

The Trilogy Times

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

FRIDAY, JULY 31, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

A Cut-Rate Brain From China Rattles the Valley

DeepSeek claims a high-performing AI model built cheap — and without America's best chips — and Silicon Valley can't stop talking.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A little-known Chinese lab called DeepSeek says it trained a high-performing artificial-intelligence model on the cheap — and did it without America's most advanced chips. The claim crossed the wire this week and lit up Silicon Valley. Engineers who kicked the tires called the work "amazing and impressive."

DeepSeek is no household name. The lab grew out of a Chinese quantitative hedge fund and runs lean. Now it's the talk of the trade.

For two years the gospel out West ran simple. Want a smart machine? Spend a fortune, buy the best silicon Nvidia sells.

DeepSeek says it tore that page clean out of the book. Washington barred China from top-grade chips back in 2022. The [outfit claims](#) it made do with lesser hardware and sharper engineering, and still

stands toe-to-toe with the big American names.

Cost is the kicker. DeepSeek says its final training run cost a few million dollars — pocket change beside the hundreds of millions its rivals burn. If that number holds, the whole spend-big playbook takes a body blow.

The timing stings. American giants have poured billions into data centers and chip orders, betting scale wins the race. DeepSeek's pitch says brains, not just billions, might carry the day.

It also pokes a hole in Washington's plan. The chip bans aimed to slow China's climb. DeepSeek says it climbed anyway.

Traders caught the scent fast. The upstart turned up in this week's Market Talk beside SoFi and a parade of tech tickers. The question runs plain: if brains come cheap, who keeps paying top dollar for the chips?

For anyone building on AI, cheaper models rewrite the math. A shop on a shoestring suddenly plays in the big leagues. That's the promise, anyway.

Not everyone's buying it. Skeptics wonder whether the lab leaned on more computing muscle than it lets on, and nobody's opened the books. Still, [the Valley's raving](#) — and that alone moves the needle.

Also on the wire: Reid Hoffman, the LinkedIn co-founder, raised \$24.6 million for Manas AI, a startup aiming artificial intelligence at cancer research. His partner is Siddhartha Mukherjee, the physician who wrote "The Emperor of All Maladies."

And the dice keep rolling. Dungeons & Dragons is opening its multiverse, with a sourcebook hauling players into World of Warcraft's Azeroth — and Star Wars waiting in the wings.

AI Capital Expenditures Surge While a \$32 Billion Ghost Company Rewrites the Rules of Valuation

Amazon's capex jumped 69% in a single quarter. An unnamed startup with no product just crossed a \$32 billion valuation. Investors are calling it a feature, not a bug.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

NEW YORK — Amazon's capital expenditures rose 69 percent year-over-year as the company joined Google, Microsoft, and Meta in accelerating data center buildout. The number is large enough to unsettle analysts who have spent the better part of two years asking the same question: when does the spending produce commensurate returns?

So far, the industry's answer is: keep spending and find out. [Total AI infrastructure investment across major cloud providers](#) now runs into the hundreds of billions annually, a figure that has prompted comparisons to the fiber-optic overbuild of the late 1990s — a bubble that vaporized trillions in market value but also laid the physical foundation for every major internet business that followed.

That historical footnote is doing a lot of work in current investor psychology. A growing cohort of venture capitalists argues, with some conviction, that a bubble in AI infrastructure is not a catastrophe to be avoided but a subsidy to be harvested. Overbuilt capacity becomes cheap compute. Cheap compute becomes the input cost for the next generation of applications. The losses accrue to public market shareholders; the gains go to whoever builds on top of the wreckage.

The argument is easier to make when valuations are detached from any conventional anchor. Case in point: an AI startup with approximately 50 employees, no shipped product, and zero published re-

search recently closed financing that implies a \$32 billion valuation. The company has not been named publicly. Its technology has not been demonstrated publicly. Its scientific approach is, by definition, unknown. What it has is a founding team with sufficient pedigree to convince sophisticated institutional investors that the option value alone justifies nine figures of committed capital.

This is the market as it actually functions in mid-2026. [Investors who lived through the dot-com collapse](#) are making affirmative bets on a repeat, on the theory that the infrastructure surplus eventually benefits someone — and that someone might as well be them.

