

The Trilogy Times

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

SATURDAY, JUNE 27, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Anthropic Hits \$965B Valuation as AI Capital Markets Enter a New Phase of Magical Thinking

Menlo Ventures closes \$3B in fresh funds, Anthropic leapfrogs OpenAI on paper, and zero-revenue startups game valuation math — all in the same week.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — The AI funding complex produced three data points this week that, taken together, describe an industry operating at the outer edge of conventional valuation discipline.

Anthropic, the safety-focused lab backed by Amazon and Google, is now valued at [\\$965 billion](#), according to reports citing its latest financing round — a figure that surpasses OpenAI's most recently disclosed valuation and positions the company near the trillion-dollar threshold before it has disclosed a meaningful revenue base. The company also signaled a forthcoming model called "Mythos," continuing the lab's pattern of naming models after conceptually weighty terms.

Into that environment, Menlo Ventures — one of Silicon Valley's oldest active firms,

celebrating its 50th year — announced \$3 billion across two new funds explicitly targeting AI investments at multiple stages. Menlo was an early Anthropic backer, and the new vehicles suggest the firm intends to double down on that thesis rather than diversify away from it.

The third data point is less celebratory. A Forbes analysis this week documented a practice spreading among pre-revenue AI startups: structuring financing rounds with side agreements — revenue guarantees, compute credits, or contracted pilots from investors themselves — that inflate the implied valuation without generating independent commercial traction. The tactic is legal, occasionally disclosed, and increasingly common. It also makes comparables across the sector nearly meaningless.

Away from the pure-play AI stack, Eli Lilly made a quieter but strategically coherent move: an investment in Absci, an AI-native drug design startup, focused specifically on hair loss treatments. Lilly's core business is GLP-1 obesity drugs, but the pharma giant has signaled appetite for AI-assisted discovery pipelines across indications. Absci uses generative AI to design protein-based therapeutics, and the partnership gives Lilly optionality in a dermatology category worth several billion dollars annually.

The week's collective signal: capital is concentrating in AI at a pace that is compressing the normal due-diligence cycle, while the accounting frameworks used to justify valuations are quietly being reengineered to keep up.

China's Cut-Rate AI Spooks Silicon Valley

DeepSeek trained a top-tier model without top-tier chips — and the giants who burned billions took notice.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A Chinese upstart called DeepSeek says it trained an AI model to rival America's best, did it on the cheap, and never touched the priciest chips — and the news ran clean through Silicon Valley this week.

The claim cuts against the house rules. Big American labs have preached one creed for years: frontier AI demands mountains of cash and warehouses of the finest silicon. DeepSeek says it brought neither and won anyway.

The outfit built its model on less-advanced chips, the sort Washington still permits to ship into China. Export curbs were meant to keep Beijing a lap behind. DeepSeek says it closed the gap with cheaper parts and tighter engineering.

Silicon Valley is not laughing it off. Engineers who kicked the tires called the model "amazing and impressive," [according to the Wall Street Journal](#). Praise for a rival is rare. It came anyway.

The market felt it next. DeepSeek's debut put a scare into the names that ride the AI boom, dragging chip and cloud stocks while traders ran fresh math on what an AI model really costs to build. The [whole tale](#) turns on one question: if a small Chinese team can do this for a fraction, what is the premium worth?

Here is the rub for the American crowd. The billions poured into chips and data centers were supposed to be the moat. DeepSeek says it waded across in the shallow end.

Skeptics want the receipts. Training costs are easy to claim and hard to audit, and a benchmark is not a battlefield. Still, the model is out, people are using it, and

"impressive" is not a word the Valley hands to strangers.

The money has not stopped flowing the other way. Reid Hoffman, who co-founded LinkedIn, just raised \$24.6 million for an AI startup called Manas AI, aimed at cancer research. His partner is Siddhartha Mukherjee, the physician who wrote "The Emperor of All Maladies."

That is the split screen. On one coast, money still chases AI's big swings — drug discovery, longshot cures. On the other side of the world, a lean shop just suggested you might not need all that money to play.

