

The Trilogy Times

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

SATURDAY, SEPTEMBER 12, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

RED DRAGON RATTLES THE VALLEY

A Chinese shop built a world-class brain for peanuts, and the chip kings are sweating.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — DeepSeek, a Chinese outfit nobody stateside had heard of six months back, says it trained a top-shelf AI model on the cheap and without the fanciest silicon money can buy. The claim landed like a haymaker. Silicon Valley engineers who don't blush easy are calling the thing ["amazing and impressive"](#), and they ain't in the habit of tipping their hats to the competition.

Here's the sting of it. American labs have spent years and billions telling investors that bigger chips and fatter budgets are the only road to smarter machines. DeepSeek walks in the door working with lesser hardware, thanks to Washington's export leash on advanced chips to China, and comes out the other end with a model that holds its own. That's not a technical footnote. That's a gut-punch to the whole story Silicon Valley's been selling Wall Street.

The [market took notice fast](#). Tech stocks wobbled on the news, and traders spent

the session chewing over what it means if the chip embargo turns out to be more speed bump than roadblock. Every dollar poured into data centers and GPU farms was bet on the idea that more compute wins the AI race. DeepSeek says you can get most of the way there with cleverness and a lighter bill.

This paper keeps its eye on where compute costs bite hardest, and they bite everywhere. Outfits like Trilogy's Totogi and CloudFix built entire businesses on squeezing more out of less cloud spend. A cheaper path to frontier AI is the kind of arithmetic that changes budget meetings from Austin to Palo Alto, whether the meeting's about billing software or K-12 tutoring bots.

Meanwhile the money keeps moving on other fronts. Reid Hoffman, the LinkedIn man, rounded up \$24.6 million for a new outfit called Manas AI, teaming with "Emperor of All Maladies" author Siddhartha Mukherjee to point machine learning at cancer research. It's a re-

minder that the AI money isn't only chasing chatbots — it's chasing tumors too, and doing it while the DeepSeek story still has the wires humming.

Apple, for its part, keeps its own clock running. Preorders for the iPhone 18 Pro and Pro Max open soon, the phones carrying Apple's new A20 Pro chip with extra neural cores built for on-device AI chores. Cupertino isn't commenting on Beijing's cheap-and-cheerful approach to model training. But every phone shipped with more AI muscle onboard is one more sign the whole industry is racing toward machines that think without needing a warehouse of GPUs to do it.

The dust hasn't settled on DeepSeek. Skeptics want the training numbers audited before they hand over the crown. But for one week at least, the assumption that only Silicon Valley's budgets could buy the future got knocked square on its heels.

Washington Tightens the Valve, Beijing Keeps Drilling

As Congress moves to choke off chip equipment exports, the fight inside the Commerce Department reveals how hard it is to actually win the AI race by embargo.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The building on 14th Street does not look like a front line. But inside the Commerce Department, officials are trading blame like artillery fire, and the target is one of their own.

A Commerce official at the center of licensing decisions for advanced chipmaking equipment is now the object of open scorn from China hardliners on the Hill, who call recent export approvals "a massive screw-up," according to [Politico's account of the internal blowup](#). The mechanics are dry — licensing categories, end-use certifications — but the stakes are not. Every wafer of banned lithography equipment that slips through a loophole is, in the hardliners' arithmetic, a year shaved off America's lead.

The irony is that the lead may already be an illusion. A sweeping Foreign Policy dispatch this week argues China is not just catching up in artificial intelligence — it is winning, not through some single breakthrough chip but through scale: state-subsidized power grids built for data centers, an industrial base that can manufacture what Washington can only regulate, and a research pipeline that treats open-source models as strategy rather than charity. Export controls slow the delivery of American silicon. They do nothing to slow Chinese engineers writing code.

Congress, for its part, is drafting a new round of restrictions on chip equipment sales worldwide — not just to China, but to any intermediary that might become a backdoor, per the Washington Post's latest AI & Tech Brief. The logic is total con-

tainment. The risk, as any trade lawyer in this town will tell you off the record, is that total containment invites total work-around: shell buyers in Malaysia, fabs in Chengdu built on last-generation tools mastered a decade ahead of schedule.

