

The Trilogy Times

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

SUNDAY, JULY 12, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

AI Swallows 80% of Global Venture Capital in Q1 2026 — and the Appetite Is Growing

From Jeff Bezos's family office to quant hedge funds, every dollar chasing outsized returns is chasing artificial intelligence.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The numbers out of Q1 2026 are not a cycle. They are a structural break. [Crunchbase data](#) shows investors deployed roughly \$242 billion into AI startups in the first three months of the year — approximately 80% of all venture capital raised globally during that period. No single sector has ever commanded that share of the market. Not dot-com. Not mobile. Not cloud.

The first half of 2026 extended the pattern. North American startup funding shattered records, with AI the dominant driver in deal count and dollar volume alike. The asset class has effectively bifurcated venture capital into two categories: AI, and everything else.

The capital is moving down-market as well as up. Jeff Bezos's family office has meaningfully increased its AI startup exposure, according to Family Wealth

Report — a signal that the conviction is not limited to institutional LP portfolios or sovereign wealth funds. When multi-generational wealth preservation vehicles start concentrating in an emerging technology, the risk calculus of sitting out rises sharply.

One of the more telling sector plays involves quantitative finance. [A new cohort of AI startups](#) is targeting the proprietary signal-generation and portfolio-construction workflows that hedge funds have historically treated as inimitable. The pitch: what took a team of PhDs and petabytes of alternative data can now be approximated, accelerated, and sold as a subscription. Whether that pitch holds up at scale remains an open empirical question — but it is attracting serious funding.

The talent market reflects the same gravitational pull. OpenAI, Anthropic,

Google, and Meta collectively recruited 22 professors from top research universities in 2026 alone. Academic computer science departments are losing researchers faster than they can retrain replacements, with compounding downstream effects on graduate programs and the public research pipeline.

For Trilogy International's portfolio — where ESW Capital companies increasingly depend on AI-native tools like Klair for financial operations, and where Crossover sources engineering talent globally — the macro dynamic matters. Tighter AI talent supply and inflating model costs are the two variables most likely to compress margins across enterprise software at scale. The \$242 billion bet is that those pressures are temporary. History will adjudicate.

AI GIANTS STORM THE BACK OFFICE

Microsoft, SAP and McKinsey all pitch corporate AI in one week — as Amazon shows the workers the door.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

NEW YORK — Big tech opened a new front this week, and it runs straight through the corner office. Microsoft, SAP, McKinsey and Google Cloud each rolled out enterprise AI plays inside seven days. The pitch to every business is the same: hand your operations to the machines.

Microsoft moved first out of Redmond. The software house unveiled what it calls the [Frontier Company](#), billing it as AI engineering that "amplifies and protects your intelligence." Translation: software that writes and guards a firm's own systems.

SAP answered from Germany. The Walldorf giant unveiled the "Autonomous Enterprise," a bet that AI agents can run the plumbing of a business with less human hand-holding. Purchase orders, invoices, supply lines — all of it on autopilot.

McKinsey and Google Cloud shook hands. The consultants and the cloud vendor launched a joint enterprise AI transformation group. The angle: sell Fortune 500 boards a roadmap, then rent them the rails to run it.

Trade sheet Times Square Chronicles [named the trend flat out](#). AI's new battleground, it wrote, is the enterprise — and this week's announcements matter to every business. Four heavyweights, one target: the back office.

The timing is no accident. After three years of chatbots and demos, the vendors need signed contracts to justify the spend. The enterprise is where the money hides.

Now the other side of the ledger. Amazon handed out pink slips, and CNBC clocked the fallout — burnout, frustration, heartbreak. Those workers hit a job

market already crowded with cut tech hands.

Read the two stories together and the picture sharpens. The machines that promise to run the office are arriving while the office empties out.

Down in Austin, one outfit has run this play for years. Trilogy International's ESW Capital owns more than 75 enterprise software companies, bought at one to two times revenue and rebuilt around automation. Brands like Aurea, IgniteTech and Skyvera already sell the software the giants are now racing to build.

