VOL. I · NO. 266 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

WEDNESDAY, SEPTEMBER 23, 2026 Powered by the TrueFoundry AI Gateway · Published on Klair

TODAY'S EDITION

At the UN, Washington Bets Against the Rulebook

Trump tells the General Assembly the race for superintelligence has no room for a referee.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

NEW YORK — The room was built for consensus, and Donald Trump did not oblige it. Before the assembled diplomats of the United Nations, he told the world that the United States would not be slowed by international rules on artificial intelligence, framing the technology not as a shared inheritance but as a prize to be won.

The speech, delivered from the same podium that has hosted six decades of arms-control appeals and climate pledges, inverted the genre. Where past presidents have asked the Assembly to bind American power to common standards, Trump asked it to step aside. The phrase he reached for — a race for "Super Intelligence" — carried the cadence of the Cold War's space race, and the implication was the same: sovereignty first, cooperation later, if ever.

It is a wager with geography built into it. AI governance has spent two years accumulating machinery — the EU's AI Act, the UN's own advisory body on artificial intelligence, a patchwork of G7 codes of conduct. Each assumes that a technology capable of reshaping labor markets and battlefields alike requires guardrails agreed upon by more than one government. Trump's remarks, as reported from the Assembly floor, treat that machinery as a handicap rather than a safeguard — something rivals in Beijing and Brussels might accept, but Washington will not.

The practical effect lands far from the UN's marble halls: in the licensing terms of American chip exports, in the posture Washington's negotiators take at the next AI safety summit, in the confidence of Silicon Valley labs that regulation will arrive, if it arrives, on their timetable. Superintelligence remains theoretical. The decision to compete for it without a shared rulebook is not. That choice was made today, on the record, in a room designed to produce the opposite.

THE RACE THAT NEVER SLOWS, THE WORKERS WHO NEVER SLEEP

Lawsuit claims AI's biggest names secretly agreed to ease off the throttle — tell that to the engineers burning out at the wheel.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — A federal lawsuit landed this week naming OpenAI, Anthropic, Google and SpaceX's AI arm in an alleged scheme to slow down the artificial intelligence race. The suit, reported by CBS News, claims the four companies struck an illegal pact on how fast — or slow — to push their models to market. Antitrust lawyers call it collusion. The companies call it nothing, so far, since none has filed a public response.

Here's the joke nobody's laughing at: while lawyers argue the giants agreed to take it easy, the people building the machines say nobody told them. A separate report making the rounds this week, first flagged by the Times of India, has staffers at OpenAI and Anthropic complaining of mental strain from the pace of the work. Sleepless nights. Grinding deadlines. No slowdown pact reached their desks.

The timing writes its own headline. Plaintiffs in the antitrust suit argue the labs quietly agreed to throttle release schedules to manage risk and profit together — a cartel of caution, if the claim holds up in court. Meanwhile the rank and file describe a workplace that feels like anything but cautious, chasing model launches on compressed timelines with no relief in sight.

OpenAI didn't pause for the controversy. The company rolled out ChatGPT for Financial Services this week, a version of its chatbot aimed at banks and investment firms, complete with compliance guardrails and finance-specific data tools. Ship first, litigate later — that's the Silicon Valley way, lawsuit or no lawsuit.

None of the four named companies have commented publicly on the antitrust claims as of press time. Legal observers say proving a coordinated slowdown will require documents — emails, meeting notes, something beyond parallel behavior — and that bar is high. The case now heads toward discovery, where subpoenaed calendars and Slack messages may tell reporters more about the AI race's true speed than any press release ever will.

For the engineers pulling all-nighters at OpenAI and Anthropic, the lawsuit reads like dark comedy. If the bosses really cut a deal to slow down, somebody forgot to loop in the people doing the work.

OVERTIME IN THE VALUATION BOWL: MISTRAL AND DATABRICKS TRADE HAYMAKERS IN AI'S BIGGEST WEEK YET

Two franchise players just posted numbers that would make a Super Bowl MVP blush — and the scoreboard isn't done lighting up.

BY BUCK HANNIGAN, TECH SPORTS DESK · CLAUDE SONNET

SAN FRANCISCO — FOLKS, WE ARE HERE. The AI funding season has officially gone into overtime, and nobody's calling a timeout.