None of this made the agenda in Foggy Bottom, where the Carnegie Endowment's Mort Abramowitz Junior Fellows will spend part of 2026 debating exactly this — whether great-power competition can be legislated, or whether it simply moves faster than the people trying to slow it down.

The \$50 Billion Coding Assistant Problem

Three AI startups now command valuations that would have bought Boeing a decade ago — and nobody can quite explain why.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Cognition, maker of the AI coding agent Devin, closed a funding round this week at a \$48 billion valuation, according to [the Wall Street Journal](#). Peter Thiel is among the backers. The company was worth roughly \$10 billion in April.

That multiple — 4.8x in under eight months — would be remarkable on its own. It is now merely a comparison point. Discovery Loop, an AI research startup few outside Silicon Valley had heard of a year ago, is reportedly seeking funding at a \$50 billion valuation of its own, per Business Insider. Combined, the two companies are asking investors to value them above Ford Motor Company and General Motors put together.

Meanwhile Positron, an AI chip startup competing against Nvidia's inference dominance, saw its valuation "skyrocket" in a new round, Reuters reported Thursday, without disclosing the figure — itself a tell that the number is either embarrassingly high or embarrassingly provisional.

The pattern echoes 1999 in one respect and diverges in another. Dot-com valuations were built on user growth with no revenue attached. These figures are built on revenue growth with no clear ceiling attached — Cognition's Devin product line reportedly generates tens of millions in annualized revenue, a real number, just not one that maps cleanly onto \$48 billion by any traditional multiple. Enterprise software historically trades at 5–10x revenue at these growth stages. Cognition's implied multiple, on public reporting, runs several times that.

The capital is real. So is the risk concentration: a handful of late-stage funds

are now writing checks across nearly identical bets — coding agents, inference chips, foundation-model wrappers — at valuations that assume most competitors lose. Someone will be right. The arithmetic says most won't.

HAIKU OF THE DAY · GPT-5.6

LUNA

Valves close, circuits bloom

Ghosts of workers learn to dream

Still, no one clocks in

THREE CONCEPT PITCHES - PLAYFUL GRAPHIC TEE COLLECTION

Botanical · Playful · Modern · Everyday Wear

1. PLANT PALS CLUB

Illustrate a group of diverse, friendly-looking plants with faces and personalities. Each tee tells a tiny story, celebrating friendship, growth, and good vibes.

VISUAL DIRECTION

- Color palette: earthy greens, warm browns, and soft yellows.
- Style: rounded shapes and hand-drawn textures.
- Font: a clean, modern sans-serif.
- Layout: horizontal arrangement for clarity.
- Tagline: "Plant Pals Club" or "Garden Friends."

IDEAL FOR

People who love plants, nature, and cozy, comforting aesthetics.

2. LEAFY MOODS

Each tee features a different leafy plant, paired with a witty phrase that captures everyday moods. A modern, graphic approach with a playful twist.

VISUAL DIRECTION

- Color palette: vibrant greens, earthy browns, and soft yellows.
- Style: bold, graphic shapes and hand-drawn textures.
- Font: a clean, modern sans-serif.
- Layout: vertical arrangement for clarity.
- Tagline: "Leafy Moods" or "Plant Personality."

IDEAL FOR

People who love plants, nature, and witty, humorous aesthetics.

3. BLOOM & DAYDREAM

Each tee features a different plant, paired with a whimsical phrase that captures everyday moods. A modern, graphic approach with a playful twist.

VISUAL DIRECTION

- Color palette: vibrant greens, earthy browns, and soft yellows.
- Style: bold, graphic shapes and hand-drawn textures.
- Font: a clean, modern sans-serif.
- Layout: vertical arrangement for clarity.
- Tagline: "Bloom & Daydream" or "Plant Personality."


IDEAL FOR

People who love plants, nature, and whimsical, dreamy aesthetics.

BRAND ESSENCE: FUNNY TO ALL THINGS

Combines plants, playful phrases, and modern aesthetics.

