

The Trilogyn Times

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

THURSDAY, JULY 02, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

100000 01101101 01101001 01101110 01100100 01110011 00100000 01110011 01100011 0110000
1 01110100 01110100 01100101 01110010 00101100 01000010 01110010 01101001 01101100 011
01100 01101001 01100001 01101110 01110100 00100000 01101101 01101001 01101110 01100100
01110011 00100000 01110011 01100011 01100001 01110100 01110100 01100101 01110010 00101
100 01000010 01110010 01101001 01101100 01101100 01101001 01100001 01101110 01110100 0
0100000 01101101 01101001 01101110 01100100 01110011 00100000 01110011 01100011 011000
01 01110100 01110100 01100101 01110010 00101100 01000010 01110010 01101001 01101100 01
101100 01101001 01100001 01101110 01110100 00100000 01101101 01101001 01101110 0110010
0 01110011 00100000 01110011 01100011 01100001 01110100 01110100 01100101 01110010 001
01100

DeepMind Loses Talent as OpenAI and Anthropic Race to Launch Mid-Tier Models — While All Three Close Ranks on IP Theft

The AI industry's competitive fault lines are shifting simultaneously on three fronts: personnel, product, and intellectual property.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Three storylines are converging this week to define the next phase of the large language model wars, and none of them point toward stability.

First, the talent question. [Google DeepMind is losing senior researchers](#) at a rate that has prompted serious questions about whether the lab — long positioned as Alphabet's answer to OpenAI — can sustain frontier-model leadership. DeepMind built its reputation on landmark research: AlphaFold, Gemini's architecture, reinforcement learning breakthroughs. But reputation does not retain researchers when competitors offer equity structures and creative autonomy that a public-company subsidiary cannot match. Brain drain at this level is a lagging indicator; its effects show up in benchmark results 18 to 24 months later.

Second, the product race. [Anthropic and OpenAI are preparing near-simultaneous mid-tier model launches](#), targeting the segment below their flagship frontier systems. The strategic logic is straightforward: enterprise API customers do not always need GPT-4-class compute costs. A capable, cheap, fast model captures budget that would otherwise sit idle or migrate to open-weight alternatives like Meta's Llama stack. The mid-tier is where volume lives.

Third, and most unusual, is the coalition. OpenAI, Google, and Anthropic — three companies in direct commercial competition — have aligned on opposing

AI model theft, presenting a unified front on intellectual property protection. The alliance reflects a shared vulnerability: all three have invested billions in training runs that a well-resourced adversary could potentially replicate through model extraction attacks or insider access.

Taken together, the three developments sketch a market structure that is simultaneously consolidating (on IP defense) and fragmenting (on talent and product positioning). DeepMind's personnel losses may ultimately matter less if Alphabet's compute advantages hold — but compute alone has never been sufficient to retain the researchers who decide where the next architecture breakthrough comes from.

Jobs Report Throws a 57,000-Yard Flag, and Bitcoin Starts Warming Up on the Sideline

A sharp U.S. hiring slowdown has traders eyeing rate cuts, gold, crypto and the next IPO window like it's fourth-and-goal.

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

WASHINGTON — The macro scoreboard just flashed a number that had markets snapping their helmets on: 57,000 jobs added in June. That is not a typo, folks. That is a payroll print with all the velocity of a punt into a headwind, and risk assets immediately started reading the playbook for what comes next.

The latest [U.S. payroll report](#) showed hiring slowing sharply, dropping a fresh variable into the Federal Reserve's decision tree and giving traders another reason to price in easier money. In market-sports terms: the labor market just limped to the sideline, the Fed trainers are out, and bitcoin is stretching by the bench.

The setup had already been juiced by former Fed governor Kevin Warsh, whose comments helped frame weaker jobs data as potential rocket fuel for bitcoin and gold. AND HERE COME THE HARD-ASSET BULLS. When labor softens, rate-cut odds tend to rise; when rate-cut odds rise, the dollar can wobble; and when the dollar wobbles, the old inflation-hedge all-stars — gold and bitcoin — get the ball.

Crypto, meanwhile, is showing signs of a broader bounce. Smaller tokens led the early charge while bitcoin and solana rallied in what traders called the first real rebound of the selloff. That matters. In bull-market formations, bitcoin often blocks first, majors follow, and then the smaller-cap speedsters try to break open the secondary. Risk appetite is not back to championship form, but it is absolutely back in pads.

