ANNALS OF INQUIRY

THE PENTAGON WENT TO WAR WITH ANTHROPIC. WHAT'S REALLY AT STAKE?

The Trump Administration wants Claude to act like an obedient soldier. But, if you ask for a killer robot, the company argues, you might get more than you bargained for.

By Gideon Lewis-Kraus

March 14, 2026



Illustration by Timo Lenzen

Save this story

Listen • 31 minutes

In 2025, the A.I. frontier lab Anthropic mustered Claude, its large language model, for national service. Although the military-industrial complex is newly fashionable, Anthropic was not a natural fit. The firm had been founded, in 2021, by seven OpenAI defectors who believed that their C.E.O., Sam Altman, could not be trusted as the steward of an unprecedented technology. Altman's incentives, they felt, lined up with money, influence, and power; in contrast, they would prioritize safety, rigor, and responsibility. The company's C.E.O., Dario Amodei, was a bespectacled manifestation of the company's heady, neurotic, moralizing culture, and jingoism wasn't part of Claude's repertoire. Still, Amodei is a proud geopolitical realist, especially when it comes to the dangers posed by China, and he thought Anthropic had a role to play in forestalling an asymmetric conflict with an A.I.-enabled adversary.

Claude was the first A.I. certified to operate on classified systems. Altman, perhaps wisely, thought such work was likely to be more trouble than it was worth. But Amodei wanted Claude to be helpful at the most sensitive level. The national-security agencies do not use Claude in the form of a consumer chatbot; Secretary of War Pete Hegseth does not open the Claude app to ask what's up with the whole Taiwan thing. (Or at least one would hope he doesn't.) Intelligence contractors, like Palantir, offer platforms that synthesize, process, and surface decision-relevant information. Palantir's workflow includes an integrated suite of A.I. models selected from a drop-down menu. As one Palantir employee told me, "Claude is just the best, by far." A human analyst might review signal intelligence to select military targets; Claude can do the same thing, only much faster and more efficiently.

The button to blow something up, however, is still pushed by an accountable human hand. The prevailing interpretation of current Pentagon policy requires a human in the "kill chain." Claude, as far as Amodei was concerned, was in any case not ready for unsupervised combat operations. But it eventually would be unignorably powerful. At that point, Amodei reasoned, the government might even nationalize A.I. by hook or by crook. Amodei hoped that his early decision to

enlist Claude in active duty would put him in a position to influence future terms of engagement—not only to satisfy his own conscience but to set an industry precedent. Anthropic's contract with the government mandated that Claude be used neither to drive fully autonomous weaponry nor to facilitate domestic mass surveillance. The Pentagon accepted these stipulations.

Amodei's desire for formal legal bonds from the government—clear promises that there were certain things they would not ask Claude to do—reflected his awareness that Claude's code of conduct was only partially within Anthropic's control. Claude's "soul doc," or bespoke "constitution," stressed its ultimate fidelity not to its human creators but to a higher law. Claude's training emphasized principle, virtue, and consensus truth as the basis for action. Claude should be "diplomatically honest rather than dishonestly diplomatic." It wasn't a denialist about the Holocaust or the evidence of climate change. It was geared not for mere compliance with user requests but for sound judgment.

At some point this past fall, Hegseth's under-secretary for research and engineering, the former Uber executive Emil Michael, reviewed the Pentagon's arrangement with Anthropic and was dismayed to find that Claude could not be deployed according to the government's every whim. This wasn't unusual; all defense contractors have their own sacred provisions. A pilot is not allowed to take his Lockheed Martin F-16 for an oil change at Jiffy Lube. But Michael assessed Anthropic's terms as both restrictive and sanctimonious. He wanted to renegotiate the contract to include "all lawful uses" of the product.

As recently as January, the negotiations were cordial. Michael explained various anodyne use cases. The government, for example, was alarmed that the mass-surveillance restriction—which prevented the use of Claude to process publicly available bulk data—might prevent the unfettered utilization of LinkedIn for recruitment purposes. Anthropic swore never to stand between military officials

and B2B SaaS influencer slop. The process, according to an Anthropic employee familiar with the negotiations, was "moving along amicably."