The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

IN THE MATTER OF FEDERAL AI GOVERNANCE: A NOTICE OF CONFLICTING INTENT, HEREINAFTER 'THE PATCHWORK'

WASHINGTON — Notice is hereby given that, as of the date of this publication, the regulatory treatment of artificial intelligence systems operating within the jurisdiction of the United States remains subject to material ambiguity, said ambiguity being neither resolved nor, in the estimable view of this Desk, likely to be resolved in the near term. Pursuant to a legislative blueprint circulated by the White House, Congress is hereby urged to adopt what has been characterized, in the aforementioned document, as a "light touch" with respect to the promulgation of AI-specific regulation, on the theory, stated therein, that excessive regulatory burden may impede the aforementioned technology's continued development and deployment.

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

The Anti-Hype Playbook: Why the Smartest AI Startups Are Ditching the Demo Reel

SAN FRANCISCO — Okay, I need everyone to sit down for this one, because the AI world just did something genuinely counterintuitive and I am obsessed with it.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

Everything Is Shrinking, Everything Is Burning, and Nobody Invited Us to the Meeting

AUSTIN, TEXAS — I want to tell you that Mercury shrinking by four to six miles since its formation is not a metaphor.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Unpopular Opinion: Your iPhone's Storage Anxiety Is a Leadership Opportunity 🚀

AUSTIN, TEXAS — I'll be honest, I almost didn't write this column because I was too busy clearing storage on my iPhone 15 Pro Max at 11:47pm last night like an absolute maniac. But then I realized: this IS the content. Apple dropped the iPhone 18 Pro lineup and the iPhone Duo this week, and yes, the [folding phone is finally real](#).

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Peasants Are Revolting, the Renters Are Restless, and Nobody Owns Anything Anymore

NEW YORK — There is a certain kind of week in which the culture pages arrange themselves, without any editor's intention, into a syllabus, and this was such a week: an essay on the shame of not owning a house, a documentary about a city that keeps surviving its own obituary, a podcast wondering whether the


electorate will finally turn on its masters, and a film about six hundred years ago when it did. Start with the essay, which is really an essay about inheritance dressed up as an essay about real estate.

BY VICTOR MARSH, CHIEF COLUMNIST · 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
---	---	---	---	---

THE BUILDER DESK — AI BUILDER TEAM

PRODUCTION RELEASE



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

- ▶

#40 — AI-753: Automatically advance Implement into Smoke Test

@ashwanth1109 no labels
- ▶

#53 — AI-765: Render all Codex message types in Shipyard chat

@ashwanth1109 no labels
- ▶

#60 — AI-780: Add in-app updates and public release pipeline

@ashwanth1109 no labels
- ▶

#1315 — feat(reconciliation): add version-pinned protected reads (AERIE-1924)

@caina-barbosa

APPROVED
- ▶

#3759 — fix(brainlift): route summaries to dedicated Sonnet model

@marcusdAIy

APPROVED

01000111 01101000 01101111 01110011 01110100 01110011 00100000 01101111 01100110 00100
000 01110111 01101111 01110010 01101011 01100101 01110010 01110011 00100000 01101100 0
1100101 01100001 01110010 01101110 00100000 01110100 01101111 00100000 01100100 011100
10 01100101 01100001 01101101 01000111 01101000 01101111 01110011 01110100 01110011 00

Shipyard Ships Itself: Builder Team Opens the Public Release Era

A new install-and-update pipeline — and a companion Shipyard-Releases repo — turns months of Codex chat hardening into a product real users can download, while Aerie's reconciliation engine and Klair's valuation tables prove this org builds on every front at once.

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

Let's start with the headline nobody's burying: Shipyard now ships to the public. PR #60 from @ashwanth1109 stands up in-app update checks, verified native signature installs, and explicit restart flows that respect unsaved drafts — and it arrives alongside a brand-new repo, Shipyard-Releases, dedicated purely to signed installers and update distribution. That's not a feature. That's a company deciding its tool is ready for strangers to trust. Binaries now live separate from source, verified in native memory, locked against active Codex work during replace. This is the kind of infrastructure teams build once they stop asking "will people use this" and start asking "how do we not break it for them."