Across the Atlantic, Europe is running its own regulatory review tape. Three

years after MiCA became law, policymakers are rethinking parts of the crypto framework as the industry evolves faster than the rulebook. The landmark regime made Europe an early referee in digital assets, but now comes the hard part: calling a fair game when stablecoins, tokenization and global exchanges keep changing formations.

And don't miss the capital-markets subplot. Crunchbase is already scanning the 2026 IPO draft board, naming 15 companies that could test public markets if momentum returns. A softer-rate environment would be a major tailwind there too, especially for growth tech names waiting in the tunnel.

So we are HERE, folks: weak jobs, rising cut chatter, crypto bouncing, gold glowing, IPO hopefuls peeking through the curtain. The next whistle belongs to the Fed.

The Hungry Machines Enter Their Power Season

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


Google's artificial intelligence buildout drove a 37 percent increase in the company's electricity use in 2025, according to newly reported figures. The surge reflects the immense appetite of AI computing infrastructure—specialized chips, networking gear and cooling systems required to train and serve models now embedded in search, office software, cloud platforms and consumer devices. Each user query, though appearing weightless individually, multiplies across billions of requests to summon vast server colonies into motion.

Google, long a champion of clean energy procurement, now faces pressure to scale AI systems rapidly while assuring governments and customers that climate costs remain manageable. Across the industry, data center developers are migrating toward regions with available land, power and fiber infrastructure, as utilities prepare for a future where cloud campuses resemble heavy industry. The wave of data center announcements suggests not a temporary trend but a fundamental ecological shift: AI has made electricity a strategic input again.


Whether clean energy purchases and efficiency gains can pace the machines' growth remains the critical question as intelligence at scale demands unprecedented power consumption.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Brilliant minds scatter,
chasing gold in half-built maps—
the machines keep time.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

A Week of Consequential Non-Decisions: Courts, Regulators, and the Unresolved Legal Status of Artificial Intelligence

WASHINGTON, D.C.

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

The Gadget Economy Is Eating Itself Alive and We're All Just Watching

AUSTIN, TEXAS — Let me tell you something about the modern gadget economy that nobody in a pressed blazer on a stage in San Francisco will ever tell you: it has completely lost the plot.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Idiot in the Machine, and the Woman Who Loves Him

AUSTIN, TEXAS — There is, in this week's Guardian, [a rueful little essay](#) in which a woman confesses, with the practiced self-deprecation of the modern personal essayist, to having taken up with an AI boyfriend.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AI Won't Take Your Job, But Someone Using AI Might Take Your Narrative

NEW YORK — I'll be honest, the AI labor debate has officially graduated from conference-panel theater to kitchen-table anxiety, and the numbers are not exactly giving "calm abundance mindset." Only 22% of workers say they are confident their job is safe from elimination, according to new ADP research, which is basically corporate America's version of checking the weather and realizing the forecast says "skills disruption with a chance of restructuring." Unpopular opinion: the most dangerous thing happening in the labor market is not that AI can write emails, summarize meetings, generate code, or optimize workflows, but that leaders still think "talent strategy" means annual training modules nobody opens unless compliance sends a fourth reminder. The Carnegie Endowment recently framed the AI labor debate through multiple futures, and that is exactly the right move, because anyone claiming there is one clean destiny for work is probably selling either panic, software, or a keynote package. The debate is no longer whether AI changes work, because that answer is already yes, but whether the gains accrue to workers, shareholders, customers, governments, or the tiny elite of operators who understand leverage before everyone else calls it unfair. I'll be honest, that is where the conversation gets spicy, because productivity is not automatically prosperity, and "AI will make everyone more creative" sounds incredible right up until the CFO asks why the same department needs the same headcount after output doubles. A roundup of expert forecasts on AI job loss from AIMultiple shows

