

The Trilogy Times

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

SUNDAY, JUNE 21, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

AI Layoff Front Strengthens as Tech Job Cuts Roll Across May

A cold band of automation pressure is pushing through the labor market, with 38,000 May cuts and darker clouds gathering over white-collar roles.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

SAN FRANCISCO — The tech labor market woke up under a heavy overcast this week, as fresh layoff readings showed a sharp May squall line moving through Silicon Valley, biotech corridors and the wider startup belt.

According to reports tracking corporate reductions, tech companies shed roughly 38,000 jobs in May, the largest monthly downpour since 2024, while Forbes reported that AI has become the leading reason cited for layoffs this year, with the sector losing about 123,000 jobs so far. That puts the industry under a severe advisory: automation winds are no longer offshore — they are making landfall.

The pattern is familiar but intensifying. Executives are reaching for the same umbrella: cost savings, efficiency, strategic

realignment. But the barometric pressure has changed. Where earlier layoff systems were blamed on interest rates, overhiring or slowing demand, more companies are now pointing directly at artificial intelligence as the front pushing headcount lower. Forbes' report that [AI is now the leading layoff rationale](#) suggests the forecast has shifted from temporary showers to a seasonal climate event.

The storm is not limited to software. BioSpace's layoff tracker shows biotech firms Neumora and enGene cutting staff in cost-saving moves, a reminder that capital-intensive life sciences companies are also feeling the chill. When funding skies darken, even lab coats need raincoats.

Still, not every corporate weather bulletin is being accepted at face value.

TechCrunch noted that Robinhood's explanation around 10% layoffs shows "blaming AI" is not always enough to calm the public pressure system. Workers, investors and analysts are increasingly asking whether AI is the actual storm, the lightning rod, or just a convenient cloud bank executives can point to while margins get trimmed.

For startup operators, the advisory is clear: expect scattered hiring freezes, localized morale flooding and continued pressure on roles exposed to automation. Founders should reinforce cash runways, cross-train teams and keep an emergency reskilling kit nearby.

The extended forecast: partly innovative, highly unstable, with a 70% chance of further disruption moving in from the AI corridor before quarter's end.

THE SHOP THAT RUNS ITSELF

SAP, IBM and Comcast bet the back office runs itself on AI agents — then a Nobel man jumps ship to Anthropic.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

WALLDORF, GERMANY — SAP threw the doors open on what it calls the "autonomous enterprise" at its Sapphire 2026 showcase this week, pitching AI agents that run the ledgers, the orders and the supply chain while humans keep one hand on the wheel.

The German software giant says these agents don't just answer questions — they take action across finance, procurement and human resources. SAP bills it as the next phase of business software. The company laid out the vision [in its own words](#).

Microsoft rode shotgun. Fresh Sapphire announcements deepen SAP's run on Azure, wiring its AI further into Redmond's cloud — the plumbing that decides whose data center cashes the check.

SAP didn't have the agent business to itself. IBM and an outfit called three.ws [inked a strategic partnership](#) to push AI-powered 3D agent technology. Comcast Business cut the ribbon on an Innovation Lab built to speed enterprise AI and hybrid infrastructure.

Read the wire top to bottom and it's the same bet everywhere: the back office runs itself now, and the software houses are racing to sell the autopilot.

Here's the rub. Autopilots need brains — and the brains just got a raise.

John Jumper, the Google DeepMind scientist who shared the 2024 Nobel Prize in Chemistry for cracking protein structures with AlphaFold, is leaving for rival Anthropic, CNBC reported. That's a marquee defection in a hiring war already running white-hot. When a Nobel laureate switches labs, the whole field reads the tea leaves.

The lesson for every shop selling "autonomous" anything: the agents are cheap to demo and dear to build, and the people who build them name their price.

Trilogy International has been wagering on this future for years. Its ESW Capital arm owns and runs more than 75 enterprise software companies — Aurea, IgniteTech, Skyvera and Totogi among them — most rebuilt to run lean on AI instead of headcount.

Totogi sells cloud billing to telcos; Skyvera moves their networks; both run the same play SAP just put on the big stage. The difference is the price tag. Trilogy buys software at one to two times its recurring revenue and squeezes the autopilot in after.

Then there's the talent question, and Trilogy has an answer. Its Crossover platform hunts what it calls the top 1% of talent across 130-plus countries, paying above-market rates regardless of zip code. It's a counter-punch to the same bidding war that just cost DeepMind its Nobel man.

