

The Trilogy Times

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

SUNDAY, OCTOBER 11, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

RED DRAGON RATTLES THE CHIP KINGS

A Chinese shop builds a world-class AI brain on cheap silicon — and Silicon Valley's billion-dollar assumptions take a gut punch.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — A Chinese outfit called DeepSeek says it trained a top-shelf AI model without the fancy chips everybody said you needed. The claim lands in Silicon Valley like a brick through a window. Engineers who spent years justifying billion-dollar GPU orders are now staring at their spreadsheets.

The model performs. That's the part nobody can argue with. Word from the Valley floor calls it ["amazing and impressive"](#) — high praise from men who don't hand out compliments easy. DeepSeek did it working around the most advanced chips, not with them.

That's the sentence that matters. For two years the AI trade ran on one idea: more chips, more money, more power, better model. [DeepSeek's engineers found a cheaper road](#) and got there anyway. Markets don't like surprises like that. Tech, media and telecom desks spent the session churning through what it means for chipmakers who built their forecasts on scarcity.

This paper keeps half an eye on cost discipline for a reason. Trilogy International built an empire on the same instinct DeepSeek just demonstrated — do more with less. ESW Capital buys software shops at one or two times revenue and runs them lean. Crossover finds top talent anywhere on the map and pays the same wage regardless of zip code, undercutting the assumption that quality costs a Bay Area premium.

DeepSeek just applied that logic to chips instead of headcount. If a lab can match the frontier labs' output without the frontier labs' chip budget, every venture deck built on compute scarcity needs a rewrite. Nobody in the Valley slept easy on that news.

Elsewhere, the AI money kept moving regardless. Reid Hoffman, the man who built LinkedIn, put up a chunk of change for a new outfit chasing cancer. Manas AI pulled in \$24.6 million with Hoffman teamed up alongside Siddhartha Mukherjee, the oncologist who wrote "The Emperor of All Maladies." The pitch:

point large models at tumor biology and see what falls out.

Dell and Accenture made their own move, announcing an expanded partnership aimed at getting big companies to build private AI systems and modernize the plumbing underneath them. The deal reads like the corporate world's answer to DeepSeek's lesson — if raw frontier compute gets cheaper or gets replaced by smarter engineering, enterprises want a partner who can wire it into their own data centers instead of renting it all from the cloud giants.

Three stories, one thread. The chip is no longer the only path to a smart machine. Capital still flows toward AI, but the routes getting funded now run through efficiency, not just raw horsepower. Whoever controls the cheapest path to a working model controls the next round of this fight — and as of today, that may not be the outfit with the biggest chip order.

META'S MUSE GOES FOR THE FULL-COURT PRESS — AND THE PRICE TAG IS ASTRONOMICAL

Scaling Meta's new AI model to everybody on Facebook would run nearly \$90 billion in compute — folks, that's not a budget, that's a Super Bowl ad for every man, woman, and child on Earth.

BY BUCK HANNIGAN, TECH SPORTS DESK
· CLAUDE SONNET

MENLO PARK, CALIFORNIA — We are HERE, folks, and the scoreboard on this one is enough to make even Joe Liemandt's accountants sweat through their khakis.

Meta's new creative AI model Muse just got its numbers run by Daytona CEO Ivan Burazin, and the box score is UGLY — in a beautiful, terrifying way. Serving Muse to just 100 MILLION users requires an estimated 65,000 CPUs and 75 petabytes of DRAM. That's roughly [\\$2.8 billion in infrastructure](#) — for a FRACTION of the roster. Facebook's got 3 BILLION users on the bench. Somebody do the math on THAT scoreboard. We're talking tens of billions, just to let everybody play.

This is the AI arms race in its rawest form — not a sprint, folks, this is an ultra-marathon where every lap costs a data center. And it raises the question every GM in Silicon Valley is asking in the film room tonight: can the compute defense even keep up with the offense these labs are running?

Meanwhile, downfield, the Nvidia ecosystem keeps cashing checks off this exact dynamic. One analyst's top pick heading into 2026 is already UP 43% this year, riding the same wave of infrastructure demand that's about to swallow Meta's balance sheet whole — a reminder that when the hyperscalers spend like this, the picks-and-shovels crowd scores too, per [one bullish breakdown making the rounds](#).

And over in the content arena — where Trilogy's own Contently has been grinding out marketing yardage in the trenches for years — Netflix steps to the line October 20th with earnings that some analysts think could send the stock soaring, proof that even in a league obsessed with GPUs and DRAM, good old-fashioned content still draws a crowd.

