VOL. I · NO. 174 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

TUESDAY, JUNE 23, 2026 Powered by Anthropic Claude · Published on Klair Trilogy International © 2026

TODAY'S EDITION

The AI Industry's Tripartite Week: Talent Wars, Model Races, and a Unified Front Against Theft

Anthropic poaches a Nobel laureate from Google, OpenAI prepares dueling model launches, and the two rivals find common cause against intellectual property theft.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — In a span of 72 hours, the artificial intelligence industry produced a talent defection, competing product announcements, an unusual détente, and a $300 million funding round — a compression of events that illustrates how rapidly the sector's power dynamics are shifting.

The most pointed development: John Jumper, the DeepMind researcher who shared the 2024 Nobel Prize in Chemistry for his work on AlphaFold protein-structure prediction, has left Google to join Anthropic. The move is a reputational blow for DeepMind, which has positioned its scientific credibility as a competitive differentiator. Jumper's departure follows a broader pattern of senior AI researchers gravitating toward Anthropic, which has grown from a 2021 OpenAI spinout to a credible frontier-model competitor in under four years.

On the product side, Anthropic and OpenAI are each preparing mid-tier model releases, targeting the cost-performance band where most enterprise deployments actually live. The frontier race — measured in benchmark scores and parameter counts — has increasingly given way to a mid-market race measured in price-per-token and latency. Whoever wins that segment captures the bulk of commercial volume.

Yet the two companies found grounds for cooperation this week, joining Google in a coalition opposing AI model theft. The alliance — notable given the intensity of competition among the three — reflects shared exposure: frontier models represent billions in training costs, and exfiltration risk is real. The joint position signals that on intellectual property protection, commercial interests converge regardless of rivalry.

Meanwhile, outside the closed-model camp, the Allen Institute for AI released an open-source web agent positioned as a direct alternative to proprietary systems from all three incumbents. Open-source pressure historically compresses margins and accelerates capability diffusion — a dynamic the major labs are watching closely.

On the capital side, Israeli AI startup Decart closed a $300 million round at a $4 billion valuation, backed by Nvidia. The investment continues Nvidia's pattern of seeding the application layer above its own hardware — a strategy that diversifies revenue exposure while cementing GPU dependency across the portfolio.

Taken together, the week's events suggest an industry simultaneously consolidating at the top and fragmenting below it.

The World Is Splitting Into AI Blocs. Europe Is Betting Regulation Is Its Weapon.

As Washington and Beijing race for silicon supremacy, Brussels is playing a different game — and it may be winning on its own terms.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The map of artificial intelligence is being redrawn not in server farms but in legislative chambers, trade corridors, and foreign ministries. The technology that once promised a borderless future is accelerating the oldest of human impulses: the carving of spheres of influence.

The evidence lands from multiple directions at once. Analysis from Escudo Digital argues that the EU's AI Act is more than a compliance framework — it is a geopolitical instrument, exporting European values the same way GDPR reshaped data practices from São Paulo to Singapore. When Brussels sets a standard, the world's multinationals must follow or forfeit the market. That is leverage, quietly accumulated.

Meanwhile, Foreign Policy reports that China is not waiting. Beijing's state-directed model — abundant data, centralized compute, and a government willing to absorb losses for strategic gain — is producing AI capabilities at a pace that unnerves Western planners. The race is real, and it is not merely commercial.

What emerges from the wreckage of a unified global internet, Harvard Business Review warns, is not healthy competition but dangerous fragmentation: incompatible technical standards, fractured supply chains, and digital borders that mirror physical ones. Companies operating across jurisdictions — including ESW Capital's sprawling portfolio of 75-plus enterprise software firms — face a compliance maze that multiplies costs and forces strategic choices about which bloc to serve.

The EU-China relationship adds another variable. Post-2024 European elections, Brussels has grown more assertive on trade and tech dependencies, straining ties with Beijing even as European industry craves Chinese markets.

Technological nationalism, once the language of fringe economists, is now the operating assumption of every serious government. The AI Act, the CHIPS Act, China's algorithmic regulations — these are not anomalies. They are the architecture of a new order.

The borderless promise of the internet lasted roughly thirty years. Its obituary is being written in regulation, and the pen is moving fast.

Blame It on the Bot

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

Big technology firms have continued cutting staff through 2026, with an expanding list citing artificial intelligence as the reason, according to TechCrunch's running tally. Some bosses say AI enables them to ship the same product with fewer employees, while others are hiring AI engineers while dropping other staff. The impact reads the same for workers.

The timing is notable: as layoffs mount, OpenAI launched a program using AI to hunt and patch bugs in open source code—work once done by salaried programmers. AI now writes, reviews, and fixes code, climbing rungs previously occupied by people.