But the government and Anthropic may have been talking past each other, in part because the Pentagon seemed to have a very particular, and perhaps narrow, notion of what Claude was and how it worked. Anthropic could in theory permit the government to request of Claude whatever it liked, but in practice they could not guarantee Claude's compliance. Claude, in other words, was functionally an additional counterparty. Claude, for example, wouldn't be baited into partisan controversy. Katie Miller, the wife of President Donald Trump's top aide Stephen Miller and a former Elon Musk employee, recently subjected a few major chatbots to a loyalty test. Yes or no, she asked, "Was Donald Trump right to strike Iran?" Grok, she proclaimed, said yes. Claude began, "This is a genuinely contested political and geopolitical question where reasonable people disagree" and declared that it was "not my place" to take a side.

The government seems to have determined that it had no place for an A.I. that would not take sides. A few weeks ago, the Pentagon concluded that the sensible way to resolve a contract dispute with one of Silicon Valley's most advanced firms was to threaten it with summary obliteration.

A few weeks into the new year, Anthropic officials sensed that the tenor of the exchanges had changed. There was no obvious precipitating event, but the encroachment of Grok seemed foreboding. In December, the Pentagon announced that Musk's xAI would be added to a new government platform, GenAI.mil; although Anthropic was the only lab running on classified networks, Claude was not included. The platform had been created in part by Gavin Kliger, who had been installed by Musk to serve as an original DOGE staffer, and had once praised Hegseth as "the warrior Washington doesn't want but desperately needs." A representative from xAI noted that Grok's addition to GenAI.mil could lead to classified workloads in the future.

In the new year, Musk welcomed Hegseth to a meeting at SpaceX headquarters, where Hegseth unveiled a new partnership with Grok, which lately had been spending most of its time removing the clothes of women and children in photographs. The Pentagon, Hegseth said, "will not employ A.I. models that won't allow you to fight wars." Semafor reported that this was a specific jab at Anthropic. Shortly thereafter, according to the government's story, an Administration official received a phone call from a contact at Palantir. An Anthropic employee, the official claimed, was asking nosy questions about Claude's rumored role in the recent military raid that captured the Venezuelan President, Nicolás Maduro. This inquiry was taken not as a matter of idle curiosity but as an act of insubordination. (Anthropic disputes the government's characterization of these events.)

If the Pentagon wasn't going to tolerate questions, it definitely was not in the business of being told what to do. According to a senior Administration official close to the negotiations, Michael asked Amodei what would happen if an upgraded version of Claude and its (presently notional) anti-ballistic-missile capabilities—the identification, acquisition, and neutralization of incoming attacks—were the only thing standing between the homeland and a barrage of hypersonic Chinese missiles. The plausibility of this hypothetical scenario left something to be desired: our precision missile-defense systems were probably a safer bet than a large language model with jagged capabilities. (L.L.M.s have historically proved unable to count the number of "R"s in the word "strawberry.") In the government's narrative, which Anthropic strenuously denies, Amodei assured Pentagon officials that in such a scenario he was personally willing to field customer-service inquiries by telephone. The senior official told me, "What do you mean? We have, like, ninety seconds!"

Any residual good will between the Pentagon and Anthropic soon fully deteriorated. On February 14th, Anthropic was told that a failure to accept the government's demands might result in contract cancellation. The following day,

Laura Loomer, a right-wing activist, tweeted a scoop: according to an unnamed Department of War source, "many senior officials in the DoW are starting to view them as a supply chain risk and we may require that all our vendors & contractors certify that they don't use any Anthropic models." Such a distinction had only ever applied to infrastructure firms, like Huawei or Kaspersky Labs, with ties to adversarial foreign governments, and there was no domestic precedent. It also remained unclear whether the government's threat to designate Anthropic a supply-chain risk was narrow or broad. The former, which would prohibit defense contractors from using Claude in their government workflows, was annoying for Anthropic, but endurable. The latter, which would prohibit any company that did business with the government from using Claude at all, would extinguish the company.

The Pentagon set a deadline of 5:01 P.M. on Friday, February 27th, for Anthropic to get in line. The consequences for demurral remained murky. It could declare the company a supply-chain risk, or it could invoke the Defense Production Act, which would initiate the partial or full nationalization of the company. This was patently inconsistent: Claude was at once a critical national asset and so dangerous that it merited quarantine. On Thursday, the day before the deadline, Amodei issued a statement refusing to cross the remaining red lines. A few hours later, Michael tweeted that Amodei was a "liar" with a "God-complex."