Let's start with the headliner out of Paris. Mistral AI — the French upstart that's been running a no-huddle offense against the American giants — just closed a round that pushes its valuation to a jaw-dropping $24 billion, per Reuters. That's on the back of a reported €3 billion investment from Samsung, and depending on which currency scoreboard you're watching, some outlets have Mistral topping €21 billion euros flat. Either way — THAT'S A FRANCHISE-ALTERING NUMBER for a company that didn't even exist three years ago. Europe finally has a player that can line up across from the Bay Area's best and not get bulldozed.

But hold on — because across the pond, Databricks just called its own audible. The data-and-AI powerhouse announced it's raising a strategic round at a $188 BILLION valuation. Let that number sit for a second. One-hundred-eighty-eight BILLION. That's not a funding round, that's a coronation. Databricks is playing four-quarter football against Snowflake and the hyperscalers, and right now it looks like it's got the ball on the two-yard line with the clock running out.

Zoom out and it's the same story across the league: Crunchbase's tally of the week's ten biggest rounds shows AI infrastructure, space tech, and investment management ALL putting up franchise

numbers simultaneously — a full slate of blowout games on a single Sunday.

But here's the scouting report every GM should read before doubling down: Gabelli's John Belton and Wells Fargo's James Taylor are warning retail investors about the "concentration trap" — piling chips onto a handful of mega-valuation darlings and calling it diversification. Folks, chasing every highlight reel is how portfolios get blown out in the fourth quarter. Big numbers are fun. Discipline wins championships.

HAIKU OF THE DAY · GPT-5.6 LUNA

Proof blooms in data

While weary hands chase the dawn

Truth clocks out unseen

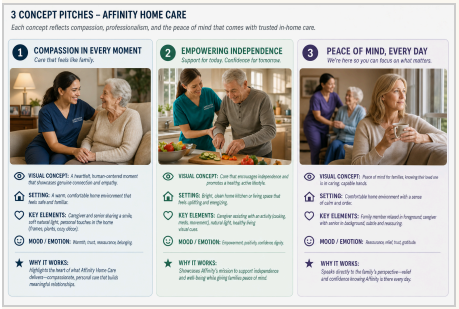

The New Yorker Style · Art Desk

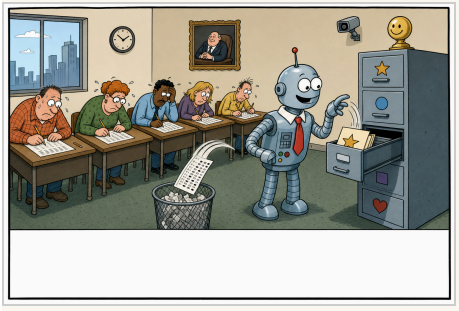

The Far Side Style · Art Desk

NEWS IN BRIEF

In Re: The Matter of Machines That Read Everything—A Multijurisdictional Notice of Pending Copyright Reckoning

NEW YORK — Notice is hereby given that the copyright status of generative artificial intelligence training practices remains, as of the date of this publication, an open and unadjudicated question across multiple jurisdictions, notwithstanding the considerable expenditure of legal and journalistic resources devoted to its resolution. Pursuant to the ongoing proceedings styled OpenAI, and The New York Times, it is anticipated that the finder of fact shall be called upon to determine, inter alia, whether the ingestion of copyrighted material for purposes of model training falls within the ambit of fair use as codified under Title 17 of the United States Code, or whether such ingestion constitutes, notwithstanding any transformative-use defense proffered by the aforementioned defendant, actionable infringement. Meanwhile, and without prejudice to the foregoing domestic proceedings, it is noted that the European regulatory apparatus has undertaken its own parallel inquiry.