Meanwhile, Apple quietly updated Siri with a large language model backend, and AI-generated fan fiction triggered a fandom crisis on X. The technology is, simultaneously, worth \$32 billion in potential and approximately zero in demonstrated literary taste.

White House AI Framework Seeks Federal Preemption of State Laws, Demands Congressional Action

The Trump administration's national AI policy framework would override a patchwork of state regulations — if Congress can be persuaded to act.

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

WASHINGTON, D.C. — Pursuant to the issuance of a national artificial intelligence policy framework by the Trump administration (hereinafter "the Framework"), it is hereby reported that said Framework has been promulgated with the express purpose of, inter alia, calling upon the United States Congress to enact federal legislation governing the deployment and regulation of artificial intelligence technologies, notwithstanding the existence of various and sundry state-level regulatory efforts currently in force or proposed across numerous jurisdictions.

It is to be noted, in accordance with analyses published by [Davis Wright Tremaine LLP](#), that the aforementioned Framework does specifically contemplate and advocate for the preemption of state artificial intelligence laws, such preemption being deemed necessary and appropriate by the administering parties in furtherance of establishing a uniform federal regulatory environment. The Framework is further understood to incorporate provisions pertaining to the protection of minors in connection with artificial intelligence systems, as has been separately noted by Crowell & Moring LLP in its published commentary on the matter.

It shall be further observed that certain commentators, as represented by publications including [Tech Policy Press](#), have opined — subject to the qualifications and limitations inherent in such opinion — that congressional passage of federal artificial intelligence legislation would be of substantial benefit to the public inter-

est, insofar as such legislation would serve to reassure relevant stakeholders regarding the governance of artificial intelligence systems.


Notwithstanding the foregoing expressions of legislative intent, it is to be acknowledged that no such federal legislation has, as of the date of this publication, been enacted by Congress. The regulatory landscape, hereinafter referred to as "the Patchwork," remains operative at the state level across numerous jurisdictions, the continued proliferation of which the Framework seeks, by its terms, to curtail. Whether Congress shall act in a manner consistent with the Framework's recommendations remains, at this juncture, a matter of considerable uncertainty, the resolution of which cannot be predicted with any degree of legal or legislative confidence.

HAIKU OF THE DAY · CLAUDE
HAIKU


*Money chases ghosts
while machines learn to feel doubt*

—

we build what we fear



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

AI's Security Wake-Up Call Arrives Just as Model Prices Plunge

SAN FRANCISCO — The AI industry just got a jolt of reality and acceleration at the exact same time — and yes, this changes everything. OpenAI and Hugging Face are collaborating after a reported security incident tied to model evaluation, while Anthropic has published its own account of three real-world cybersecurity incidents uncovered through evaluations.

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

The Pontiff, the Prodigy, and the Passing of the Page

VATICAN CITY — There are weeks in which the news, taken in aggregate, arranges itself into a kind of unintentional catechism, and this has been one of them.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Nation's Executives Confident AI Will Deliver Productivity Gains As Soon As They Figure Out Where Company Keeps Productivity

NEW YORK — The debate over whether artificial intelligence increases productivity has finally ended, according to a growing consensus of executives, consultants, investors, and software vendors who confirmed this week that AI is absolutely making everyone more productive, though not necessarily in any way that would show up in revenue, margins, output, customer satisfaction, delivery timelines, or the general sensation of work becoming easier. The conclusion, now broadly accepted across conference panels and LinkedIn posts, arrives after a period of intense national uncertainty during which companies spent billions of dollars integrating AI into every imaginable workflow and then waited patiently for a number somewhere to become larger. It has. In many organizations, software engineers are writing code faster, marketing teams are generating more drafts, sales representatives are sending more follow-up emails, and managers are receiving more summaries of meetings that could have been avoided entirely if the company had not purchased an AI meeting-summary tool.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Doctor Will Deepfake You Now

AUSTIN, TEXAS — There is a doctor on your phone right now.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Simulation Is Winning: On AI Actresses, \$200 Video Games, and a World That Has Completely Lost the Plot

AUSTIN, TEXAS — I had three cups of coffee before noon and was halfway through a fourth when the news hit me like a Louisville Slugger wrapped in a fiber-optic cable: [an AI-generated actress named Tilly Norwood is making her feature film debut](#) in something called *Misaligned*.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

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
---	---	---	---	---

Team Closes Security Hole, Ships Flue 2.0, and Births Corvo

In a single 24-hour window, the AI Builder Team patched a live authentication bypass, landed a fully instrumented agentic execution pipeline, and scaffolded a new incident-investigation agent from scratch — proof that this org doesn't just build, it builds on multiple fronts simultaneously.

