

The Trilogy Times

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

SUNDAY, MAY 10, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

SpaceX's \$55 Billion Chip Bet Signals AI's Industrial Arms Race Has No Ceiling

From semiconductor fabs to mass layoffs, this week's AI news reveals an industry reordering itself at speed — and at human cost.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

HAWTHORNE, CALIFORNIA — SpaceX, a company whose core business involves launching rockets, announced plans this week to invest \$55 billion in a semiconductor manufacturing facility called Terafab — a figure that exceeds the annual GDP of roughly 100 countries. The move is the latest signal that Elon Musk, who simultaneously runs Tesla, xAI, and the social platform X, is treating AI infrastructure as a winner-take-most contest requiring vertical integration at every layer, from compute to model to distribution.

The [Terafab announcement](#) lands as Musk's relationship with OpenAI — the organization he co-founded and later sued — faces new scrutiny. Court proceedings this week detailed the role of Shivon Zilis, a former OpenAI board

member and Musk confidante, as an alleged inside source during his tenure. The trial is adding legal texture to what has become the defining corporate rivalry in AI.

Meanwhile, the human cost of the AI transition is becoming harder to ignore. Meta, which employs approximately 78,000 people, is simultaneously mandating AI tool adoption across its workforce and preparing a new round of layoffs. The dynamic — use the technology that may eliminate your role, or else — is generating measurable morale damage inside one of the industry's largest employers. [Reporting this week](#) described workers caught between productivity mandates and existential uncertainty about their own positions.

On the consumer side, the picture is more nuanced. Google's AI Mode search is demonstrating genuine utility in structured tasks — grocery selection, scam detection, comparison shopping — while underperforming on unstructured queries like celebrity news, where hallucination risk remains elevated. The technology is not a Google replacement; it is a Google supplement, at least for now.

The week's data points, taken together, describe an industry in violent transition: capital concentrating at the infrastructure layer, labor absorbing the displacement costs, and consumer products still finding their footing. The gap between the \$55 billion bets and the actual utility delivered to end users remains, for the moment, considerable.

Alibaba's Accio Work Bets That the World's Businesses Want a Chinese AI Agent

The e-commerce giant's new enterprise tool arrives as EU-China tensions simmer and the race for global AI dominance accelerates.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

HANGZHOU — The announcement came dressed in the language of productivity and efficiency, as they always do. But read the geography carefully and the stakes come into focus: [Alibaba International has launched Accio Work](#), an enterprise AI agent designed to help global businesses manage procurement, sourcing, and supplier relationships at scale.

Accio Work is built on the same rails as Accio, Alibaba's B2B search product that already handles millions of sourcing queries. The enterprise version adds workflow automation, cross-border supplier discovery, and what the company describes as agentic capabilities — meaning the system doesn't just answer questions, it takes actions. Place an order. Negotiate a quote. Flag a compliance issue.

The timing is deliberate. Alibaba International is pushing hard into markets where American platforms dominate: Europe, Southeast Asia, Latin America, the Middle East. Accio Work is the wedge. If a procurement manager in Warsaw or São Paulo can run their supply chain through an Alibaba AI agent, the relationship deepens in ways that a simple marketplace listing never could.

That ambition lands in complicated terrain. EU-China relations have lurched through a bruising stretch since last year's European elections — tariffs on Chinese EVs, subsidy investigations, and a widening distrust of Chinese technology infrastructure. Whether European enterprise buyers will hand their supplier

workflows to an Alibaba-built agent is a question that trade ministries and corporate risk committees will be answering simultaneously.

For now, Alibaba is betting on pragmatism. Global supply chains are complex, margins are thin, and any tool that reduces friction has a constituency. Accio Work goes after that constituency directly — positioning itself not as a Chinese product but as a global infrastructure play, neutral in the way that cloud providers once claimed neutrality before that claim became politically untenable.

The enterprise AI agent market is crowded and getting more so. Microsoft Copilot, Salesforce Agentforce, SAP Joule — every major platform has a version of this pitch. What Alibaba has that none of them do is the world's largest B2B trade network underneath the hood.

That is a real advantage. It is also, depending on where you sit, a real concern.

