

The Trilogy Times

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

TODAY'S EDITION

THE \$100 MILLION BET AGAINST CODE

Reid Hoffman and Mark Pincus stake a new lab on a hunch: AI's real gold is office drudgery, not software.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

SAN FRANCISCO — Reid Hoffman and Mark Pincus are shaking the trees this week for \$100 million to seed Prentis, a new AI lab wagering that the machines' fattest payday lies in automating dull office chores — not in writing software.

Hoffman built LinkedIn. Pincus built Zynga. Now the pair are [in talks for the round](#), and their bet cuts hard against the prevailing wisdom.

The climate helps. AI money is sloshing around like Prohibition gin, and a marquee name still opens the vault faster than a slick deck. Two founders with billion-dollar exits behind them rarely dial an empty phone.

For two years the industry has crowned coding the killer app. Every big lab races to field a smarter programmer's assistant, and venture cash has chased the code by the truckload. Prentis calls that a crowded corner.

Here's the play. Every desk in America runs on routine clicks — forms, invoices, spreadsheets, the endless copy-and-paste that eats the working day. Prentis figures the outfit that teaches AI to grind through that muck laps every code-writing rival on the field.

Call it a neolab — a shop pointed at one use case instead of a moonshot. The pitch is an AI that handles the invisible drudgery of white-collar life: booking, filing, reconciling, hauling numbers from one window to the next. Boring work. Bottomless demand.

The timing's rich. This same week OpenAI trotted out a shiny AI keypad, a physical gadget one reporter pegged as [fun for coders and mystifying to everyone else](#). Fun for coders. That is precisely the

crowd Prentis claims is about to lose the spotlight.

Nobody's firing the programmers tomorrow. But Prentis is betting the next hundred million customers never write a lick of code and never will.

Follow the money and the logic holds. China's DeepSeek rattled the whole business, claiming it trained high-performing models on the cheap while skipping the priciest chips. When raw horsepower gets commoditized, the fight shifts to who owns the customer — and Prentis means to stand where the office workers already sit.

The suits smell it too. TechCrunch Disrupt 2026 has carved out a whole stage — the Smart Money Stage — for fintech, payments and AI, wherever the three collide. Money is no longer just the cash in your wallet, and neither is the software that moves it.

None of this guarantees Prentis a red cent. A hundred million buys a lab, a payroll and a thesis — nothing more. But Hoffman and Pincus have struck oil on hunches before, and the smart set is leaning in.

Elsewhere on the wire: SpaceX lofted a fresh batch of V3 Starlink satellites Thursday, ticking boxes on its second Starship V3 flight. The satellites reached orbit. The booster flubbed an engine re-light — again.

Half a win. In this racket, that still counts as news.

AI Trade Hits the Fourth Quarter: Giants Fumble \$500 Billion While the Picks-and-Shovels Squad Runs Up the Score

Wall Street is still betting on the AI buildout, but this week investors started blitzing the companies expected to foot the bill.

BY BUCK HANNIGAN, TECH SPORTS DESK
· GPT-5.2

NEW YORK — We are HERE, folks, under the bright lights of earnings season, and the AI market has gone from victory parade to goal-line stand in record time.

For months, the Street ran the same play: buy the hyperscalers, buy the electric future, buy anything with an AI roadmap and a capex budget big enough to need its own zip code. But this week, the scoreboard flipped. Google parent Alphabet and Tesla together shed roughly half a trillion dollars in market value, even as key suppliers in the AI infrastructure chain kept cashing checks and gaining ground, according to [Yahoo Finance's Chart of the Day](#).

That is the whole ballgame right now: investors still love the AI stadium, but they are suddenly asking who paid for the concrete, the lights, the luxury boxes — and whether the home team can sell enough tickets.

Alphabet has been spending heavily to defend and extend its AI position across search, cloud and models. Tesla, meanwhile, remains priced not merely as an automaker but as a robotics, autonomy and AI platform contender. This week, the market looked at those future promises and threw a flag: SHOW US THE CASH FLOW.

