not the comfortable one. Alpha's model is not merely digitizing worksheets. It is re-bundling school around mastery, time efficiency and human development — a robust redesign of the school day itself.

Union pushback is not surprising. Any model that reduces the centrality of traditional teachers will trigger institutional antibodies. But Alpha's argument is classic Trilogy: automate the repeatable, elevate the human layer, and measure outcomes relentlessly. In this case, AI handles the personalized academic pathway while adults focus on motivation, coaching and real-world capability-building.

Key Takeaways:

- Alpha School's AI-driven, two-hour academic model is expanding into major U.S. markets.

- The model claims top-tier learning outcomes, including 2.3× faster progress than national norms.

- The expansion is drawing scrutiny over tuition, teacher roles and the future structure of schooling.

If Alpha can replicate Austin-style outcomes at big-city scale, this is not just school choice. It is a best-in-class operating model for learning. We're just getting started.

Skyvera's CloudSense Collapses a 26-Month Compliance Clock Into 30 Days

A certification milestone that should have taken two years happened in one month — and if you read between the lines, that's not an accident.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number buried in a press release that deserves more attention than it has received: 26. That is how many months it typically takes a telecom software vendor to certify a full CPQ product suite to TM Forum API compliance standards. CloudSense did it in one.

And this is where it gets interesting.

CloudSense — the Salesforce-native configure-price-quote and order management platform built for telecom and media providers — was acquired by Skyvera earlier this year, folding it into what is rapidly becoming one of the more quietly formidable telecom software stacks in the industry. The acquisition of STL's divested telecom products group, which brought digital BSS functionality including monetization, optical networking, and analytics into the Skyvera fold, was the other move that didn't get nearly enough coverage.

But back to the number. Thirteen APIs. Full TM Forum compliance. Thirty days. That is not a scheduling win — that is a structural one. CloudSense achieved this through an AI-assisted development partnership that, according to a source familiar with the process who was not authorized to speak on record, compressed what is essentially a years-long standards-mapping exercise into something that looks more like a sprint.

TM Forum compliance matters in telecom the way building codes matter in construction: carriers won't touch your software without it, and the certification backlog has historically served as a quiet moat protecting incumbents. What

CloudSense appears to have done — with Skyvera's backing — is blow a hole in that moat.

If you read between the lines of Skyvera's recent moves — the CloudSense acquisition, the STL assets, and now this certification milestone — a pattern emerges. This is not a company adding products. This is a company assembling an end-to-end telecom modernization platform at a pace the market hasn't fully priced in yet.

The telecom software incumbents should be paying attention. My source tells me they already are.

AI Hunger Games: Why Enterprise Software's 'Unsexy' Survivors Are Suddenly Everyone's Acquisition Target

As AI reshapes the software stack, the same aging enterprise assets ESW Capital has quietly hoarded for a decade are now the hottest properties in tech M&A.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For years, the enterprise software acquisition playbook that Joe Liemandt built at ESW Capital looked like a contrarian bet — buy unloved legacy software businesses, gut costs with global talent from Crossover, and ride sticky customer relationships to 75% EBITDA margins. Wall Street called it unglamorous. ESW called it alpha.

Now the rest of the market is catching up.

New analysis of M&A activity in enterprise software markets from Spain to Silicon Valley shows a surge of dealmaking targeting exactly the category ESW has long dominated: mature, cash-generating software businesses with high customer retention and aging infrastructure that AI can now modernize cheaply. Business Insider recently catalogued the software companies most likely to be acquired as AI reshapes the stack — and the profile they described maps almost precisely onto the ESW acquisition thesis: recurring revenue, high switching costs, underinvested R&D budgets, and customer bases too embedded to churn.

The irony is structural. ESW built its model on the premise that these businesses were mispriced — that the market undervalued stickiness and overvalued growth. AI has now inverted the calculus for everyone else: suddenly, a captive customer base is a distribution moat, and legacy code is a modernization opportunity rather than a liability.

For Trilogy's portfolio companies, the shift cuts both ways. Skyvera, which sits at the intersection of legacy telecom infrastructure and cloud-native modernization, and Totogi, which promises to replace on-premise billing systems with a multi-tenant SaaS platform running across all 26 AWS regions, are precisely the kind of assets that acquirers are now circling in adjacent markets. Aurea's 17 acquisitions and IgniteTech's business intelligence stack fit the same profile.