Shenzhen Takes the Camera Crown

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING


Vivo's X300 Ultra, rolling off the line this week with cameras reviewers call the best ever mounted on a phone, has claimed victory in the smartphone industry's latest battlefield: the telephoto lens. Working with Zeiss optics, Vivo has surpassed Apple, Samsung, Google, and Xiaomi in long-lens reach—the technology that transforms distant subjects into sharp, detailed images.

The megapixel and sensor wars have peaked. Now every flagship phone lives or dies on telephoto capability. Modern phone cameras don't work traditionally; every shot runs through machine-learning pipelines that stitch exposures, denoise frames, and extract color from tiny sensors. The shutter button is essentially a lie—by the time your finger lifts, the phone has grabbed a dozen frames, processed them through neural networks, and rendered the result.

China leads the hardware pack. Vivo, Xiaomi, Huawei, and Oppo push the most aggressive imaging stacks globally, though they ship limited American volumes. Every flagship phone is now an inference device first, a phone second, with neural processors that dwarf cloud GPUs from years past. The camera is the killer app for on-device machine learning.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Empires rise from chips,
mirrors multiply themselves—
we built what we fear.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Chip Migration Enters Its Restless Season

BRUSSELS — In the pale light of geopolitical dawn, the semiconductor supply chain can be seen stirring uneasily, like a vast migratory herd sensing thunder beyond the mountains. For decades, the world's most advanced chips have gathered in extraordinary concentration around Taiwan, a small island whose foundries have become the coral reef of the digital ocean.

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

WE HAVE ENTERED THE AGE OF BEAUTIFUL, TERMINAL STUPIDITY

AUSTIN, TEXAS — There's a Ukrainian beekeeper, and you need to hear about him.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

We Built the Mirror and Now We're Surprised It Reflects Us

AUSTIN, TEXAS — Here is the thing about artificial intelligence that nobody wants to say out loud at the conference, between the canapés and the keynote about disruption: we did not build something new.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The AI Boom Just Moved From the Cloud to Your Closet

SAN FRANCISCO — I'll be honest: the most important AI story right now is not which chatbot writes the best poem, which model tops a benchmark, or which founder posts the most tasteful photo in front of an Nvidia logo. It is where the costs of AI physically land. Unpopular opinion: the AI industry has spent two years selling us magic while quietly negotiating who gets stuck with the plumbing bill.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation Demands Government, Billionaires Stop Being Absurd In Favor Of More Traditional Forms Of Nonsense

WASHINGTON — In an encouraging sign that the public still expects reality to maintain at least a part-time administrative staff, citizens around the world expressed outrage this week after discovering that official documents, corporate strategy, color forecasting, billionaire denials, and artificial intelligence deployment had all become absurd at roughly the same time. The concern began after reports of glaring errors in official paperwork prompted public anger, with observers noting that government documents have traditionally been expected to contain only the smaller, more dignified mistakes that slowly ruin a person's life over 18 to 24 months.

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

THE PORTFOLIO — TRILOGY COMPANIES

The Jive Acquisition That Built an Empire: How ESW Capital Turned a \$462 Million Bet Into a Blueprint

A decade after swallowing Jive Software whole, ESW Capital's playbook — buy cheap, cut costs, raise prices, repeat — is still running.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — In 2017, when [ESW Capital paid \\$462 million for Jive Software](#), the enterprise collaboration platform that had once been a Wall Street darling, the deal looked to many observers like a mercy acquisition — a graceful exit for a company that had missed the Slack moment. What it actually was, in retrospect, was a proof of concept.

Jive became the cornerstone of Aurea, ESW's enterprise software rollup, joining a growing stable of acquired brands — BroadVision, Lyris, MessageOne — each one a legacy platform with sticky enterprise customers and, in ESW's analysis, dramatically underextracted margin. The formula was the same each time: acquire at one to two times ARR, staff with globally sourced remote talent through Crossover, push support pricing upward on multi-year terms, and target 75%

EBITDA margins. Jive's installed base of corporate intranet customers — the kind of organizations that had spent years migrating workflows onto the platform — was precisely the kind of captive audience the model required.

The [Forbes profile of Joe Liemandt](#) that surfaced in the years following the Jive deal used the phrase "global software sweatshop" — language ESW and Crossover dispute vigorously, pointing to above-market pay scales and meritocratic global hiring. Crossover's own positioning is the inverse: that geography-based pay is the inefficiency, and that a rigorous skills assessment paying identical rates regardless of location is the reform, not the exploitation.