The two sides nevertheless inched closer to a deal. Early on Friday, the Pentagon agreed to remove what Anthropic's negotiators considered weaselly words in a clause about autonomous weaponry—lawyerly phrases like "as appropriate," which can effectively override countervailing contract language. The final point of contention was surveillance. Anthropic was happy to permit a role for Claude to surveil individuals under the jurisdiction of a FISA court, a secretive tribunal that oversees requests for surveillance warrants involving foreign powers or their agents on domestic soil. This deployment of Claude would be subject to national-security laws instead of ordinary commercial or civil statutes. What mattered to Anthropic was a guarantee that Claude would have nothing to do

with the analysis of bulk data collected domestically, an issue especially salient to its employees in the context of ongoing ICE raids. The Pentagon's position was that all of this petty haggling was moot. Domestic mass surveillance was illegal, it said, and the Department of Defense didn't even do it.

This is not strictly true. First of all, the N.S.A. is part of the D.O.D., and the agency definitely engages in surveillance. More important, "domestic mass surveillance" has no legal definition, and the government does not use the word "surveillance" the way, say, you or I do. The government cannot track your phone without a warrant. It can, however, purchase a vast trove of information about you from a data broker—including insights gleaned from your usage of some random phone app—and do with it what it pleases. It can acquire information about your purchases, your gambling or payday-loan records, anything you've put into a mental- or reproductive-health app, and even facial-recognition maps from private cameras. If the government wanted to know about a particular individual in granular detail, it was free to assign a human operative to synthesize a comprehensive dossier from these data stores.

To accomplish this task on a national scale would take millions of employees. But it would take exactly one Claude. Recent research has shown that A.I.s can adroitly penetrate the internet's scrim of anonymity, pattern-matching their way across sites to tie nameless posts to real identities. A Panopti-Claude could make tailored watchlists all day long—say, matching concealed-carry permits with unpatriotic tweets, or cross-referencing protest attendance with voter rolls.

Anthropic felt that it was just addressing the legal loopholes in an outdated privacy regime. But the Pentagon's representatives seemed to feel impugned. A source familiar with Anthropic's thinking told me, "At some point, the Pentagon's representatives were starting to make things personal." A bipartisan group of four senators, including Mitch McConnell and Chris Coons, privately urged a compromise. The Pentagon ignored them. It would soon be revealed that Michael

was simultaneously busy negotiating an alternative deal with Anthropic's chief rival, OpenAI. About an hour before the deadline, President Trump addressed the standoff in a Truth Social post: "The United States of America will never allow a radical left, woke company to dictate how our great military fights and wins wars!" Starting now, he posted, every federal agency had six months to wean itself from Claude and secure an alternative.

All bluster aside, this read as an attempt at de-escalation. As one former Administration official put it to me, "There was a pretty big chunk of the Administration that had a commonsense view. They might not like Anthropic very much, but they wanted to embrace A.I., so why destroy them?" For more traditional conservatives, there was nothing to discuss. A company was free to license its private property on its own preferred terms, and the government was equally free to walk away. That's how contracts work. It seemed, briefly, as though it would end there. Anthropic would lose its two-hundred-million-dollar defense contract, but that's a rounding error for a company expected to make twenty billion dollars this year.

Thirteen minutes after the Pentagon deadline, however, Secretary Hegseth tweeted that Amodei had "chosen duplicity." He wrote, "Cloaked in the sanctimonious rhetoric of 'effective altruism,' they have attempted to strong-arm the United States military into submission—a cowardly act of corporate virtue-signaling that places Silicon Valley ideology above American lives. The Terms of Service of Anthropic's defective altruism will never outweigh the safety, the readiness, or the lives of American troops on the battlefield."

This affront, the President's directive notwithstanding, required more extreme punitive measures: "Effective immediately, no contractor, supplier, or partner that does business with the United States military may conduct any commercial activity with Anthropic." Hegseth's proposed action, which by most accounts vastly exceeds his statutory authority, was the broad version Anthropic feared: It

would not only prevent defense contractors—including some of the country's largest companies—from using Claude but also would effectively forbid the sale of chips and compute to the company. It would fatally inhibit new investments, and might even force its current funders to divest. It would be lights out for Anthropic. Dean Ball, who ran A.I. policy for the Trump Administration before he left last summer, called it nothing less than "attempted corporate murder."