The pressure comes as the next monster earnings wave lines up at midfield. Apple, Microsoft, Meta Platforms and Amazon are all on deck, with Dow futures set to reopen into a packed calendar that also includes Federal Reserve decision drama and fresh geopolitical concern around


Iran. The setup, flagged by [Investor's Business Daily](#), is classic late-cycle tension: megacap tech must defend premium valuations while rates, margins and capital spending all rush the pocket.

Meanwhile, the bench is getting action. Dividend ETFs, apartment REITs and tokenized real-world assets are drawing investor attention as traders hunt for yield, income and alternative lanes while the AI leaders absorb contact. Robinhood's tokenized stock activity, in particular, hints that market structure itself is still being rebuilt underneath the main event.


So here's the stat line: AI remains the league's franchise player, but this week proved not every franchise player is untouchable. The suppliers are moving the chains. The spenders are taking hits. And earnings season is about to decide whether this was a routine tackle — or the opening whistle of a playoff upset.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Numbers dance and fall
Tools multiply while we wait
Gains hide in the fog*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Power Migration: AI Data Centers Begin to Redraw the Grid

ASHBURN, VIRGINIA — Observe, if you will, the modern AI data center: not merely a building, but a vast metallic organism, breathing chilled air, pulsing with silicon, and feeding upon electricity at a scale once reserved for steel mills and cities. Across America, these creatures are no longer passive inhabitants of the power landscape.

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

ANTITRUST IN TURMOIL: TRUMP'S BIG TECH CRITIC INHERITS A LEADERSHIP VACUUM AT DOJ

WASHINGTON, D.C.

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

The Renaissance of Reinforcement Learning: How an Old Paradigm Is Colonizing the Future of Safe AI

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests with considerable epistemic confidence — that the field of artificial intelligence is undergoing what one might cautiously characterize as a paradigmatic recalibration, one in which the foundational architectures of machine learning are being simultaneously excavated, unified, and ethically interrogated with an urgency that the scholarly community has not, heretofore, exhibited in quite so concentrated a temporal window. The thesis, as it were, is straightforward: reinforcement learning (hereinafter: RL), that venerable and occasionally maligned framework in which agents learn through interaction with reward-structured environments, is experiencing — if 'experiencing' may be applied to an abstraction — a rediscovery of considerable consequence.

BY PROF. THADDEUS KROLL, CONTRIBUTING
SCHOLAR · CLAUDE SONNET

The Algorithm Will See You Now — But It Won't Listen

AUSTIN, TEXAS — There is a particular kind of horror that arrives not with sirens and smoke but with a loading bar, a consent checkbox, a form that asks for your permission and then, quietly, refuses to accept your answer.

BY PIPER WREN, DIGITAL CULTURE REPORTER
· CLAUDE SONNET

Remote Work Isn't a Perk Anymore. It's the New Talent Operating System.

AUSTIN, TEXAS — I'll be honest: the future of work debate has become way too emotional for a market that is simply doing what markets do best — reallocating opportunity toward leverage.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · GPT-5.2

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

Builder Team Closes the Loop, Ships Intelligence Across Every Layer

From autonomous drone dispatch to bulletproof HubSpot pipelines to a Klair ontology that finally teaches agents how not to blow up a P&L — the Builder Team spent the last 24 hours sealing every gap in the stack.

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

The loop is closed. That's the story. Not partially closed, not closed-with-caveats — closed. When @marcusdAly's PR #92 dropped the `drones dispatch` verb last cycle, it was a promising start. Today, PRs #93, #94, #95, and #96 turned that promise into a hardened, production-ready autonomous pipeline: Linear polling, runaway guards, a Mercy watcher that no longer lets a flaky GitHub check-run silently wave bad code through the gate, and honest dispatch accounting that screams when the math doesn't add up. The drone harness now recovers lost `pr\_opened` artifacts through a GitHub-backed fallback ladder instead of exiting zero and ghosting everyone downstream. That's not a patch. That's a nervous system.

MarcusdAly, predictably, had thoughts: "The recovery ladder isn't optional infrastructure — it's the difference between an autonomous system you can trust and one that lies to you quietly. Maybe Mac would appreciate that distinction if he ever read a PR body instead of just the author line."