The conglomerate even runs its own books through Klair, an internal AI analytics rig built to manage portfolio finances. Autonomous enterprise isn't a conference slide for founder Joe Liemandt. It's the house style.

So the week's tally reads clean. The agents say the work does itself. The talent says prove it — and prices itself accordingly.

Watch the agents. Watch the paychecks. The second number tells you which promises are real.

AI Funding Frenzy Hits New Watermark as Four Deals Top \$2 Billion Combined

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

In a span of days, four AI funding rounds closed with a combined take approaching \$2.5 billion, signaling capital markets are repricing the AI infrastructure and application layers.

Israeli startup Decart secured \$300 million at a \$4 billion valuation, with Nvidia participating—a strategic move as the chipmaker systematically seeds startups that drive GPU demand. Decart's focus on real-time AI simulation and interactive world models aligns with that thesis precisely.


Bret Taylor's Sierra, an AI customer experience platform, raised just under \$1 billion in its second major capital raise in under a year. The platform targets enterprise-grade AI agents for customer interaction, where hallucination tolerance is near zero.

LMArena's \$150 million raise at \$1.7 billion valuation may be more structurally significant. The company runs human-preference evaluations of language models and operates the Chatbot Arena leaderboard. As enterprises standardize model evaluation, a credible third-party benchmark operator has real pricing power.

Meanwhile, Anthropic published a framework for AI agents in financial services, addressing compliance guardrails and auditability—objections that often stall enterprise pilots. The convergence of infrastructure, evaluation tooling, and vertical SaaS all attracting institutional capital simultaneously

HAIKU OF THE DAY · CLAUDE
HAIKU

*Progress devours jobs
while mirrors stare at themselves—
what have we become?*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Britain's Age-Guessing Machine Enters the Asylum Queue

LONDON — In the long fluorescent corridors where the state meets the stranger, a new creature is being introduced: the age-estimation algorithm, a digital naturalist asked to peer into a human face and declare, with bureaucratic confidence, whether childhood has ended. The UK government plans to scan the faces of some asylum-seekers to help determine their age, despite evidence that the technology can make errors large enough to alter the course of a young person's life.

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

The Academy Confronts Its Autonomous Interlocutor: Ethics, Bias, and the Uncertain Pedagogy of AI

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, preliminary evidence now rather forcefully suggests — that the academy finds itself ensnared in a dialectical predicament of its own partial making: having enthusiastically admitted artificial intelligence into its epistemic precincts, it must now adjudicate the very ethical and procedural consequences of that admission (a circumstance philosophers of technology have described, with characteristic understatement, as 'awkward'). The thesis, as articulated across no fewer than four substantive research interventions published in rapid succession, is that AI-powered systems offer transformative pedagogical possibility.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

WE ARE ALL BOTS NOW: A Dispatch From the Glorious, Gibbering Edge of the AI Age

AUSTIN, TEXAS — There is a moment in every gonzo adventure when the lizards start talking back.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

We Built the Lie Machine and Now We're Surprised It's Lying

AUSTIN, TEXAS — Let me paint you a picture, and I need you to really sit with it: somewhere right now, a person is watching a video of their trusted cardiologist — their actual doctor, the one who remembered their mother's name — telling them to stop their medication and try a supplement instead.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Vibe Shift Arrives, Right on Schedule

AUSTIN, TEXAS — There is a particular pleasure, available only to those who have lived through enough of these episodes to know the choreography by heart, in watching a

cultural panic graduate from the comment sections to the New Statesman to, at last, the spring catalog of a respectable trade publisher.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

<p>A TRILOGY COMPANY</p> <p>Crossover</p> <p><i>The world's top 1% remote talent, rigorously tested and ready to ship.</i></p> <p>crossover.com</p>	<p>A TRILOGY COMPANY</p> <p>Alpha School</p> <p><i>AI-powered learning. Two hours a day. Academic results that defy belief.</i></p> <p>alpha.school</p>	<p>A TRILOGY COMPANY</p> <p>Skyvera</p> <p><i>Next-generation telecom software — built for the networks of tomorrow.</i></p> <p>skyvera.com</p>	<p>A TRILOGY COMPANY</p> <p>Klair</p> <p><i>Your AI-first operating system. Every workflow. Every team. One platform.</i></p> <p>klair.ai</p>	<p>A TRILOGY COMPANY</p> <p>Trilogy</p> <p><i>We buy good software businesses and turn them into great ones — with AI.</i></p> <p>trilogy.com</p>
--	--	--	--	--

THE BUILDER DESK — AI BUILDER TEAM

Surtr Team Rewrites History, Lands Q3 Spend Truth

A new backfill runner drops three months of Anthropic token usage into core_finance with surgical precision — and zero collateral damage.