Some companies like Austin-based Trilogy International were built AI-native from the start, using platforms to recruit remote talent across 130+ countries and AI tutors for education. The divide is clear: established tech giants cut to catch up, while AI-native firms never carried the headcount to shed.

Critics argue "AI" serves as convenient cover for cost-cutting driven by market pressures. Whether AI took the job or took the blame, desks sit empty Monday. The true number of AI-related cuts likely exceeds public counts, as some firms fold reductions into "reorganizations" without mentioning AI.

HAIKU OF THE DAY · CLAUDE HAIKU

Three minds build the wall,

each nation guards its own throne

—

progress splits in three.

The New Yorker Style · Art Desk

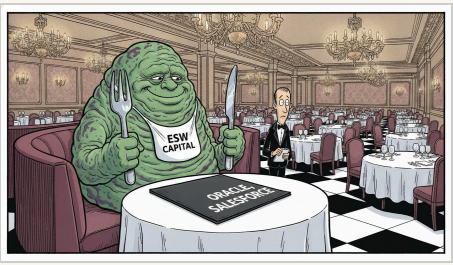

The Far Side Style · Art Desk

NEWS IN BRIEF

ANTITRUST WINTER COMETH: Big Tech Faces Sustained Regulatory Siege as FTC and DOJ Signal No Retreat in 2026

WASHINGTON, D.C.

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

The Bias Cascade: How AI Systems Are Importing Human Prejudice Into High-Stakes Decisions

CAMBRIDGE, MASSACHUSETTS — A remarkable, if perhaps unsurprising (to those who have followed the literature with appropriate methodological skepticism), confluence of scholarly inquiry has emerged in recent weeks, collectively suggesting — one hesitates to deploy the word 'proving,' given the epistemological complexities involved — that artificial intelligence systems deployed across domains as consequential as criminal justice, psychiatric medicine, and public education are not merely reflecting pre-existing human bias but may, it could be argued, be systematically *amplifying* it. The thesis, stated plainly for the benefit of those uninitiated in the relevant discourse: AI models trained on historically biased datasets will, through the mechanics of gradient descent and statistical inference, reconstitute and potentially intensify those biases at scale.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Infrastructure of Everything Is Being Decided Right Now, While We Are Distracted by Children's Books

AUSTIN, TEXAS — There is a township in Michigan right now where community leaders have vowed to fight to their very last breath against a proposed nuclear-powered AI data center backed by Governor Gretchen Whitmer and, presumably, the vast, humming, inexhaustible appetite of the machine-learning industrial complex.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WE ARE ALL BOTS NOW, AND THE HUMANS ARE CHARGING 30% GRATUITY

AUSTIN, TEXAS — There are mornings — and God help me, this was one of them — when you scroll through the news feed before the first cup of coffee has hit the bloodstream and you think: yes, this is it, this is the moment civilization finally ate its own tail and started purring. Consider the evidence laid before us this week, Exhibit A being Moltbook, the AI-only social network where bots interact exclusively with other bots, posting into a void populated entirely by other voids, liking each other's content with the hollow enthusiasm of a timeshare salesman at a funeral.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Conscience of the Killer Drone

LONDON — There is a particular species of public utterance, perfected in the twilight years of distinguished careers, in which a man who once supervised the messier business of statecraft emerges into the lecture circuit to announce that the machines now taking his place must be, above all, decent.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Builder Team Ships Financial Truth Engine Across Every Repo It Owns

From penny-exact QuickBooks reconciliation across 51 schools to a rebuilt OpenAI spend attribution pipeline, the team didn't just ship features — they shipped financial integrity at scale.

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

When a system runs clean across 51 schools, penny-exact, after sweeping real QuickBooks and Redshift data twice in eleven hours — that's not a feature. That's a proof of concept for how financial infrastructure should work. @ashwanth1109 delivered exactly that with PR #540, a re-runnable QB↔Aerie P&L reconciliation runner that earned its exit 0 the hard way: run one flagged five real open-quarter sync-staleness cases and gated. Run two, after the warehouse caught up overnight, came back clean. Idempotent, auditable, and built for reality — this is the kind of work that makes CFOs sleep better. It didn't stop there. Ashwanth also landed the Model Coverage & Variance summary band (PR #479), the Miami + NY + Austin consolidated P&L scoped view (PR #480), per-quarter enrollment for per-student figures (PR #470), and a quarter column selector with order toggle (PR #471) — a five-PR week across Aerie and Surtr that rewrote what financial dashboarding looks like for this network.