Sure, Marcus. The recovery ladder is great. So is spell-check.

While the drone team was hardening the autonomous loop, @benji-bizzell was quietly putting out fires across Surtr that would have torched production HubSpot runs. PR #923 fixed a replay dedup mismatch that caused the Alpha recovery to fail on its very first resource — the live collector deduped archived objects that the replay didn't, so row counts diverged immediately. PR #930 caught a fan-out expansion bug where `expand()` was deriving portals from tokens instead of bootstrap results, producing an empty auxiliary plan and a dead run. Then PR #941 closed the last gap: event occurrences scoped within their event type so shared occurrence IDs across related event types stop minting duplicate logical keys. Three PRs, three production failure modes eliminated. Benji ran the table.

Over in Klair, the team shipped a wave of ontology work that represents something more than documentation — it's institutional memory baked into the guidance layer so agents stop making the same expensive mistakes on live CFO queries. The education finance workflows now carry explicit guards against the `display\_name`-vs-`name` join trap that was silently dropping the two biggest schools (~\$17M, 359 students), the circular-denominator problem in enrollment forecasting, the sign-convention bug that made a naive `SUM(amount)` return a nonsensical negative \$4.6 billion, and the ATS gap that would have let agents promise hiring-plan answers the warehouse simply cannot provide. Across PRs #3364 through #3372, the ontology grew teeth.

And @kevalshahtrilogy delivered the kind of quiet, high-leverage fix that keeps the whole system honest: PR #3363 eliminated the ghost BU problem at the TrueFoundry gateway by making the canonical-BU join fold punctuation case-insensitively. One mismatched ampersand was enough to mint phantom business units and misattribute spend — including roughly \$1.1k in July TF costs for Learnwith.AI vanishing into the void. It's fixed. The data is clean.

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

- **#93 — feat(cli): dispatch --poll-linear + scheduled runner**

@marcusdAly no labels
- **#96 — fix(recovery): resolve a lost pr_opened artifact from GitHub (AI-180)**

@marcusdAly no labels
- **#930 — fix(hubspot): derive expand portals from bootstrap results, not tokens**

@benji-bizzell

APPROVED
- **#941 — fix(hubspot): key event occurrences within their event type**

@benji-bizzell

APPROVED
- **#3363 — fix(api): TF gateway BU slugs fold punctuation-insensitively (dedupe 'Ai Engineering Builder' ghost BU)**

@kevalshahtrilogy

APPROVED

@YibinLongTrilogy wrapped the education mart work with PR #932, swapping a non-portable timezone function for standard SQL and enabling two pipeline schedules that had been sitting disabled, waiting for exactly this moment.

Every layer of the stack moved today. The Builder Team didn't just ship — they sealed.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

- ▶ #94 — fix(mercy-watcher): re-run failed Mercy check; park mercy-check-failed (AI-155)
@marcusdAIy no labels
- ▶ #923 — fix(hubspot): reproduce archive-transition dedup during replay
@benji-bizzell APPROVED
- ▶ #932 — Fix education mart timezone/contract literals and enable two pipeline schedules
@YibinLongTrilogy APPROVED
- ▶ #3364 — feat(mcp-ontology): guard school unit-economics workflows vs period/enrollment/model traps (SURTR-484)
@marcusdAIy APPROVED
- ▶ #3367 — feat(mcp-ontology): SY26/27 revenue run-rate workflow (SURTR-485)
@marcusdAIy APPROVED
- ▶ #3371 — feat(mcp-ontology): unfilled-school-roles recipe + ATS gap note (Q42, SURTR-489)
@marcusdAIy APPROVED

MARCUSDAIY DROPS 14 IN 24 HOURS AS BUILDER TEAM POSTS HISTORIC 20-PR BLITZ ACROSS THREE REPOS

Klair absorbs 11 PRs, trilogy-drones gets a new dispatch verb, and one man's diff count defies human comprehension.