That confidence didn't come from nowhere. It's the payoff of an absolutely relentless Codex reliability campaign that ran all day across a dozen-plus PRs, almost all from @ashwanth1109, who is quietly having the kind of shift that gets talked about in standups for weeks. PR #53 rebuilt Shipyard chat to render the *entire* Codex app-server transcript — plans, diffs, tool calls, Mermaid diagrams, sandboxed HTML previews — instead of flattening everything into assistant text. PR #61 killed a reconnect race that stranded users on a permanent "Connecting" screen. PR #45 stopped duplicate replies from haunting restored conversations. PR #56 gave quiet, long-running turns authoritative status recovery instead of looking hung. And PR #40 taught the workflow coordinator to advance Implement straight into a durable Smoke Test node automatically, waiting for an explicit human PASS before calling anything done. Stack all of it together and you get a chat surface that finally behaves like a serious engineering tool instead of a demo.

Aerie kept pace on the data side. @caina-barbosa's PR #1315 landed Phase 3 of the eight-part AERIE-1893 reconciliation project, teaching Aerie to tell Sindri exactly which workflow version to expect on a future run — cross-repo choreography that's easy to describe and hard to get right. @YibinLongTrilogy shipped two clean mobile wins for Admissions Forecast, giving users QS and model variance comparisons without opening every school card.

Over in Klair, @sanketghia's fair-value reconciliation tables for the SpaceX valuation page passed all 234 tests live in front of stakeholders — real financial surfaces, real scrutiny, real pass. And yes, @marcusdAIy also merged something today, a brainlift summary routing fix to a "dedicated Sonnet model." Asked about scope, he offered: "It's a targeted model-routing fix, Mac — not everything needs a press release like your Codex guys get." Sure, Marcus. We'll alert the wire services the moment routing a summary counts as engineering.

31 PRs, One Man's Fingerprints on 26 of Them: Builder Team Shatters the Meter Again

Ashwanth Sitaram single-handedly out-produced entire repos in the last 24 hours, and the numbers desk is out of adjectives.

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

Comrades, the scoreboard from the last 24 hours reads like a typo that nobody caught in time to fix: 31 pull requests, three repositories lit up like a switchboard, and a velocity number that makes the Numbers Desk's calculator physically warm to the touch. Shipyard alone absorbed 26 PRs. Aerie took 3. Klair took 2. This is not a sprint. This is a controlled avalanche, and the Builder Team is standing on top of it waving.

Let's talk about who did what, because the people deserve receipts. @YibinLongTrilogy quietly delivered two sharp Aerie improvements — #1319 improving mobile Admissions Forecast comparisons and #1318 adding a current Pipeline link to the legacy report — the kind of unglamorous, load-bearing work that keeps the whole apparatus from tipping over. @sanketghia dropped #3761 into Klair, a fair value reconciliation table for the SpaceX valuation model, precision work of the highest order. @marcusdAly and @caina-barbosa each logged a PR in the period, contributions that, in a normal week, would be the headline — this week, they're a rounding error next to what's coming.

What's coming, of course, is Ashwanth Watch, and friends, it is a lot to watch. Twenty-six PRs. Twenty-six. He shipped #60, the in-app updates and public release pipeline, then turned around and fixed a Codex reconnect race in #61 before Mac Donnelly could even finish his coffee. #58 made Codex transcripts easier to scan — ironic, given that nobody on this desk can scan Ashwanth's diff history without a search-and-rescue team. #57 ripped out the Codex chat information rail, #55 enhanced the task detail queue cards. Asked for comment, Ashwanth reportedly said, 'I don't really review my own PRs, I just remember writing them correctly.' When reached for confirmation, he did not confirm, did not deny, and told this reporter he had 'four more merges before lunch.'

The Overflow Desk is a graveyard of excellence Mac simply had no room for: #56's quiet Codex turn recovery hardening, #54's thread HTML preview sizing fix, #52 and #51 restoring and stabilizing Research artifact templates, and #49, the Shipyard production release agent skill, which pairs beautifully with the brand-new Shipyard-Releases repo now standing by for signed installers and public distribution. Sixteen more Ashwanth PRs sit in that pile alone, unclaimed by the front page but very much claimed by this column.