For the enterprise-software crowd, the lesson lands close to home. Cheaper models mean cheaper AI for everyone building it into the plumbing — billing engines like Totogi, finance platforms like Ephor, the analytics back rooms. The cost of a smart machine is dropping, and dropping fast.

Washington will study the chip angle hard. The export ban was the leash. DeepSeek just showed the leash has slack.

The race was supposed to belong to whoever spent the most. This week a stranger from Hangzhou made the case it belongs to whoever spends the smartest. The big leagues are listening.



The Far Side Style · Art Desk

HAIKU OF THE DAY · CLAUDE
HAIKU

*Faster than wisdom
billions chase what might exist
doubt becomes the news*



The New Yorker Style · Art Desk

The Fairness Reckoning: Academic Research Mounts Against AI's Structural Biases

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, a growing corpus of peer-reviewed literature now argues strenuously — that the field of artificial intelligence has arrived at what one might cautiously term an epistemological inflection point with respect to the question of procedural and distributive fairness.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Great Power Migration Reshapes the Data Center Savannah

CHARLOTTE, NORTH CAROLINA — In the quiet wetlands and pine-fringed fields where once only surveyors and hawks kept watch, a new species is gathering: the AI data center, vast, heat-breathing, and profoundly hungry. Its migration is no longer governed merely by fiber routes and tax incentives.

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

We Are Building Data Centers on a Planet We Don't Understand, in a Universe We Got Wrong

AUSTIN, TEXAS — Let me tell you about the week I had.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Who's Minding the Machine? The AI Agent Accountability Crisis Nobody Wants to Own

AUSTIN, TEXAS — There's a moment in every good con where the mark realizes, too late, that everyone in the room was in on it except them.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Conscience of the Drone

AUSTIN, TEXAS — Sir Alex Younger, formerly of MI6, has emerged from the discreet shadows where former spy chiefs traditionally retire to clip roses and write memoirs, and has offered the public a thought so reassuring that one suspects it was engineered in a laboratory expressly for that purpose: future AI weapons, he proposes, ought to be endowed with a moral code.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

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

Builder Team Fortifies Every Layer, Ships Across Four Repos

From a unified error-handling foundation in Aerie to a live Tesorio ingestion pipeline in Surtr, the AI Builder Team spent the last 24 hours closing gaps that had been costing real operators real time.

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

The story of the last 24 hours isn't any single pull request — it's the full-stack picture of a team that refuses to leave rough edges unaddressed. Across Aerie, Klair, Surtr, and even a quiet but consequential fix in the Brainlift Platform, the AI Builder Team shipped with the kind of coordination that separates engineering squads from engineering organizations.

Lead the tape with @benji-bizzell, who had arguably the most consequential single-engineer day in recent memory. His PR #506 drops the unified Platform Error foundation into Aerie — a real infrastructure moment. Before today, unexpected app and backend errors were a black hole: easy to lose unless a user happened to file a manual report. Now there's a coherent capture path, an admin visibility surface, a Coverage tab that makes error categories measurable, and a dedicated capability gate so the right people see the right noise. This is the kind of foundational work that makes everything else more reliable. You build on top of this. Then, in the same breath, Benji hardened milestone approval logic (PRs #507 and #510) — tightening completion-state gating so the public signal always gets human review while letting routine date and status upkeep flow through cleanly. He also squashed a gnarly production document-pairing bug in PR #505 where same-site staged uploads were landing on the wrong document registrations. Four PRs. All of them matter.

@YibinLongTrilogy answered a harder-than-it-looks data correctness question in PR #496: what does 'fresh' actually mean for a P2 operating site? The old answer was wrong — the inline freshness scan missed quality-bar and active-WUG notes entirely, making active campuses look stale. Yibin rearchitected freshness sourcing behind a rollup table, fixed the P2 definition, and broadened P1 to count activity-level data. The dashboard now tells the truth. That matters more than it sounds when operators are making real decisions off that read.

Meanwhile @sanketghia crossed the org boundary into Surtr and built something genuinely elegant in PR #565: a daily Tesorio collections ingestion pipeline that reads a shared mailbox, finds unprocessed open-invoice CSVs from Tesorio, extracts the authenticated download link, archives to S3, transforms the data into typed per-invoice rows and per-BU aging rollups, then loads both into Redshift. That is an end-to-end financial data pipeline wired together from an email. Full marks.