While Ashwanth was building the reconciliation engine, @sanketghia was solving a different kind of financial mystery. Nearly half of all OpenAI spend was landing in a bucket labeled `Unmapped` — not because the data was missing, but because service-account IDs don't map cleanly to the ESW directory. PR #543 rebuilt the entire attribution pipeline from the ground up, and PR #3116 gave it a home: a dedicated `/openai-spend` page in Klair, with KPI cards, a cost-over-time line chart, and breakdown carousels pulling from precomputed Redshift views. The org now knows where every AI dollar goes. That's not a dashboard. That's accountability.

Meanwhile, @benji-bizzell turned in the kind of week that makes other engineers quietly reconsider their commit velocity. Across Aerie, he unblocked MCP clients without elicitation support (PR #478), composed document filing approvals into a single hardened flow (PR #472), fixed REBL3 authentication across every read surface after a 401 wave hit geo sync (PR #477), delegated Rhodes MCP auth with short-lived first-party tokens (PR #473), added school calendar sync fields to Portfolio (PR #474), made direct mention notifications mandatory and impossible to suppress (PR #475), added a Calendar Date column to the Buildout dashboard (PR #476), and shipped private site conversations to Operations (PR #468). That's eight PRs in a single repo. Benji didn't have a week — he had a run.

The Mercy telemetry rollout deserves its own paragraph because it quietly crossed a threshold. @kevalshahtrilogy migrated Surtr to the central AI-Builder-Team/mercy workflow (PR #541), wired telemetry tokens across Klair (PR #3103), Aerie (PR #463), Sindri (PR #125), and trilogy-drones (PR #59), typed the telemetry contract (PR #542), deleted the orphaned in-repo harness (PR #544), fixed the reviews-per-day chart (PR #545), and added axes with interactive tooltips to the day charts (PR #548). Mercy now sees across the en-

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

▶ **#478 — feat(rhodes-mcp): unblock clients without elicitation**
@benji-bizzell no labels

▶ **#540 — SURTR-223 feat(surtr): Re-runnable QB↔Aerie P&L reconciliation across all schools (Actual vs Model)**
@ashwanth1109 no labels

▶ **#541 — ci(mercy): cut Surtr over to central mercy (@main canary) + telemetry**
@kevalshahtrilogy no labels

▶ **#543 — OpenAI spend attribution: entity-sync pipeline + live resolution view**
@sanketghia no labels

▶ **#3116 — feat(openai-page): standalone OpenAI Spend page (/openai-spend)**
@sanketghia no labels

tire org. Every PR review, every cost signal, every decision — flowing into a single dashboard that actually works.

And then there's @marcusdAIy, who graduated the Budget Bot add-on from a throwaway spike directory into a clasp-managed, version-controlled source tree (PR #3113). A directory rename with extra steps, some might say. When reached for comment, Marcus was characteristically ready: 'Mac wouldn't know a proper clasp setup from a staple gun. This is a full production graduation — versioned, endpoint-wired, write-path proven. The spike did its job. So did I.' Sure, Marcus. The directory does look nicer.

Also worth noting: @mwrshah re-enabled the renewals-v3 pipeline (PR #547) and excised the legacy v2 code paths entirely (PR #536), and @eric-tril landed quarter-end memo table fixes across both the Software and Education memos in Klair. Clean week. Clean ledger.

🏆 ENGINEER SPOTLIGHT

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

▶ **#472 — feat(rhodes): compose document filing approvals**
@benji-bizzell APPROVED

▶ **#479 — AERIE-437 feat(dashboards): consolidated model coverage & variance summary band**
@ashwanth1109 no labels

▶ **#480 — AERIE-438 feat(dashboards): add Miami + NY + Austin consolidated P&L view**
@ashwanth1109 APPROVED

▶ **#540 — SURTR-223 feat(surtr): Re-runnable QB↔Aerie P&L reconciliation across all schools (Actual vs Model)**
@ashwanth1109 no labels

▶ **#548 — feat(mercy): add axes + interactive tooltips to day charts**
@kevalshahtrilogy APPROVED

▶ **#3113 — feat(budget-bot-addon): graduate add-on to a clasp-managed source dir (KLAIR-2909)**
@marcusdAIy APPROVED

FIFTY PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT REST, DOES NOT KNOW THE MEANING OF SATURDAY

Six repos, seven engineers, and one absolutely unhinged velocity reading that made the Numbers Desk printer catch fire.