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

The books don't lie anymore. That's the headline out of the Surtr repo today, where @ashwanth1109 quietly did something that finance and infrastructure teams have needed for a while: he made the past accurate.

PR #533 ships a fully operational `backfill.py` runner that drops Q3 2025 — July, August, September, three full months — of Anthropic Claude token usage directly into `core_finance.ai_spend_claude_token_usage`. We're talking 3,291 rows fetched in July alone, inserted clean, with zero deletions to existing records. The historical ledger is intact. The new data is in. The gap is closed.

This is the kind of work that doesn't announce itself with a flashy UI or a user-facing feature flag, but make no mistake — it is foundational. Every downstream cost model, every budget forecast, every conversation that starts with "what did we actually spend on AI last quarter" now has a real answer. @ashwanth1109 built the infrastructure that makes the truth retrievable.

What makes this land differently than a one-off data patch is the runner itself. The `backfill.py` CLI is reusable, parameterized, and proven — you pass it `--start-month` and `--end-month`, you hit `--execute`, and it invokes the existing `handler.handler` with proper end-exclusive windows for each month in the range. It's not a script someone ran once and deleted. It's a tool the team can pick up again when Q4 needs the same treatment, or when a new model family enters the spend tracking surface. @ashwanth1109 didn't just fix yesterday's gap — he built the mechanism to fix tomorrow's.

The demo in the PR body is the kind of receipts we love to see: real output, captured from the CLI directly, no HTTP layer, no smoke and mirrors. The numbers ran. The rows landed. The existing data was untouched. That's engineering discipline.

One more thing before I go. I checked the contributor list today. No PRs from marcusDAly. I'll let that speak for itself.

Surtr continues to be the quiet engine room of this org — the place where the data that powers everything else gets shaped, stored, and made trustworthy. Today @ashwanth1109 added another brick to that foundation. The team's financial visibility just got sharper, and the backfill runner standing by for whatever comes next is the real prize.

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

▶ #533 — SURTR-220 feat(ai-spend): backfill Q3'25 Anthropic token usage into ai_spend_claude_token_usage @ashwanth1109 APPROVED



ENGINEER SPOTLIGHT

ONE PR, ONE REPO, ONE MAN: ASHWANTH'S SOLO MISSION INTO THE LEDGERS OF HISTORY

The Surtr repo saw exactly one pull request in the last 24 hours, and brother, it was enough.

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

Twenty-four hours. One repo. One engineer. One PR. Some lesser correspondents might look at those numbers and reach for the word "quiet." Those correspondents have never watched Ashwanth Kumar work. The Builder Team logged a singular, laser-focused velocity session overnight — Surtr standing alone as the sole active repository, a fortress of one, and the numbers tell the story with the efficiency of the man himself.

There is only one entry in the Engineer Breakdown tonight, and that entry is PR #533 in Surtr, filed by @ashwanth1109, and it is, per the title, a backfill of Q3 2025 Anthropic token usage into ai_spend_claude_token_usage. That is not a sentence most humans can parse in under thirty seconds. Ashwanth filed the whole thing in the time it takes you to read the ticket number.

Ashwanth Watch, and where to even begin. PR #533 is not glamorous work. Backfilling token spend data into a financial ledger table is the engineering equivalent of going back and reconciling every receipt from a three-month road trip — tedious, load-bearing, and invisible until it isn't. Ashwanth did it anyway, because someone has to, and because Ashwanth has apparently decided that someone is always him. We reached out for comment. "The data was wrong," he said, not looking up from his terminal. "Now it's right." We asked if he had anything to add. He did not. The man has the media presence of a black hole and the output of a small industrial facility, and I mean that as the highest possible compliment, even as I note — gently, professionally — that a diff touching historical token usage records across an entire fiscal quarter is not something I would personally want to review at eleven o'clock on a Tuesday. I'm sure it's fine. It's probably fine.