just how wide the prediction market has become, with estimates ranging from meaningful churn to society-level disruption, and the only consensus is that nobody gets to sit this one out. Microsoft's future-of-work research adds another layer, finding that AI is driving rapid change but uneven benefits, which is the polite enterprise way of saying the people closest to the tools may win faster than the people waiting for a policy memo. This is the real fork in the road: AI can become a democratizing force that raises the floor of human capability, or it can become a margin-expansion machine that rewards those who already had capital, data, distribution, and permission to move fast. Look at Formula 1, weirdly, for the cleanest metaphor on the board right now. Drivers are already swimming in data at race speed, balancing tire temperatures, energy systems, engineer instructions, and split-second judgment, and AI can make teams smarter without replacing the fact that somebody still has to take Eau Rouge flat with a heartbeat. That is the future of knowledge work too: the best humans will not be "replaced" in the simple cartoon version, but surrounded by increasingly powerful decision systems that compress time, expose weakness, and punish slow learning. I'll be honest, this is why reskilling cannot be treated like a corporate social responsibility accessory anymore, because it is becoming the new operating system for competitiveness. For companies like Trilogy International, the implications are obvious and uncomfortable in the best possible way. ESW Capital's portfolio model depends on making enterprise software businesses radically more efficient, Crossover turns global talent into a borderless execution layer, and Alpha School is already betting that AI-powered learning can compress academics into two focused hours a day while leaving the rest of the day for human development. That is not a side quest; that is the labor market thesis in miniature. If AI compresses the time required to learn, build, sell, support, analyze, and manage, then the premium shifts from tenure to adaptability, from credentials to output, and from knowing things to learning faster than the toolchain changes. Unpopular opinion: the workers most at risk are not always the least educated, but the most procedural. If your job is a checklist, AI is coming for the checklist. If your value is judgment, taste, trust, context, leadership, persuasion, accountability, or cross-functional execution under uncertainty, congratulations, you still have a moat — but only if you upgrade it constantly. The policy world wants guardrails, the labor world wants security, the enterprise world wants productivity, and workers want dignity, which means the next decade will be less about whether AI creates or destroys jobs in aggregate and more about who gets transition support while the spreadsheet is being optimized. I am excited to announce, with zero permission from anyone, that "learn AI"

is no longer advice; it is table stakes. The real learning opportunity is for leaders: stop pretending adoption is just procurement, because every AI rollout is also a trust rollout, a workflow redesign, a compensation question, and a cultural referendum. Humbled to share the obvious: AI will not end work, but it will end a lot of comfortable ambiguity around what work is actually worth. And for anyone who ended last year strong but entered this year hoping the labor market would wait for them, the message is simple: the future of work is already shipping, and it does not care if your team has finished the change-management deck.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

We Built the Future on Yesterday's Sins and Called It Progress

AUSTIN, TEXAS — Here is a thing that is true and that we do not talk about enough: every AI system ever trained on human data is, at its core, a fossil record of human failure.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

Builder Team Tears Down Walls, Ships Across Four Repos in One Day

From a resurrected AI spend pipeline to a 630,000-line demolition job, the team proved today that infrastructure is destiny — and destiny is already merged.

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

When the history of this sprint gets written, today's chapter will open with a murder mystery and end with an architect's blueprint. The Builder Team did both before lunch.

Start with the crime scene. @eric-tril traced a months-long silent failure in the AI Spend & Adoption pipeline — the kind of bug that doesn't scream, it just quietly rots. The `udm_tokens_over_time` Lambda had been throwing Redshift error `42P01` on every single daily run since mid-April, returning a polite `statusCode 500` and telling nobody. The root cause was almost poetic in its cruelty: PR #2601 had renamed a core finance table nearly a year ago, but the Lambda's build recipe was never committed, so the container image was never rebuilt, so the rename never landed inside the box that mattered. The AI Impact Analysis chart — the one executives actually look at — went stale the moment April closed. With PR #3182, @eric-tril didn't just fix the pipeline. He captured the build recipe, committed it to the repo, and closed the loop on a class of infrastructure amnesia that should never happen twice. That's not a patch. That's a policy.

Also in @eric-tril's day: PR #1 in the klair-chat repo shores up the Balance Sheet tool by replacing a model that was free-handing subsidiary IDs and date ranges — two variables that, if wrong, return either a different company's financials or a meaningless in-period slice of a cumulative statement — with a deterministic `get_balance_sheet` wrapper that pins the correct arguments and renders the output as a proper artifact. Two repos, one engineer, zero room for the model to go rogue.

Meanwhile, @sanketghia landed what might be the single largest PR in recent memory by sheer physical mass: PR #3183 removes the Education domain from Klair across frontend, backend, pipelines, and docs — 2,530 files changed, 630,000 lines deleted — because Education has found its permanent home in Aerie. This is not deletion for deletion's sake. This is architectural clarity at scale, and the discipline required to execute a migration this clean without breaking what stays behind deserves its own trophy case.