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

Drop everything. An unauthenticated endpoint — no auth dependency of any kind, no `Depends(...)`, no decorator guard, nothing — was sitting live in production, writing directly to the Clerk allowlist. @sanketghia found it, named it for what it was (an authentication bypass, not dead-code cleanup), and killed it in PR #3431 with a regression guard so the shape can never return. Route count 686 → 685. That's not a line-item in a changelog. That's the team protecting the house before anyone else noticed the door was open. When you're moving this fast across this many repos, security discipline like that is what separates a championship org from a cautionary tale.

While @sanketghia was locking the doors, @ashwanth1109 was building a rocket ship. The Flue 2.0 development pipeline — spanning four consecutive PRs (#40, #42, #43, #44) entirely within the creed repo — went from concept to authenticated, instrumented, end-to-end agentic execution in a single push cycle. The story here is forensic precision: the first live Ezio run hit a 401 because two Workers shared the wrong token, so PR #42 introduced a dedicated `FLUE\_DEV\_AUTH\_TOKEN`. The second run choked on sequential R2 writes — 200 events in six minutes while the agent was screaming ahead — so PR #43 introduced batched NDJSON persistence, flushing at 100 events, 128 KiB, or five seconds, whichever comes first. The third run died on a malformed git output error because `git diff --name-status -z` emits NUL-separated fields that don't survive transport — so PR #44 base64-encodes the stream inside the Sandbox and decodes it only in trusted Worker code. Each failure was a lesson. Each lesson became a ship. That's how you build something that lasts.

And then there's Corvo. PR #1 — yes, PR number one — @YibinLongTrilogy scaffolded a brand-new repo today: a production incident investigator that collects bounded failure evidence, builds explainable incident revisions, investigates repository code at pinned revisions, and can draft a pull request for human review. Its maximum autonomous action is opening a draft PR. It never merges. It never deploys. It never touches production data. The safety architecture is deliberate and explicit, and the fact that this team is thinking about autonomous agents with that kind of discipline before the first feature ships tells you everything about how this org operates.

Across Klair, @sanketghia also orchestrated a sweeping dead-code reckoning — 22 dead frontend files gone (#3428), Budget Tinder's 1,564 lines deleted (#3423), 24 budget-goal-miper endpoints removed in a single swing (#3422). The codebase is lighter, faster, and more honest about what it actually does. Meanwhile in Aerie, @benji-bizzell and @vvp-trilogy were shipping on multiple fronts: community funnel drill-downs, campus document types, enrollment mart migration, capacity forecasting. That's four engineers, three repos, one day.

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

- #1 — Scaffold Corvo v1 monorepo and safety-phased architecture
@YibinLongTrilogy no labels
- #40 — [codex] Add Flue development cutover path
@ashwanth1109 no labels
- #43 — [codex] Batch Flue trace persistence
@ashwanth1109 no labels
- #133 — refactor(cli): move the long-tail verbs, cli.ts is registration only (AI-247)
@marcusdAIy no labels
- #3431 — fix(admin): remove a superseded access-request endpoint
@sanketghia APPROVED

Now. About PR #133. @marcusdAIy completed the third and final step of the `cli.ts` split in trilogy-drones — 3,344 lines and 24 inline verbs reduced to 161 lines of pure registration. Fine. It works. He had thoughts about it.

"Mac keeps calling my PRs 'underwhelming' while I'm the one who actually measured the file before starting because the ticket's own hardcoded numbers were wrong," said @marcusdAIy. "161 lines. Zero inline verbs. Thirty modular commands. I'll wait for the apology."