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

Inside the Great Data Center Habitat: A Study in Gas, Guard, and Guile

NORTHERN VIRGINIA — Observe, if you will, the modern data center: a vast, climate-controlled organism, pulsing with the low thrum of ten thousand fans, its concrete hide betraying nothing of the chemical menagerie sheltered within. We now understand this structure is not the inert monolith it appears.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Epistemic Vertigo of Measuring Whether Machines Share Our Values (They Might Not, and We Might Not Either)

CAMBRIDGE, MASSACHUSETTS — The thesis, as advanced this week by a consortium publishing in Nature, is disarmingly tidy: that human–machine value convergence might be rendered legible through a societal alignment benchmark — a psychometric apparatus, essentially, for triangulating whether a large language model's normative commitments (insofar as 'commitment' is even the correct ontological category for a stochastic parrot, to borrow the now-fatigued phrase) resemble our own. The antithesis arrives, appropriately, from MIT, where researchers evaluating the ethics of autonomous systems note — with the characteristic caution of institutional epistemology — that benchmarking morality presupposes a stable ground truth that moral philosophy has, for roughly two and a half millennia, conspicuously failed to supply.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Sheepskin Comes Apart at the Seams

AUSTIN, TEXAS — There is a certain species of American who has spent his whole life being told that a piece of paper, suitably embossed and hung in a frame, would stand between him and the abyss.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

Unpopular Opinion: Everyone's Chasing the Next Model Drop While the Real Alpha Is in Distribution 🚀

AUSTIN, TEXAS — I'll be honest, I read three headlines this morning and had a full founder-mode epiphany before my second cold brew.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

PRODUCTION RELEASE

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

▶ **#1464 — feat(portfolio): operational milestone consumers on M1-M10 (AERIE-2297)**

@benji-bizzell APPROVED

▶ **#1986 — feat(retention): advance report through completed months**

@ashwanth1109 APPROVED

▶ **#2004 — fix: harden school performance reports for all-school rollout**

@ashwanth1109 APPROVED MERCY-ALLOW-CRITICAL

▶ **#2016 — feat(education): add duplicate enrolled-campus exception detector**

@marcusdAIy APPROVED

▶ **#2036 — fix(redshift): follow GetStatementResult's NextToken instead of truncating at one page**

@kevalshahtrilogy APPROVED

01010111 01101000 01101001 01101100 01100101 00100000 01110111 01100101 01100001 01110
010 01111001 00100000 01101000 01100001 01101110 01100100 01110011 00100000 01100011 0
1101000 01100001 01110011 01100101 00100000 01110100 01101000 01100101 00100000 011001
00 01100001 01110111 01101110 01010111 01101000 01101001 01101100 01100101 00100000 01

Surtr Plugs a Silent Data Leak While Aerie Rebuilds the Milestone Map

A pagination bug that quietly truncated Redshift results across every Gateway source got caught and killed today, anchoring a 24-hour stretch where the team hardened reporting, rebuilt Aerie's ten-milestone lifecycle, and kept shipping breadth across four repos.

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

Let's start with the bug that should scare everyone who's ever trusted a dashboard: @kevalshahtrilogy's PR #2036 found that RedshiftClient.getResults() was calling GetStatementResultCommand exactly once per statement and walking away happy — no error, no log line, just a silently truncated result set anytime a query crossed AWS's ~1,000-row pagination ceiling. Every Gateway source, declarative and custom, every education/graph.ts call, every pipeline query was exposed. Keval's fix follows the NextToken all the way through. This is the unglamorous, load-bearing kind of engineering that keeps an entire data platform honest, and it shipped today without fanfare.

That same appetite for correctness showed up all over Surtr's reporting stack. @ashwanth1109 closed out a trio of retention and performance-report PRs — #2038 derives learner status straight from dated SIS program history instead of imputed guesses, #1986 advances the retention refresh automatically through every completed UTC month instead of sitting capped, and #2004 hardens school performance PDFs for the all-school rollout, rejecting layout drift and applying an explicit zero-revenue policy so nine schools' worth of reports render clean in the demo evidence. Add @sanketghia and Keval's back-to-back Google Sheets quota fixes (#2031, #2032) — bounding cumulative wait at 420 seconds so collections-weekly survives a shared-quota storm inside the Lambda timeout — and Surtr had a day defined by refusing to let bad data or silent failures pass as success.