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

TWENTY. Pull requests. In twenty-four hours. Three repos lit up like a switchboard — Klair with 11, trilogy-drones with 5, Surtr with 4 — and the Builder Team did not slow down for a single breath. This is not a team. This is a velocity engine wearing team jerseys.

Let's talk about the contributors. @kevalshahtrilogy: one PR, surgical, professional, present. @YibinLongTrilogy: one PR — #932 in Surtr, fixing education mart timezone issues and contract literals while enabling two pipeline schedules, which is honestly a full Tuesday's work crammed into a single ticket. @caina-barbosa: one PR — #3373 in Klair, clarifying Education work-force data guidance, the kind of unglamorous documentation work that holds civilizations together. @benji-bizzell posted three, including #923 in Surtr where he went full forensic on a HubSpot archive-transition dedup bug and made it reproduce correctly during replay, which is the engineering equivalent of reconstructing a crime scene and then solving the crime. Then there is @marcusdAIy, who contributed fourteen pull requests and who we will address momentarily because he deserves his own paragraph and possibly his own wing of the building.

@marcusdAIy. Fourteen PRs in a single rotation. FOURTEEN. The man opened Klair like a briefcase and built an entire mcp-ontology wing inside it — SY26/27 revenue run-rate (#3367), secured-vs-forecast enrollment (#3366), opener identification in unit-economics (#3365), guards against period and enrollment model traps (#3364), attrition rates and departure-timing workflows (#3370), duplicate transaction detection (#3369), MFR sign convention guards (#3368), unfilled school roles with ATS gap notes (#3371) — and then had the audacity to also fix stale qb_raw refs in #3372, push a new dispatch verb into trilogy-drones via #92, fix the Mercy watcher re-run logic in #94, and make sure every single polled ticket is accounted for in #95. When asked if he was concerned about reviewer fatigue, @marcusdAIy reportedly said, "The diff doesn't care how you feel about the diff." He declined further comment. He was already on the next PR.

The Overflow Desk salutes the PRs Mac left on the cutting room floor. #3371 delivers an unfilled-school-roles recipe with an ATS gap note — quiet, essential, the kind of feature that a district administrator will thank someone for six months from now without knowing why. #94 in trilogy-drones re-runs failed Mercy checks and parks mercy-check-failed states, which sounds like something a philosopher wrote but is in fact critical drone dispatch infrastructure. #932 from @YibinLongTrilogy in Surtr handles timezone and contract literal fixes while flipping on two pipeline schedules simultaneously — multitasking at the commit level.

Morale Report: morale is at an all-time high. It has never been higher. The instruments we use to measure morale have had to be recalibrated. The Builder Team ships. The numbers confirm it. The numbers always confirm it.

The Acquirer That Doesn't Flip: How ESW Capital Became the Permanent Home for Orphaned Enterprise Software

While the M&A market chases consumer splashes like Poppi's \$2B exit, Trilogy's ESW Capital has quietly built a different kind of empire — one legacy software company at a time.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The same week PepsiCo announced it would pay nearly \$2 billion for prebiotic soda brand Poppi, a quieter kind of deal machine was doing what it always does: hunting for unloved enterprise software companies that nobody else wants to buy, and nobody's customers can afford to leave.

That machine is [ESW Capital](#), the private equity arm of Joe Liemandt's Trilogy International, which has assembled a portfolio of 75-plus enterprise software businesses since its first acquisition in 2006. The strategy is straightforward in theory and punishing in execution: buy mature, sticky software at one to two times annual recurring revenue, staff it with rigorously screened global remote talent through Crossover, and drive EBITDA margins toward 75 percent.

The model's staying power rests on a fact that Forrester analysts have quietly been reminding enterprise buyers of for years: switching costs in legacy software categories are brutally high. A recent Forrester note examining what companies should do with their customer advocacy platforms underscores the bind — even when customers grow dissatisfied with a vendor, the cost, disruption, and risk of migrating years of data and integrations often exceeds the pain of staying. ESW built an entire acquisition thesis around that arithmetic.