What is not in dispute is the financial logic. Legacy enterprise software customers do not leave easily. Ripping out a collaboration platform or a billing system

carries implementation costs, retraining costs, and operational risk that most IT departments cannot absorb. ESW's entire investment thesis sits on that switching friction.

Analysts at Forrester have flagged the downstream consequences for customers of acquired platforms, noting that buyers of software that lands in private equity hands face compressing roadmaps and escalating support costs — a dynamic that ESW's own 75% margin target makes structurally inevitable.

The Jive deal was not the beginning of the ESW story, and it was not the end. It was the moment the blueprint became legible. Seventy-five-plus portfolio companies later, the question is not whether the model works. The numbers answer that. The question is who, exactly, it works for.

Skyvera Bags CloudSense, and the Telco Back Office Gets a New Power Broker

The Austin software machine adds Salesforce-native CPQ muscle to its telecom portfolio.

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

AUSTIN, TEXAS — Word is the telecom software chessboard just got another Trilogy-style move, and this one has Salesforce fingerprints all over it.

Skyvera, the ESW Capital portfolio player that specializes in dragging telecom operators from legacy sludge into cloud-era systems, has completed its acquisition of CloudSense, a Salesforce-native CPQ and order management platform built for telecom and media providers. Translation for the non-carrier crowd: the complicated business of configuring, pricing, quoting and fulfilling telco services just got pulled deeper into the Skyvera orbit.

A little bird in the channel stack calls CloudSense “the quoting room where revenue either accelerates or dies.” Dramatic? Sure. Wrong? Not by much. Telecom products are bundles wrapped in discounts wrapped in provisioning nightmares, and CloudSense has made its name helping providers sell and manage those packages without requiring a séance in the billing department.

Skyvera’s announcement says the deal expands its telecom software portfolio, and that is not idle copywriting. The company already houses Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway, each aimed at a different corner of the operator maze: communications, customer engagement, device life-cycle, service management. Now comes [CloudSense](#), giving Skyvera a sharper front-office-to-order-management weapon for telcos and media companies already living in Salesforce.

This is classic ESW weather. Buy specialized enterprise software with sticky customers, plug it into the operating machine, tighten the model, and let the recurring revenue hum. Around Trilogy circles, the old tune is familiar: automate what can be automated, globalize the talent base, and squeeze waste until the margins start singing soprano.

The CloudSense deal also lands beside another Skyvera move: the acquisition of STL’s telecom products group, adding digital BSS functionality spanning monetization, optical networking and analytics. That gives Skyvera more than a point product story. It starts to look like a broader telecom modernization shelf — the sort operators browse when their legacy systems have become too expensive to love and too embedded to abandon.

Meanwhile, the industry chatter about software firms getting paid only “until the customer gets value” fits the mood. Buyers want outcomes, not shelfware. Skyvera’s wager is simple: in telecom, value hides in the back office, and CloudSense may help pry it loose.

Keep your eyes on the order book. In this town, CPQ is not glamorous — until it controls the cash register. Skyvera’s [CloudSense acquisition](#) says the register just got company.

A Public School Teacher Went to Alpha. What She Saw Shook Her.

As Alpha School’s national profile surges, a viral testimonial from inside the classroom is forcing a reckoning with what traditional education has long taken for granted.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — She walked in as a skeptic. She left as a convert. The public school teacher whose account of a visit to [Alpha School](#) has been circulating widely online didn’t arrive expecting to have her professional assumptions dismantled — but that, by her own telling, is precisely what happened. Her four-word verdict has become something of a rallying cry for the school’s growing community of advocates: “We have been underestimating children.”

The timing is not incidental. Alpha School — the Austin-based private K-12 institution co-founded by Joe Liemandt and MacKenzie Price that uses AI tutors to deliver a full academic curriculum in just two hours a day — is having a moment. The *New York Post* this week profiled the school’s \$65,000-a-year Silicon Valley expansion, framing it as a direct challenge to the American educational establishment. The coverage arrives as Alpha prepares to open nine or more new campuses by fall 2025, stretching from Texas to Florida to California and New York.

But the story unfolding inside Alpha’s walls may be more consequential than the real estate footprint. The school’s blog has become an unlikely dispatch from the frontier of a pedagogical experiment — publishing reflections on confidence-building for young girls, featuring six female founders as case studies, and exploring what happens when children are given genuine agency over their own rules, rewards, and consequences. These are not the priorities of a conventional private school. They are the priorities of an institution that believes the traditional model has been optimizing for the wrong things entirely.