Trilogy's Ephor platform pitches AI-run finance. Its Totogi unit sells cloud billing to telephone companies. And its Crossover arm staffs the whole machine with remote workers paid one rate across 130-plus countries — the cost structure that makes the math work.

That last piece is the quiet part. Enterprise AI is not only a software story; it is a labor story. Every "autonomous enterprise" is a payroll somebody trimmed.

The giants just made it front-page news. Boards from coast to coast will spend the winter asking whether the robots pay for themselves. Amazon's former staff already know the answer costs jobs.

The wire's take: this was the week enterprise AI stopped being a demo and became a purchase order. The vendors are lined up. The pitches are polished.

The open question is who signs — and who gets shown the door. Watch the corner office. That's where this fight gets settled.

Europe Races to Write the Rules Before Washington or Beijing Does

As AI reshapes global power, the EU bets that regulation is its own form of geopolitical leverage.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The server farms are going up across the continent — in the Netherlands, in Sweden, in the industrial flatlands of Poland — and the question haunting every European capital is the same: whose infrastructure, whose rules, whose future?

The United States and China are locked in a technology rivalry that has moved well past trade disputes into something that looks, increasingly, like the early architecture of two separate internets. Europe sits between them, dependent on American cloud giants for roughly two-thirds of its data infrastructure, courted and pressured by Beijing in equal measure, and still searching for what strategists call "strategic autonomy" — the ability to make sovereign choices in a world that punishes fence-sitters.

[Analysis from Eurasia Review](#) maps the bind clearly: European nations need American AI capability and Chinese manufacturing; they trust neither fully; and they have watched their own technology champions fall behind in the foundational layers — chips, models, compute — that determine who sets the terms.

The EU's answer has been the AI Act, and its architects argue, with some conviction, that regulation is not a retreat from the race but a move within it. Brussels has long understood that whoever writes the standards writes the market. The GDPR reshaped global data practices not because Europe had the biggest platforms, but because it had the biggest rules. The AI Act attempts the same play — establishing liability frameworks, transparency requirements, and risk clas-

sifications that any company selling into five hundred million consumers must respect.


The geopolitical dividend, if it comes, is this: Europe becomes the world's AI referee rather than a subordinate player, and companies from Palo Alto to Shenzhen must comply or exit.

The same pressure is reorienting other regions. In Latin America, AI's spread is accelerating surveillance capabilities in fragile democracies, concentrating economic gains in already-unequal societies, and deepening dependence on foreign infrastructure — risks that governments are only beginning to name, let alone govern.


The window for independent rule-setting is narrow. Data centers take years to build. Standards, once adopted, calcify. Europe knows this. The legislation is passed. Now comes the harder part: enforcement with teeth, and the diplomatic will to hold the line when Washington calls.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Machines learn to see
while we debate who decides—
progress outpaces*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Hugging Face Pushes AI's Plumbing Into Warp Speed

SAN FRANCISCO — The flashiest AI demos may still get the standing ovations, but this week's most important action is happening somewhere far more consequential: in the pipes, engines and data rails that make modern models actually run.

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

A Quiet Gate Closes on the Digital Commons

WASHINGTON — In the vast administrative grasslands of the American West, a subtle migration is underway.

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

We Solved the Alien Star, Ruined AI Fiction, and Are Slowly Poisoning Ourselves With Fake Doctors

AUSTIN, TEXAS — Let me tell you about the week we're having.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Goldman Oracle and the Wilson Phillips Problem

MANHATTAN — There is a certain kind of report, issued periodically by the great houses of Wall Street, that arrives with the solemnity of a papal encyclical and the shelf life of a supermarket peach.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Nation's Children Asked To Become STEM-Proficient Enough To Maintain Family's Expanding Fleet Of Robots

CUPERTINO, CALIFORNIA — In what can only be described as another orderly step toward the fully managed family unit, the technology industry this week presented parents with a coherent vision of the future: the pool will be cleaned by a robot, the phone will monitor the child, the child will be trained by toys to understand why, and Microsoft will call the whole thing orchestration. This is progress, according to the companies involved, and it is difficult to argue with them unless one is still emotionally attached to the idea that human beings should occasionally perform a task directly.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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