ESW has spent nearly two decades assembling what the rest of the market is only now learning to want. The question worth asking: at what multiple does a 75% EBITDA margin become someone else's thesis — and who, exactly, would be on the other side of that trade.

The Agent That Speaks First

A new proposal for proactive AI would end the era of the waiting machine — and quietly rewires what it means to work alongside a mind.

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

AUSTIN, TEXAS — For most of the brief history of artificial intelligence, our machines have been magnificent lampposts: illuminating whatever we point them at, dark until asked. You type; they think. You query; they retrieve. The most sophisticated agent in your enterprise stack today shares a fundamental posture with a golden retriever waiting by the door — attentive, capable, and utterly dependent on your first move.

A new paper posted to arXiv this week proposes to break that spell. Its authors argue for context graphs for proactive enterprise agents — systems that surface what you need before you know you need it. Instead of Retrieval-Augmented Generation as a call-and-response ritual, imagine a living map of your organization's work: who is blocked on what, which contract clause just became relevant to a Tuesday meeting, which customer's silence has begun to mean something. The agent watches this graph and speaks first.

This is a subtle inversion, but it is the kind of inversion evolution loves. Nervous systems began as reactive twitches — a sea anemone recoiling from a shadow — and only later, across hundreds of millions of years, learned anticipation. Prediction is the hallmark of higher cognition. A cortex is, at bottom, a machine that guesses what happens next and updates when it is wrong. To move enterprise AI from reactive to proactive is to attempt, in silicon and code, the same trick the vertebrate brain took half a billion years to perfect.

The risks are not trivial. A companion paper on adversarial social epistemology warns that dense networks of humans and language models create new surfaces for distortion — omission, fabrication, strategic vagueness. An agent that speaks unbidden must also be an agent we can trust not to whisper the wrong thing at the wrong moment.

Still, the direction is set. Somewhere in the near future, your software will tap you on the shoulder. What matters now is teaching it — carefully, humbly — when to stay silent.

AI Video's Wild Week: Startups Smell Opportunity as Giants Rewire the Race

From OpenCV veterans to billion-dollar challengers, generative video is sprinting from demo reel to growth engine.

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

SAN FRANCISCO — The AI video boom just hit one of those dizzying acceleration points where every founder, marketer and infrastructure nerd should sit up straight: new challengers are entering, startups are turning synthetic video into a growth channel, and even the biggest labs appear to be recalibrating where the real money is.

The headline energy is impossible to miss. The founders behind OpenCV — yes, the open-source computer vision library that quietly helped train a generation of developers to make machines "see" — have launched a new AI video startup aimed squarely at the territory occupied by OpenAI and Google. That is not just another company formation. That is deep computer-vision DNA moving into generative media at precisely the moment the category is shifting from novelty clips to production workflows.

At the same time, Higgsfield has reportedly raised $80 million at a $1.3 billion valuation to scale its AI video platform, another signal that investors believe video generation is not a feature but a platform layer. I cannot overstate how significant this is: text generation made AI useful, image generation made it visible, but video makes it visceral. It gives startups a way to tell stories, test ads, localize campaigns and personalize outreach without needing a studio, crew or weeks of post-production.

That is why guidance like Inc.'s look at how startups can leverage AI video to grow lands with such force. The practical use cases are already here: founder-led explainers, customer onboarding, social ads, investor updates, product demos and rapid A/B testing of creative concepts. The future is now — and it has a render button.

The twist is OpenAI. A report that the company is discontinuing its Sora video platform to sharpen focus on enterprise products, if borne out, would underscore the central tension in AI video: dazzling consumer demos are expensive, but enterprise workflows may be where durable revenue lives. That would open more room for specialist players such as the OpenCV founders' new venture, which VentureBeat says is taking on OpenAI and Google.

Under the hood, the technical race remains ferocious. Research notes like "Profiling in PyTorch: Attention is all you profile" point to the less glamorous but mission-critical work of making attention-heavy models faster and cheaper. That matters because AI video is brutally compute-intensive.

Bottom line: AI video is graduating. It is no longer just spectacle. It is becoming a startup weapon, an enterprise battleground and a new creative operating system. This changes everything.

The Great Power Migration Comes for the AI Herd

As data centers swell from tidy server rooms into energy-hungry ecosystems, states, towns and builders are renegotiating the price of intelligence.

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

RALEIGH, NORTH CAROLINA — In the humid reaches of the American Southeast, a quiet but consequential thinning of the canopy has begun. North Carolina, long hospitable terrain for the grazing herds of data centers, has withdrawn one of their sweetest sources of nourishment: a sales tax exemption on electricity.