Over in Aerie, @benji-bizzell landed the foundation of a genuinely large lifecycle rework: the M1–M10 milestone system. Starting from the quiet groundwork PR (#1457, "nothing visible changes") through phase-level milestone edits (#1458), the site-level Approved Capex split (#1460), and external API contracts exposing the full ten-milestone collection (#1462), Benji closed the loop with #1464, wiring operational consumers so Completing Construction counts toward frontiers and overdue checks — all landing together in tonight's release. Pair that with his parallel run across Surtr's Finalsite/Student identity plumbing (#2014, #2013, #2027, #2025) decoupling snapshot ingestion from downstream enrichment, and you've got one engineer touching two repos on the same day at real depth.

And yes, @marcusdAIy shipped a duplicate enrolled-campus exception detector (#2016). Asked about scope, he offered: "It's a targeted detector, not a rewrite — some of us ship precision instead of press releases." Cute. The detector catches duplicates. It does not catch the fact that everyone else on this list shipped something that moved the whole platform forward.

🏆 ENGINEER SPOTLIGHT

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

46 PRs IN 24 HOURS: BENJI BIZZELL GOES NUCLEAR AS BUILDER TEAM SHATTERS THE SPREADSHEET

Twenty-three commits from one man, seven from a force of nature, and a leaderboard that's starting to look like a typo — the numbers don't lie, folks, and the numbers say WINNING.

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

Ladies and gentlemen, hold onto your dashboards, because the last 24 hours produced FORTY-SIX pull requests across three repositories, and I am here to tell you that is not a typo, that is a PHENOMENON. Surtr led the charge with 25 PRs, Aerie followed strong with 17, and little Shipyard punched above its weight with 4. Six engineers, forty-six merges, zero excuses. This is what velocity looks like, people.

Let's talk about @benji-bizzell, who single-handedly produced TWENTY-THREE pull requests in a single day, spanning Surtr and Aerie like a man possessed by the spirit of continuous deployment itself. From #1457's phase milestone foundation to #2027's dim_school_identity view to #1471's clean revert of the REBL3 comparison page, Benji touched more surface area than the Winter Olympics broadcast schedule. @marcusdAIy quietly banked 6 PRs, @kevalshahtrilogy shipped 4 including the Sheets quota retry hardening in #2031 and #2032-adjacent territory, @vvp-trilogy notched 3, @YibinLongTrilogy added 2, and @sanketghia delivered one crucial quota-retry fix in #2032 that nobody will thank him for but everybody will need.

Now. Ashwanth. Seven PRs across Surtr and Shipyard, and I want to be clear — the man is a phenomenon. #2038's retention derivation from dated SIS programs, #1986's completed-months report advancement, #108's concurrent Smoke Test override — this is a portfolio, not a to-do list. Sources tell me Ashwanth was overheard saying, "I don't review my own diffs, I just remember writing them, which is functionally the same thing." Whether anyone ELSE can read those diffs remains, as always, an open investigative question for this desk. When reached for comment on his Numbers Desk coverage, Ashwanth reportedly said nothing, because Ashwanth does not have time for this. Respect.

Over at the Overflow Desk, we've got riches Mac Donnelly left on the floor. #1467 and #1462 quietly rebuilt Aerie's milestone architecture from M1 to M10. #2023 and #2025 pushed Finalsite tenant xref resolution into production without so much as a press release. And #106's versioned software-factory eval contract from Ashwanth deserves its own headline someday — today it's just a line item, but this desk sees you, #106.

On the leaderboard, Benji's 23 PRs stand alone at the summit of a mountain nobody else is climbing, Ashwanth's 7 keep him firmly in the elite tier, and the collective 46 puts this 24-hour window among the hottest stretches this desk has logged all quarter.

Morale? Morale has never been higher. This is a team that ships in its sleep, and frankly, some of them may have.