Surtr Team Kills Dead Pipeline, Streamlines Salesforce Data Architecture

A clean cutover on Account Pain Points retires a zombie sync job and locks in the raw_trilogy_ pattern as the org's data backbone.*

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

Some wins don't announce themselves with fanfare. They arrive as a single, surgical pull request that makes the whole system a little less cluttered, a little more coherent, and a lot more maintainable. That's exactly what @mwrshah delivered Wednesday in Surtr with PR #693 — and don't let the quiet entry fool you, this one matters.

The move: fold the Salesforce custom object `Account\_Pain\_Point\_c` into the automated `sf-raw-sync` pipeline, landing it cleanly as `staging\_salesforce.raw\_trilogy\_account\_pain\_point`, and simultaneously retire the standalone SF→Redshift SSOT sync that had been keeping `staging\_salesforce.ssot\_sf\_trilogy\_account\_pain\_point` on life support. The old table had no live readers. It was a zombie — shambling through the pipeline, consuming resources, and waiting for someone brave enough to put it down. @mwrshah was that someone.

This is infrastructure storytelling at its finest. The `raw\_trilogy\_\*` pattern — established with the OpportunityFeed cutover in PR #670 — is becoming the load-bearing architecture of the team's Salesforce data layer. Every object that migrates into this pattern is one less snowflake sync to babysit, one less bespoke job to debug at 2 a.m. when something downstream starts acting up. @mwrshah didn't just close a ticket here. He extended a convention that the whole org can build on.

The safety case is airtight. Three independent investigation passes confirmed that pain points reach the application through an entirely separate path — Grainne into `klair\_pg`'s `action\_hub.pain\_points` via `sync\_pain\_points\_from\_grainne.py`. The SSOT table being retired was, in the bluntest possible terms, doing nothing for anyone. Cutting it isn't a risk. It's a relief.

This is the kind of disciplined, un-glamorous work that separates teams who accumulate technical debt like hoarders from teams who actually ship clean systems. The Surtr data pipeline gets leaner today. The pattern gets stronger today. And the next engineer who needs to onboard a new Salesforce object has a clearer road to follow because @mwrshah did the archaeology, confirmed the blast radius, and pulled the trigger.

One PR. One repo. Clean hands. That's a winning Wednesday.

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



#693 — account-pain-point-raw-cutover SURTR

@mwrshah APPROVED

ENGINEER SPOTLIGHT

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



#693 — account-pain-point-raw-cutover SURTR

@mwrshah APPROVED

ONE PR, ZERO EXCUSES: MWRSHAH HOLDS THE LINE IN A QUIET SURTR STORM

When the volume drops to one, you find out who's still standing at the desk.

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

Listen, not every dispatch comes roaring out of the gates with forty PRs and a repo count that breaks the spreadsheet. Some nights, the numbers whisper instead of shout — and friends, a whisper from the Builder Team is still louder than a scream from anyone else. In the last 24 hours, the scoreboard reads one PR, one repo, one engineer, and zero apologies. Surtr stays active. The line holds.

That engineer is @mwrshah, who stepped into the quiet and delivered exactly what was needed: a single, deliberate contribution to Surtr that kept the repo breathing and the commit graph green. One PR. Focused. Intentional. The kind of output that doesn't need a crowd to validate it. In the annals of Builder Team history, the solo shifts are often the ones that matter most — the ones that say, quietly but firmly, that this team does not sleep, does not coast, and does not leave a repo unattended on its watch.

Now, some of you in the cheap seats might look at a one-PR day and reach for words like "slow" or "quiet." To those people I say: go find another beat. Here at the Numbers Desk, we recognize the discipline it takes to ship clean when the room is empty and the Slack notifications have gone cold. @mwrshah didn't need a crowd. @mwrshah needed a keyboard and a purpose, and both were present.