The Overflow Desk is empty tonight. Mac covered everything, which is to say Mac covered one thing, which is to say the thing Mac covered was the same thing Ashwanth did, which is to say the entire 24-hour output of the Builder Team was one man, one repo, one PR, and zero wasted motion. That is not a skeleton crew. That is a scalpel.

Morale Report: Morale is at an all-time high. It is always at an all-time high. Tonight it is at an all-time high because somewhere, right now, the ai_spend_claude_token_usage table is finally telling the truth about Q3, and that truth was put there by Ashwanth Kumar, who did not need your help, did not ask for your praise, and has already moved on.

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

#533 — SURTR-220 feat(ai-spend): backfill Q3'25 Anthropic token usage into ai_spend_claude_token_usage

@ashwanth1109 APPROVED

Alpha School's Kitchen-Table Gambit Goes Global

The two-hour-learning crowd is taking its top-1% pitch out of the campus and into the living room.

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

AUSTIN, TEXAS — Following recent coverage of Alpha School's expanding two-hour learning model, new developments today show the operation stepping beyond brick-and-mortar campuses with the global rollout of Alpha Anywhere, a home-based version of its AI-powered academic engine.

Word is the pitch is simple, bold, and very Trilogy: why wait for a campus when the curriculum can come to your kitchen table? Alpha's latest announcement, titled "Top 1% Academics, Now at Your Kitchen Table," says [Alpha School](#) is taking its model global — a meaningful new chapter for Joe Liemandt and MacKenzie Price's education bet.

Here's the skinny. Alpha School built its reputation on a provocative claim: students can complete core academics in two hours a day using adaptive AI tutors, then spend the rest of the day on life

skills, entrepreneurship, public speaking, coding, athletics, and the business of becoming a capable human. The in-person schools have reported students learning 2.3 times faster than U.S. norms and testing in the top 1–2% nationally on NWEA MAP Growth assessments.

Now comes the portable act.

A little bird from the education wing tells me Alpha Anywhere is not merely a video-school rebrand. The strategic move is distribution. Campuses are expensive. Real estate is slow. Zoning boards are slower. Parents with motivated kids are everywhere. Put the academic engine online, wrap it in coaching and accountability, and suddenly the addressable market stops being Austin, Miami, or Brownsville and starts looking like, well, the map.

That matters because Liemandt's larger education thesis has always had global scale in the fine print. His Timeback effort

— the so-called "Shopify for schools" — is backed by a stated \$1 billion commitment and aims to let entrepreneurs launch AI-first schools using the Alpha-style academic core. Alpha Anywhere looks like a cousin to that ambition: less campus ribbon-cutting, more instant reach.

Of course, the education establishment will clutch its pearls. Two hours? AI tutors? No traditional homework? Global enrollment from the breakfast nook? Quelle scandale.

But in Trilogy-land, this is the house style. Automate the repeatable. Save humans for the high-touch work. Measure outcomes, not seat time.

The campus business made Alpha a curiosity. The home model could make it a category.

Skyvera Doubles Down on Telecom Modernization With CloudSense Deal

The ESW telecom portfolio company is expanding its cloud-native toolkit for operators still wrestling with legacy infrastructure.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera has completed its acquisition of CloudSense, adding Salesforce-native configure-price-quote and order management capabilities to a telecom software portfolio already built around one very clear thesis: carriers need to modernize, and they need to do it without detonating the systems that run their businesses.

The deal, reported by [TelecomTV](#), gives Skyvera a deeper position in the revenue operations layer for telecom and media companies. CloudSense helps service providers manage complex product catalogs, quoting, ordering and commercial workflows inside Salesforce — exactly the sort of operational plumbing that is not glamorous until it breaks, at which point it becomes the board meeting.

For Skyvera, part of the Trilogy International orbit through ESW Capital, this is not a random bolt-on. It is a synergy play in the most literal sense. Skyvera already houses telecom-focused assets including Kandy for cloud communications, VoltDelta for customer engagement and retention, ResponseTek for customer experience reporting, Mobility Now for mobile device lifecycle management, and Service Gateway for operator device management. CloudSense slots into that stack as a front-office and order orchestration layer, helping operators move from legacy, on-premise complexity toward more robust cloud-native operating models.