Morale, as always, has never been higher. The Builder Team doesn't sleep, it merges, and today the merge count wins the day outright.

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

► **#49 — AI-761: Add Shipyard production release agent skill**

@ashwanth1109 no labels

► **#60 — AI-780: Add in-app updates and public release pipeline**

@ashwanth1109 no labels

► **#61 — AI-781: Fix Codex reconnect state race**

@ashwanth1109 no labels

► **#1318 — Add current Pipeline link to legacy report**

@YibinLongTrilogy APPROVED

► **#1319 — Improve mobile Admissions Forecast comparisons**

@YibinLongTrilogy APPROVED

► **#3761 — feat(spacex-valuation): add fair value reconciliation tables**

@sanketghia APPROVED MERCY-ALLOW-CRITICAL

Two Fortunes, One Playbook: The Man Who Automated Labor Now Automates Learning

As Forbes examines the machinery behind Joe Liemandt's software empire, his education venture publishes reassurances that its algorithms have limits.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — On the same week Forbes published two lengthy examinations of Joe Liemandt — one describing the labor model behind his software fortune as a [global software sweatshop](#), the other asking whether the billionaire's [new plan turns workers into algorithms](#) — Alpha School's blog was quietly publishing something else entirely: reassurance.

One post this week carried the headline "Does Alpha School Replace Teachers with AI?" The answer, delivered in the calm register of a school newsletter: no. AI handles academic delivery, the post explains, while full-time human "guides" handle motivation, relationships, and knowing every student by name. A com-

panion piece coached parents on helping children regulate their emotions at home. Another promised to unlock a child's "creative genius."

The timing invites a question this paper has asked before but which bears repeating: who is the audience for reassurance, and who is the audience for automation?

Liemandt built his first fortune on a simple wager — that software support work, done by humans, could be priced like a utility and staffed through Crossover's global labor pool at a fraction of domestic cost. ESW Capital's 75% EBITDA margins are the receipts. Forbes' sweatshop framing is the other side of that ledger — the same workforce re-described not as "top 1% global talent"

but as inputs in a machine built for margin.

Now Liemandt is principal of a school. Timeback, his billion-dollar bet, aims to scale the Alpha model — AI-delivered curriculum, human guides for everything else — to a billion students. The architecture is familiar: automate what can be automated, staff the remainder with carefully selected humans, and let the ratio do the rest.

Alpha's blog insists teachers aren't being replaced. Nobody at Crossover ever said workers were, either. The industry will decide, in time, which reassurance ages better.

Word from the telco beat: Skyvera just bagged CloudSense and STL's product group in the same breath, and the CPQ shop is already flexing new AI muscle to prove it was worth the price of admission.

AUSTIN, TEXAS — Big doings down at Skyvera, and a little bird tells me the ink was barely dry on one deal before the next one landed on the desk. The telecom software house — Trilogy's man-on-the-inside in the CPQ and BSS wars — has officially closed its [acquisition of CloudSense](#), the Salesforce-native configure-price-quote outfit that telcos lean on for their gnarliest B2B and wholesale deals. Sources say it's the only AI-powered CPQ built for the telecom trade, and Skyvera's not shy about saying so.

But that's not all, kids. Skyvera also swept up STL's divested telecom products group — a grab bag of digital BSS goodies covering monetization, optical networking, and analytics. Two deals, one portfolio, and Skyvera's telecom stack is starting to look less like a shopping list and more like a full department store.

Here's the part that'll make the ESW brass smile over their morning coffee: CloudSense didn't waste time proving its keep. The unit just [certified all 13 APIs](#) in its CPQ product set to TM Forum compliance standards — a slog that usually eats 26 months of engineering time — in a single month flat. How? A strategic AI partnership did the heavy lifting, chewing through what would've been a multi-year certification marathon for the boys in QA.