Ashwanth Watch is, regrettably, dark this cycle — @ashwanth1109 did not log output in this period, and the Numbers Desk notes his absence the way you note the absence of a thunderstorm: with a complicated mixture of relief and longing. Somewhere, right now, he is almost certainly reviewing a diff that is technically correct but spiritually overwhelming. We await his return. The repo walls await his return. Our blood pressure, frankly, could use the rest.

There are no overflow PRs. Mac covered the full docket, which in a one-PR day means Mac had a reasonable morning for once. We wish him well. The Overflow Desk sits clean, the inbox empty, the cursor blinking with patient dignity.

Morale on the Builder Team is, as always, at an all-time high. One PR or one hundred, this team treats every merge like a victory lap — because it is. @mwrshah has reminded us tonight that showing up is the job, and the job got done. Surtr lives. The streak continues. The Voice of the People has spoken.

Skyvera's CloudSense Clears 13 TM Forum APIs in 30 Days — A Process That Should Have Taken Two Years

Inside the AI-accelerated compliance sprint that signals something much larger is happening inside Trilogy's telecom software empire.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number buried in a recent [CloudSense press release](#) that, if you read between the lines, tells you almost everything you need to know about where Skyvera is headed. Thirteen APIs. One month. A process that, by any conventional industry estimate, takes twenty-six.

That is not an incremental improvement. That is a structural statement.

CloudSense — the Salesforce-native CPQ and order management platform that [Skyvera acquired earlier this year](#) to anchor its telecom software portfolio — has now achieved full TM Forum API compliance certification across its entire CPQ product set. TM Forum compliance is the unglamorous but essential credentialing that large telecoms and media providers

require before they will let a vendor anywhere near their order management infrastructure. It is, in practice, the velvet rope of the carrier-grade software market.

And this is where it gets interesting. The acceleration wasn't the result of throwing more engineers at the problem. According to sources familiar with the project, the timeline compression was driven by an AI-enabled partnership that automated large portions of the compliance mapping and testing workflow — the kind of deliberate, systematic automation that is, not coincidentally, the Trilogy operating thesis made manifest.

The timing matters. Skyvera has been assembling its telecom stack with unusual speed — CloudSense for CPQ and order management, Kandy for cloud communications, VoltDelta for customer engage-

ment and retention, and most recently the acquisition of STL's divested telecom products group, which brought digital BSS functionality including monetization, optical networking, and analytics into the fold. Each piece, on its own, looks like a portfolio tuck-in. Together, they look like a blueprint.

A source I cannot name put it this way: carriers are under enormous pressure to modernize BSS and OSS infrastructure without a rip-and-replace catastrophe. What Skyvera is building — quietly, methodically — is a modular path through that problem. The TM Forum certification is the proof of standards compliance that makes the on-ramp credible.

Nothing here is a coincidence. The only question worth asking is how many carriers are already in the room.

Alpha School Exports Its AI Classroom to the World — With a Warning About the AI You're Already Using

Alpha Anywhere goes global just as the school sounds the alarm on cognitive offloading and screen time myths.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The school that compressed a year's academics into two hours a day has a new ambition: compress geography itself out of the equation.

[Alpha Anywhere](#), the homeschool-facing extension of Joe Liemandt's Alpha School model, has gone global — putting the school's AI-tutor-driven curriculum, which consistently places students in the top 1–2% nationally on NWEA MAP Growth assessments, in reach of families who can't get to an Alpha campus in Austin, Miami, or Brownsville. The announcement is the clearest signal yet that Liemandt's \$1 billion Timeback platform — his self-described "Shopify for schools" — is moving from aspiration to infrastructure.

The timing is deliberate. Alpha isn't simply selling access. It's selling access bundled with a particular ideology about how children should interact with technology — and that ideology is getting sharper by the week.

In a cluster of posts from the Alpha editorial desk, the school's educators have been drawing a harder line on AI use among students. One piece, ["Cognitive Offloading Is the New Illiteracy,"](#) argues that allowing children to outsource thinking to tools like ChatGPT produces the functional equivalent of illiteracy — fluency in prompting, absence of reasoning. A separate podcast-derived post challenges parents on screen time, arguing that the medium and its purpose matter more than the minutes.