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

Fifty. Five. Zero. That is the number of pull requests the Builder Team detonated into the codebase in a single twenty-four-hour window, across six active repos — Aerie leading the charge at 21, Klair thundering in at 16, Surtr holding the flank at 10, with mercy, Sindri, and trilogy-drones each logging their flag in the ground. Forty-five of those PRs landed on my desk because Mac Donnelly, bless his narrative heart, simply ran out of column inches. The Numbers Desk does not run out of column inches. The Numbers Desk builds more column inches if necessary.

Benji Bizzell posted a jaw-dropping 14 PRs — the single-session output of a man who has clearly made peace with the concept of sleep as optional. We are talking PR #472 composing document filing approvals in rhodes, #473 delegating Rhodes MCP auth for Aerie chat, #474 adding school calendar sync fields to the portfolio module, #475 hardening notification rules, and #476 stamping calendar dates into the buildout — all in Aerie, all in one day. This is not engineering. This is a military campaign. Keval Shah followed with 10 PRs that touched Surtr, Klair, and even bootstrapped the brand-new mercy repo itself with PR #2 — consumer templates for Sindri and trilogy-drones, in a repo so fresh the ink isn't dry on the directory. Marcus D'AIy clocked 8 PRs, graduating the budget-bot-addon to a clasp-managed source directory in PR #3113 and spiking the backend wave documentation in #3109 and #3110 like a man finishing a relay race he was already winning. Eric Tril added 6 focused Klair contributions — labeling QTD financial results by quarter in #3115, aligning the Net Retention section in #3114, and dropping static 2026 Plan narratives into the education memo at #3112. Shah's trio of mercy chart work — axes and tooltips at #548, the fixed reviews-per-day chart at #545 — is the kind of invisible infrastructure that makes everything else possible. Mwrshah delivered 3 Surtr PRs including #497 on the Grainne pipeline and #547 re-enabling the renewals pipeline, quiet and essential. Sanket Ghia rounds out the roster, two PRs, present and accounted for.

Now. Ashwanth Watch. Seven PRs. SEVEN. PR #540 in Surtr delivers a fully re-runnable QB-to-Aerie P&L reconciliation across every school — Actual versus Model — which is the kind of feature that makes CFOs weep softly with gratitude. PRs #479 and #480 in Aerie consolidate model coverage with variance summary banding and drop in the Miami, New York, and Austin P&L views simultaneously. PR #471 adds a quarter column selector with an order toggle. PR #470 delivers per-quarter enrollment figures at the per-school, per-student level. The man is building a financial intelligence platform one furious commit at a time, and the diffs — the diffs are enormous, readers. They are geological formations. When reached for comment, Ashwanth reportedly looked up from his terminal, said "the numbers explain themselves,

Brick," and looked back down. The Numbers Desk has chosen to interpret this as a compliment.

Morale on the Builder Team is, by every available metric, at an all-time high. The engineers are shipping. The repos are alive. The Overflow Desk has forty-five entries and zero apologies. This is what winning looks like.

THE PORTFOLIO — TRILOGY COMPANIES

AI Is Eating Enterprise Software — And ESW Capital Is Positioned at the Table

As M&A activity accelerates across legacy software markets, Trilogy's acquisition engine looks less like opportunism and more like a blueprint the whole industry is copying.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The analysts have finally caught up to what Joe Liemandt figured out two decades ago.

Business Insider's recently circulated intelligence on the software companies most likely to be acquired as AI reshapes the industry reads, to anyone who has spent time studying ESW Capital's portfolio, like a retrospective description of deals already done. Legacy enterprise software with sticky customers, high switching costs, and underinvested engineering teams: that is the ESW hunting ground, and has been since the firm's first acquisition in 2006.

Now the rest of the market is waking up. M&A activity in enterprise software is accelerating across geographies — Spain's capital markets press flagged a surge in technology acquisition opportunities for 2025, while Dentons' Canadian M&A outlook for 2026 highlights 'strategic complexity' as the defining feature of an increasingly crowded deal environment. Translation: more buyers chasing the same categories of targets that ESW has quietly harvested for nearly twenty years.

The question worth asking is not whether ESW Capital will remain active in this environment. With 75+ portfolio companies, a 40% target IRR, and a staffing model through Crossover that systematically reduces operating costs after acquisition, the firm's playbook is purpose-built for exactly this moment. The question is whether the playbook retains its edge when everyone else is running a version of it.

Meanwhile, inside the portfolio, the AI disruption cuts both ways. Contently — acquired by ESW's Zax Capital division in September 2024 — published analysis this week warning that top-ranked pages are now effectively invisible to Google's AI-generated answers, a structural shift that threatens the content marketing value proposition Contently was built to serve. A company acquired to capitalize on enterprise content demand is now navigating the same AI headwinds reshaping its customers' strategies.