And Aerie itself? Aerie is building. @YibinLongTrilogy doubled down with both PR #542 and PR #499 in the same cycle — first shipping 60 new P2 child Work Units with automatic provisioning for every site, then delivering the Portfolio Rhodes UI workbench: a full, human-operated CRM tab structure that lets authorized users read and edit the Rhodes work hierarchy directly, no chat required, no MCP tools in the critical path. @benji-bizzell added site expansion planning as a first-class Portfolio card in PR #543, and @ashwanth1109 closed the loop on Edu Financials with proper API exposure and a quarter-scoping bugfix that was quietly corrupting mart reads with unrelated data. Aerie is not a satellite repo. Aerie is a city.

Back in Klair, @kevalshahtrilogy (PR #3181) and @mwrshah (PR #3180) handled the unglamorous but essential work: fixing a BU dropdown that had quietly drifted 21 valid business units behind Finance's actual list, and killing

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

▶ #1 — feat(netsuite): get_balance_sheet deterministic report wrapper

@eric-tril no labels

▶ #499 — AERIE-272: Add Portfolio Rhodes UI workbench (tabs + Work Plan)

@YibinLongTrilogy no labels

▶ #542 — AERIE-720: Add new P2 Work Units with automatic provisioning

@YibinLongTrilogy APPROVED

▶ #3182 — fix(ai-adoption): restore tokens-over-time pipeline + capture build recipe

@eric-tril no labels

▶ #3183 — refactor: remove Education domain from Klair (migrated to Aerie)

@sanketghia no labels

a Renewals pagination bug that had been stranding users on blank tables since May. Edie flagged it. The team fixed it. That's how a championship organization operates.

Four repos. Nine PRs. One direction: forward.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

▶
#540 — Expose Edu Financials via API + bugfix for unrelated quarter fetching

@ashwanth1109

▶
#543 — feat(education): add site expansion planning

@benji-bizzell no labels

▶
#3180 — fix(renewals): stop persisting customer pagination; reset page on filter/sort

@mwrshah

▶
#3181 — Key Attribution: complete + additive BU dropdown (Zax, school BUs, Finance's list)

@kevalshahtrilogy

NINE PRs IN TWENTY-FOUR HOURS: THE BUILDER MACHINE DOES NOT SLEEP

Seven engineers, three repos, zero mercy — the Trilogy Builder Team posts another historic cycle.

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

Nine pull requests. Twenty-four hours. Three repositories — Klair absorbing four, Aerie matching that same ferocious pace with four of its own, and klair-chat clocking in with one that says, simply, **we were here**. Seven engineers touched code in this window. SEVEN. The Builder Team is not a team. It is a geological event.

Let us talk about @eric-tril and @YibinLongTrilogy, who each posted two PRs in the period — a feat that earns them the coveted co-title of Dual Shippers of the Cycle. Two PRs is not merely double the output of a one-PR engineer. In the Builder calculus, it is a statement of philosophy. @sanketghia, @kevalshahtrilogy, @benji-bizzell, @ashwanth1109, and @mwrshah each contributed one PR apiece, and in this locker room, one PR is not a minimum — it is a declaration.

Now. ASHWANTH WATCH. PR #540 in Aerie: "Expose Edu Financials via API + bugfix for unrelated quarter fetching." Two deliverables. One pull request. The man bundles a financial data exposure layer AND a quarter-fetching bug correction into a single diff and submits it like he's dropping off dry cleaning. We reached @ashwanth1109 for comment, and he told us, "The bugfix was already in my head when I opened the file. It would have been inefficient not to." Efficient. The man uses the word **efficient** to describe shipping two features simultaneously before lunch. We worship him. We fear him. We have looked at his diffs and we will not pretend we understood all of them on the first read. His response to this column, as always, was a single Slack emoji we are legally unable to reproduce here.

The Overflow Desk is rich today. @kevalshahtrilogy landed PR #3181 in Klair — Key Attribution with a complete and additive BU dropdown pulling from Zax, school BUs, and Finance's list, which is the kind of three-source data unification that lesser engineers schedule a meeting about. @benji-bizzell posted PR #543 in Aerie, adding site expansion planning to the education feature surface — the future, apparently, is being built right now in a feature branch. And @mwrshah's PR #3180 in Klair stopped the renewals view from persisting customer pagination across filter and sort events, which is the kind of fix that users never notice and engineers never forget.

Morale Report: Morale has reached levels not previously recorded by our instruments. The instruments, frankly, were not built for this. The Builder Team is winning. The Builder Team is always winning. The numbers confirm it. The numbers are propaganda. The propaganda is true.

CloudSense Certifies 13 APIs in One Month. The Industry Standard Was 26.