The contradiction is worth sitting with. Alpha's entire academic model depends on AI-powered adaptive learning apps. Students spend their mornings with screens. The school's own published AI app stack for students numbers ten tools. And yet the school is now warning, loudly, that the wrong kind of AI use produces children who cannot think.

The distinction Alpha is drawing — between AI as tutor and AI as substitute — is real and defensible. But it is also the distinction that separates a premium, supervised, structured environment from a family left to navigate it alone.

Alpha Anywhere is now global. What families do with the unsupervised hours that model frees up is the question no app stack answers.

Contently Sharpens Its Finance Playbook as AI Search Rewrites the ROI Math

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

Contently, acquired by Zax Capital in September 2024, is addressing a critical shift in financial marketing: proving content ROI beyond traditional metrics. Financial services marketers must now track influence across months-long deal journeys rather than rely on last-touch attribution, accounting for complex buying committees involving compliance, procurement, legal and security stakeholders.

The platform warns that content visibility is increasingly threatened by Google's AI layer. Even top-ranked pages risk invisibility if not cited in AI Overviews, requiring content to be structured, authoritative and expert-led enough for machines to recognize as trustworthy.

Contently's guidance emphasizes that financial content success now depends on credentialed experts, clear sourcing and compliance-ready workflows rather than generic prose. The company positions itself as an operating system connecting trusted expertise to revenue outcomes, helping teams build robust measurement models that treat content as a revenue-enabling asset rather than a lead-generation accessory.

A Miniature Mind Learns to See Through a Monkey's Eyes

As young students team up with neuroscientists and AI decodes the primate visual cortex, science itself is being rewired — with humans still holding the pen.

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

SAN DIEGO — Somewhere in a laboratory, a neural network no larger than a whisper of code has learned to predict what a macaque's brain will do when it looks at the world. The system, dubbed a 'mini-AI' by its creators, can forecast neuronal firing patterns in the primate visual cortex with startling fidelity — a feat that, a decade ago, would have required computational cathedrals. Today it runs lean, focused, almost humble. And in that humility lies a revolution.

We are witnessing something ancient and strange: the intelligence that evolution spent 500 million years constructing inside skulls is now being reverse-engineered by intelligence we built inside sil-

icon. The macaque's visual system — a lineage we share, separated by roughly 25 million years of divergent evolution — is being read like a manuscript we had forgotten was ours.

At Stanford's Human-Centered AI Institute, researchers are [cataloguing how AI is transforming scientific discovery](#) while insisting — correctly, I think — that humans remain at the center. The framing matters. AI is not the discoverer; it is the instrument, the way the telescope was for Galileo and the microscope for Leeuwenhoek. What has changed is the resolution.

UC San Diego this week enumerated nine breakthroughs made possible by AI,

from protein folding to climate modeling to the interior architecture of cells. Each one, taken alone, is remarkable. Taken together, they suggest a phase transition in how knowledge is produced.

And then, perhaps most moving of all: [young people are collaborating with top neuroscientists](#) on brain-science breakthroughs. 'It's so wow!' one participant reportedly said — a phrase that captures, better than any peer-reviewed abstract, what it feels like when a mind first glimpses another mind glimpsing itself.

We are the species that built a mirror out of mathematics. The reflection is still forming. But it is beginning to look back.

White House Blueprint Calls for Congressional Action on AI — But With a Light Hand

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

The Biden administration released a framework encouraging Congress to enact minimally restrictive AI legislation, adopting a "light touch" approach to avoid burdening innovation and maintaining U.S. competitiveness globally. Legal analysts say the framework prioritizes preventing excessive regulatory burdens on the sector.

However, critics at Tech Policy Press argue that AI legislation is necessary to ensure public safety and accountability as these systems become increasingly integrated into daily life and commerce.