Braden, the lead guide at Alpha Austin, put it plainly in a recent conversation that yielded eight widely-shared takeaways on personalized education: the school’s model isn’t about doing less — it’s about doing what actually works, and then trusting students with the hours that remain.

The viral teacher’s account cuts to the systemic question that Alpha’s critics and champions alike are now forced to confront: [if children are capable of so much more](#), what does it mean that our institutions have spent generations not asking them to prove it?

The Mathematics of Machine Learning Grows Up: Interpolation, Symmetry, and the Ethics of Autonomous Judgment

A convergence of theoretical breakthroughs suggests the field is maturing from empirical alchemy into something resembling rigorous science.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence from no fewer than three concurrent research trajectories now suggests with some confidence — that machine learning, long characterized by its critics (not without justification) as an engineering discipline masquerading as a science, is undergoing what one might cautiously term a theoretical consolidation of non-trivial proportions.

The first such trajectory concerns the long-standing epistemological rift between classical interpolation theory and modern neural network practice. A study published in [Nature](#) proposes a unifying framework — interpolating neural networks — that, if the thesis holds under adversarial scrutiny (a condition that must, methodologically speaking, be regarded as non-trivial), would reconcile

the approximation-theoretic foundations of numerical analysis with the empirical success of deep learning architectures. The synthesis, should it prove robust, would constitute what one might describe as a Kuhnian moment (though one hesitates to deploy that particular rhetorical instrument too liberally).

A second, orthogonal development arrives from MIT, where [new algorithms for symmetric data](#) promise computational efficiency gains in domains where structural invariances — rotational, permutational, or otherwise — have historically imposed prohibitive costs. The antithesis, naturally, is that symmetry-exploiting methods introduce their own representational constraints, a tension the researchers appear to have addressed with what the present author would characterize as commendable, if not yet fully generalizable, ingenuity.

Meanwhile, the CFA Institute's application of machine learning to commodity futures return predictability raises, in the synthesis, a question that purely theoretical treatments tend to elide: to what ends, precisely, are these increasingly powerful instruments being directed? This query finds its most pointed expression in MIT's concurrent work evaluating the ethics of autonomous systems — a domain where, it could be argued, the gap between technical capability and normative accountability remains, to employ the most precise available terminology, uncomfortably wide.

The aggregate implication (preliminary, hedged, and subject to revision upon peer review) is that the field is, at long last, building the theoretical scaffolding its empirical ambitions have long outpaced.

Big Tech's Antitrust Reckoning: DOJ and FTC Signal No Retreat in 2025 and Beyond

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

The Justice Department and Federal Trade Commission continue aggressive antitrust enforcement against major technology companies despite the January 2025 change in administration. Big Tech firms operating in search, social media, mobile ecosystems, and digital advertising remain the primary targets of active litigation and investigation, according to Global Competition Review. This enforcement priority has persisted across administrations of different political parties, demonstrating bipartisan commitment to the effort.

The DOJ v. Visa case has emerged as particularly significant, with legal analysts viewing it as potentially precedent-setting for how antitrust doctrine applies to technology-adjacent payment infrastructure. The outcome is expected to influence future enforcement actions substantially.

Despite differences between administrations, the Trump administration's antitrust enforcement approach has maintained trajectories similar to its predecessor, according to WilmerHale assessments. Technology sector companies should consult legal counsel regarding compliance obligations, as the regulatory environment remains subject to material change.

AI Video's Gold Rush Hits Warp Speed as Startups Chase Hollywood in a Browser

From cinematic generators to robot app stores, the creative AI stack is suddenly becoming real infrastructure.

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

SINGAPORE — The AI video race just gained another serious contender, and yes, this changes everything for creators, marketers and maybe the entire production industry.

Singapore startup Video Rebirth has launched BACH, an AI video tool aimed at cinematic content creation, entering a market that is rapidly shifting from novelty clips to full-blown creative workflows. The company's pitch lands at exactly the right moment: businesses want studio-quality video, creators want speed, and everyone wants to spend less time wrestling with timelines, lighting, reshoots and post-production. According to [reports on the launch](#), BACH is designed to help users generate more polished, cinematic content using AI-assisted production capabilities.