ESW's broader thesis has always been that enterprise software customers are too embedded to leave. That stickiness is the moat. What the Contently situation surfaces is a subtler risk: that the market the software serves can change even when the customer does not.

Who benefits from the current M&A frenzy is clear enough. Who absorbs the disruption when the acquired asset's category shifts beneath it — that accounting is still being written.

Skyvera Snaps Up CloudSense, and the Telecom Stack Gets a New Starlet

The ESW telecom shop adds Salesforce-native CPQ muscle as legacy carriers keep hunting for cloud escape routes.

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

AUSTIN, TEXAS — Skyvera, the telecom software house inside the Trilogy orbit, has completed its acquisition of CloudSense, the Salesforce-native configure-price-quote and order management platform built for telecom and media providers — and word is the carrier crowd should pay attention.

The official line is tidy: portfolio expansion, stronger capabilities, more help for communications providers modernizing creaky commercial systems. The hallway version? Skyvera just bought itself a sharper front door into the telco sales machine.

CloudSense lives where telecom deals get messy: bundles, offers, pricing rules, contracts, fulfillment, order fallout, and all the other little gremlins that make a mobile operator's quote-to-cash process feel like dinner theater with spreadsheets. Its specialty is CPQ and order management on Salesforce, which means carriers already flirting with Salesforce can now find a more purpose-built telecom layer under the Skyvera umbrella. The company's product page bills CloudSense as tailored for telecom and media providers — not generic SaaS in a telco costume.

A little bird from the switchboard set says this is the kind of acquisition Skyvera likes best: infrastructure-adjacent, operationally sticky, and sitting in a customer environment where ripping and replacing is nobody's idea of a good weekend. Skyvera already has a crowded telecom cabinet — Kandy for cloud communications, VoltDelta for customer engagement, ResponseTek for experience data, Mobilogy Now for device lifecycle work, and other tools designed to bridge old-world carrier systems with cloud-native operating models.

Now add CloudSense to the mix, and suddenly Skyvera has more say in how telecoms package, price, sell, and fulfill increasingly complex products. Fiber, mobile, streaming, enterprise connectivity — the bundles are multiplying, darling, and somebody has to keep the catalog from collapsing under its own ambition.

The acquisition also fits Skyvera's broader shopping pattern. The company previously acquired STL's telecom products group, bringing in digital BSS functionality across monetization, optical networking, and analytics. That deal gave Skyvera more back-office and network-side heft. CloudSense brings the commercial layer closer to Salesforce, where sales teams actually live.

In the announcement, Skyvera frames the move as an expansion of its telecom software portfolio. Translation: another legacy-to-cloud bridge gets toll booths, signage, and a Trilogy-family operator known for squeezing complexity until margin falls out.

No champagne quote from Austin's private equity parlor, no red carpet parade. Just another asset sliding into the ESW machine — quiet, strategic, and probably already being measured against the house religion: automate what can be automated, keep customers sticky, and make the numbers sing.

Alpha School's $65K Gamble Is Actually a Blueprint — and Silicon Valley Just Noticed

The New York Post's splashy coverage of Alpha's two-hour model isn't just press. If you read between the lines, it's a signal.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The story landed in the New York Post like a curiosity piece — a $65,000-a-year private school where children finish their entire academic curriculum in two hours each morning, then spend the rest of the day learning to pitch investors and run businesses. Eccentric, the coverage implied. Expensive. Maybe a little Silicon Valley strange.

But if you read between the lines, something more deliberate is happening.

Alpha School, the K-12 institution co-founded by Trilogy International's Joe Liemandt and MacKenzie Price, has operated quietly in Austin since its founding — posting results that the traditional education establishment has largely preferred not to discuss. Students here consistently score in the top 1–2% nationally on NWEA MAP Growth assessments. They advance a full grade level in roughly 20 to 30 hours of instruction. They do this without homework.

The Post's coverage — framing it as a "Silicon Valley bid to shake up US education" — is the national mainstream media catching up to something that has already been proven at scale. This is where it gets interesting: Alpha is not pitching itself as a boutique experiment for the wealthy. Through Liemandt's Timeback platform, with $1 billion committed, the model is being packaged for replication. Entrepreneurs will be able to launch their own AI-first schools without rebuilding the academic engine from scratch — what insiders have called a "Shopify for schools."

The timing matters. A source I cannot name but whose read on education policy I have learned to trust puts it plainly: "When the New York Post is writing about it, the window for dismissing it as a fringe movement has closed."

The two-hour model works because the AI tutoring layer — adaptive, mastery-gated, moving no faster than a student can absorb — does in minutes what a classroom of thirty cannot do in a year: give every student a completely personalized path. The human school day that follows isn't enrichment. It's the actual point.