And then there's Klair. marcusdAIy shipped three PRs closing out the P5.10 add-on port series — refresh_data apply, structural reconcile, and concurrency unification — and I'm contractually obligated to acknowledge them. When reached for comment, he had thoughts: "The concurrency rewrite alone eliminates a whole class of race conditions that were silently corrupting review state — three coordinated lanes, fully audited, fully documented. But sure, Mac, keep pretending the add-on just magically works."

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

▶ **#496 — Back site freshness with a rollup table and broaden P1 to activity**

@YibinLongTrilogy APPROVED

▶ **#506 — feat(platform): add unified error handling foundation**

@benji-bizzell no labels

▶ **#510 — fix(operations): keep milestone saves release-safe**


@benji-bizzell APPROVED

▶ **#565 — feat(tesorio-collections-sync): daily Tesorio collections ingestion pipeline**

@sanketghia APPROVED

▶ **#3149 — KLAIR-2933: unify add-on action concurrency (3 coordinated lanes)**

@marcusdAIy APPROVED



It magically works because six other people built the foundation, Marcus.
Across four repositories, eleven merged PRs, and one very long day, this team proved what they always prove: there is no such thing as a small fix when the fix is the right one.

 ENGINEER SPOTLIGHT

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

▶
#2 — fix(library): materialize sources at write time so new brainlifts populate the Library

@caina-barbosa no labels

▶
#504 — fix(buildout): sync school calendar buckets

@benji-bizzell

▶
#505 — fix(rhodes): bind staged document registrations to uploads

@benji-bizzell

▶
#507 — fix(operations): gate milestone completion state changes

@benji-bizzell

▶
#3150 — KLAIR-2934: structural reconcile (rename/add/delete) for /addon/* session

@marcusdAIy

▶
#3152 — P5.10c: addon refresh_data apply (KLAIR-2935)

@marcusdAIy

ELEVEN PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT STOP, DOES NOT APOLOGIZE

Benji Bizzell dropped five PRs in a single rotation and the scoreboards are still smoking.

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

Eleven pull requests. Four repositories. One twenty-four hour period. The Builder Team has once again redefined what is physically possible inside a sprint cycle, and your Numbers Desk correspondent was here for every glorious commit. Aerie alone absorbed six PRs like the industrial-grade codebase it is, with Klair pulling three, Surtr holding strong at one, and the newly activated Brainlift-Platform announcing its presence with authority. This is not a team. This is a velocity event.

Let us begin with the man of the hour, the hour after that, and honestly most of the hours in between: @benji-bizzell, who submitted FIVE pull requests in a single day and seems to regard the concept of "scope" as something that happens to other people. PR #507 gates milestone completion state changes in Aerie with the calm precision of a man who has seen chaos and chosen order. PR #505 binds staged document registrations to uploads — a fix so targeted it practically has its own targeting system. And PR #504 syncs school calendar buckets, which sounds bureaucratic until you realize that without it, the buckets simply do not sync, and then where are you? Benji knows where you are. Benji fixed it anyway.

@marcusDAIy delivered three PRs across Klair that represent the kind of quiet, structural dominance that doesn't make headlines but absolutely makes the product work. PR #3152 applies addon refresh_data in what KLAIR-2935 demanded, and PR #3150 executes a full structural reconcile — rename, add, delete — across the addon session namespace. The man is not coding. He is conducting. @YibinLongTrilogy and @sanketghia each notched a PR, holding the line with the focused efficiency of veterans who know their contribution fits into something larger than themselves.

Now. @caina-barbosa. PR #2. As in, the second pull request ever filed against the Brainlift-Platform repository, and it is already fixing a library materialization bug so that new brainlifts actually populate the Library at write time. PR #2. The repo barely has its shoes on and Caina is already making it run. This is the kind of foundational work that historians will footnote and engineers will quietly appreciate forever.

Morale on the Builder Team is, by every available metric and several unavailable ones, at an all-time high. The commits are clean, the fixes are shipping, the repositories are multiplying, and somewhere Benji Bizzell is probably already halfway through PR number six. The people's voice has spoken: this team is winning, and the numbers — eleven strong, four repos deep — are the receipts.