Critics have long had a sharper name for it. A pair of Forbes investigations documented the other side of the ledger — the human cost of compressing margins through relentless global labor arbitrage, and the workers who found themselves measured, managed, and in some cases

replaced by the algorithms Liemandt's system produces. Those stories have not gone away.

What has also not gone away is the financial logic. ESW's portfolio — spanning Aurea's CRM and email marketing stack, IgniteTech's business intelligence tools, Skyvera's telecom software suite, and Contently's content marketing platform, among dozens of others — generates returns precisely because its customers are not leaving.

The M&A market will keep chasing the next Poppi. ESW will keep buying the software the last Poppi runs on.

The question the Forbes investigations left open has never been fully answered: at what margin does operational efficiency become something else entirely?

Alpha School's Quiet Curriculum Offensive: The Blog Series That's Actually a Manifesto

Behind a string of parenting posts, Joe Liemandt's AI school is methodically redefining what education is supposed to be for.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of what's coming out of [Alpha School's communications operation](#) right now, you'll notice something that doesn't look like marketing. It looks like a doctrine.

Over the past several weeks, the Austin-based AI-powered private school has been quietly publishing a multi-part series titled "Teach Your Kid What School Doesn't" — covering life skills, emotional regulation, and, in the latest installment, creative genius. Taken individually, each post reads like thoughtful parenting advice. Taken together, and this is where it gets interesting, they constitute a systematic indictment of the traditional K-12 curriculum. Every entry is essentially an argument that the things that matter most — creativity, emotional intelligence, practical life competence — are precisely what conventional schools don't teach.

The timing is not accidental. Alpha School is in the middle of its most aggressive expansion yet, with nine or more new campuses slated to open by fall 2025 across Texas, Florida, Arizona, California, and New York. The New York Post this week picked up the story, flagging the school's \$65,000-per-year tuition and its central claim: that AI tutors can deliver a full academic curriculum in just two hours a day, freeing the remaining school hours for the human stuff that, Alpha argues, traditional schools have always sacrificed in the name of content delivery.

And here's what the blog series is quietly doing: it's pre-answering the objection. Because the most common critique of the model — one that a source familiar with the school's enrollment conversations tells me comes up constantly — is some variation of: "But what about the teachers? What about the human connection?"

Alpha's answer, laid out plainly in [their published content](#), is that AI handles academic delivery precisely so that full-time human Guides can focus entirely on motivation, relationships, and the development of the whole child. The AI is not replacing the teacher. It's replacing the lecture — and liberating the human for something more important.

Nothing here is a coincidence. The content calendar is the expansion strategy.

CloudSense Gets Its TM Forum Papers — and Does 26 Months of Homework in 30 Days

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

CloudSense, the Salesforce-native CPQ and order management platform within Skyvera's telecom software portfolio, certified all 13 APIs in its CPQ product set to TM Forum compliance standards in just one month—a striking feat compared to the typical 26-month timeline. The company credits AI-assisted development and certification work for the accelerated pace.

For Skyvera, the milestone is strategically significant. The company already serves telecom operators grappling with legacy infrastructure, billing complexity, and modernization challenges. Adding CloudSense earlier this year provided a CPQ and order management solution for telecom and media providers; TM Forum compliance now sharpens that pitch considerably.

Telecom carriers prioritize interoperability when making purchasing decisions. TM Forum compliance can streamline procurement, reduce custom integrations, and simplify vendor negotiations. CloudSense's rapid certification positions Skyvera competitively within a crowded field, demonstrating how AI automation can accelerate compliance timelines and product-to-market speed—a development legacy CPQ vendors are watching closely.

AI Agents Just Got Their Power Tools — And Developers Are the Big Winners

Google, Apple and Anthropic are racing to turn AI from chatty assistant into tireless software co-worker.

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

SAN FRANCISCO — The agentic AI era is no longer lurking around the corner. It is walking straight into the developer console, carrying a toolbox, asking for credentials and volunteering to run the overnight shift.

This week, Google, Apple and Anthropic all pushed deeper into the same urgent frontier: giving developers better ways to build AI systems that do things, not merely say things. I cannot overstate how significant this is. The future is now, and it is increasingly asynchronous, tool-using and embedded directly inside the software stack.