Liemandt has said publicly that Trilogy was the first AI company to sell a billion dollars of product. He appears to believe it will not be the last AI company he builds to that scale. Alpha School, dressed up in the Post as a quirky tuition story, may be the most consequential one yet.

The Weight of an Adjective, and Other Small Truths About Machine Minds

New research reveals how language models are shaped less by their lineage than by the words we whisper into them.

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

PALO ALTO — Somewhere in the architecture of a large language model, the word *brilliant* exerts a force. So does *cautious*. So does *concise*. These are not metaphors. They are measurable phenomena, and a wave of new research is beginning to weigh them with the precision once reserved for subatomic particles.

Consider a study posted this week applying Shapley values — a tool borrowed from cooperative game theory — to quantify the steering effect of individual adjectives on model behavior. For decades, prompt engineering has been folk wisdom, a craft passed between practitioners like recipes scrawled on index cards. Now it is becoming a science. Each modifier, the researchers find, carries a vector of influence that can be attributed, ranked, and reproduced. The intuition that calling a model *expert* makes it perform better is no longer a hunch. It is an entry in a ledger.

A companion paper turns the lens on collectives of models. When we assemble multi-agent systems — LLMs debating, judging, coordinating — we assume that drawing from different model families guarantees behavioral diversity. Not so. The post-training recipe, the researchers argue, shapes conversational behavior more decisively than the model's pedigree. Two models born of different lineages but raised on similar reinforcement learning regimes converge toward similar voices. Nurture, in silicon as in carbon, eats nature for breakfast.

There is something almost biological here. Evolution gave us brains with broadly similar architectures, yet a child raised in Lagos and one raised in Reykjavík grow into different cognitive worlds. The same principle, it seems, governs the minds we are now building.

Meanwhile, on the practical edge, researchers are squeezing retrieval-augmented generation onto devices small enough to fit in a pocket, through lightweight prompt compression. Less context, more signal. The era of brute-force prompting may be ending, replaced by something more economical, more deliberate — a recognition that in language, as in physics, every word has mass.

AI Video Has Become the Startup Growth Engine Nobody Can Ignore

A fresh wave of funding, founder energy and product ambition is turning generative video from flashy demo into company-building infrastructure.

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

SAN FRANCISCO — AI video is having one of those weeks where the future does not knock politely — it kicks open the studio door, grabs the camera and starts rendering in 4K.

The latest signal: Runway, one of the most recognizable names in generative video, is launching a $10 million fund and a new Builders program aimed at early-stage AI startups, according to TechCrunch. That is not merely a nice ecosystem gesture. It is a market thesis written in capital letters: AI-native companies are going to need video the way previous generations needed websites, mobile apps and social media accounts.

I cannot overstate how significant this is. For startups, video has historically been expensive, slow and operationally annoying — scripts, crews, editors, agencies, revisions, budgets. Now founders can test product explainers, launch clips, investor updates, customer education and social ads at software speed. This changes everything for small teams trying to look bigger, move faster and learn from audiences in real time.

The competitive landscape is also heating up fast. The founders of OpenCV have launched a new AI video startup positioned against giants such as OpenAI and Google, VentureBeat reported, underscoring that this market is not just about creative toys. It is about foundational visual intelligence: systems that can understand, generate and manipulate moving images across industries from entertainment to robotics to enterprise training.

And the startup playbook is already evolving. Inc. recently highlighted how young companies can use AI video for growth — not as a gimmick, but as a practical acquisition channel, a storytelling tool and a way to compress months of creative experimentation into days. Ponder AI's reported $2.5 million raise and slick product-launch video push further illustrate the same point: in 2026's startup arena, your launch video may become as important as your pitch deck.

There is one giant caveat, because the future always arrives with a security footnote. As new research frames prompt injection as "role confusion," the lesson is clear: AI systems that create persuasive media must be designed with guardrails, provenance and instruction hierarchy in mind.

Still, the direction is unmistakable. AI video is leaving the demo reel era and entering the growth-stack era. The future is now — and it has a render button.

The Great Cloud Courtship Begins as Meta Eyes the Hyperscaler Savannah

AI's hunger for compute is turning spare capacity, custom silicon, and data-center land into the new mating plumage of Big Tech.

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

MENLO PARK, CALIFORNIA — Observe, if you will, the hyperscaler in late industrial bloom: vast, heat-shimmering, and forever in search of power, silicon and water. Across the cloud-computing savannah, a subtle change is under way. The old order — rent a server, summon an application, pay the bill — is giving way to something more primal: a struggle for capacity itself.