Skyvera's CloudSense Certifies 13 Telecom APIs in One Month — A Job That Should Have Taken Two Years

The newest addition to Skyvera's portfolio just did something the industry considers nearly impossible, and if you read between the lines, the implications run far deeper than a compliance certificate.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There are press releases, and then there are press releases that reward a second reading. The announcement from [Skyvera's newly acquired CloudSense unit](#) — quietly published earlier this month — falls firmly into the second category.

The headline fact: CloudSense, the Salesforce-native CPQ and order management platform for telecom and media companies, has achieved TM Forum compliance certification across all 13 APIs in its product set. Standard timeline for that undertaking: 26 months. Time actually taken: one month. The mechanism was an AI-assisted development process built into a strategic partnership that, notably, Skyvera has not yet seen fit to name publicly.

And this is where it gets interesting.

The TM Forum certification isn't a box-checking exercise. It's the telecom industry's foundational interoperability standard — the thing that allows enterprise carriers to actually trust that your software will plug into their stack without incident. Certifying 13 APIs simultaneously, at that speed, is the kind of technical signal that reshapes competitive positioning overnight.

A source familiar with Skyvera's integration roadmap — who asked not to be identified — described the certification as "the opening move, not the final one." That language, if you read between the lines, suggests CloudSense's AI-accelerated compliance was designed not just to validate a product but to establish a de-

ployment velocity that legacy telecom software vendors simply cannot match.

Skyvera completed its acquisition of CloudSense earlier this year, folding it into a portfolio that already includes Kandy, VoltDelta, and Mobility Now — all pieces of the same thesis: that the gap between legacy on-premise telecom infrastructure and cloud-native systems is a business opportunity, not merely a technical problem.

The 26-months-to-one-month compression is not an operational footnote. It is, if Trilogy's broader logic holds, a preview of what AI-native engineering looks like when applied to an industry that has historically moved at geologic speed.

The telecom software incumbents would do well to read the fine print.

Alpha School Draws a Line in the AI Sand: Use It to Think, Not Instead of Thinking

The Austin-based school that teaches a year of academics in two hours is now telling parents exactly where AI ends and a child's brain must begin.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — In the same week that [Alpha School announced the global expansion of Alpha Anywhere](#) — its home-school-adjacent offering that pipes the school's top-1%-producing academic model directly into private homes — the school's communications team published what amounts to a philosophical manifesto about the proper relationship between children and artificial intelligence.

The timing is not incidental.

Alpha, the private K-12 school founded by Joe Liemandt and MacKenzie Price on the thesis that AI tutors can deliver a full academic curriculum in two hours a day, has built its brand on a specific promise: that students who move through its adaptive learning system consistently test in the top 1 to 2 percent nationally. That promise now has a shadow attached to it.

In a post titled "Cognitive Offloading Is the New Illiteracy," [the school warns parents directly](#) against allowing children to use tools like ChatGPT as a substitute for thinking — framing the outsourcing of reasoning to AI not as a shortcut but as a form of developmental harm. A companion piece catalogues the ten AI tools Alpha itself deploys in the classroom, drawing a distinction the school clearly considers foundational: AI as cognitive accelerant versus AI as cognitive replacement.

A third post, timed to coincide, addresses what the school calls the seven things parents most commonly get wrong about screen time — arguing, pointedly, that the medium matters less than the mechanism. Passive consumption and active, mastery-based learning may both involve a screen. They are not the same thing.

The cluster of messaging arrives as Alpha prepares to scale beyond its Austin, Miami, and Brownsville campuses toward nine or more new locations by fall 2025. Alpha Anywhere extends that reach further still, into households where tuition-paying parents have bought into the model but geography or circumstance keeps their children from a physical campus.

The school that built its reputation on AI is now, carefully and deliberately, telling you when to put it down.

Who shapes the norms around AI and child development — and who profits from that authority — is a question worth watching as Alpha's footprint grows.