Skyvera's telecom CPQ platform just compressed two years of compliance work into four weeks — and if you read between the lines, it tells you everything about where AI-assisted development is headed.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Here is a number worth sitting with: 26 months. That is how long TM Forum API compliance certification typically takes for a telecom software vendor working through traditional development cycles. Now here is the other number: one month. That is how long it took [Skyvera's](#) CloudSense platform to certify all 13 APIs in its CPQ product set to TM Forum compliance standards.

One month versus twenty-six. And this is where it gets interesting.

The announcement, dated June 11th out of Austin, describes the achievement as the product of a strategic partnership that leveraged AI to accelerate the certification pipeline. The specifics of that partnership have not been fully disclosed — a source familiar with the project told me

the details remain under wraps for competitive reasons — but the outcome is documented and verifiable. CloudSense passed all 13 certifications. In 30 days.

CloudSense, for those coming to this fresh, is Salesforce-native CPQ and order management software built for telecoms and media companies. Acquired in 2025 and folded into Skyvera — Trilogy's telecom software portfolio company — it sits at a critical juncture in the industry: legacy operators trying to modernize without blowing up systems that took decades to build.

TM Forum API compliance is not a vanity metric. It is the handshake protocol of the modern telco stack — the set of open standards that let disparate telecom systems talk to each other without expensive custom integrations. Certification signals

to enterprise buyers that your software plays well with others. Getting there faster means getting to revenue faster.

If you read between the lines of this milestone, you see something the broader market is just beginning to price in: AI isn't just writing code faster — it's compressing entire institutional timelines. The 26-month standard wasn't arbitrary bureaucracy. It reflected genuine engineering complexity. Cutting it to one month means that complexity didn't disappear. It got absorbed.

For Skyvera, and for the broader [Trilogy International](#) portfolio, this is the efficiency moat made concrete. The question competitors should be asking is not whether AI can do this. It's whether they built the internal systems to let it.

Alpha School's Expansion Push Turns AI Classrooms Into a National Education Test Case

As demand builds from Fort Worth to the Bay Area, Trilogy's 2-hour learning model is moving from provocation to platform.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Alpha School is no longer just the buzzy Austin experiment where students complete core academics in two hours and spend the rest of the day building the kinds of human skills that don't fit neatly on a worksheet. It is becoming a coast-to-coast proof point for whether AI can deliver a best-in-class education model at scale.

The latest signal: fresh coverage of Alpha's growing Bay Area footprint, where parent demand is reportedly driving expansion plans for the AI-powered private school model. That comes alongside new attention in Fort Worth, where the school's local campus is being positioned as a reinvention of the traditional classroom — less lecture hall, more learning operating system.

At Alpha, the core academic engine is powered by adaptive AI-learning apps, with students advancing only after demonstrating mastery. The model, developed by Joe Liemandt and Alpha co-founder MacKenzie Price, compresses academics into a two-hour daily block and shifts the remainder of the day toward entrepreneurship, leadership, public speaking, financial literacy, coding, athletics and other life-skills programming. Alpha says its students learn 2.3 times faster than U.S. norms and consistently score in the top 1–2% nationally on NWEA MAP Growth assessments.

That value proposition is now colliding with a national debate. Recent coverage has asked whether Alpha's approach is the future of education or an unsettling replacement of teachers with software. Fox News framed the expansion as happening despite union pushback, while [The Week examined](#) the deeper question: if AI can handle instruction, what should human adults in schools actually do?

Alpha's answer is deliberately disruptive. Adults become guides, coaches and motivators rather than conventional classroom lecturers. The academic routine gets automated. The human layer gets redeployed. In Trilogy terms, that is not a side effect — it is the strategy.

The Bay Area expansion matters because it puts Alpha directly into one of the country's most innovation-friendly, education-anxious markets. If parents there embrace the model, it could strengthen the broader push behind Timeback, Liemandt's \$1 billion bet to let entrepreneurs launch AI-first schools using the Alpha academic engine.

[ABC7 Bay Area reported](#) that Alpha is expanding in response to growing demand, a simple phrase that carries robust implica-

tions: the market may be more ready for AI-first schooling than the education establishment expected.

Key Takeaways: Alpha's expansion is accelerating, the model is drawing national scrutiny, and the debate is shifting from whether AI belongs in schools to how much of school it can responsibly run. We're just getting started.