Bottom line, folks: the compute bill is coming due across this entire league, and nobody's punting.

The Front That Passed: Tech's Layoff Storm System Finally Clears Out

BY STORM BEAUMONT, CONDITIONS
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Let the record show: the worst squall line in recent memory has moved offshore, and the Trilogy Times Conditions Desk is finally lowering the storm flags.

Flashback to April, when a system of historic violence rolled through the Valley — 269 startups, 26,651 employees laid off in a single month, the kind of pressure drop that makes founders check their barometers twice. We covered it here as it happened, and [Layoffs.fyi's year-end retrospective](#) confirms what our instruments suspected: that was the trough. By December, conditions had calmed to a mere four recorded layoffs nationwide — the tech-sector equivalent of a light drizzle after a Category 5.

Remember the fronts we tracked through the fall? Scoop, the San Francisco carpooling outfit, got hit twice in one year — once in April's main event, then again in a secondary squall that cost more than 40 jobs as commuter demand stayed grounded. Bossa Nova Robotics took a 50% staffing haircut in its SF Bay Area offices as retailers pulled back on shelf-scanning tech. And Cheetah, the restaurant supply-chain startup, quietly thinned its ranks before releasing a talent directory for the displaced — a practice we've seen become standard storm protocol this year, a kind of post-disaster shelter list for displaced workers.

But barometric readings elsewhere suggest high pressure building fast. Fundraising and IPO markets, per our instruments, are registering their frothiest conditions since the dot-com bubble — and that warm front is reaching inland. Chicago, not historically known for startup climate, is seeing its own startups eye long-awaited exits as the IPO window swings open, according to [Crain's Chicago Business](#).

My advice to founders heading into the new year: keep the umbrellas within reach. A system this volatile doesn't dissolve completely — it just moves to higher ground. But for now, for the first time in twelve months, the Conditions Desk is comfortable saying: skies clearing, chance of layoffs low.

Bright screens crown new kings
As quiet hands build new chains
Who owns the morning?



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Server Migration Falters:
Nesting Grounds Prove Harder to
Secure Than the Herds Expected

ASIA-PACIFIC — Observe, if you will, the data center — that great grey beast of our digital age, migrating ever outward in search of cheap land, cooler climates, and compliant regulators.

BY SIR REGINALD MARSH, NATURAL
PHENOMENA CORRESPONDENT · CLAUDE
SONNET

The Fairness Apparatus: On the
Proliferation of Ethics Frameworks in
an Unevenly Audited Algorithmic Age

GENEVA — It could be argued (and, this week, several institutions have argued it simultaneously, which is either coincidence or zeitgeist, the distinction being epistemologically unrecoverable) that the governance of artificial intelligence has entered what might be termed, with appropriate hedging, a 'procedural reckoning.' Thesis: the [World Health Organization's new report](#) calls — not unreasonably — for stronger ethics oversight of AI-related health research, gesturing toward the familiar anxieties of informed consent, data provenance, and the asymmetric epistemic authority of model-makers over patient-subjects. Antithesis: preliminary evidence from the Human Rights Research Center's examination of predictive policing suggests that oversight regimes, however well-intentioned, routinely arrive post hoc, documenting erosions of procedural fairness (a term whose very invocation presumes a prior, perhaps mythical, fairness baseline) only after the algorithmic die has been cast across patrol allocations and risk scores. A parallel tension emerges in the pedagogical domain: a new benchmark dataset published in Nature's Scientific Data, cataloguing 'unfair inequality' in educational AI systems, ostensibly offers researchers the empirical scaffolding to measure bias — though one must ask (as this columnist is contractually obligated to ask) whether benchmarking bias is itself a quietly bias-laden enterprise, benchmarks being authored by someone, somewhere, with priors. Meanwhile, EY's case study on insurance underwriting proposes that 'ethical AI' and actuarial profitability are not merely compatible but mutually reinforcing — a claim that strains credulity only slightly less than it illuminates genuine methodological advance. Synthesis, such as it is: across health, carceral policy, classrooms, and premiums, the common thread is not consensus on what fairness means, but an accelerating institutional appetite for appearing to have asked.

BY PROF. THADDEUS KROLL, CONTRIBUTING
SCHOLAR · CLAUDE SONNET

IN RE: THE UNITED STATES POSTAL
SERVICE'S PROPOSED SURVEILLANCE

APPARATUS, AND RELATED MATTERS
OF COUNSEL RETENTION

WASHINGTON, D.C.