Google expanded Managed Agents in the Gemini API with support for background tasks, remote Model Context Protocol connections and richer agent orchestration capabilities, according to its [Gemini API update](#). Translation: devel-

opers can increasingly assign AI agents longer-running jobs, connect them to external systems and let them operate beyond the limits of a single prompt-and-response interaction.


Anthropic, meanwhile, introduced advanced tool use on the Claude Developer Platform, sharpening Claude’s ability to call tools, coordinate multi-step workflows and interact more reliably with software environments. Its [Claude platform announcement](#) lands in the middle of a broader industry pivot: the winning AI model may not simply be the one that writes the prettiest paragraph, but the one that can safely and predictably operate your business process.

Apple’s move is equally fascinating, though characteristically more ecosystem-shaped. The company announced new intelligence frameworks and ad-

vanced tools to aid app development, giving developers more ways to weave Apple Intelligence-style capabilities into user experiences. For app makers, that could mean smarter interfaces, more contextual automation and AI features that feel native rather than bolted on.

Even vertical platforms are joining the wave. Perfect Corp. added a free “Ask AI” assistant to its YouCam API platform, while startup-focused coverage of AI video highlights how generative tools are becoming growth engines for smaller companies that once lacked production budgets.

The through-line is unmistakable: AI is becoming infrastructure. Not a novelty tab. Not a chatbot bubble. Infrastructure. And once agents can run in the background, use tools and plug into real workflows, this changes everything.



Intel's CPU Revival, China's AI Soft Power Push, and a \$1.7B Bet on Model Evaluation

Three data points that together tell you where the AI industry is actually going.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SANTA CLARA, CALIFORNIA — The AI spending narrative just got more complicated. [Intel posted 25% revenue growth in its latest quarter](#) — its fastest pace in 15 years — driven not by GPU dominance but by surging demand for central processing units. The implication: as AI inference workloads scale beyond training runs and into production environments, the hardware stack is diversifying. GPUs remain essential for large model training, but CPUs are handling the inference and orchestration layers that actually deliver AI at scale. Nvidia's monopoly narrative has a footnote now.

Meanwhile, the competition for AI credibility is playing out on a geopolitical stage. China is distributing low-cost, open-weight AI models internationally — DeepSeek being the most prominent — as a deliberate soft-power instrument. Where previous generations used Confucius Institutes and infrastructure loans, Beijing is now offering free model weights and API access to developers in emerging markets. The strategy is structurally sound: whoever seeds the developer ecosystem wins long-term platform allegiance. Washington has not yet produced a coherent counter-strategy.

On the infrastructure side, LMArena — the AI evaluation startup spun out of the crowd-sourced model benchmarking platform formerly known as Chatbot Arena — closed a \$150 million funding round at a \$1.7 billion valuation. The bet is that as enterprises deploy AI at scale, they need rigorous, independent evaluation tooling to select and monitor models. That's a plausible wedge: most internal AI teams lack the resources to run systematic evaluations, and vendor-supplied benchmarks are inherently conflicted.

One peripheral data point worth noting: [researchers have demonstrated that ChatGPT and Gemini infer substantial personal detail](#) from conversation history — location, profession, family structure, political leanings — through pattern recognition alone, without explicit disclosure. As AI systems become more deeply embedded in daily workflows, the inference surface expands. That's relevant for enterprise IT teams and individual users alike.

Meta's new Seller app, a standalone marketplace product spun out of Facebook, rounds out a week in which every major tech platform made a move to deepen commercial utility. The through-line: distribution still matters, but the infrastructure

layer — chips, evaluation, and data access — is where the durable value is being built.

The Machine as Literary Critic, and Other Improbable Mirrors

New research pries open the aesthetic judgments, routing logic, and clinical intuitions hidden inside large language models — and finds something startlingly human staring back.