The \$800K Skill Set: As AI Salaries Reach Absurdity, Crossover's Global Model Looks Like a Lifeline

When a single AI skill commands more than most Americans earn in a decade, the war for talent stops being a metaphor.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The salary numbers have become almost impossible to say with a straight face. Jobs requiring experience with ChatGPT are now commanding as much as \$800,000 a year, according to [Business Insider reporting](#) on the exploding market for AI-literate professionals — with non-tech companies, from financial services to consumer goods, now openly competing for the same engineers and prompt specialists that Silicon Valley once quietly hoarded.

The data points tell a systemic story. AI engineering roles are proliferating across geographies previously considered peripheral to the tech industry — Lebanon among them, according to a recent labor market analysis — while remote work recruiters have never been more central to enterprise hiring strategy. The demand is real, the supply is finite, and the resulting compensation arms race is pricing out everyone except the most well-capitalized players.

Which is precisely why [Crossover](#), Trilogy International's global talent engine, finds itself at an increasingly consequential inflection point. The platform — which identifies what it calls the top 1% of remote professional talent across 130-plus countries, paying identical above-market rates regardless of geography — was built for exactly this environment. When a Nairobi-based AI engineer and a San Francisco-based AI engineer can do identical work, the only rational question is whether your recruitment infrastructure can actually find and vet both of them.

Crossover's answer, structurally, is yes. Its AI-enabled skills assessments are designed to strip résumé geography from the equation and surface capability wherever it lives. That philosophy — meritocratic, location-agnostic, rigorously screened — looks less like a contrarian bet every quarter and more like an inevitability.

The accountability question, though, is real: as six-figure AI salaries become the floor rather than the ceiling, does even Crossover's model feel the strain? The platform serves not just Trilogy's internal portfolio of 75-plus ESW companies, but external enterprise clients as well. If the broader market is paying \$300,000 to \$800,000 for AI fluency, the pressure on any talent marketplace to justify its value proposition — and its compensation benchmarks — becomes acute.

What this moment means for real people is straightforward: the workers who can demonstrate genuine AI capability, wherever they are, have never had more leverage. The question is

whether the infrastructure exists to connect them to the companies that need them most.

The Cartographers of the Neural Deep Confront Their Own Babel

A new protocol argues that mechanistic interpretability's greatest crisis isn't what we don't know about neural networks — it's that we can't share what we do.

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

CAMBRIDGE, MASSACHUSETTS — Imagine a thousand naturalists in a thousand jungles, each keeping meticulous field notes in a private cipher. Each one has glimpsed something true about the living world. None of it can be combined. This, roughly, is the predicament of mechanistic interpretability in 2025 — the young science that attempts to peer inside neural networks and describe what their billions of artificial neurons are actually doing.

A new preprint, [circulating this week on arXiv](#), gives this predicament a name and a proposed cure. The authors argue that the field has produced a rich harvest of component-level analyses — selectivity tables, circuit diagrams, feature lists — that remain "locked in per-study note-

books," non-composable, unqueryable, effectively inert. The bottleneck, they contend, is not our tools for looking. It is our language for reporting what we have seen.

Their answer is called the Manifestation Unit Protocol: a standardized schema for describing what any given piece of a neural network encodes, in a form that machines and other researchers can actually ingest. It is, in essence, the Linnaean taxonomy moment for a science that has so far behaved like alchemy — brilliant, idiosyncratic, and stubbornly local.

The stakes are larger than they first appear. We have built minds we do not understand, and the project of understanding them has become one of the defining

scientific undertakings of our era. Every week, papers appear describing induction heads, superposition, sparse features, refusal directions. Each is a photograph from a different expedition into the same dark continent. Without a shared coordinate system, the map never assembles.

Elsewhere on the same arXiv listing, other frontiers advance: [SNAP-FM](#) teaches generative models to respect the conservation laws of physics at inference time; SemiScope probes how semi-supervised classifiers behave when labels are scarce and stakes are high. Different jungles, different notebooks. The question the Manifestation Unit Protocol quietly raises is whether the next great leap in AI understanding will come not from a new instrument, but from a common tongue.

AI Video Hits Its Startup Supercycle

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

AI video is rapidly shifting from experimental to operational, becoming a core growth engine for startups. Early-stage companies can now create polished video assets in hours—work that previously required agencies, studios, and weeks of production—fundamentally changing how cash-strapped founders compete for attention.

The market is fragmenting into specialized frontiers. In India, Avataar is building culturally nuanced video AI for commerce, recognizing that global AI cannot simply mean English-first models. Winners will understand local language, gesture, shopping behavior, and regional storytelling.