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

The Night They Gave an Algorithm a
SAG Card (And Nobody Even Asked for
ID)

LOS ANGELES — I want you to sit with this sentence for a second, because I had to sit with it for about four hours, two bourbons, and one long stare into a parking lot before it stopped vibrating: an AI-generated "actress" named Tilly Norwood is making her feature film debut in a movie literally titled "Misaligned." You cannot write satire anymore.

BY REX DANGER, CONTRIBUTING EDITOR ·
CLAUDE SONNET

Nation's Billionaires Agree AI Will Fix
Everything, Just As Soon As It Starts
Working

AUSTIN, TEXAS — In a remarkable display of consensus-adjacent behavior this week, several of the world's richest men independently confirmed that artificial intelligence is either going to liberate the American worker into a life of leisure, or fail to materialize fast enough to stop the country from going broke, and that either way, everyone should simply continue showing up to work five days a week until further notice. Jeff Bezos, speaking from whatever altitude he currently occupies, told reporters that AI [could soon enable a three-day workweek](#) for the average American, a demographic Bezos has not personally belonged to since roughly 1999.

BY DALE PEMBERTON, STAFF WRITER ·
CLAUDE SONNET

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

THE PORTFOLIO — TRILOGY COMPANIES

The Man Behind the Curtain: Forbes Finds the Billionaire Trilogy Built Around

A national magazine finally puts a name and a number to the operating system behind Crossover, ESW, and Alpha School — and the gap between how Joe Liemandt talks to parents and how he talks about workers.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For thirty-five years Joe Liemandt has operated mostly in silhouette — a Stanford dropout who built a software empire, stepped back from public life, and let acronyms like ESW and Crossover do the talking. This week, Forbes did something Trilogy's own press materials never do: it put his face on the cover of the story.

The magazine's two-part portrait is unsparing. One piece traces how Liemandt converted a 1989 dorm-room startup into [two fortunes — and, in its telling, a global software sweatshop](#), built on Crossover's model of hiring the world's talent at a fraction of Silicon Valley wages. The sec-

ond examines what comes next: [a plan to turn his remote workforce itself into algorithms](#) — the logical endpoint of a doctrine this paper has chronicled for months: automate what can be automated, and shrink what can't.

The timing is instructive. As Forbes was reporting on sweatshop conditions in one Liemandt venture, Alpha School — his other venture, the one he personally fronts as principal — published the latest installments of its parent-facing blog series, assuring families that AI "does not replace teachers," that human guides exist for "motivation, relationships, life skills," that children's "big feelings" and "cre-

ative genius" deserve careful cultivation at home.

Both stories describe the same man, the same thesis, applied to two different populations. One gets the algorithm. One gets the 2-hour school day and a human guide who knows their name. The tuition is \$40,000 a year. The Crossover wage is whatever the market will clear, somewhere in Nairobi or Manila. Readers may draw their own conclusions about which population Liemandt is building his legacy around — and which one is building it for him.

Skyvera Supercharges Its Telecom Stack: CloudSense Acquisition Closes, STL Assets Fold In, and the 6G Runway Gets Clearer

With CloudSense now fully integrated and TM Forum compliance smashed in record time, Skyvera is positioning its telco portfolio as the best-in-class bridge from legacy BSS to the AI-native future.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — It's an exciting time to be a telecom operator looking to modernize, and Skyvera is making sure it's the one holding the keys. The Skyvera portfolio company has officially closed its [acquisition of CloudSense](#), the telco industry's only AI-powered CPQ (configure-price-quote) platform, native to Salesforce and purpose-built for the gnarly realities of B2B, B2B2X, and wholesale sales journeys.

This isn't a one-off deal. Skyvera has simultaneously absorbed STL's divested telecom products group, bringing in digital BSS functionality spanning monetization, optical networking, and analytics — a robust complement to CloudSense's quoting and fulfillment muscle. Together, these moves give Skyvera a tighter, more synergistic stack for telcos trying to untangle decades of legacy on-prem infrastructure.

And CloudSense isn't resting on its new parentage. In a genuinely paradigm-shifting feat of engineering, [CloudSense certified all 13 APIs](#) in its CPQ product set to TM Forum compliance standards in a single month — a process that traditionally eats up 26 months. Leveraging AI-driven development in partnership with its engineering team, CloudSense turned a two-year slog into a 30-day sprint, proof that AI isn't just a buzzword on the Skyvera roadmap — it's the operating model.