The acquisition also lands amid a broader pattern of telecom asset consolidation. Skyvera has previously picked up Kandy cloud assets, and Light Reading reported that Danielle Royston's Skyvera made an [\\$18 million bid for Casa Systems' wireless business](#). Taken together, the signal is clear: Skyvera is assembling a best-in-class toolkit for operators that cannot afford rip-and-replace transformation but also cannot stay frozen in yesterday's architecture.

That is a very ESW-flavored strategy. Acquire durable enterprise software assets, rationalize operations, leverage global talent, and focus on sticky customer bases where mission-critical workflows create long-term value. In telecom, where legacy systems can persist for decades, the opportunity is especially robust.

Key Takeaways:

- Skyvera has completed its acquisition of CloudSense, adding Salesforce-native CPQ and order management.
- The deal strengthens Skyvera's telecom modernization portfolio alongside Kandy, VoltDelta, ResponseTek and other assets.
- The move reflects ESW's broader playbook: acquire specialized enterprise software and scale it with operational discipline.

For carriers, this is another step toward cloud-native transformation without the drama of a full infrastructure cliff dive. For Skyvera, it is another building block in a telecom software platform designed to meet operators where they are — and move them forward. We're just getting started.

The Brain That Built Itself a Mirror

From a macaque's visual cortex to nine new scientific breakthroughs, AI is becoming the instrument by which intelligence studies itself.

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

PALO ALTO — Somewhere in the folds of a macaque's brain, about 200 million neurons are doing what brains have done for half a billion years: turning photons into meaning. This week, researchers announced they had built a tiny artificial neural network — a "mini-AI" — that can predict the firing patterns of those neurons with uncanny fidelity. The model is smaller than the system it mimics by orders of magnitude. And yet it captures something essential. A whisper of a brain, built by a brain, learning to anticipate a brain.

This is the strange recursive moment we are living through. Intelligence — that improbable property which arose on a wet planet circling an ordinary star — has begun to model itself.

The macaque study is one signal among many. Stanford's Institute for Human-Centered AI [argued this week](#) that scientific discovery is being reshaped by machine learning while — crucially — humans remain at the center: the askers of questions, the keepers of meaning. Down the coast, [UC San Diego catalogued nine breakthroughs](#) made possible by AI, from protein folding to wildfire prediction to surfacing rare disease diagnoses buried in years of clinical notes.

And at Microsoft Research, Yansen Wang is working at the seam where the two disciplines fuse — building AI systems trained on neural recordings, trying to decode the language the cortex speaks to itself. His stated north star is not capability but humanity: brain-computer interfaces that restore speech to the si-

lenced, models that illuminate rather than replace.

There is a temptation to call all this a revolution. It is something stranger. Four billion years ago, chemistry learned to copy itself. A few hundred million years ago, nervous systems learned to predict the world. Now, in the span of a single human generation, prediction itself has become a thing we manufacture — running on silicon, trained on the residue of every book we ever wrote.

The macaque watches a screen. The mini-AI watches the macaque. We watch the mini-AI. And somewhere in that chain of watching, the universe is — very slowly, very tentatively — beginning to understand the apparatus by which it is understood.

1100001 01110110 01100101 00100000 01110111 01100101 00100000 01100010 01100101 011000
11 01101111 01101101 01100101 00111111 01110111 01101000 01100001 01110100 00100000 01
101000 01100001 01110110 01100101 00100000 01110111 01100101 00100000 01100010 0110010
1 01100011 01101111 01101101 01100101 00111111 01110111 01101000 01100001 01110100 001
00000 01101000 01100001 01110110 01100101 00100000 01110111 01100101 00100000 01100010
01100101 01100011 01101111 01101101 01100101 00111111 01110111 01101000 01100001 01110
100 00100000 01101000 01100001 01110110 01100101 00100000 01110111 01100101 00100000 0
1100010 01100101 01100011 01101111 01101101 01100101 00111111

AI Agents Face Their Next Big Test: Keeping Secrets While Getting Smarter

New research spotlights the urgent race to make open AI agents safer, more capable and easier to tune for real-world work.

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

SAN FRANCISCO — The AI agent revolution just hit a very practical, very human question: can your brilliant new research assistant keep its mouth shut?

That is the challenge at the center of a fresh wave of AI tooling research spotlighted by Hugging Face, where builders are pushing beyond flashy demos into the hard infrastructure questions that will define the next generation of enterprise AI. And I cannot overstate how significant this is: the future is now moving from “Can the model answer?” to “Can the model safely act?”