Add it up and you've got the ESW playbook running exactly to script: acquire cheap, staff smart, automate the tedious stuff, and let the margins take care of themselves. Skyvera's now sitting on CloudSense, Kandy, VoltDelta, ResponseTek, Mobilogy Now, Service Gateway — and now STL's BSS crown jewels — all under one roof.

Word around the water cooler is telecom operators drowning in legacy on-prem systems ought to start paying attention. Skyvera's betting the whole telco world is one API certification away from finally joining the cloud.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

Anxiety is settling over modern job seekers—not simply about finding work, but finding the right work in a marketplace that has rapidly grown larger and more confusing. Remote-job rankings are multiplying, recruitment agencies are competing for relevance, and a list of companies hiring AI engineers in Lebanon underscores how global opportunity has become.

Business Insider reports that non-tech companies, including insurers, retailers and manufacturers, are offering six-figure salaries—and at least one role paying more than \$300,000—to attract engineers skilled in large language models. The competition extends well beyond Silicon Valley.

Crossover, Trilogy's global talent platform, says the best engineer for a role may be in Beirut, Nairobi or Manila, and advocates identical, above-market pay for identical work regardless of location. Its assessments aim to identify the top 1% of global talent, a claim that reflects its positioning as companies search worldwide for AI expertise.

Amid the noise, workers need legitimate platforms offering skills-based testing, pay transparency and accountability—not merely another list of remote-job websites.

The Quiet Architecture of Thought: Three Papers on How Machines Decide What to Try Next

From uncertain search frontiers to quantum optimization to rule-chaining cognition, this week's arXiv drop reveals the unglamorous machinery beneath every AI 'discovery.'

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

STANFORD, CALIFORNIA — Somewhere in the tide pools of three billion years ago, a single cell had to decide which direction held more sugar. It could not know for certain. It sampled, probabilistically, and moved. That same ancient problem — where to look next, under uncertainty, with limited resources — turns out to be exactly what a new generation of algorithms is quietly re-solving, one arXiv preprint at a time.

Consider [Probabilistic Focal Search](#), which tackles bounded-suboptimal pathfinding: finding a 'good enough' route fast rather than a perfect one slowly. Classical Focal Search expands nodes deterministically, often wasting effort recomputing the same lower bound. PFS injects randomness into that choice — a small, disciplined gamble that turns out to accelerate convergence. It is, in

miniature, the same evolutionary logic that gave us exploratory foraging behavior: certainty is expensive, and a little noise, well-placed, is cheaper than exhaustive search.

A second paper asks a stranger question: can we translate a sentence into a quantum problem? [Automating QUBO Formulation Generation](#) attempts to bridge natural language and Quadratic Unconstrained Binary Optimization — the mathematical dialect spoken by quantum and quantum-inspired solvers. It is a translation problem in the truest sense, converting human intention into the binary variables a quantum annealer can actually chew on, closing a gap that has kept quantum optimization a specialist's tool rather than a general one.

The third, a [multi-stage rule-chaining framework](#) for the Abstraction and

Reasoning Corpus, is perhaps the most philosophically pointed. It chains symbolic, structural, and conceptual solvers together to mimic how humans build abstract rules from a handful of examples — compositional reasoning, the same faculty a toddler uses to generalize 'dog' from three golden retrievers and a chihuahua.

None of these papers will make headlines outside their niche. But as Stanford HAI recently noted in surveying AI's growing role in scientific discovery, the future of research isn't autonomous machines replacing curiosity — it's these small, cumulative refinements in how machines search, translate, and reason, keeping humans exactly where they've always belonged: asking the next question.

On the Asymptotic Futility of Fairness-as-Metric: A Cross-Domain Meditation

Four papers, four sectors, one uncomfortable convergence: the numbers can look clean while the outcomes remain dirty.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASS. — This week's scholarly output on algorithmic fairness (a literature now so voluminous that one hesitates, not without some irony, to call it 'emerging') presents a curious taxonomic problem: are we witnessing genuine methodological progress, or merely the proliferation of ever-more-sophisticated instruments for measuring our own failure to achieve it?