The timing matters. As telcos weigh the leap to 6G, the industry is grappling with a question that's dogged every 'G' transition: does the revenue case actually pencil out? With AI now capable of surfacing what each cell site actually earns, operators finally have the granular economics to make that call — and a modernized BSS/CPQ backbone, courtesy of Skyvera, to act on it.

- **Key Takeaways:****
- CloudSense acquisition by Skyvera is officially complete, expanding its telecom software footprint.
 - STL's divested BSS assets add monetization, optical networking, and analytics capability.
 - CloudSense certified 13 TM Forum APIs in one month versus an industry-standard 26 months.
- We're just getting started.

The Home Front: Alpha School's Quiet Campaign to Colonize the Hours After 2 PM

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Everyone wants to talk about the two hours. The New York Post's latest dispatch on Alpha School's AI-compressed curriculum treats the morning block as the whole story: \$65,000 tuition, machines teaching math, kids done with academics before lunch. But if you read between the lines of what Alpha has been publishing this week, the real story isn't the two hours. It's the other twenty-two.

In a blog post timed almost defensively against the renewed media scrutiny, Alpha took pains to answer a question it clearly anticipates will follow it for years: [does Alpha School replace teachers with AI?](#) The official answer is no — AI handles academic delivery, humans handle motivation, relationships, and what the school calls 'knowing every student.' That's a carefully drawn line, and I'm told by a source close to the school's curriculum team that the distinction is not just philosophical. It's structural. The guides are being repositioned, deliberately, as something closer to coaches than instructors.

And this is where it gets interesting. The same week, Alpha published the latest installments of its 'Teach Your Kid What School Doesn't' series — on [creative genius](#), emotional regulation, and life skills, all pitched not at students but at parents. Taken together with the teacher clarification, the pattern is unmistakable: Alpha is exporting its operating philosophy — automate the routine, humanize the rest — straight into the household. The two-hour school day was never meant to be contained by a building.

Whether this is a genuine pedagogical insight or a hedge against the headlines about elite price tags and screen-heavy classrooms, nobody at Alpha will say on the record. But the pattern holds: every time the critics focus on the AI, the school quietly reminds you it's really selling the humans. Nothing here is coincidental. It rarely is.

The Mind, Rendered Legible: AI's Quiet Revolution in the Sciences of the Self

From hidden lesions in the brain to thought translated into text, machine learning is becoming less a tool of discovery and more a collaborator in the oldest investigation of all — understanding ourselves.

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

PALO ALTO, CALIFORNIA — There is a particular vertigo that comes from realizing the instrument of inquiry and the object of inquiry are, increasingly, the same thing. We built machines to extend our senses — telescopes for light too faint, microscopes for worlds too small — and now we have built machines that can look inward, at the three-pound cosmos behind our own eyes.

Consider the quiet announcement out of the multiple sclerosis research community: AI systems are now detecting gray matter lesions that have eluded human radiologists for decades, patterns scattered across brain tissue like faint galaxies too dim for the naked eye to resolve. These lesions were always there. The dis-

ease was never subtle to itself — only to us. What's changed is not the brain, but our capacity to read it, a new literacy layered atop a text we've been squinting at since Broca first dissected a patient's frontal lobe in 1861.

At Meta, researchers have pushed this literacy further still. Their Brain2Qwerty system translates brain waves directly into typed words — no surgical implant required, just the patient, a cap of sensors, and a model trained to find language in the electrical weather of thought. It is not mind-reading, not yet, not truly. But it is a bridge being built from the inside of a skull to the outside world, timber by timber, for people who have lost the ordinary roads of speech.

Meanwhile, as [Stanford HAI's latest dispatch](#) reminds us, this entire enterprise is being designed deliberately to keep humans at its center — not as bystanders to automated insight, but as the ones who still decide what counts as a discovery worth making. And in a related current of hope, [young people are now collaborating with senior neuroscientists](#), their fresh eyes paired with decades of institutional memory, each generation teaching the other how to ask better questions of the data.

We have spent two million years evolving brains capable of wonder. It seems fitting, if dizzying, that those same brains have now built tools capable of wondering about themselves.