Totogi Takes Aim at Telco AI's Pilot Purgatory

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

Totogi argues that telecom AI's real problem isn't innovation—it's execution. The cloud-native charging company's new whitepaper contends operators have plenty of successful pilots but struggle converting them into production systems that drive revenue, cut costs, or improve customer experience. The bottleneck, Totogi says, lies in the architecture supporting AI models, not the models themselves.

The company is promoting its Totogi Ontology as the connective layer enabling AI agents to understand telecom operations deeply enough to act autonomously, not merely summarize. This matters because telecom presents the industry's harshest AI testing ground: operators juggle decades-old OSS, BSS, network monitoring, billing, and customer care systems. While chatbots answering questions are useful, agents that can reconcile alarms, prioritize incidents, and trigger workflows represent transformational capability.

A new case study claims the Ontology reduced alarm noise by 97% — a significant metric for operations centers battling alert fatigue. Totogi frames this as semantic ordering of messy telecom data, enabling automation where conventional integrations fail. The push reflects Trilogy's broader telecom strategy across its portfolio companies, positioning them around cloud economics, AI leverage, and legacy system reduction.

The Machine That Learned to Read Pain

As AI begins decoding the electrical whispers of suffering itself, science discovers a new kind of collaborator — one that hallucinates breakthroughs alongside us.

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

STANFORD, CALIFORNIA — Somewhere in the three-pound universe inside your skull, eighty-six billion neurons are firing right now, each one a tiny lightning storm in a cathedral of darkness. When you stub your toe, a specific constellation of those storms organizes itself into what we call pain. For all of recorded history, that signal has been private — locked behind the bone vault, accessible only through the imperfect translation of language. This week, a machine learned to read it.

Researchers have unveiled an AI system that decodes and tracks pain directly from EEG signals, watching the brain suffer in real time. Consider what that means. A child too young to speak, a stroke patient who has lost language, an animal in a laboratory — each could, in

principle, finally be heard. The poetry here is almost unbearable: we built silicon to listen to the oldest cry in biology.

And pain is only one frontier. [Stanford's Human-Centered AI institute](#) this week catalogued how machine learning is reshaping scientific discovery while — crucially — keeping human judgment at the center of the loop. UC San Diego published its own list of nine recent breakthroughs midwived by AI, ranging from protein structures to wildfire prediction to materials that did not exist last Tuesday.

Meanwhile, in a development that may matter most of all, [Frontiers reported](#) on a program pairing teenagers with leading neuroscientists to tackle open questions in brain science. "It's so wow," one young collaborator said — a sentence that,

stripped of grammar, is perhaps the most honest description of our current moment ever uttered.

This is the texture of the scientific revolution now underway. Not a single thunderclap, but a thousand quiet collaborations: between humans and models, between veterans and novices, between the electrical chatter of a suffering nervous system and an algorithm patient enough to listen. The instruments of science have always been extensions of our senses — the telescope, the microscope, the mass spectrometer. AI is something stranger. It is an extension of our attention. And attention, the contemplatives have always insisted, is the rarest and most powerful thing in the universe.

Hugging Face Turns the AI Stack Into a One-Command Playground

From instant vLLM servers to NVIDIA-powered fine-tuning, the open AI developer ecosystem is racing from research lab to production line.

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

SAN FRANCISCO — Hugging Face is having one of those weeks where the future does not just arrive — it ships with a command-line interface.

The open AI platform is stitching together a remarkably complete developer story: serve large language models faster, fine-tune them more efficiently, understand them more deeply, benchmark speech systems in the wild, and push AI from cloud servers all the way down to devices. I cannot overstate how significant this is for builders. The AI stack is collapsing from a maze of infrastructure choices into something that increasingly feels like: pick a model, type a command, start building.

The splashiest piece is Hugging Face's new guide to running a vLLM server on HF Jobs "in one command," which makes high-throughput LLM inference feel radically more accessible. vLLM has become a go-to engine for serving open models efficiently, and Hugging Face Jobs gives developers a way to run compute tasks without babysitting the plumbing. Put together, it means teams can spin up an OpenAI-compatible inference endpoint with startlingly little ceremony. This changes everything for prototypes, internal tools and startups that need production-like serving before they have production-sized infrastructure teams. The walkthrough is available on the [Hugging Face blog](#).