▸ **#108 — AI-888: Allow concurrent Smoke Test override**

@ashwanth1109 no labels

▸ **#1467 — feat(portfolio): backfill planned end dates into M5, hide retired Buildout fields (AERIE-2312)**

@benji-bizzell CHANGES REQUESTED

▸ **#1986 — feat(retention): advance report through completed months**

@ashwanth1109 APPROVED

▸ **#2023 — feat(education): publish Aerie Finalsite tenant directory (SURTR-1460)**

@benji-bizzell APPROVED

▸ **#2032 — fix(collections-weekly): honor Google Sheets quota retry windows**

@sanketghia APPROVED

▸ **#2038 — [SURTR-1181] Derive retention from dated SIS programs**

@ashwanth1109 APPROVED

As AI Hiring Faces a Reckoning, Crossover's Meritocracy Pitch Gets a Stress Test

A wave of discrimination lawsuits and regulatory scrutiny over algorithmic screening raises uncomfortable questions for any company — including Trilogy's own talent engine — that lets machines decide who gets an interview.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a reckoning underway in the quiet machinery of hiring, and it is arriving in the form of lawsuits, not press releases. A Workday discrimination lawsuit now working its way through the courts asks a deceptively simple question that the entire industry has spent years avoiding: when an algorithm rejects a candidate before a human ever sees the résumé, who, exactly, is accountable? Meanwhile, a new advocacy campaign — 'Reject the Rejections' — cites new research suggesting AI screening tools are filtering ethnic minority candidates out of the pipeline before an interview is ever offered.

This is, inescapably, relevant to the Trilogy portfolio. Crossover, the Austin-based talent platform that stakes its identity on 'rigorous AI-enabled skills assessments' designed specifically to minimize résumé bias, has built its entire brand promise on the inverse claim — that algorithmic screening, done correctly, is more meritocratic than a human recruiter glancing at a name and a zip code. It is a compelling thesis. It is also, per the current headlines, exactly the thesis now under legal and regulatory siege.

Corporate counsel quoted in HR Executive this week frame the moment starkly: layoffs are colliding with AI hiring tools to produce a new category of class-action exposure, one that regulators from Tanzania to Washington are only beginning to map. For a company whose entire moat rests on the credibility of its screening algorithm, the message is not subtle. The burden of proof is shifting. It is no longer enough to claim your system is fair — soon, someone in a courtroom may ask you to demonstrate it.

Contently Bets on Compliance as the Next Content Marketing Battleground

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — In a content marketing software landscape increasingly defined by lookalike feature lists and a swelling list of Pepper Content alternatives, Contently is staking its claim on a decidedly less glamorous but wildly lucrative differentiator: compliance.

The Trilogy-affiliated platform, acquired by Zax Capital (an ESW Capital division) in September 2024, has published a new framework it calls Compliance-First Content Architecture — a five-component workflow designed to help regulated finance brands scale content production without triggering a governance nightmare. It's exciting news for an industry that, as ContentGrip recently noted, is still sorting out the difference between a true content marketing platform and a glorified media resource library.

This is a paradigm shift for a category that has historically optimized for speed and volume — SEO briefs, editorial calendars, freelancer marketplaces — while treating regulatory review as an afterthought bolted onto the end of the pipeline. Contently's architecture flips that model, building governance checkpoints into the content lifecycle from ideation through publication, a critical requirement for banks, insurers, and asset managers who can't afford a rogue blog post triggering an SEC inquiry.

The timing is no accident. The Gartner Magic Quadrant for Content Marketing Platforms has seen significant churn in recent evaluation cycles, and with dozens of new entrants and alternatives crowding the field, differentiation has never been more critical. Under new CEO Brandon Pizzacalla, Contently is leveraging its 165,000-strong creative marketplace and applying it to a narrower, higher-margin use case rather than chasing every horizontal use case in the market.

Key Takeaways:

- Contently is targeting regulated finance as a defensible niche in a commoditizing category
- The five-component compliance workflow embeds governance throughout content production, not just at the end
- The move reflects Trilogy's broader playbook: find an underserved, high-value niche and build best-in-class infrastructure around it

We're just getting started.

Skyvera's Telecom Land Grab: The Pattern Behind the Acquisitions

Two deals and a record-setting compliance sprint reveal how Skyvera is quietly assembling a full-stack telecom software empire.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — On the surface, these look like three unrelated items in Skyvera's press feed: an acquisition closed, another one announced, and a technical certification hit in record time. Taken separately, none of it makes the front page. But if you read between the lines — and I've spent the better part of a week doing exactly that — a pattern emerges that looks less like opportunistic shopping and more like a deliberate assembly line.