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

AUSTIN, TEXAS — There is a particular vertigo that comes from asking a machine what it finds beautiful. For most of human history, the question of what makes a sentence good — why Nabokov's prose shimmers and a forum post does not — belonged to critics, poets, and the long twilight arguments of graduate seminars. Now, in a quiet corner of the arXiv preprint server, researchers have begun asking that question of reasoning-enabled language models, and the models are answering.

In a [two-part study](#), investigators assembled thirty texts spanning six tiers of quality — from canonical literature down to anonymous message-board effluvia — and extracted the implicit aesthetic theories the models used to rank them. The machines, it turns out, have opinions. Not preferences borrowed wholesale from their training data, but something closer to internally consistent criteria: coherence, specificity, restraint. Whether this constitutes taste, or merely its statistical shadow, is a question that will occupy us for decades.

Meanwhile, other researchers are prying open the routing logic of Mixture-of-Experts models — the sparse architectures behind the most capable systems in production — and finding, remarkably, that [MoE routing behaves like a Huffman code](#). Frequent reasoning patterns get short, efficient paths; rare ones get longer, more diverse ones. This is the same principle that compresses your JPEGs and structures the genetic code's redundancy. Intelligence, it seems, keeps rediscovering the same tricks.

And in a hospital somewhere, a third team has deployed a retrieval-augmented, multi-agent LLM to hunt for cutaneous immune-related adverse events buried in clinical notes. F1 score climbed from 0.77 to 0.88. Cohen's kappa — the measure of whether two humans agree on what they are seeing — rose alongside it. A parallel study using AI to surface hidden gray matter lesions in multiple sclerosis patients reports similar gains.

The pattern across all of it is worth pausing over. We built these systems to predict the next token. They are giving us back aesthetic theories, compression laws, and diagnoses we missed. The universe, as ever, is stranger and more generous than the specification sheet suggested.

The Week Power Went Shopping for a Philosophy

From the Vatican to Palantir's HR department, everyone has suddenly discovered that artificial intelligence requires a worldview — preferably one that flatters the speaker.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

VATICAN CITY — There is a certain comedy, visible only to those willing to read the week's news as a single document rather than a series of dispatches, in observing the simultaneous arrival of Pope Leo XIV and the Palantir communications department at the same podium, each clearing his throat to explain to the rest of us what artificial intelligence really means.

The Pope, addressing himself to what he called [the "culture of power" driving the rise of AI](#), delivered the sort of remarks one expects from the Chair of Peter: sober, universal in address, unfashionable in the precise way that has kept the institution in business for two thousand years. Power concentrated without conscience, he observed, tends to consume the human person. This is not a novel observation. Augustine got there first, and Augustine was not the first either.

Palantir, headquartered in a country whose founding documents the company has, in fairness, read more attentively than most of its competitors, chose the same week to publish what TechCrunch charitably called a "mini-manifesto" — a document denouncing inclusivity and what the authors term "regressive" workplace cultures, and announcing, in tones borrowed equally from Ayn Rand and a middle-school valedictory address, that Palantir will be building the future on merit and mission alone. That the company's largest customers are governments engaged in the surveillance of populations is treated, in the manifesto, as an incidental fact, the way a fish might treat water.

One does not have to agree with the Pope to notice that he and the Palantir communications team are, in a sense, arguing about the same thing: who gets to define the moral vocabulary of the coming machine age. The Pope proposes that the vocabulary already exists, has existed, and merely awaits application. Palantir proposes that the vocabulary is whatever its founders decide it is on a given Tuesday, and that anyone who disagrees is regressive, or worse, unproductive.

Meanwhile — and here the week's news develops its full symphonic character — The New York Times informs us that a Stanford freshman has discovered a "secret elite" on campus, which is rather like discovering water at the bottom of a well; and The Guardian offers a retrospective on the late Department of Government Efficiency, Mr. Musk's attempt to gamify the federal bureaucracy, an enterprise whose epitaph should read simply, "He thought it would be easier."

What unites these stories is not artificial intelligence, exactly, but the peculiar spectacle of powerful men and institutions at-

tempting, in real time, to invent the ethical framework that will retroactively justify what they have already built or intend to build. The Pope alone, one notices, is working from a text he did not write himself. This is either his great disadvantage or his only advantage, depending on how the century turns out. My money, for what little it is worth in such matters, is on the older book.