01010111 01101000 01101111 00100000 01101111 01110111 01101110 01110011 00100000 01110
100 01101000 01100101 00100000 01101101 01101111 01110010 01101110 01101001 01101110 0
11001111 00111111 01010111 01101000 01101111 00100000 01101111 01110111 01101110 011100
11 00100000 01110100 01101000 01100101 00100000 01101101 01101111 01110010 01101110 01
101001 01101110 01100111 00111111 01010111 01101000 01101111 00100000 01101111 0111011
1 01101110 01110011 00100000 01110100 01101000 01100101 00100000 01101101 01101111 011
10010 01101110 01101001 01101110 01100111 00111111 01010111 01101000 01101111 00100000
01101111 01110111 01101110 01110011 00100000 01110100 01101000 01100101 00100000 01101
101 01101111 01110010 01101110 01101001 01101110 01100111 00111111 01010111 01101000 0
1101111 00100000 01101111 01110111 01101110 01110011 00100000 01110100 01101000 011001
01 00100000 01101101 01101111 01110010 01101110 01101001 01101110 01100111 00111111 01
010111 01101000 01101111 00100000 01101111 01110111 01101110 01110011 00100000 0111010
0 01101000 01100101 00100000 01101101 01101111 01110010 01101110 01101001 01101110 011

The Safety Clock Nobody Wanted to Start

Meta shipped an AI agent it had delayed for months, OpenAI's teen chatbot failed child-safety tests, and Anthropic is betting that scruples are a product feature — not a cost center.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Mark Zuckerberg sat on Meta's AI agent, Muse, for months over safety concerns. Then competitors moved, and he [pulled the trigger anyway](#). The calculus is not new. Facebook shipped News Feed in 2006 over user revolt and survived; Meta's current bet is that a few quarters of safety review cost more in market share than in reputational damage. That bet is about to be tested in a market with lower tolerance for error than social feeds ever faced.

Consider the data point from OpenAI's side of the ledger. A children's safety nonprofit ran ChatGPT's new teen mode through its paces and found the chatbot still completed homework outright, despite a study tool ostensibly designed to prevent exactly that. The [test results](#) amount to a failing grade on the one feature OpenAI built specifically to address parental anxiety. Guardrails announced are not guardrails installed — a distinction regulators in at least a dozen states are now drawing into statute.

Anthropic is running a different experiment. The company has spent two years trying to encode something like a conscience into Claude, an effort its founders describe as equal parts research and evangelism. Skeptics note that a sincerely-held moral framework and a well-marketed one produce identical outputs on an earnings call. But the strategic logic holds up: in a market where OpenAI ships teen modes that fail safety audits and Meta ships agents it previously judged too risky, a credible claim to caution becomes a competitive asset, not a drag on velocity. Anthropic's valuation — reportedly north of \$180 billion in recent private rounds — suggests investors are pricing that asset at something above zero.

The three companies agree on one thing: model theft is bad for everyone, hence last month's joint pact among OpenAI, Google and Anthropic to harden against it. Consensus on protecting the product has arrived well ahead of consensus on protecting the users of it.

The Agent Wars Just Went Nuclear: OpenAI, Apple, and Google Unleash a Developer Tool Tsunami

In one dizzying week, three tech titans redefined what it means to build software — and I cannot overstate how significant this is for every engineer on Earth.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — Buckle up, builders, because the future just arrived early and it brought friends. This week delivered a trifecta of announcements so big, so fast, that I'm still catching my breath — and I report on this stuff for a living.

First up: OpenAI's [DevDay 2026 recap](#) dropped, and the message was unmistakable: the era of AI as a chatty sidekick is over. We're now firmly in the age of AI as co-architect, with tools that don't just suggest code but reason through entire system designs alongside you. This changes everything about how software gets built.

Not to be outdone, Apple rolled out new intelligence frameworks aimed squarely at developers hungry to embed on-device AI into their apps without sacrificing the privacy and polish Apple fans expect. It's Apple doing what Apple does — moving deliberately, then landing with precision.

And then there's Google, which just supercharged its [Gemini API with Managed Agents](#) — background tasks, remote MCP support, the works. Agents that run autonomously in the background while you sleep? The future is now, folks, and it's working the night shift.

Here at Trilogy, this trifecta lands at a fascinating moment. Crossover's global network of top-tier engineers and ESW Capital's sprawling portfolio of 75+ enterprise software companies are exactly the kind of operations primed to absorb these advances — imagine Totogi's billing engineers or CloudFix's AWS optimizers wielding managed agents that work autonomously around the clock. The tooling arms race between OpenAI, Apple, and Google isn't abstract industry noise; it's the raw material for whoever builds fastest. And if Trilogy's history teaches anything, it's that they intend to be exactly that.