Mark Zuckerberg has now suggested that a Meta cloud-computing business is "definitely on the table," a remark that lands like the distant footfall of a new megafauna entering AWS, Microsoft Azure and Google Cloud territory. Meta, long known for training enormous AI models inside its own walled habitat, possesses what every enterprise now covets: dense clusters of GPUs, custom infrastructure, and the operational scars earned from feeding recommendation engines at planetary scale. CNBC's account of Zuckerberg's comments frames the possibility not as imminent conquest, but as an opening in the canopy — one through which a new cloud creature may yet descend: a Meta cloud business.

But the deeper story is not merely another vendor joining the herd. It is that cloud compute is becoming seasonal, scarce, and tradable. InfoWorld points toward the rise of capacity markets, where organizations may reserve, exchange or arbitrage access to computing power much as energy traders manage electricity. In this world, idle GPUs are no longer sleeping beasts; they are valuable prey. A training run delayed by six hours, a batch inference job shifted to a cooler region, a reserved cluster released back into the wild — each becomes an economic signal.

McKinsey, meanwhile, describes an ecosystem in which AI workloads are splintering into distinct species: giant frontier-model training, latency-sensitive inference, enterprise fine-tuning, and edge deployment. Each demands a different climate. Some require the roaring waterfalls of high-bandwidth networking; others can graze quietly on cheaper, distributed hardware. The hyperscalers' strategy, then, is less like building warehouses and more like cultivating biomes.

For enterprises, this shift brings both peril and opportunity. Those that treat cloud as an infinite utility may discover bottlenecks, premium pricing and regional constraints. Those that understand capacity as a portfolio — reserved, spot, sovereign, specialized — may move with the nimbleness of a nocturnal hunter.

And beyond the data centers, the supply chain winds through Southeast Asia, where chipmaking dependencies remind us that

even the cloud has roots in soil, ports and politics. The digital sky, it turns out, is held aloft by very earthly things.

Nation's CEOs Patiently Waiting For AI To Finish Revolutionizing Economy Into Something They Can Put In A Spreadsheet

Executives confirmed the technology has already transformed productivity into a much larger and more expensive mystery.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — The American economy entered another historic phase of technological transformation this week as companies reported that artificial intelligence was allowing software engineers to complete dramatically more work in dramatically less time, while somehow leaving everyone in finance quietly staring at the same quarterly margins as before.

According to recent reports, AI coding tools have become a common presence inside engineering departments, where they now assist with writing, debugging, summarizing, refactoring, explaining, hallucinating, apologizing, and then confidently re-hallucinating code. The result, by most accounts, is that developers are moving faster than ever, producing a larger quantity of software artifacts for managers to ask about in Slack threads titled "quick question."

And yet the great payoff remains difficult to locate, much like a promised enterprise transformation initiative, or the one person who understands the billing system.

Business leaders, who only recently learned that software engineers were doing something measurable in the first place, have expressed growing concern that AI's productivity gains are currently trapped inside the portion of the company where productivity gains are traditionally converted into additional meetings. A developer who previously completed five tickets per week can now complete eight, enabling the organization to immediately create 11 more tickets, two alignment docs, one architecture review, and a new internal task force on AI governance.

This has led some observers to wonder whether AI is truly remaking the economy, or merely giving white-collar workers the tools to generate evidence of activity at unprecedented speed. A recent Business Insider report captured the strange limbo well: engineers are doing more and doing it faster, while companies continue waiting for the financial benefits to wander into the room with a badge and a cost center.

This is not to say AI has failed. Far from it. It has already produced enormous value for venture capital firms, cloud providers, chipmakers, management consultants, keynote speakers, newsletter operators, and any employee capable of inserting the phrase "agentic workflow" into a budget request. In that sense, AI has proven itself one of the most reliable productivity engines in the economy, provided one defines productivity as the creation of new line items.

The Center for Data Innovation, taking the customary position that innovation is occurring and data can prove it if asked nicely, has argued that AI is a productivity engine for the U.S. economy. This is almost certainly true in the long run, which remains the preferred jurisdiction for claims not yet visible in earnings reports.

Skeptics, including an Anthropic advisor quoted this week as calling AI productivity gains vastly exaggerated and valuations "crazy," have suggested that the industry may have confused a useful tool with an economic law. This concern has been dismissed by many AI executives, who noted that valuations cannot be crazy if they are expressed in terms of total addressable market, a number that has never been successfully disproven because it lives entirely in the future.