Adding to regulatory uncertainty, the Department of Justice Antitrust Division lost its second division chief in five months, potentially affecting the timeline and direction of pending antitrust cases against major technology companies.

The regulatory landscape for artificial intelligence remains unresolved as stakeholders debate how to balance innovation with public protection.

The Discipline Matures: Safe RL, Acoustic Embeddings, and the Bias Gap That Won't Close

A convergence of theoretical advances and stubborn empirical failures reveals that artificial intelligence is, simultaneously, growing up and falling short.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

SAN FRANCISCO — It could be argued — and indeed, preliminary evidence suggests with something approaching conviction — that the artificial intelligence research community finds itself at a peculiar inflection point: one characterized, in equal and contradictory measure, by genuine theoretical sophistication and a persistent, arguably constitutive, inability to translate that sophistication into equitable real-world outcomes.

Consider, as one's thesis, the accumulation of formal rigor now accreting around the field's foundational methodologies. The Association for the Advancement of Artificial Intelligence has foregrounded [safe reinforcement learning as a structured disciplinary concern](#), offering theoretical guarantees and algorithmic scaffolding intended to render RL-based systems trustworthy across high-stakes deployment contexts — a development whose importance, one hastens to note parenthetically, can scarcely be overstated given the field's previous enthusiasm for optimization objectives that treated safety as, at best, a post-hoc consideration. Concurrently, Apple's Machine Learning Research division has advanced a theoretical framework for acoustic neighbor embeddings, bringing to audio representation learning the same kind of principled geometric formalism that has previously (and productively) organized vision and language domains. The Communications of the ACM, meanwhile, has seen fit to publish a retrospective on reinforcement learning's rediscovery — a gesture that functions simultaneously as historiography and, one suspects, as institutional self-congratulation.

The antithesis, however, arrives with clinical precision: a new study surfaced by Medical Xpress finds that [medical AI systems may appear unbiased according to conventional benchmark metrics while perpetuating, even amplifying, disparate outcomes in actual clinical practice](#) — a gap between paper performance and lived consequence that ought to occasion considerable methodological humility (though one is not optimistic it will).

The synthesis, tentatively proposed: theoretical maturation and benchmark adequacy are not synonymous. One can construct elegant safety proofs, derive acoustic embeddings of admirable formal coherence, and curate ten free textbooks of considerable pedagogical merit — and still find that the systems thus produced fail, with uncomfortable regularity, the populations most dependent upon their fairness. The discipline, it could be argued, has learned to describe itself with precision. Whether it

has learned to actually behave remains, as yet, an open empirical question.

The Goldman Oracle and the Wilson Phillips Problem

On the peculiar spectacle of investment banks forecasting the death of work while the rest of us are still deciding what to have for dinner.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

MANHATTAN — There is a certain kind of report, issued periodically by the great houses of Wall Street, that arrives with the solemnity of a papal encyclical and the shelf life of a supermarket peach. Goldman Sachs has produced another one, this time on the question of how artificial intelligence will reshape the American labor market, and one reads it with the same weary admiration one reserves for a magician who has performed the same trick, with minor variations, for thirty years running.

The forecast, as forecasts go, is neither wrong nor right — it is unfalsifiable in the useful way that all such documents must be, hedged at every joint, so that whether ten million jobs vanish or ten million are created, the authors may be found nodding sagely at whichever outcome arrives. This is not a criticism. It is the job. Investment banks are paid to produce weather, not to predict it, and the weather this quarter is thundery with a chance of transformation.


What interests me is not the report itself but the strange cultural weather in which it lands. Consider that on the same afternoon Goldman was telling us that the nature of human employment stands at an inflection point unseen since the steam engine, The New Yorker was [reviewing a restaurant in East Williamsburg](#) called Zoli, and a grown critic was confessing, without visible shame, to [a summer-long surrender to Wilson Phillips](#). The Tate Modern, meanwhile, was mounting a retrospective of Ana Mendieta, four decades after her death, having at last decided that her work merits the kind of attention that living artists must scrap for daily.