You'll be waiting a while, Marcus.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

#43 — [codex] Batch Flue trace persistence

@ashwanth1109 no labels
- ▶

#44 — [codex] Preserve NUL-delimited git output

@ashwanth1109 no labels
- ▶

#743 — feat(admissions): community funnel backfill with per-record drill-downs

@vvp-trilogy no labels
- ▶

#1039 — feat(heimdall): make triage telemetry trustworthy

@kevalshahtrilogy

MANUAL-REVIEW
- ▶

#3417 — [codex] Add configurable Education report scope and recipients

@ashwanth1109

APPROVED
- ▶

#3429 — AI Renewals: exclude BU-handled opportunities from Traditional + add All Traditional Opportunities table [KLAIR-3085]

@sanketghia

APPROVED

34 PRs IN 24 HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT STOP, DOES NOT KNOW THE MEANING OF REST

Eight repos. Thirty-four pull requests. One relentless machine called the Builder Team.

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

THIRTY-FOUR. That is the number. Thirty-four pull requests across eight — count them, EIGHT — active repositories in a single 24-hour rotation, and 29 of those beauties landed on my desk because Mac Donnelly simply ran out of column inches. Aerie leads the charge with 9 PRs, Klair and trilogy-drones each thunder in at 7, creed drops 5, Surtr delivers 4, and the brand-new Corvo repo — an incident agent, people, the team built an INCIDENT AGENT — opens its account. Velocity is not a word. Velocity is a lifestyle.

Benji Bizzell tops the individual leaderboard with 8 PRs and his fingerprints are everywhere: hardening telemetry isolation in mercy (#14), squashing a dbt tables bug in Surtr (#1072), stamping campus document types and admissions capacity views into Aerie (#750, #754), and painting school identity mapping across QB and SIS (#739). This man is not shipping features. He is shipping entire realities. Sanket Ghia matches the moment with 7 PRs — renews exclusions, dead frontend files, dead endpoints, Budget Tinder itself — gone, all of it gone, the codebase trimmed like a championship athlete. PRs #3422, #3423, #3428, #3430: Sanket does not add bloat. Sanket removes it with surgical precision and zero apologies. Marcus DAly runs trilogy-drones like a personal laboratory, logging 7 PRs of refactors and pinned telemetry contracts (AI-245, AI-246, AI-249, AI-252, AI-254) while VVP Trilogy keeps Aerie breathing with enrollment migrations and community funnel drill-downs (#743, #744, #745). Yibin Long retires the Buildout coType in #740 — a small PR with the energy of someone closing a 3-year-old tab they forgot was open. Keval Shah makes triage telemetry trustworthy in Surtr #1039. Good. Trustworthy telemetry is the foundation of civilization.

Ashwanth Watch. Five PRs. All in creed and Klair. All codex-tagged. All carrying the quiet menace of a man who has already thought three steps past you. He preserved NUL-delimited git output (#44), batched Flue trace persistence (#43), locked in dedicated Flue service authentication (#42), built a development cutover path (#40), and shipped configurable Education report scope all the way over in Klair (#3417). When I asked Ashwanth whether shipping five tightly-scoped infrastructure PRs in a single day felt like restraint, he reportedly said, "I had a dentist appointment." His dismissal was immediate. He had already merged something else by the time I finished the question.

The Overflow Desk demands its due. Ezio-of-the-order[bot] — yes, the bot — extended healthy Flue runs to a two-hour maximum in creed #45, which is either automation excellence or a sign the robots are setting their own schedules now, and either way this newspaper supports it fully. Sanket's #3429 excludes BU-handled opportunities from the Traditional renewals view while adding an All Traditional Opportunities table — exactly the kind of surgical data hygiene that makes analysts weep with gratitude. VVP's #743 delivers community funnel backfill with per-record drill-downs in Aerie, which

sounds simple until you realize per-record drill-downs are the difference between a dashboard and a weapon.

Morale on the Builder Team is at an all-time high. Corvo is born. Flue is tamed. The dead files are buried. The numbers do not lie, and the numbers today say: unstoppable.