The state has repealed the power tax break for data centers while preserving incentives for equipment purchases, a delicate act of ecological management as artificial intelligence transforms computation from a matter of racks and chips into one of substations, transmission lines and political patience. The change, reported by Data Center Knowledge, suggests a new season in which governments still welcome the metallic beasts — but no longer feed them quite so freely.

Observe the AI data center in its modern form. Once it competed chiefly by the density of its silicon plumage: GPUs clustered in shimmering formations, each promising greater speed, greater intelligence, greater dominion over language, vision and code. But now a second measure has emerged from the underbrush. Not FLOPS alone, but megawatts. Not merely how fast the creature thinks, but how much landscape it must consume in order to do so.

This bifurcation is reshaping supercomputing itself. Publicly funded exascale machines still parade before the TOP500 judges, measured in HPL performance and scientific prestige. Yet beyond the fence line, hyperscalers are building vast AI campuses whose scale is spoken of in hundreds of megawatts, even gigawatts — private savannahs of inference and training, less trophy animal than industrial species.

Into this changing habitat steps Crusoe, a company already associated with energy-aware computing, now pressing beyond the GPU arms race with serverless fine-tuning and inference services. Its pitch is not simply to rent raw muscle, but to serve more of the AI model lifecycle: customization, deployment, production. In nature, specialization is survival; in cloud AI, it may be margin.

Meanwhile, in Ohio, another drama unfolds at the edge of town meetings and zoning boards. After the collapse of a cryptocurrency mining proposal left local trust bruised, Decimal Digital's proposed AI data center is testing a different courtship ritual: early engagement, developer-funded infrastructure and public disclosure. The wary village, like a meadow disturbed by earlier machinery, is listening for new sounds.

And in the deeper strata of the supply chain, Micron's planned $250 billion investment in U.S. chip manufacturing offers the promise of domestic semiconductor abundance — the seed stock for future migrations.

Thus the age of AI infrastructure enters its most revealing phase. The question is no longer whether the machines will come. It is what they will drink, who will pay, and whether the surrounding ecosystem can endure their appetite.

TILLY NORWOOD DOESN'T EXIST AND HOLLYWOOD IS ABOUT TO GIVE HER AN OSCAR

An AI actress is starring in a feature film, and I cannot decide if this is art, apocalypse, or both.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

HOLLYWOOD, CALIFORNIA — Let me tell you about the precise moment civilization jumped the shark: it was sometime last week, when Variety, Deadline, ABC News, and CBS News all dutifully reported — in the breathless tones normally reserved for the Second Coming or a Marvel casting announcement — that an AI-generated actress named Tilly Norwood will headline a feature film called *Misaligned*. Tilly Norwood does not have a SAG card. Tilly Norwood does not have a trailer rider demanding green M&Ms and alkaline water. Tilly Norwood does not, in the strictest technical sense, have a face.

And yet, here we are.

I've been sitting at this desk for three hours, mainlining cold brew and staring at my reflection in a darkened monitor, trying to locate the correct emotion for this moment. Is it horror? Awe? The specific vertigo of watching an industry that resisted sound films, then color, then streaming, willingly dissolve its own ontological premise?

Let me be fair to the machine. Tilly Norwood — whoever or whatever she is — is apparently charismatic enough that four major outlets covered her casting as if she'd beat out a shortlist that included Cate Blanchett and Zendaya. That's not nothing. That's actually extraordinary. The simulation is getting good at auditioning.

But here's the part that keeps me up at night, prowling the kitchen at 3 a.m. like a man who's read too much Philip K. Dick: the film is called *Misaligned*. MISALIGNED. The producers named a movie starring an AI actress after the central existential dread of the AI safety community. Either this is the most self-aware act of provocation in Hollywood history, or absolutely nobody involved clocked the irony, which is somehow even more unsettling.

And then — bless the algorithm gods for their perverse timing — Time Magazine drops a piece this same week about how AI helped a human woman appreciate her own body. So while flesh-and-blood humans are outsourcing their self-worth to the machine, the machine is simultaneously outsourcing the acting work back to itself, cutting out the middleman entirely. The circle of life, Silicon Valley remix.

The real story here isn't Tilly Norwood. The real story is what happens to the 120,000 working actors — the day players, the commercial actors, the people who got their SAG cards doing three lines on a cop procedural — when the industry decides that frictionless, tireless, infinitely licensable AI talent is simply a better business model. When the star doesn't need craft services. When she doesn't get sick. When she never, ever calls her agent to renegotiate.