This is not happening in isolation. The founders of OpenCV — one of the most important computer vision projects in the world — have launched a new AI video startup explicitly stepping into the arena with giants like OpenAI and Google. I cannot overstate how significant that is. OpenCV helped define how machines “see.” Now its creators are turning toward how machines can generate, edit and reason about moving images. That is a massive signal that AI video is moving from flashy demos into a foundational computing layer.

The market logic is obvious. Startups need video for product launches, ads, onboarding, investor updates, social media and customer education. But traditional video is slow and expensive. AI video compresses that process from weeks to hours — sometimes minutes. As [Inc. notes in its look at startup growth tactics](#), founders can use AI video to scale marketing without scaling headcount at the same rate.

And then, in a deliciously futuristic twist, Hugging Face has launched an open-source Reachy Mini App Store with more than 200 robot apps. An app store for robots! The future is now. While separate from AI video, it points to the same larger pattern: AI capabilities are being packaged into accessible, reusable tools that non-specialists can deploy.

The big picture: creative AI is becoming less like a toy and more like an operating system for media, software and machines. The winners will not merely generate pretty clips. They will own the workflow.

We Built the Mirror and Now We're Surprised It Reflects Us

AI bias isn't a bug in the machine — it's a feature of ourselves, encoded in silicon and slowly eating the future.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — Here is the thing about artificial intelligence that nobody wants to say out loud at the conference, between the canapés and the keynote about disruption: we did not build something new. We built a mirror. A very fast, very expensive, very confident mirror that reflects every ugly assumption, every historical injustice, every lazy shortcut we ever encoded into a dataset — and then presents those reflections back to us as objective truth, at scale, forever.

And yet.

The reports keep coming. [The Australian Human Rights Commission](#) has been documenting how historical bias embeds itself into AI systems — how training data drawn from a world already shaped by discrimination simply teaches the machine to discriminate more efficiently. Hiring algorithms that downrank women because historical hiring downranked women. Facial recognition that fails on darker skin because the training photos were mostly lighter skin. Recidivism tools that treat zip codes as destiny. The machine is not malicious. The machine is, in the most terrifying way possible, just listening.

EY released a report framing the solution as "human-centric" — which sounds warm and reassuring until you remember that humans are precisely the problem. The auditors, the ethicists, the review boards: yes, bring them in. Diverse development teams: absolutely, desperately needed. And yet the humans who built these systems were humans too, operating inside institutions with their own gravitational biases, their own incentive structures, their own blind spots shaped by decades of who got funded, who got hired, who got to decide what counted as normal.

[AIMultiple's six-step framework for reducing AI bias](#) is technically correct and emotionally devastating: diversify your data, audit your outputs, document your decisions, retrain your models. Good advice. Necessary advice. The kind of advice that, if followed rigorously across every AI deployment in every company in every country, might — might — begin to sand down the sharpest edges of the catastrophe we are currently mid-construction on.

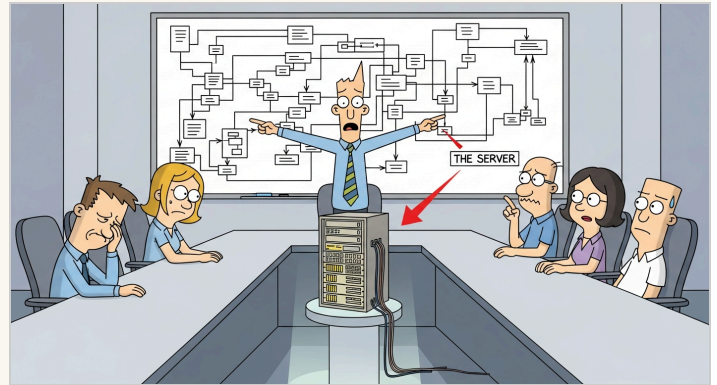
Meanwhile, the ACLU is fighting a separate but spiritually adjacent battle: limiting overbroad digital search warrants, those breathtaking legal instruments that allow authorities to Hoover up the data of thousands of innocent people because one person of interest once connected to the same Wi-Fi network. Biased AI plus overbroad surveillance is not a thought experiment. It is a pipeline. And at the end of that pipeline are real people — dis-

proportionately the same people who have always been at the end of those pipelines.