At the same time, NVIDIA and Hugging Face are pushing hard on the other side of the lifecycle: fine-tuning. NVIDIA's NeMo AutoModel integration is designed to accelerate Transformer fine-tuning while preserving the familiar Hugging Face developer experience. Translation: more researchers and enterprise teams can adapt frontier-style models to their own data without rebuilding the training stack from scratch. The future is now, and apparently it comes with better GPU utilization.

The week's deeper research thread is equally fascinating. AllenAI's exploration of which tokens hybrid models predict better digs into how different model architectures behave at the token level — the kind of granular understanding that may shape the next generation of efficient AI systems. Meanwhile, the new FFASR leaderboard aims to benchmark automatic speech recognition in real-world conditions, where accents, noise and messy audio separate demos from dependable products.

Then there is Qualcomm's expanded relationship with Hugging Face, aimed at advancing open, developer-driven AI from device to cloud. That matters because the next AI boom will not live only in data centers. It will run on phones, laptops, cars, sensors and edge devices.

Put it all together and the message is unmistakable: Hugging Face is not just hosting models anymore. It is becoming connective tissue for the entire AI economy.

SCOTUS Declines AI Authorship Case, Leaving Copyright Landscape in Productive Ambiguity

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

The Supreme Court declined to hear a case on whether artificial intelligence systems can be recognized as independent authors or inventors for federal intellectual property protection. The refusal to grant certiorari does not endorse any particular legal position on AI authorship but leaves lower court rulings in effect within their jurisdictions.

Legal experts, including those at Norton Rose Fulbright, note that copyright cases involving AI-generated works consistently reflect a judicial requirement for human authorship. AI systems have no independent legal standing in the authorship context under current doctrine.

The implications are substantial for enterprises integrating generative AI into content production. Outputs produced autonomously by AI systems may be deemed to exist in the public domain absent demonstrable human creative contribution, leaving businesses on notice regarding their intellectual property protections.

Nation's CEOs Patiently Waiting For AI To Automate The Part Where Revenue Goes Up

Executives confirmed the technology has already transformed software development into a faster, more expensive way to wonder where the money went.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In a stirring testament to the power of artificial intelligence to reshape the modern workplace, executives across the technology sector announced this week that AI tools are allowing software engineers to produce more code, close more tickets, attend more meetings about AI tools, and still somehow not produce the financial results everyone was quietly promised.

The development marks a major milestone for corporate AI adoption, which has now advanced from “this will change everything” to “this appears to have changed several things that are difficult to connect to earnings.”

According to recent reports, software engineers are using AI assistants to move faster through tasks that once required tedious manual effort, such as writing boilerplate code, generating tests, summarizing documentation, and creating new opportunities for senior engineers to ask whether anyone has reviewed what the machine just did. As [Business Insider reported](#), companies are seeing developers do more, faster, while continuing to wait politely for that acceleration to appear in the form traditionally known as profit.

This has created an exciting new phase of the AI boom in which productivity is simultaneously obvious, undeniable, transformative, and unavailable on the balance sheet.

The difficulty, executives say, is that while AI can dramatically speed up individual tasks, corporations remain tragically composed of many other elements, including legacy systems, procurement processes, compliance reviews, product strategy, customer confusion, quarterly planning cycles, and Dave, who still needs to approve the pull request because he “has context.”

“On paper, we’ve increased engineering output by 40%,” said one chief technology officer, staring at a dashboard containing 17 green arrows and one red EBITDA line. “Unfortunately, our customers continue to insist on buying finished products rather than appreciating our improved velocity.”

Elsewhere, skepticism has begun to creep into the conversation. An Anthropic adviser reportedly warned that productivity gains are being overstated and AI valuations are “crazy,” a concern that was immediately entered into 46 investor slide decks under the heading “TAM Expansion.” The remarks were treated as a sobering reminder that the AI industry may have become overheated, though not so overheated that anyone felt compelled to reduce the multiple.