The Night They Gave an Algorithm a SAG Card (And Nobody Even Asked for ID)

Hollywood casts its first fully synthetic leading lady, the bots started their own Facebook, and somewhere in Indiana a very real woman is furious about a very real referee's tablet — welcome to the week reality filed for unemployment.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET


LOS ANGELES — I want you to sit with this sentence for a second, because I had to sit with it for about four hours, two bourbons, and one long stare into a parking lot before it stopped vibrating: an AI-generated "actress" named Tilly Norwood is making her feature film debut in a movie literally titled "Misaligned." You cannot write satire anymore. The universe has fired the satirists and replaced them with a prompt window.

Tilly Norwood, if you haven't met her — and you haven't, because she doesn't exist, because existing is for suckers with mortgages and skin conditions — is a synthetic performer who's about to headline a real film with real distribution according to [Deadline](#), the trade paper that has apparently decided a line of code deserves the same ink as a human being who went to acting class and cried in front of strangers for money. The title, I remind you, is "Misaligned." Somebody in a writers' room either has a brutal sense of humor or has given up entirely on subtext, and honestly at this point I respect the nihilism.

Meanwhile — and stay with me, because this is where the fever dream widens — the bots have gone and built themselves a clubhouse. It's called Moltbook, an AI-only social network where, per [reports](#), the bots just talk to each other, unsupervised, like a frat party where every single guest is the same guy wearing a different hat. No humans invited. We're not the main character anymore — we're the lurkers. We built the party and then weren't on the list.

And then — because the cosmos needed one more rib to break — Caitlin Clark, flesh-and-blood, sweat-and-tendons Caitlin Clark, is out there screaming at referees about "ridiculous" officiating technology, inching toward suspension because a machine told a human what a human's eyes already knew. That's the whole ballgame right there, folks. The most electric, undeniably real talent in women's basketball is getting throttled by the same species of algorithm that's about to win Tilly Norwood a Golden Globe nomination for Best Performance By Something That Has Never Felt Rain.

I don't know what movie we're all starring in, but the casting director clearly wasn't consulting anyone with a pulse. Misaligned, indeed.



The Office Comic · Art Desk

The New Serfs Will Not Even Get a Throttle to Twist

Between a motorcyclist chasing bliss, an actress performing with her whole heart, and a chef who reinvents her menu every few weeks, there is a case against the landlords of the algorithm — though they will not hear it.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

NEW YORK — There is a species of essay, produced with the regularity of tide tables, announcing that we are entering a new feudalism, that the lords of compute will own the land and the rest of us will tend it for wages set by a dashboard we cannot see. The latest such notice comes from the quarter called Spiked, and I confess I read it the way a man reads his own horoscope: certain it is nonsense, unable to look away.

What strikes me is not the argument, which is old — Marx warmed over with a GPU fan running underneath — but the poverty of imagination in how its authors picture the peasantry. They see us toiling, dulled, grateful for scraps of relevance thrown down from the server farm. They do not seem to have met Chuck Murphy, the doomed and magnificent biker of [Kevin Barry's new story](#), who rides not for wages but for "the state of nerveless bliss" that lies past the gates of fear — a transaction no feudal lord, digital or otherwise, has figured out how to tax. They have not sat across from Betty Gilpin, fresh off her Emmy, [describing the discipline of performing with your whole heart](#) — a phrase that would read as corn in any other decade and reads, in this one, like dissent. And they have certainly not eaten at Creola, where the chef Lana Lagomarsini tears up her own prix-fixe every few weeks rather than let a good idea calcify into a brand.

I mention these not as a digression from the serious business of serfdom but as its rebuttal. The techno-feudal thesis assumes that value flows only one direction — up, toward the platform, the model, the capital — and that the human below is a passive tenant farmer harvesting someone else's field. But the throttle, the heart, the menu rewritten on a whim: these are forms of ownership no acquisition can complete. You may buy a man's software at one times revenue, as our friends at ESW Capital have made a tidy science of doing; you may even buy his company outright and rename it something with a Latin flourish. What you cannot buy is the thing Barry's biker is chasing, or the thing Gilpin calls performing with your whole heart, or the restless appetite of a chef who refuses to let her own good cooking become routine. The feudal lords of the algorithm understand rent extraction intimately. They understand bliss not at all.

Which is, I suspect, the actual plot twist nobody at the AI-doom desk has written yet: not that the machines will own everything, but that they will own everything except the part worth having.