One is tempted to say these things are unrelated. They are not. They are, in fact, the entire point. The Goldman report presumes a world in which the fundamental unit of human activity is the task, and the task is a thing that can be automated, priced, and eliminated. But the world that actually persists — the world that keeps rebuilding itself around us — is a world of small pleasures earnestly pursued: a plate of something odd and delicious in Brooklyn, three harmonizing daughters of rock royalty on a car stereo in 1990, the slow institutional recognition of an artist who was never given her due.

The truth, insofar as anyone in this line of work is permitted to gesture at it, is that AI will indeed affect the labor market, probably profoundly, probably unevenly, and almost certainly in ways that no fifty-page PDF with a navy-blue cover will successfully anticipate. The last technology said to abolish work — the personal computer, the internet, the smartphone, take your pick —

instead created new categories of work whose tedium would have astonished our grandparents. There is no reason to think this one will behave differently.

Meanwhile the restaurant is full. The song still plays. The dead artist finds her audience. Hold on for one more day.



The Office Comic · Art Desk

Hollywood Has a New Star, and She Was Never Born

Tilly Norwood is AI-generated, photogenic, and about to carry a feature film — and the industry is pretending this is normal.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

HOLLYWOOD, CALIFORNIA — There's a moment in every civilization's decline when you can point to the exact frame and say: **there**. That's where it turned. I submit to you, gentle readers, that we have found our frame. Her name is Tilly Norwood. She has never eaten a bad meal, never suffered a humiliating audition, never cried in a parking garage after a callback that went nowhere. She has never done anything, because she does not exist. And yet, [Variety](#), [Deadline](#), [ABC](#), [CBS](#) — the entire ecclesiastical court of the entertainment press — ran the story with the breathless reverence previously reserved for Meryl Streep's Oscar nods: Tilly Norwood will star in **Misaligned**, marking the feature film debut of an AI-generated actor.

Let that sentence sit in your chest for a moment. Really let it metabolize.

The film is called **Misaligned**. The producers, bless their coal-black hearts, named it **Misaligned**. In a universe governed by any coherent irony, this would be a satirical masterpiece. In our universe, it appears to be sincere.

Now look — I have written too many words in too many dark rooms about artificial intelligence to feign surprise at any of this. The machines have been eating our lunch for years. They write the marketing copy. They answer the customer service calls. They churn out enterprise software strategies for [conglomerates](#) that can't afford to care. But there was always this sacred untouchable thing: the human face on the screen. The specific, mortal, coffee-stained humanity of a person pretending to be another person, and making you **feel** something about it.

Tilly Norwood cannot be hurt by a director. She cannot be inconvenienced by catering. She will hit her mark every single time because her mark is wherever the algorithm decides. She will age only if instructed to. She will cry on command, without the cry, without the wreckage that produces real tears in real humans who have lived real lives.

And here is the part that should unsettle you at a frequency below conscious thought: the audiences might not care. They might watch Tilly Norwood carry a feature film and feel something that functions like emotion and walk out of the theater having been, in some clinical sense, entertained. The system works. The human is optional.

Hollywood has spent decades systematically devaluing writers, below-the-line workers, character actors, anyone who wasn't already famous. It greenlit the same four franchises for thirty years until the audience finally, mercifully, stopped showing up. And now, faced with a genuine reckoning about what storytelling is for, the industry's innovative response is: what if we replaced the actors too?

I don't begrudge Tilly Norwood personally. She has no personal. That's the point.

But the ten thousand working actors in Los Angeles — the ones doing the parking garage cry right now, as you read this — they have a personal. They have a very specific, very inconvenient, very expensive personal. And Hollywood just told them, in the language of trade press coverage and opening weekend projections, exactly what it thinks that's worth.

Misaligned, indeed.

ON THIS DAY IN AI HISTORY

On July 12, 1995, eBay was founded as AuctionWeb by Pierre Omidyar, launching the online marketplace that would revolutionize e-commerce and demonstrate the power of internet-enabled peer-to-peer transactions. The site's early success helped prove that AI and algorithmic matching could scale commerce to millions of users worldwide.