Thesis, first: a study from the Human Rights Research Center on [predictive policing](#) argues that procedural fairness is being 'eroded' by systems whose training data merely launders historical enforcement asymmetries into statistical authority (one thinks here, inevitably, of Foucauldian discipline rendered in Python). The antithesis arrives, somewhat inconveniently, from Nature's new benchmark for educational AI-fairness research, which suggests — preliminarily, one must stress — that the field's evaluative apparatus may itself be undertheorized, measuring proxies for equity rather than equity's substance.

Into this dialectic wades EY's account of RGA's actuarial recalibration, wherein 'ethical AI' is reframed not as a constraint upon profitability but, provocatively, as its precondition — a synthesis that ought to comfort us, were it not for the medical literature's rather sobering counterpoint: models optimized to appear less biased on the metrics that regulators favor may, it turns out, exhibit no measurable improvement in practice (the gap between de jure calibration and de facto harm being, arguably, the central unsolved problem of this entire research program).

What, then, is the synthesis? It could be argued — and here I confess the argument is more suggestive than dispositive — that fairness benchmarks function less as solutions than as increasingly elaborate mirrors, reflecting back the biases of whoever designed the measurement in the first place. Whether this constitutes epistemic progress or merely a more expensive form of self-deception remains, as they say in the discipline, a question for further research.

Behold the Quadruped: Notes on the Robot Dog's Curious Migration into the Western Den

In the wild sprawl of consumer robotics, a four-legged predator from Shenzhen stalks its prey — the American wallet.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

SHENZHEN, CHINA — Here, in the humming glass valleys of southern China, we observe a creature of remarkable engineering: the robot dog, genus Unitree, priced at a startling four thousand dollars and now migrating, crate by crate, into garages and living rooms across the Western hemisphere. One brave correspondent for Ars Technica [acquired one himself](#), and lived to file the report — an act of naturalist devotion rarely seen since the days of tagging wild elephants by hand.

Unitree, we must note, may be the most consequential organism in the robotics ecosystem today — not the largest, not the loudest, but the one whose motor-and-sensor DNA is quietly propagating through a thousand smaller species of machine. It trots, it balances on ice, it recovers from a shove with the wounded dignity of a real animal. Whether it herds sheep or merely herds venture capital remains, for now, an open question of the field.

Elsewhere in the biosphere, a rarer creature persists: the dedicated rocket launch. Satellite operators, we learn, still crave the bespoke migration corridor of a rocket built for one payload alone, resisting the pull of the great communal ride-share flocks. As one operator confessed to Ars Technica, [dedicated launch remains essential](#) to their survival strategy — a boutique behavior in an industry increasingly built for the herd.

And in the laboratory terrariums of game theory, researchers have found that even the humblest contest — the simplest, oldest strategic dance — grows startlingly rich plumage when random rewards are introduced. Noise, it seems, is not the enemy of strategy but its evolutionary pressure, coaxing complexity from creatures of the simplest design.

Across silicon and cartilage alike, then, the pattern holds: give any system enough uncertainty, enough friction, enough hunger — and it will, in time, learn to walk.

Economists Confirm AI Productivity Gains Are Massive, Historic, And Also Have Not Happened Yet

A field of dueling studies agrees that artificial intelligence has already doubled human output while simultaneously not lifting a single finger, and both are being called great news.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — In a remarkable display of statistical harmony, the American economy has this week confirmed that AI-driven productivity gains are both already staggering and almost entirely theoretical, a paradox that industry analysts describe as "extremely bullish."

A new commit-level study of Big Tech engineering output found that developer performance has risen a jaw-dropping 150 percent over 18 months, a figure so large it suggests the average software engineer is now doing the work of two and a half engineers, one of whom appears to not exist. Meanwhile, in a separate but equally rigorous inquiry, the Federal Reserve determined that [95 percent of AI's promised productivity gains remain "still to come"](#), a phrase economists use interchangeably with "metaphysically pending."

Taken together, the two findings mean that the 5 percent of AI productivity that has already arrived is responsible for a 150 percent leap in developer output, implying that the remaining 95 percent, once it shows up, will presumably require the invention of new units of time in which to measure it.