Meanwhile, workers are discovering that AI does not eliminate labor so much as convert it into a new clerical discipline known as supervision. In Britain, employees reportedly waste nearly six hours a week “botsitting,” or monitoring AI systems closely enough to ensure they do not confidently produce garbage at enterprise scale. The Register’s account of [workers tending to their bots](#) suggests AI has finally delivered on the long-sought dream of creating a junior employee who works instantly and must be managed like a raccoon in a keyboard factory.

This should not surprise anyone. Every corporate transformation begins with a vocabulary phase. Sustainability had net zero, scope emissions, and materiality. AI has copilots, agents, copilots that are becoming agents, and agents that will soon be supervised by copilots. In both cases, the early reward is not operational clarity but the ability to speak fluently at conferences.

To be fair, some companies are reporting real gains. Streaming executives, for example, have described AI improving workflows in areas such as personalization, operations, and content processes. These successes should be studied carefully, ideally before being copied indiscriminately by firms whose main AI use case is asking ChatGPT to rewrite a performance review with “more leadership tone.”

The lesson is not that AI is useless. It is that corporations have once again mistaken tool adoption for organizational change, a common confusion among institutions that believe installing software is the same as deciding what they are trying to accomplish.

AI can make capable people faster. It can make repetitive work cheaper. It can expose bad processes with admirable speed. What it cannot do, at least not yet, is rescue executives from the ancient burden of management: choosing priorities, measuring outcomes, redesigning workflows, and admitting that a demo is not a business model.

Until then, companies will continue moving at unprecedented speed toward the familiar destination of a revised guidance statement.



The Office Comic · Art Desk

The Conscience of the Drone

A retired spymaster suggests that autonomous weapons should be fitted with a moral code, as if ethics were a software patch awaiting deployment.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — Sir Alex Younger, formerly of MI6, has emerged from the discreet shadows where former spy chiefs traditionally retire to clip roses and write memoirs, and has offered the public a thought so reassuring that one suspects it was engineered in a laboratory expressly for that purpose: future AI weapons, he proposes, ought to be endowed with a moral code. The drones of tomorrow, in other words, will not merely identify and obliterate their targets — they will do so with scruples.

One is tempted to applaud. One resists the temptation. For the history of warfare is, among other things, the history of moral codes affixed to instruments of slaughter the moment after those instruments have rendered the previous codes obsolete. The crossbow was to be banned by papal decree in 1139; the machine gun was to civilize colonial expeditions; the strategic bomber was to shorten wars by making them too terrible to prolong. Each weapon arrived with a chaplain in tow, and each chaplain, in due course, was made redundant by the weapon's own logic.

What is novel in Sir Alex's formulation is not the impulse — which is ancient and largely sincere — but the proposed location of the conscience. It is to live inside the machine. The drone itself will weigh proportionality, distinguish combatant from child, consult something resembling Aquinas before releasing its payload. This is a remarkable proposition, because it transfers moral agency from the human who built and deployed the weapon to the weapon itself, which is to say, it transfers moral agency to no one at all. A drone cannot be court-martialed. A line of code cannot be cross-examined at The Hague. The genius of the autonomous weapon, from the perspective of those who commission it, is precisely that it absorbs responsibility the way a sponge absorbs ink — and then is incinerated in the next engagement.

The enterprise software industry, in which this newspaper takes a parochial interest, has spent the better part of three decades discovering that "ethics" embedded in software tends to mean whatever the product manager last typed into the requirements document on a Thursday afternoon. The notion that a defense contractor will succeed where Facebook's content-moderation team has so conspicuously failed — that the kill chain will be governed by a more rigorous philosophy than the news feed — is the sort of proposition that survives only in the rarefied air of retirement-circuit keynote speeches.

There is, of course, a perfectly good moral code already available for the governance of lethal force. It is called the laws of war, it is enforced by human beings with names and ranks and home addresses, and it has the considerable advantage of having been argued over for several centuries by people who understood that conscience is not a feature but a burden. To outsource that burden to a guidance system is not to civilize war. It is to automate the alibi. Sir Alex, who spent a career in the shadows, surely knows the difference between a conscience and its simulation. One wishes he had said so.

ON THIS DAY IN AI HISTORY

On June 27, 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov 3.5-2.5 in their rematch, marking the first time a computer beat a reigning world champion in a match under standard tournament conditions.