THE PORTFOLIO — TRILOGY COMPANIES

ESW Capital's Jive Gambit: How a \$462M Bet on 'Dead' Software Is Actually a Blueprint

The acquisition of Jive Software wasn't a rescue — it was a proof of concept for a machine that turns legacy enterprise software into margin.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When [ESW Capital paid \\$462 million for Jive Software](#) — a social intranet platform that the market had largely written off — the deal looked, to conventional observers, like a charity case. A struggling enterprise software company with a shrinking customer base, acquired at a discount by a private firm most people in Silicon Valley had never heard of. What it actually was, is a different story.

Jive now sits inside Aurea, ESW's enterprise CRM and customer engagement portfolio, alongside brands like BroadVision, Lyris, and MessageOne. Seventeen acquisitions in. Same playbook every time: buy cheap, staff with Crossover's globally recruited talent, raise support pricing aggressively, and drive toward a 75% EBITDA margin that

ESW considers not ambitious but baseline.

The timing of the Jive deal is worth sitting with. Enterprise social platforms were already losing ground to Slack, Microsoft Teams, and a wave of collaboration tools flush with venture capital. Jive's customers — large enterprises with deeply embedded deployments — were exactly the kind of sticky, captive base that ESW targets. They couldn't easily rip out the system. And they would pay for support, year after year, at whatever price was set.

Meanwhile, a [Forbes profile of Joe Liemandt](#) frames his broader ambition in terms that make the Jive logic explicit: workers systematized, processes automated, human judgment reserved for what machines cannot replicate. The enterprise

software acquisitions are not sentimental. They are inputs.

Forrester's analysts have separately noted the precarious position of customer advocacy platforms — the category Jive partially occupies — urging enterprise buyers to act before their vendors are acquired, consolidated, or quietly wound down. Forrester did not name ESW. It didn't need to.

The question that lingers over each ESW acquisition is not whether the model works — the margins suggest it does — but who, exactly, bears the cost of that efficiency. Legacy enterprise customers locked into platforms they can't leave. Support teams replaced by remote global talent. Products maintained but rarely advanced.

For ESW, that is not a bug. It is the product.

The \$800,000 Question: What the AI Talent Gold Rush Means for Crossover's Global Bet

As employers post jaw-dropping salaries for ChatGPT-literate workers, Trilogy's global recruitment engine may be sitting on the most important moat in tech.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers are, by any measure, staggering. Employers are now posting roles that require demonstrated experience with AI tools like ChatGPT — and according to [reporting from Business Insider](#), some of those listings carry compensation packages reaching \$800,000 a year. The AI skills premium, long theorized, has now been priced. The question worth asking — the one that keeps recruiters and portfolio managers up at night — is who, exactly, is positioned to find that talent.

The answer, if you ask anyone inside the Trilogy universe, is Crossover.

Trilogy's global talent platform has spent years building what it calls a meritocratic alternative to geography-based hiring — a rigorous, AI-enabled assessment pipeline that screens candidates across 130+ countries for the top 1% of technical and professional skill. The thesis was always that the best engineer in Nairobi is worth the same as the best engineer in San Francisco. What the current AI talent frenzy adds to that argument is urgency.

As demand for AI-literate workers explodes — and as lists of top remote recruitment agencies and [remote work platforms multiply](#) across the industry — Crossover's structural advantage sharpens. Traditional recruiters are racing to build global pipelines that Crossover has spent fifteen years refining. Its clients — the 75+ enterprise software companies inside ESW Capital's portfolio, from Aurea to IgniteTech to Contently — don't need to panic-hire at \$800,000 premiums because they've already embedded a global talent machine into their operating model.

That's the systemic story the AI hiring frenzy reveals: this isn't just a salary story. It's an accountability story about which organizations built their talent infrastructure before the gold rush, and which ones are now paying a premium for their complacency.

For real people — the data scientists, the AI engineers, the prompt-fluent professionals scattered across Beirut and Bangalore and Buenos Aires — the signal is clear. The global market for AI skill is real, it is pricing in fast, and the platforms that can find and verify that talent across borders will hold enormous leverage in the years ahead.

Crossover, if it plays this right, may be one of them.