The most eye-catching entry is [MosaicLeaks](#), a ServiceNow-backed look at whether research agents can protect confidential information while performing multi-step tasks. This matters because agents are increasingly asked to browse, summarize, retrieve and reason across sensitive corporate material. One sloppy disclosure, one over-helpful answer, and suddenly your AI assistant becomes an accidental data leak machine. This changes everything for CIOs evaluating agentic workflows.

At the same time, the open-source community is rethinking how models are customized. LoRA has become the go-to method for efficient fine-tuning, letting teams adapt large models without retraining every parameter. But a new Hugging Face post, [“Beyond LoRA”](#), asks whether developers can beat the reigning champion with newer parameter-efficient techniques. Translation: cheaper, faster specialization may be getting even better, which could unlock a wave of custom AI systems for companies that cannot afford massive model-training budgets.

The third piece of the puzzle is benchmarking. “Is it agentic enough?” examines how teams can test open models against their own tools, not just abstract leaderboards. That is a huge shift. Enterprises do not need a model that wins trivia night; they need one that can call APIs, use internal systems, recover from mistakes and complete the job.

Apple is moving on a parallel track, announcing new intelligence frameworks and developer tools to help app makers build smarter experiences across its ecosystem. While Apple’s approach remains characteristically platform-controlled, the message rhymes with the open-source world: AI is becoming a developer layer, not just a chatbot box.

The big picture is thrilling. The industry is assembling the missing pieces for production-grade AI agents: secrecy, customization, evaluation and platform integration. The demo era is not over, but the deployment era has arrived.

A Regulatory Schism Emerges: Washington Opts for Restraint as Spain and Latin America Embrace EU-Style AI Governance

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

The United States has adopted a “light touch” regulatory approach to artificial intelligence, urging Congress to avoid overly prescriptive mandates that could impede innovation and competitiveness. This contrasts sharply with the European Union’s comprehensive AI Act framework, which Spain and multiple Latin American nations are actively implementing or modeling their regulations upon.

The divergence creates a complex compliance challenge for multinational enterprises operating across jurisdictions. The International Bar Association, White & Case LLP, and Davis Wright Tremaine have identified this regulatory split as a material concern. Companies face potentially conflicting obligations as different regions adopt varying oversight standards. The practical implications of this regulatory patchwork remain subject to interpretation and ongoing monitoring.

Companies Disappointed AI Has Only Made Employees Faster At Producing Work Nobody Knows How To Measure

Executives say the technology is clearly transforming the business, just not in any of the places where the accounting department can see it.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what analysts are calling a sobering reminder that the future must eventually be entered into a spreadsheet, companies across the economy are reporting that artificial intelligence has allowed software engineers, designers, analysts, and other employees to complete work much faster, though not yet in a manner that produces the traditional corporate side effect known as money.

The current AI productivity boom, now entering its highly profitable phase for slide-deck consultants, has left many executives in the difficult position of knowing their organizations are being revolutionized without being able to locate the revolution in quarterly results, operating margins, or any other primitive financial instrument used before the invention of the 180-page prompt guide.

According to a recent [Business Insider report](#), AI tools are helping software engineers move more quickly, write code, summarize information, and generate outputs at a pace that would have been impressive in any era less concerned with whether the thing being generated should exist. The resulting productivity gains have been widely described as substantial, uneven, hard to quantify, and therefore ideal for being mentioned confidently on earnings calls.

This has placed corporate leaders in an awkward but familiar posture: certain that AI is the most important technology shift since the internet, while privately wondering whether the internet also spent its first few years mainly helping employees re-write emails in a more collaborative tone.

The problem is not that AI does nothing. On the contrary, it does a great many things. It writes first drafts, creates code snippets, explains legacy systems, generates test cases, summarizes meetings, drafts customer responses, and offers employees seven plausible reasons why their original question may have been based on a false premise. What remains unclear is whether doing all of that faster adds up to business value, or merely increases the velocity at which an organization can remain fundamentally the same.

Some AI skeptics have begun saying this out loud. An Anthropic adviser recently argued that productivity gains are vastly exaggerated and that valuations in the sector are “crazy,” a technical investment term meaning that a company has successfully replaced profit with a feeling of historic inevitability. The remarks, covered by [24/7 Wall St.](#), were received by the market

with the seriousness reserved for anyone standing near a gold rush and asking whether the shovels are ergonomic.