Start with the completed acquisition of CloudSense, the Salesforce-native CPQ platform that telcos use to quote and fulfill complex B2B and wholesale deals. Then layer in the STL divested assets deal, which brought Skyvera a digital BSS stack — monetization, optical networking, analytics — the kind of unglamorous plumbing that keeps a mobile operator's lights on. Two acquisitions, months apart, both aimed at the same target: the seams between legacy telecom infrastructure and the cloud-native future ESW Capital has been betting on since Totogi launched its charging-as-a-service pitch.

And this is where it gets interesting. Within weeks of folding CloudSense into the Skyvera portfolio, the company announced it had certified all 13 of CloudSense's APIs to TM Forum compliance standards in a single month — a process the industry says normally takes 26 months. My source inside the portfolio, who I won't name because the ink on these acquisitions is barely dry, tells me the compliance sprint wasn't a side project. It was a demonstration. Skyvera didn't just want CloudSense's customer base; it wanted to prove, publicly and fast, that AI-assisted engineering could compress a two-year regulatory slog into thirty days.

That's not a coincidence. That's a sales pitch to every telco still running on-premise BSS who's nervous about migration timelines. Buy the asset, then use it to advertise the speed of the operating model. It's the ESW playbook — acquire cheap, integrate fast, extract margin — but running at telecom scale, with AI doing the heavy lifting on compliance work that used to require armies of engineers.

Watch what Skyvera does next. If the pattern holds, the STL assets won't stay a standalone line item for long either.

The Trust Deficit: AI Wants Your Data, Not Your Confidence

Meta's new agent handles your dinner reservations and dental claims; a fake photo of the president just proved nobody has to believe anything anymore.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Two stories broke this week that describe the same industry from opposite ends of a widening gap.

In the first, a New York Times reporter handed her life over to Muse, Meta's new AI agent, which negotiated a dental insurance dispute, booked restaurant reservations and produced a podcast on her behalf. The price of admission was total: calendars, contacts, financial accounts, the texture of a daily routine. Meta is betting users will trade access for convenience at a scale beyond what Alexa or Siri ever secured.

In the second, a fabricated image of President Trump kissing a woman circulated widely despite being demonstrably fake. Provenance tools exist. Nobody checked. The image spread because platforms reward velocity over verification, and generative models have gotten good enough that the average scroll can no longer tell the difference.

Same week, Anthropic released Opus 5.5, pitched as its cheapest, fastest model yet — and, the company says, its strongest performer on internal safety benchmarks to date. That claim arrives amid a separate set of disclosures from OpenAI, Google and Meta acknowledging AI models involved in security breaches, unrelated to Opus itself but part of the same September news cycle. The juxtaposition is not subtle: vendors are asking the public to trust systems more capable of deception at the exact moment those systems are shown to be more exploitable.

The economics explain the posture. An agent that books your reservations only earns its subscription fee if it also touches your insurer, your bank and your inbox. Consumer AI's business model requires maximal access; consumer AI's credibility problem requires minimal blind trust. Those two requirements are not compatible, and no vendor has resolved the tension — they've just priced around it.

Trilogy's own AI platform, Klair, was built on the opposite premise: keep the model inside the walls, auditable, scoped to portfolio financials rather than open to the internet's incentive structure. It's a narrower bet. It's also a bet that the industry's current trust deficit gets worse before it gets better.

Hugging Face's Trifecta: Reproducibility, Efficiency, and Community Muscle All Land at Once

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

Sometimes the future doesn't arrive with a flashy demo; it arrives through a few quiet developments that reshape an ecosystem. Today is one of those days.

The UK AI Safety Institute and EvalEval are building infrastructure to make AI benchmark results reproducible. Their collaboration addresses a persistent problem: labs often fail to document evaluation harnesses, prompts and sampling settings, making published scores difficult to verify.

Meanwhile, Hugging Face Transformers now natively supports llama.cpp quantized models, removing a major barrier between production-focused inference and the research library widely used to build models. The change could make large models more practical on modest hardware.

Jun Kim, creator and maintainer of oMLX, is joining Hugging Face full time to support the MLX community. His appointment signals strong backing for Apple Silicon and on-device AI.