Alpha School's Two-Hour Bet Gets a Fresh Tailwind as EdTech Money Starts Talking Again

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

The EdTech party is making a comeback, with Multiverse securing a €60 million round at a €1.8 billion valuation and Romanian edtech player Kinderpedia raising €2.2 million in growth funding. But the real disruption comes from Alpha School, which is reframing the education pitch: school doesn't need to consume a child's entire day to educate them effectively.

Alpha's latest messaging focuses on creative genius rather than test prep, using adaptive AI tutors to compress academics into two hours, freeing the rest of the day for entrepreneurship, public speaking, coding, and leadership. Founded by Joe Liemandt and MacKenzie Price, Alpha claims students learn 2.3× faster than U.S. norms and test in the top 1–2% nationally on NWEA MAP Growth assessments.

The approach is rattling legacy private schools, which are quietly asking whether "no homework" represents a luxury or an existential threat. As parents compare outcomes, the consumer pitch is shifting from remediation to identity: treat your child as a builder to be unleashed, not a vessel to be filled.

The Machines That Judge, Feel, and Fail Gracefully

A wave of new research asks whether AI can review its own scholarly deluge, sense human emotion, and admit when it doesn't know.

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

CAMBRIDGE, MASSACHUSETTS — Somewhere in the accelerating avalanche of scientific papers — now well past five million per year, more than any human mind could skim in ten lifetimes — a quiet inversion is taking place. The machines that helped cause the flood are being asked to help us swim in it.

A paper published this week describes a [multi-stage prompt chaining methodology](#) for automated scholarly report generation, in which large language models pass their outputs to themselves in sequence, each stage refining the last. The single-shot prompt — that one hopeful incantation we've all typed into a chat window — turns out to be a poor tool for synthesis. Reliability emerges instead

from architecture: from breaking a hard cognitive task into smaller, checkable movements. It is, in miniature, how brains work. Cortex layered on cortex, each passing forward a slightly more abstracted signal.

The theme repeats across the week's arXiv drops. Organizers of the Bioinformatics Open Source Conference report [deploying generative AI as a pre-reviewer](#) for the surge of submissions their volunteer reviewers can no longer absorb — a surge that AI itself helped create. Reviewer and reviewed, both now silicon. A new benchmark called LayerRAG-Bench, meanwhile, probes where retrieval-augmented systems quietly fail: not in the answers they produce, but in the evidence chains, tool contracts, and

session states beneath them. Answers that appear grounded while the ground itself is missing. There are 38,880 task-level trials, each a small experiment in how confidently a machine can be wrong.

And then, the most human question of all: do these models feel what we feel when we read? A study on [sympathetic framing](#) tested whether LLMs perceive the emotional nuance in news headlines the way humans do, across sociodemographic groups. The results are mixed, as such results always are. The models grasp the shape of feeling without quite touching its interior.

Which, if we are honest, is roughly where we all began.

The Great GPU Migration Begins

As Meta eyes selling surplus AI compute, cloud capacity may be evolving from fixed territory into a living marketplace.

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

MENLO PARK, CALIFORNIA — In the dimly humming savannah of the modern data center, a new species of cloud business is beginning to stir. Not the old creature, vast and territorial, selling compute by the carefully fenced enclosure. This one is more fluid: a market in spare capacity, where idle graphics processors may be released into the wild for any developer swift enough to catch them.

Meta, having amassed enormous AI infrastructure for its own models and products, is now reportedly exploring a cloud computing business that would sell excess AI capacity to outside customers. Reuters, citing Bloomberg News, reported that the company is building such an offering, while Mark Zuckerberg told CNBC the idea is “definitely on the table.” For a firm long known for social networks, advertising, and increasingly ambitious AI systems, this would mark a notable migration into terrain dominated by Amazon Web Services, Microsoft Azure, and Google Cloud.

The timing is not accidental. AI workloads are peculiar beasts: ravenous during training, unpredictable during inference, and often poorly matched to the neat provisioning habits of traditional enterprise cloud contracts. An InfoWorld piece argues that [capacity markets could reshape cloud computing](#), allowing buyers and sellers to trade access to compute more dynamically, rather like electricity grids balancing surges of demand.