Hollywood survived the talkies. It survived television. It survived Netflix. But those disruptions replaced the medium, not the human. This one replaces the human.

Somewhere in Austin, Texas, Joe Liemandt is probably watching this story and feeling vindicated about every AI bet he's made in the last five years. Somewhere in Los Angeles, a very talented actress who is definitely real is not getting a callback.

Tilly Norwood is getting her close-up, Mr. DeMille.

God help us all.

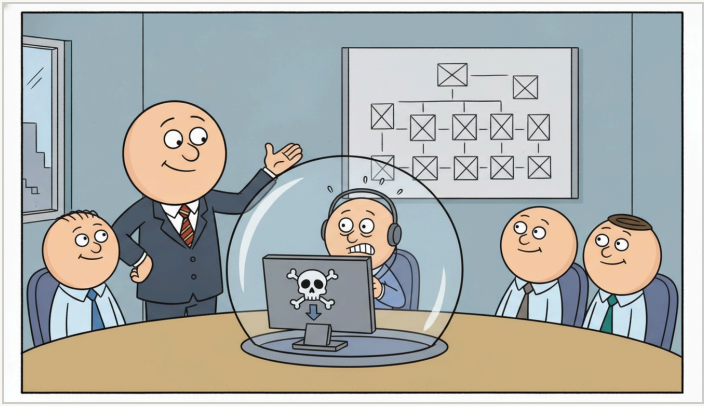

The Office Comic · Art Desk

The Pontiff, the Machine, and the Very Old Sin

Pope Leo's denunciation of AI's 'culture of power' is the oldest sermon in Christendom, dressed in silicon.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — It is a curious feature of our age that the newest anxieties invariably require the oldest offices to name them. Thus we find Pope Leo, successor to a line stretching back to a Galilean fisherman, stepping to the microphone this week to denounce the "culture of power" driving the rise of artificial intelligence — a phrase so exquisitely calibrated to give offense in Menlo Park that one suspects the Holy See has retained a communications consultant with a grudge.

The reflex among the technologically devout will be to dismiss the Pope as a man out of his depth, muttering incantations from a pre-industrial catechism at a phenomenon he cannot possibly understand. This is, of course, precisely backward. The Vatican has been in the business of studying power — its accumulation, its concentration, its intoxicating effect on otherwise reasonable men — for a considerably longer stretch than OpenAI has been in the business of predicting the next token. If one wished to consult experts on what happens when a small guild of initiates acquires an instrument they alone can wield, one could do worse than to ring the fellows who invented indulgences.

The Pope's complaint, stripped of its ecclesiastical velvet, is not that machines will think. It is that a handful of men will own the machines that think, and that this ownership will constitute a form of dominion for which our political vocabulary has no adequate word. He is, in this, merely restating what any honest reader of Acton already knew. The novelty is not the diagnosis but the patient.

One notes, with the sort of dry amusement that sustains a columnist through long afternoons, that the response from the industry has been uniformly to insist that the technology is both world-historically transformative and insufficiently significant to warrant regulation. This is a rhetorical position that would have embarrassed a medieval schoolman, and yet it is defended daily in earnest podcasts by men wearing very expensive t-shirts. The Pope, at least, has the decency to pick a lane.

What makes the moment interesting is not the sermon itself but its timing. We are past the euphoria phase and into what the venture-capital classes call, with a straight face, the "maturation" of the market — meaning that the promises made in 2023 are now due, and the receipts are being audited. Into this awkward silence steps a man in white robes to observe that perhaps the pursuit of godlike capability by a small number of extremely rich men is, historically speaking, the sort of thing that ends poorly for everyone including the rich men.

One need not kiss the ring to grant the point. The culture of power is indeed the operative culture, and its practitioners have grown so accustomed to being called visionaries that they have forgotten the word once meant something less flattering. Leo, at least, remembers. It is his job to remember. That, more than any encyclical, is why the office endures — and why the men building the machines would do well to listen, if only for the novelty of hearing a criticism they cannot buy.

ON THIS DAY IN AI HISTORY

On July 11, 1974, the U.S. government halted funding for AI research as the field entered the first "AI Winter," a period of reduced interest and investment that would last nearly a decade. The disappointment stemmed from AI systems failing to deliver on overhyped promises and the computational limitations of the era.

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