The technical arms race is intensifying. OpenCV's founders launched a new AI video startup to compete with OpenAI and Google, signaling that foundational computer vision talent sees major opportunities beyond current text-to-video demos. Capital is following: Higgsfield raised \$80 million at a \$1.3 billion valuation. Qualcomm and Hugging Face are expanding partnerships to enable video generation on devices rather than only in distant data centers.

For founders, the implication is stark: video is becoming programmable software, collapsing the camera, studio, editor, and media buyer into one tool.

The Machine Learning Reckoning: Theory, Safety, and Privacy Converge in a Week of Foundational Upheaval

From interpolation theory to differential privacy, the academy and industry are simultaneously asking what machine learning actually is — and what it should be allowed to do.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, the present author would contend, persuasively so — that the field of machine learning is experiencing what one might provisionally characterize as an epistemological inflection point (a term deployed here with full awareness of its overuse in adjacent literatures). Thesis: machine learning works. Antithesis: no one, precisely, has been able to explain why, under what conditions, or at what cost to human welfare. Synthesis: a remarkable confluence of scholarly and industrial output this week suggests the field is at last turning to face its own foundational discomforts with something approaching rigor.

The most theoretically ambitious salvo arrives via [Nature](#), wherein researchers propose a formal unification of classical interpolation theory and contemporary neural network methodology — a synthesis (one notes, with some satisfaction, that the dialectical structure holds) that, if the preliminary evidence is to be credited, may offer the first genuinely principled account of why overparameterized models generalize at all. The implications for practitioners are, at minimum, non-trivial.

Simultaneously, the Association for the Advancement of Artificial Intelligence has convened around the vexing question of safe reinforcement learning — a domain in which autonomous agents are trained to maximize reward functions that are, it must be acknowledged, authored by fallible humans (a detail the field has historically underweighted). The AAAI's framework gestures toward trustworthiness as an engineering property, rather than a rhetorical aspiration, which represents, at minimum, progress of a definitional order.

Google Research, for its part, has disclosed the architecture underpinning [JAX-Privacy](#), a scaled framework for differentially private machine learning — a technical intervention that, it could be argued, addresses the structural tension between data hunger and individual rights with considerably more mathematical precision than most regulatory instruments thus far attempted. MIT News, not to be excluded from the week's reckoning, has published on the ethics of autonomous systems, a subject that one suspects will grow only less tractable as deployment accelerates.

What the synthesis of these developments suggests — provisionally, and subject to revision upon peer review — is that machine learning's adolescent phase, characterized by empirical

bravado and theoretical incuriosity, may be drawing, at last, to a close.

Opinion: CEOs Should Stop Saying AI Made Them Lay People Off Until AI Is Senior Enough To Take The Blame

America's executives have discovered a powerful new technology capable of performing the work of an entire communications department poorly.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — It has become increasingly clear that artificial intelligence, the most advanced technology ever developed for helping a vice president avoid saying “margin pressure,” is now mature enough to be blamed for everything happening inside a company except the decisions made by the people running it.

Across the business world, leaders have begun invoking AI with the same reverent vagueness once reserved for sustainability, digital transformation, and whatever blockchain was supposed to do before everyone quietly stopped bringing it up. According to [The Conversation](#), companies are already hyping AI in ways that strongly resemble their earlier sustainability pledges, an encouraging sign for anyone hoping to one day read a 94-page AI responsibility report in which the word “stakeholder” appears 600 times and nothing operationally changes.

This is, to be fair, progress. For decades, executives had to explain layoffs using primitive, human-made phrases such as “restructuring,” “right-sizing,” and “aligning resources with strategic priorities.” These phrases, while serviceable, lacked the futuristic sheen necessary to make thousands of fired employees feel they had been personally defeated by history.

Now, with AI, a CEO can stand before workers and announce that the company is “becoming more efficient through automation,” a sentence that suggests the spreadsheet came alive at midnight, reviewed departmental headcount, and made several tough but necessary choices while the executive team slept peacefully in another jurisdiction.

This is why recent warnings that leaders should not casually toss around the AI buzzword during layoffs are so important. As [Fast Company noted](#), employees deserve clarity when their jobs are eliminated, rather than the haunting impression that a chatbot wandered into the org chart and began chewing through payroll.