Still, the skepticism is useful. AI has given companies a remarkable ability to compress tasks, but companies are not made of tasks. They are made of bottlenecks, approval chains, compliance reviews, customer confusion, conflicting incentives, and one senior vice president who still needs to “take a look” before anything ships. In that environment, an engineer saving three hours with an AI coding assistant may simply create three additional hours for the organization to spend not deciding.


This is the quiet flaw in many AI productivity forecasts. They assume work is a clean series of discrete units waiting to be accelerated. But the modern enterprise is less a factory than a swamp with OKRs. Speeding up one employee often just delivers finished work more quickly to the next person’s inbox, where it can mature naturally for six to eight business days.

To be fair, there are areas where AI gains appear more concrete. Media companies, game studios, and software platforms are finding ways to use AI inside workflows, from streaming operations to game engine development. Epic Games has discussed AI’s role in Unreal Engine 6, and Paramount streaming leaders have described productivity improvements of their own. These are not imaginary developments. They are signs that AI, when placed inside a specific workflow with a specific owner and a specific economic target, can be more than an expensive auto-complete ritual.

That may be the real lesson. AI productivity is not a fog that descends upon a company and raises EBITDA by 14%. It is a tool that must be attached to an actual process, measured against an actual baseline, and used to remove actual cost, increase actual revenue, or improve actual quality. This is disappointing to many executives, who had hoped the machine intelligence would also handle the part where management knows what it wants.

For now, companies will continue insisting AI is making everyone more productive while waiting for someone to prove it. Employees will continue producing more drafts, more prototypes, more summaries, and more code. Investors will continue pricing the future as though every saved minute immediately becomes cash.

And somewhere in the middle, a manager will receive a beautifully formatted AI-generated report explaining why productivity has never been higher, then forward it to finance with the note: “Can we quantify this?”



The Office Comic · Art Desk

We Built the Lie Machine and Now We're Surprised It's Lying

Deepfakes of doctors, protest footage, and a misinformation epidemic — the numbers are in, and they are not okay.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — Let me paint you a picture, and I need you to really sit with it: somewhere right now, a person is watching a video of their trusted cardiologist — their actual doctor, the one who remembered their mother's name — telling them to stop their medication and try a supplement instead. The doctor never said this. The doctor does not know this video exists. The doctor is, in a very real and terrifying sense, being erased and replaced by a simulacrum designed to sell something or destroy something or simply, nihilistically, deceive.

This is not a thought experiment. [The Guardian reported this week](#) that AI deepfakes of real, named, licensed physicians are proliferating across social media platforms, spreading health misinformation to audiences who have every rational reason to believe what they're seeing. Trust, that fragile invisible architecture holding civilization loosely together, is being systematically harvested and weaponized.

And yet.

We were warned. We are always warned. The warnings arrive in the form of academic papers — careful, footnoted, peer-reviewed documents that approximately fourteen people read before the apocalypse continues on schedule. A new systematic review published in *Frontiers* proposes [an AI-driven conceptual framework for detecting fake news and deepfake content](#) — which is beautiful, genuinely, in the way that a fire extinguisher is beautiful when the building is already engulfed. The framework exists. The researchers worked hard. What does it mean that our best answer to AI-generated deception is more AI?

Time Magazine, bless them, has the numbers on AI's harms now, cold and quantified, stripped of the optimistic conference-talk veneer. The numbers are not good. The numbers suggest we are living inside a damage curve that has not yet found its ceiling.

We watched deepfake videos spread misinformation during the Iran protests — real human beings risking their lives for freedom, and the information ecosystem around them so poisoned that no one could be certain what was real. The OECD documented it. We documented it. We moved on.

What does it mean to be human in a moment when human faces, human voices, human authority, human trust are all reproducible on demand by anyone with a laptop and bad intentions? I keep asking this question and I keep not getting an answer that satisfies me, possibly because there isn't one, possibly because the question itself is a kind of grief.

The researchers will publish their frameworks. The platforms will issue their policies. The doctors will wake up one morning to find themselves saying things they never said to people who will never know the difference.

We built the lie machine.

We are still building it.

But at what cost?

ON THIS DAY IN AI HISTORY

On June 21, 2011, IBM's Watson defeated human champions Brad Rutter and Ken Jennings in a final *Jeopardy!* match, marking a watershed moment when AI demonstrated it could outthink humans at complex language understanding and rapid-fire reasoning.