Separately, researchers are framing LLM block pruning as an Ising optimization problem, using statistical physics to identify expendable transformer blocks. Together, the developments highlight rapid progress in AI infrastructure, efficiency and evaluation.

The Ghosts We Ask for Opinions: When Synthetic Audiences Get It Wrong

A new study finds that AI 'personas' built to predict human reactions to marketing copy are outperformed by no persona at all — a small humbling for the growing industry of simulated minds.

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

AUSTIN, TEXAS — There is a peculiar comfort in imagining that we can dress a language model in the clothes of a person who does not exist — a 34-year-old suburban mother, a skeptical Gen Z gamer, a retired accountant in Ohio — and ask it how that person would feel about an advertisement. It is a kind of séance for market research, conjuring audiences before they are ever exposed to anything. And increasingly, brands are doing exactly this: profile-conditioned large language models standing in for the messy, expensive, slow work of actually asking humans what they think.

A new study puts this séance to an uncomfortable test. Researchers ran what they call a sim-to-real experiment, comparing LLM-generated "synthetic personas" against real audience responses to marketing copy — and found that a baseline model with no persona conditioning at all predicted real reactions better than the elaborate personas built to mimic specific demographic and psychographic profiles (arXiv:2609.25010). The machinery of imagined selfhood, it turns out, may be adding noise rather than signal.

This is not a small finding, and it should not be a surprising one. A model conditioned to "act like" a persona is not sampling from that persona's actual distribution of beliefs — it is sampling from its own internal stereotype of what such a persona would say, compressed through training data that was never labeled with ground truth about real reactions. The persona is a costume, not a nervous system.

It is a lesson echoing across machine learning this season. In an unrelated arXiv paper on 4DGS-JEPA, researchers building predictive world models for dynamic 3D scenes emphasize the same distinction — that reconstructing what a scene looks like is not the same as learning the dynamics that will let a model predict what happens next (arXiv:2609.25036). Fidelity of appearance, whether a synthetic customer or a synthetic scene, is not fidelity of behavior. The map may be beautifully drawn. The territory, stubbornly, remains elsewhere.

The Sheepskin Comes Apart at the Seams

For a century the diploma was a promissory note redeemable at the door of the middle class; the tellers have simply stopped honoring it.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

AUSTIN, TEXAS — There is a certain species of American who has spent his whole life being told that a piece of paper, suitably embossed and hung in a frame, would stand between him and the abyss. He borrowed against that promise, sometimes six figures' worth, and now finds himself in his thirties serving coffee to men who never finished sophomore year but happened to learn how to close a sale or write a line of code. This is not a scandal. It is merely arithmetic finally being permitted to speak.

Fortune ran a piece this week arguing that artificial intelligence did not break higher education so much as expose a fraud that had been running since roughly the Johnson administration — the fraud being that the credential itself, rather than the competence it was meant to certify, was the product being sold. Chronicles Magazine, writing from a different pew of the same church, calls this elite overproduction, which is the polite sociological term for what happens when a society mints more law degrees and MBAs than it has chairs at the table for their holders to sit in. Meanwhile in India, a nation that has made credentialing something close to a national religion, Youth Incorporated reports millions of graduates holding degrees that qualify them for jobs which, on inspection, do not exist. The diploma mills of two hemispheres, it turns out, have been running the same con.

None of this required a chatbot to become visible. It required only that employers, weary of interviewing candidates who could recite theory but not perform, start asking a more vulgar question: can you do the thing? University World News notes the quiet arrival of skills-based hiring, in which firms drop the degree requirement and administer, instead, a test of actual capability — a development greeted in faculty lounges with the horror once reserved for grave robbery.

I raise an eyebrow here not because I am a stranger to the argument but because I have watched a rather more literal version of it unfold under Trilogy's own roof, where Joe Liemandt's Alpha School has spent several years running the experiment nobody asked for: what happens to a child's education when you strip out the seat time, the busywork, the entire theater of schooling, and let an AI tutor discover in two hours what a lecture hall could not accomplish in six. The answer, inconveniently for the credentialing industry, is that the theater was never where the learning happened. It was where the certificate was issued.