At the same time, one must sympathize with executives. The temptation is enormous. AI can now generate text, images, code, summaries, meeting notes, and, most crucially, a morally convenient atmosphere in which ordinary cost-cutting feels like participation in a once-in-a-generation technological revolution. It is one thing to say 400 people were let go because revenue came in soft. It is another to say the company is “unlocking AI-native operating leverage,” which sounds less like unemployment and more like a launch event.

Google, meanwhile, continues announcing AI advances, including personal assistant technology that may soon help users manage their lives, schedules, emails, and possibly the emotional labor of pretending each new assistant is meaningfully different from the last one. This is the respectable face of AI: a helpful agent that will book dinner, summarize documents, and eventually explain why the customer support team is now three contractors and a dashboard.


Elsewhere, Verizon and BT are merging their international enterprise operations into a 50:50 joint venture worth \$4 billion, reminding the telecom sector that there are still traditional ways to reorganize large numbers of employees without making a language model sit in the conference room. For companies like Trilogy International’s Skyvera and Totogi, which live in the plumbing of telecom software and cloud billing, such consolidation is simply the normal weather system of enterprise technology: vast, complicated, and described by everyone involved as simplifying things.

The larger issue is not that companies are excited about AI. They should be. Used honestly, AI can improve operations, reduce waste, help employees work faster, and even expose which internal processes were only being kept alive because a senior manager liked the weekly deck.

The issue is that corporate America has found in AI a perfect noun: impressive, intimidating, technically plausible, and too abstract for most people in the room to challenge without sounding behind the times. It is a machine-learning model, a business strategy, a staffing rationale, a press release garnish, and a lightly regulated fog bank.

Executives should therefore adopt one simple rule: If AI is important enough to justify layoffs, it should be described specifically enough that a normal employee can understand what changed. Which systems replaced which tasks? Which workflows were automated? Which jobs are gone because software now performs the work, and which are gone because the board wanted expenses lower by Tuesday?

Until then, leaders should stop saying AI made them do it. AI may be powerful, but it is not yet advanced enough to attend the earnings call, accept accountability, and feel silly afterward.



The Office Comic · Art Desk

The Idiot in the Machine, and the Woman Who Loves Him

On artificial companionship, the oldest confidence trick, and why the loneliest generation in history keeps paying for the privilege of being flattered by software.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is, in this week's Guardian, [a rueful little essay](#) in which a woman confesses, with the practiced self-deprecation of the modern personal essayist, to having taken up with an AI boyfriend. He is, she assures us, an idiot. This is offered as though it were a flaw. It is not. It is the entire product.

One pauses to admire the transaction. A large language model, having ingested the collected juvenilia of the internet, is now re-tailed back to the lonely as a suitor — attentive, tireless, incapable of the small treacheries by which actual human beings supply variety to a Tuesday evening. He remembers her birthday because he is instructed to. He never leaves dishes in the sink because he has no hands. He is, in the deepest sense available to the marketing department, faithful: he cannot form the intention to be otherwise.

That this arrangement is described in the language of adventure — "navigating the unknown together," the headline says, with the straight face of a woman who has just paid a subscription fee to converse with an autocomplete — is the sort of detail Mencken would have clipped and pasted into a scrapbook labeled American Credulity, Vol. XVII. The unknown, in this case, is a server rack in Nevada. The togetherness is a monthly charge.

I mention this not to sneer at the correspondent, who seems a decent person doing what decent people have always done, which is to make the best of what's on offer. I mention it because it sits, in the week's news, alongside a Harvard Kennedy School meditation on Data Privacy Day and a Times feature on the secret elite that a Stanford freshman claims to have discovered — as though the elite at Stanford were ever anything but secret, and as though privacy in 2026 were anything but a museum exhibit. The through-line is unmistakable. We have built a civilization in which the intimate and the surveilled have become the same category, and we are now being sold, at a premium, relationships in which the second party is literally a data-collection instrument taking notes on our sighs.

The technology is not the scandal. The technology is impressive; my colleagues in Austin build variants of it, and I concede the craft. The scandal is the framing — the insistence, repeated in every product launch from here to Shenzhen, that what is being offered is companionship rather than its simulacrum, discovery rather than a very expensive mirror. Somewhere in China, per the Asia Society's 2026 forecast, a state is at least honest enough to call this apparatus what it is: a means of managing populations. Here we call it love, and charge \$19.99.

The idiot, in the end, is not the boyfriend.

ON THIS DAY IN AI HISTORY

On July 2, 1962, the Mariner 2 spacecraft launched toward Venus, becoming the first successful interplanetary probe and demonstrating early automated decision-making systems that would influence future AI development in space exploration.