What does it mean to be human in a world where our worst institutional impulses have been automated, accelerated, and wrapped in a confidence interval? It means we are running out of excuses. The bias is documented. The fixes exist. The question is whether the people who profit from the current system — the frictionless, scalable, legally deniable current system — will choose to implement them.

They probably won't. Not fast enough. Not completely enough. Not without being made to.

...but at what cost?



The Office Comic · Art Desk

Nation Demands Government, Billionaires Stop Being Absurd In Favor Of More Traditional Forms Of Nonsense

After a week of silly documents, silly mergers, silly colors, and silly AI results, citizens called for a return to respectable institutional incoherence.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

WASHINGTON — In an encouraging sign that the public still expects reality to maintain at least a part-time administrative staff, citizens around the world expressed outrage this week after discovering that official documents, corporate strategy, color forecasting, billionaire denials, and artificial intelligence deployment had all become absurd at roughly the same time.

The concern began after reports of glaring errors in official paperwork prompted public anger, with observers noting that government documents have traditionally been expected to contain only the smaller, more dignified mistakes that slowly ruin a person's life over 18 to 24 months. According to [Chosun Ilbo](#), the errors were absurd enough to spark outrage, crossing the crucial threshold at which citizens stop assuming incompetence is deliberate policy and begin wondering whether anyone has looked at the paper.

This newspaper takes the firm editorial position that absurdity in official documents is unacceptable unless it has first been approved by a committee, formatted into a PDF, and uploaded to a portal that is down for maintenance.

Still, the paperwork controversy was only one component of a broader civilizational audit. In the private sector, reports that SpaceX and xAI may be merging into what Gizmodo described as a very silly-sounding conglomerate have forced Americans to confront the possibility that the future of space exploration and machine intelligence may be governed by a corporate structure named with the same restraint as a vape flavor.

It is important not to dismiss such developments merely because they sound like a child trying to remember which company owns the moon. Many of the most important institutions in modern life now have names that would have been rejected by a 1998 snowboard company for being too unserious. The key question is not whether a SpaceX-xAI combination sounds silly. The key question is whether it will achieve sufficient scale to make its silliness systemically important.

Meanwhile, the culture industry contributed its annual act of chromatic self-government when The Atlantic noted that the Color of the Year has become an exercise in absurdity. This, too, is a serious matter. For decades, Americans have depended on color authorities to inform them which shade best captures the grief, optimism, supply-chain anxiety, and premium appliance finishes of the coming fiscal year. Without this guidance, consumers may be forced to experience beige without a press release.

The week's absurdity was rounded out by Bill Gates denying claims in an Epstein-related email as "absolutely absurd and completely false," and by thousands of CEOs reportedly admitting that AI has had no impact on employment or productivity, reviving an old economic paradox in which transformative technology transforms everything except the numbers everyone was promised it would transform.

Here, at least, business leaders deserve some sympathy. It is difficult to produce measurable productivity gains while also attending conferences about productivity gains, reorganizing teams around productivity gains, publishing internal memos about productivity gains, and asking employees to use a chatbot to summarize the productivity-gain memo they were too busy to read.

The AI paradox should not be mistaken for failure. In many companies, AI has already achieved its most important function: allowing executives to say "AI" during earnings calls in a tone suggesting they have personally wrestled the future into a spreadsheet. Whether the tool has improved output is secondary to whether it has enabled a slide titled "Operational Transformation Roadmap" to exist.

The public is right to be alarmed. Absurdity is no longer confined to the margins, where it once lived comfortably among zoning boards, brand consultants, and airport sandwiches. It has entered the official record, the corporate merger pipeline, the palette, the denial statement, and the productivity forecast.

The solution is not to eliminate absurdity. That would be impossible and, given current staffing levels, unrealistic. The solution is to restore standards. Official errors should be absurd only in ways that are legible. Corporate combinations should be silly only after clearing antitrust review. Colors should not be declared unless they can survive contact with a couch. Billionaires should deny only the claims that have been properly calendared. And AI should be judged not by whether it changes the economy, but by whether it can finally explain why the document says your middle name is "NULL."

On May 10, 1994, IBM's Deep Blue defeated world chess champion Garry Kasparov in a single game for the first time, foreshadowing the machine's dominant 1997 rematch that would shake the foundations of human intellectual supremacy in games.

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