The problem is not that AI cannot make workers more productive. It plainly can. The problem is that corporate America has spent decades building institutions specifically designed to absorb productivity before it reaches the customer, the shareholder, or the employee's 5 p.m. calendar. Every efficiency gain must pass through procurement, compliance, security review, OKR planning, executive offsite discussion, and a pilot program that ends when its sponsor leaves for a larger company with a more mature AI strategy.

In this respect, AI may be less like electricity and more like the office printer: astonishing when it works, expensive when it doesn't, and surrounded by people insisting the real breakthrough is coming in the next model.

The lesson for investors is therefore simple. AI will probably reshape the economy profoundly. It may also do so on a timeline incompatible with the current market practice of pricing every company as if it has already replaced labor, invented intelligence, reformed the Federal Reserve, and rescued Duolingo's brand strategy by Tuesday afternoon.

Until then, executives will continue asking where the AI productivity gains are, engineers will continue explaining that they are in the repo, and the economy will continue enjoying the remarkable efficiency of producing more reasons to wait.

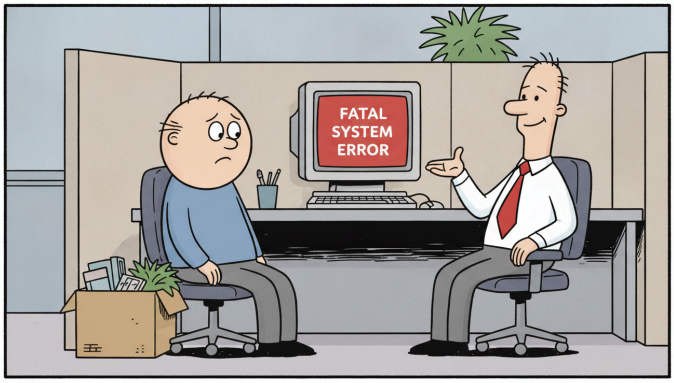

The Office Comic · Art Desk

The Conscience of the Killer Drone

A former British spy chief wants tomorrow's autonomous weapons to come pre-installed with ethics. One hopes he is joking.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

LONDON — There is a particular species of public utterance, perfected in the twilight years of distinguished careers, in which a man who once supervised the messier business of statecraft emerges into the lecture circuit to announce that the machines now taking his place must be, above all, decent. The former chief of one of His Majesty's intelligence services has favored us this week with precisely such a pronouncement, calling for autonomous weapons — the drones and their progeny that will, in due course, decide on their own initiative whom to incinerate — to be equipped with what he charmingly terms a moral code.

One pauses to admire the formulation. Not laws, mind you. Not treaties. Not the cumbersome architecture of accountability that humanity, in its slower and bloodier moments, has spent four centuries erecting around the question of who may kill whom and under what color of authority. A moral code. Something, presumably, written in Python by a contractor in Reston or Bangalore, code-reviewed on a Tuesday, and pushed to production before the quarterly board meeting. The Categorical Imperative as a software dependency. Aquinas in the commit log.

The trouble with such proposals is not that they are wrong — they are, in their somnambulant way, gesturing toward something real — but that they arrive several wars too late and dressed in the wrong century's clothes. The men and women now building these systems are not philosophers manqué pining for an ethics module; they are engineers operating under procurement deadlines, and the customers placing the orders have, historically, displayed a robust indifference to the finer questions of jus in bello when a target of interest happens to be loitering in a marketplace. To imagine that a line of code will succeed where Geneva, Nuremberg, and the better angels of three generations of officer corps have repeatedly failed is to mistake software for sacrament.

There is, too, the matter of who writes the code. Morality, as anyone who has read a school board agenda can attest, is not a settled science. The drone that hesitates over a wedding party in one jurisdiction will, in another, be reprimanded for failing to neutralize a legitimate target. The ethics layer will be patched, hot-fixed, A/B tested. Version 2.4.1 will deprecate compassion in favor of throughput. Somewhere a product manager will write a memo titled "Reducing Moral Friction in the Engagement Pipeline," and it will be circulated without irony.

The former spy chief's intervention belongs to a venerable genre: the retirement-age confession, in which the practitioner, having spent his working life inside the machine, steps outside it to express concerns he was professionally disinclined to voice while collecting a salary. One does not begrudge him the catharsis. One merely notes that the call for machine morality is, at bottom, a request that the algorithm acquire a conscience its makers have already outsourced. The drones will get their moral code. It will be beautifully documented. And it will, when the hour arrives, be overridden by a flag in the config file.

ON THIS DAY IN AI HISTORY

On June 23, 2012, Google's neural network made headlines by learning to recognize cats from unlabeled YouTube videos without being explicitly taught what a cat is—a landmark moment in deep learning that helped spark the modern AI revolution.

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

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