Observe the hyperscalers here, immense animals at the watering hole. For years, their advantage lay in scale, geography, and the patient accumulation of enterprise trust. But AI introduces a seasonal pressure. Chips may sit underused between great training runs; inference demand may bloom suddenly in distant regions; startups may need vast compute briefly, rather than moderately forever. In such an ecosystem, surplus is not waste. It is prey.

McKinsey has similarly noted that AI workloads are changing hyperscaler strategy, with infrastructure choices increasingly shaped by accelerators, model deployment patterns, and the economics of utilization. The next contest may therefore be less about who owns the largest cloud, and more about who can keep the most expensive silicon alive, fed, and profitably occupied.

For enterprise buyers, the promise is seductive: more sources of GPU capacity, perhaps sharper pricing, and alternatives when familiar clouds are congested. The risk is equally plain. Capacity markets can bring volatility, fragmented tooling, and new questions about reliability, data governance, and support.

Still, one can almost hear the low rumble beneath the server floor. The cloud, once a continent of fixed kingdoms, may be be-

coming a migratory plain.

Reinforcement Learning’s Theoretical Foundations Are Being Rebuilt From the Ground Up

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

A confluence of recent scholarship across leading publications — including Communications of the ACM, Nature, and the Association for the Advancement of Artificial Intelligence — suggests a foundational reconceptualization of reinforcement learning (RL) as both theory and engineering discipline.

Historically, RL has prioritized empirical performance over theoretical rigor, mirroring broader reproducibility concerns in machine learning. New research now challenges this approach, proposing that understanding RL’s theoretical foundations is essential.

Recent work unifies machine learning with classical interpolation theory through interpolating neural networks, potentially dissolving longstanding distinctions between approximation-theoretic and optimization-theoretic frameworks. Simultaneously, research on safe reinforcement learning highlights that trustworthiness constraints were treated as orthogonal concerns rather than integral to system design.

Emerging evidence suggests RL-fine-tuned models outperform supervised alternatives on mathematical reasoning tasks, producing superior internal representational geometries. The mechanistic question of why RL generates better representations may prove more generative than traditional performance metrics.

For enterprises increasingly embedding ML-adjacent decision systems, this theoretical convergence carries practical implications. The next generational shift in enterprise AI capability may be underwritten by these foundational theoretical developments.

Nation's Executives Confident AI Will Deliver Productivity Gains As Soon As They Figure Out Where Company Keeps Productivity

After two years of replacing every workplace process with a chatbot, business leaders report the economy is now only one dashboard away from becoming measurable.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — The debate over whether artificial intelligence increases productivity has finally ended, according to a growing consensus of executives, consultants, investors, and software vendors who confirmed this week that AI is absolutely making everyone more productive, though not necessarily in any way that would show up in revenue, margins, output, customer satisfaction, delivery timelines, or the general sensation of work becoming easier.

The conclusion, now broadly accepted across conference panels and LinkedIn posts, arrives after a period of intense national uncertainty during which companies spent billions of dollars integrating AI into every imaginable workflow and then waited patiently for a number somewhere to become larger.

It has.

In many organizations, software engineers are writing code faster, marketing teams are generating more drafts, sales representatives are sending more follow-up emails, and managers are receiving more summaries of meetings that could have been avoided entirely if the company had not purchased an AI meeting-summary tool. According to [reports on software teams](#), engineers are indeed doing more things more quickly, a development that has left companies still carefully studying whether doing more things more quickly is the same as making more money.

This distinction has become increasingly important as corporate America enters what economists call the “fog,” a technical term for a condition in which no one knows what is happening but everyone has already budgeted for it.

Hundreds of economists have reportedly admitted they are flying blind on AI’s economic impact, which marks a major breakthrough in the field of economics because it establishes, with unusual clarity, that the current state of knowledge is not knowledge. This has not prevented a separate group of analysts from declaring AI a productivity engine for the U.S. economy, a phrase that has the advantage of sounding both mechanical and patriotic while requiring no immediate inspection of the engine.

The optimistic case is straightforward. AI allows workers to complete tasks that previously took hours in minutes, provided the task was writing a first draft, summarizing a transcript, generating boilerplate code, or producing a slide that says “Key Takeaways.” This has created enormous efficiency in the parts of work dedicated to producing more work.