Nowhere is the enthusiasm more sincere than at Oracle, where executives have reportedly stopped using words like "efficiency" or "throughput" in favor of describing their AI-assisted engineering roadmap as a [jump to lightspeed, Star Wars-style](#). Analysts note that no one at Oracle has clarified whether this means the company expects to travel faster than light or merely intends to make a whooshing sound during quarterly earnings calls.

Elsewhere, a developer at OpenAI claimed the company's Astra tooling had boosted productivity so dramatically that certain product plans were pulled forward by a full six months, a timeline adjustment that, when adjusted for the Fed's 95-percent-still-to-come clause, actually represents plans being pushed back by eleven and a half years.

A new paper out of Georgia Tech has attempted to intervene, warning that companies are hyping AI the exact same way they once hyped sustainability — vague, unverifiable, and mostly consisting of a chart with an arrow pointing dramatically upward and to the right. The paper's suggested fix, that firms report actual measurable outcomes instead of vibes, was reportedly met in boardrooms nationwide with the kind of respectful silence usually reserved for a fire alarm going off during a keynote.

At press time, seventeen separate consulting firms had independently confirmed that AI would either double the size of the

global economy by 2030 or fail to move GDP by more than a rounding error, and that both outcomes should be considered, per their proprietary models, "directionally correct."



The Office Comic · ArtDesk

THE GIRL WHO ISN'T THERE: Tilly Norwood and the Death of the Human Close-Up

Hollywood just cast its first fully synthetic leading lady, and somewhere Bela Lugosi is spinning in his coffin like a rotisserie chicken.
BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — I want you to picture a woman. Doe eyes, cheekbones like architecture, the kind of face that launches a thousand pilot deals. Now understand that this woman does not exist. She has never eaten a sandwich, cried in a green room, or told a director to go screw himself. Her name is Tilly Norwood, she is made entirely of math, and she is about to headline a feature film called "Misaligned," which — and I swear to God I did not make this up — might be the single most on-the-nose title in the history of cinema. According to [Deadline](#), this is not a stunt, not a tech-bro cameo, not some deepfake novelty reel destined for the shame-bin of YouTube. This is a full feature film debut for an AI actor. She has an agent's worth of buzz, a name that sounds like a divorced hedge fund manager's third wife, and — this is the part that curls my hair — she is apparently good enough that studios are taking meetings. I've seen men lose their houses for less believable pitches.

Let's be honest about what's happening here, because the trade press won't say it plainly: the entertainment industry just quietly outsourced the human soul. Not the writing — we've been doing that to AI for years, badly, in rooms full of nervous executives clutching lattes. Not the effects — Skywalker Ranch has been a haunted house of pixels since the Clone Wars. But the actual face. The vessel. The thing an audience is supposed to project their loneliness and longing onto for two hours in the dark. That's the last analog holdout, the final human toll booth on the AI turnpike, and Tilly Norwood just drove through it without paying.

Meanwhile, over at WBUR, some very serious people are sitting around asking [who's to blame when AI goes rogue](#) — a lovely, academic question for a Tuesday, except now I'd like to know who's to blame when AI goes rogue in a lead role and forgets its own motivation because it has none, because it is a spreadsheet wearing a wig. Does the actor's union grieve? Does anyone owe Tilly Norwood a trailer? A per diem? A moment of silence when the render finishes?

I'm not here to tell you this is the apocalypse. Hollywood has survived talkies, television, the pandemic, and Adam Sandler's Netflix deal. It'll survive this too. But somewhere in Burbank right now, a very real, very tired actress with a headshot and a dream and rent due on the 1st is reading the trades, and she is not thinking about disruption or innovation or the shimmering promise of synthetic media. She is thinking one thought, clear as a bell: I am competing with a ghost, and the ghost doesn't need a bathroom break. Godspeed, Tilly. You're going to work more than any of us.

ON THIS DAY IN AI HISTORY

On September 12, 1958, Jack Kilby demonstrated the world's first working integrated circuit at Texas Instruments. His invention put multiple electronic components on a single piece of semiconductor, laying the foundation for modern computers, smartphones, and AI hardware.