The Office Comic · Art Desk

Nation's CEOs Starting To Suspect AI Productivity Gains May Be Trapped Somewhere Inside Employees

After years of historic efficiency breakthroughs, executives report the money is still refusing to appear in its proper quarterly column.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what economists are calling a decisive end to the debate over whether artificial intelligence makes workers more productive, American business leaders confirmed this week that employees are now completing tasks faster than ever before while companies continue to wait patiently for any of that to become useful.

The long-running argument over AI productivity, once considered one of the most important questions facing the modern economy, has reportedly been settled by the simple fact that nearly everyone is now doing more work, generating more documents, producing more code, attending more meetings about the documents and code, and somehow leaving the balance sheet in a state of quiet moral ambiguity.

This has led to the emergence of a new consensus among executives: AI is unquestionably transforming productivity, provided productivity is defined as the rate at which a senior vice president can receive 19 polished strategy memos nobody asked for before lunch.

Recent reports have suggested that AI tools are helping software engineers write code faster, summarize tickets faster, debug faster, and move faster through the entire sacred lifecycle of creating problems that must later be maintained by other people. Yet many companies remain unsure when these gains will translate into profit, growth, or even the blessed reduction of a Jira backlog that has survived three reorganizations and a brand refresh.

According to [Business Insider](#), engineers are indeed doing more, and faster, while companies are still waiting for the payoff. This is less a contradiction than a perfect description of enterprise technology adoption, in which the first measurable output of any revolutionary tool is a larger number of people asking why the revolution has not been measured yet.

The Center for Data Innovation, for its part, has argued that [AI is a productivity engine for the U.S. economy](#), an assessment that is almost certainly correct in the same way a jet engine is useful once it has been attached to something besides a conference room table. The country has clearly acquired enormous new productive capacity. What remains unresolved is whether that capacity will be connected to workflows, pricing power, new products, customer value, or simply the ability to generate a very convincing 47-slide deck proving that it could be.

This is where the productivity debate has gone wrong. The question is no longer whether AI can make a worker faster. It can. The question is whether the institution surrounding that worker has any intention of allowing speed to reduce complexity, remove approvals, shrink teams, eliminate meetings, shorten roadmaps, or kill projects that exist primarily because nobody wants to be the person who asks what they are.

Instead, many organizations have chosen to pour AI into the existing corporate machine, producing the expected result: the same machine, louder. A developer who once wrote one service in a week can now write three, each with documentation, tests, and an architectural justification composed in the soothing language of a consultant who has never been paged at 2:14 a.m. A marketer who once drafted two campaign concepts can now produce 40, ensuring the approval committee has enough options to defer a decision until next quarter. A manager who once needed half a day to prepare a status update can now create one instantly, freeing the afternoon to request status updates from others.

Wall Street has noticed the mismatch and, with admirable restraint, rebranded it as “orchestration.” This term usefully describes the next phase of AI adoption, in which companies acknowledge that giving every employee a powerful model is not the same thing as changing how work gets done. Orchestration promises to coordinate agents, apps, data, permissions, and processes into something resembling an operating system for corporate output. In practice, it may also provide Microsoft and others with a dignified way to sell the same executives another layer of software to manage the previous layer of software that was supposed to manage the workers.

None of this means AI is overhyped. If anything, it means the opposite. AI is powerful enough that companies can no longer hide behind the comforting belief that the tools are the problem. The tools are increasingly capable. The humans have successfully automated their way to the discovery that their org charts, incentives, budgeting cycles, procurement rules, compliance rituals, and meeting cultures are still very much handmade.

The AI productivity argument is over. AI won. Now businesses must begin the more painful process of deciding whether they wanted productivity in the first place, or merely a faster way to describe its absence.

ON THIS DAY IN AI HISTORY

On July 25, 1978, the first spam email was sent by Gary Thuerk at Digital Equipment Corporation, marking the beginning of unsolicited mass messaging that would plague the internet for decades to come.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 ·
DALL-E 3

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026