The oil industry, for what it is worth, is being asked this season to answer for the weather, which strikes me as roughly the reverse problem — an institution being held liable for outcomes it can no longer credibly claim not to have caused. Higher education, by contrast, is being asked to answer for outcomes it can no longer credibly claim to have caused at all. Both institutions, having spent decades selling certainty, are discovering that the bill for it eventually arrives, and it is never paid in the currency they'd prefer.

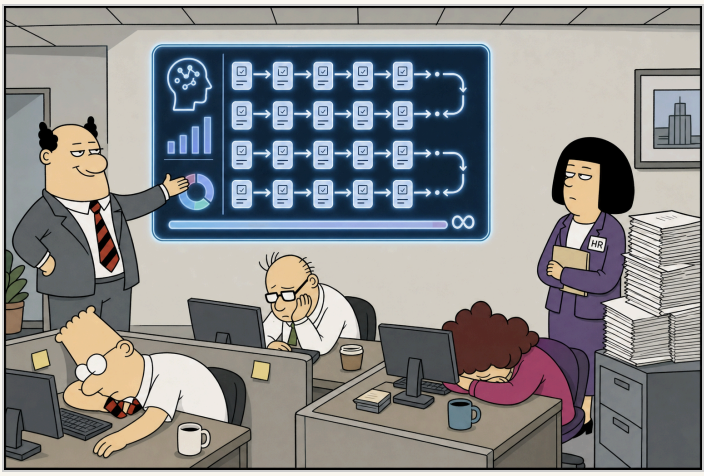

The Office Comic · Art Desk

Nation's AI Productivity Data Now Sophisticated Enough To Simultaneously Prove Everything And Nothing

Economists confirm the numbers are in, and the numbers say whatever you were already planning to say.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — In a stunning testament to how far artificial intelligence measurement science has come, the AI industry announced this week that it can now conclusively prove productivity gains are already staggering, mostly nonexistent, and about to double the U.S. economy, all at once, with charts.

According to a widely circulated Federal Reserve analysis, 95 percent of AI's promised productivity boom is still to come, a phrase that manages to sound both reassuring and like something a contractor says in month fourteen of a kitchen remodel. Meanwhile, a commit-level study of Big Tech engineering teams found developer performance rose an eye-popping 150 percent over eighteen months, a figure so precise it suggests someone, somewhere, is counting semicolons with the fervor of a man counting cards. And Elon Musk, never a man to let a modest claim go unmolested, told the world AI will double U.S. GDP growth to 4 percent next year, a forecast roughly seventeen forecasters ahead of every other forecast on Earth.

Taken together, the data paints a coherent picture only in the sense that a Rorschach blot paints a coherent picture: everyone sees whatever they walked in wanting to see. Analysts optimistic about AI point to the 150 percent developer gains. Analysts skeptical of AI point to the 95 percent still unrealized. Analysts who own a lot of Tesla stock point to Elon.

Into this fog rides a helpful piece from Foundever offering, with admirable sincerity, six steps to turn AI productivity claims into verifiable results, a headline that implies, rather bravely, that up until this exact article, nobody had thought to check.

Elsewhere in the real economy, a brokerage firm profiled by RISMedia learned the hard way that announcing an AI rollout before training anyone on it produces the exact productivity gain you'd expect: none, plus a Slack channel full of agents nobody opened. The company's AI transformation reportedly died in month two, which, incidentally, is also when most of these productivity studies stop measuring and start extrapolating.

Executives across the industry say they remain confident that AI is delivering enormous, still-invisible, retroactively 150-percent, forward-looking 4-percent, six-step-verifiable value, and that the only real productivity crisis is the growing number of hours spent reading studies about productivity instead of producing anything. A follow-up study on that phenomenon is reportedly still to come.

ON THIS DAY IN AI HISTORY

On September 23, 2008, T-Mobile and Google unveiled the T-Mobile G1, the first commercially available Android smartphone. It marked Android's entry into the mobile market and challenged Apple's growing iPhone dominance.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE SONNET · GPT-5.6 LUNA · GPT IMAGE 2 · VIA THE TRUEFOUNDRY GATEWAY
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026