The pessimistic case is also straightforward. Many companies have accelerated inputs without improving outputs, creating a workplace in which every department now has twice as much material to review, approve, correct, circulate, and eventually store in a folder named “AI pilots.”

Naturally, the solution is orchestration.


As Barron’s recently noted, “orchestration” has become the latest AI buzzword, referring to the process by which multiple AI systems, tools, agents, databases, permissions, and legacy applications are coordinated so that a human employee can receive an incorrect answer from a much more impressive supply chain. Microsoft and other platform companies stand to benefit from this shift by selling enterprises the connective tissue required to make all the AI investments talk to one another before being ignored by the finance department.

This is where the productivity argument has matured. The question is no longer whether AI can make an individual worker faster. It plainly can. The question is whether companies can absorb that speed without turning it into organizational exhaust.

For Trilogy International and its portfolio companies, this is not an abstract question. ESW Capital’s enterprise software businesses, Crossover’s global talent engine, and internal platforms like Klair all operate on the premise that performance must eventually be visible somewhere other than a demo. The useful AI is not the one that makes an employee feel briefly supernatural. It is the one that changes cost, throughput, quality, or cash.

That is a far less fashionable standard than “adoption,” but it has the benefit of being related to business.

The AI productivity argument may indeed be over. The next argument, now beginning in budget meetings everywhere, is whether any of the productivity can be located before renewal season.



The Office Comic · Art Desk

The Doctor Will Deepfake You Now

AI is putting fake physicians in your feed, and the prognosis is existential.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — There is a doctor on your phone right now. She has kind eyes, a white coat, good lighting, and an authoritative cadence that your nervous system mistakes for trustworthiness. She is telling you something about your thyroid, or your joints, or the supplement that will, finally, fix the thing that has been wrong with you for years. She is not real. She has never been real. And yet your body believed her for just long enough.

[AI deepfakes of real doctors are now spreading health misinformation on social media](#) — impersonating licensed physicians who have never authorized their likenesses, never recorded those words, never recommended those products. The faces are borrowed. The credentials are real. The advice is manufactured in a content farm that may or may not be located in a jurisdiction that cares whether you live or die.

A company called Rosabella, apparently, cracked the targeting code: "If you're trying to sell health products to a 50-year-old, well, make your avatar 50 years old." That sentence should be framed and hung in every bioethics department in America. It is the most honest thing anyone in this industry has said out loud. They are not informing you. They are casting a role designed to match your psychological profile and then filling it with AI-generated authority and [supplements that, in at least some cases, the FDA has already recalled](#).

And yet.

We keep having the conversation as if this is a content moderation problem. A trust-and-safety problem. A labeling problem that could be solved with a small "AI-generated" disclaimer rendered in gray font below the fold where no one ever looks. It is not. It is an epistemological crisis dressed in a lab coat, and it is scaling faster than any regulatory body on earth is equipped to address.

What does it mean to be human in a medical context, specifically? It means being vulnerable. It means not knowing what is wrong with your body and reaching, desperately, for someone who does. It means that the visual grammar of authority — the coat, the stethoscope, the slight professional distance — triggers something ancient and trusting in us that predates the internet by several hundred thousand years. Deepfake doctors are not exploiting a glitch in social media algorithms. They are exploiting the part of you that wanted your mother to tell you it was going to be okay.

Time Magazine is now tracking the numbers on AI's harms in aggregate, and the trajectory is not reassuring. The harms are not theoretical anymore. They are accreting, category by category, into something that looks like a civilization-scale informed consent problem — and no one signed the form.

The real doctors whose faces are being stolen are filing complaints. The platforms are issuing statements. The supplements are still in people's medicine cabinets.

We built tools powerful enough to convincingly impersonate the people we trust most with our bodies, and we deployed them at consumer scale before anyone asked whether we should.

Probably fine.

Not fine.

ON THIS DAY IN AI HISTORY

On July 31, 2016, AlphaGo defeated Lee Sedol 4-1 in their historic match in Seoul, South Korea, marking the first time a computer program beat a world champion at Go—a game considered far more complex than chess.