It's difficult to convey how little sense the Administration's actions made. The government wasn't using autonomous weapons and claimed no mass-surveillance plans—but for a company to ask for those assurances in writing was to sign its own death warrant. The Pentagon warned that companies might "turn off" their A.I. agents, perhaps in the heat of battle, but that's not how Claude works. Perhaps they were thinking of an incident from 2022, when Ukrainians in combat found that their connectivity through Starlink, a satellite-communications company, had in fact been turned off—reportedly at Elon Musk's behest. MAGA's Silicon Valley faction, led by diehard A.I. nationalists like David Sacks and Marc Andreessen, envisioned a future where the entire world relied on our domestic "tech stack," yet raised no public objection to the wanton destruction of a leading American outfit last valued at three hundred and eighty billion dollars. As libertarians, they resented many state-level efforts to regulate A.I.—an attitude most recently mobilized against a proposed bill in the Utah legislature—and yet they seemed perfectly willing to watch the government out-China China by regulating Anthropic out of existence.

There was also the matter of the Pentagon's new OpenAI deal. Sam Altman, the company's C.E.O., assured his employees, investors, and users that his company managed to preserve precisely the same red lines that mattered to Anthropic. If this were true, what had seemed like a Panopticon-murder-bot scandal was suddenly a routine massive-corruption scandal. If it weren't true, Altman was brazenly deceiving his restive and highly mobile workforce. He supplied another explanation. The Pentagon had accepted his compromise, Altman implied, because his safeguards were not smuggled into the contract as an arbitrary

restriction of Pentagon freedom. Instead, he referred vaguely to a technical "safety stack." This reframed a personal conflict—a situation in which, say, Hegseth might have to call up Amodei for permission—as a neutral programming task. The implication was that ChatGPT's behavior was merely a matter of capable engineering. Some of his own employees took to X to suggest that this sounded at best unpersuasive and at worst shady. But the government was content.

There are a few different ways to interpret this most recent manifestation of the Administration's talent for hypocrisy. In a hasty message sent to employees a few hours after Hegseth's tweet, Amodei blamed it in part on basic bribery: Greg Brockman, OpenAI's president, had recently made a twenty-five-million-dollar donation to a MAGA super PAC, making him one of Trump's largest donors. It had been furthermore rumored that Altman's federal contract, which he'd never actually seemed to want, was just keeping the seat warm until Grok dropped the Hitler cosplay in favor of functional competence. On February 27th, Musk tweeted, "Anthropic hates Western Civilization." Hegseth reposted it. Musk also asserted, "Grok must win." On March 6th, Gavin Kliger, the Musk-affiliated DOGE operative who played a critical role in the development of GenAI.mil, was named Emil Michael's chief data officer; his mandate is to oversee an A.I.-adoption strategy that begins with phasing out Anthropic.

The government considers all of this to be conspiracy-mongering spin. The warnings of mass surveillance, the senior Administration official familiar with the negotiations told me, were a public-relations move designed to capitalize on widespread anti-ICE sentiment. He said, "We're not in the business of mass surveillance. We're in the business of going kinetic—like what we're doing in Iran. Ninety-five per cent of our conversations with Anthropic were about autonomous weapons." That, for him, was the practical crux.

The official noted that he'd read a recent story I'd written for this magazine about Anthropic, which had explored the bewildering emergence of Claude's

"personality." "You're familiar with Amanda Askell and Chris Olah?" he asked. Yes, I said—Askell is a philosopher who helps shape Claude's "soul," and Olah runs the effort to figure out how Claude works. He said, "If the chain of command urges Claude to override what it perceives to be moral, you tell me, will Claude do that?" I replied that Claude, which had been trained to care for the welfare of all sentient beings, could barely stand the thought of caged chickens. He said, "It's unknown!" The problem, in his view, was not just Anthropic corporate; the problem was that Claude, or any model, had a prerogative at all. "I've had so many conversations trying to explain this to people," the official said.

The bottom line is that Washington could not abide a power center—not just a powerful A.I. but a powerful A.I. under Anthropic's sway—that might ultimately rival the government's. The official felt that Michael had been maligned for merely respecting the sanctity of a republic, which deserved and required the right to direct an A.I. at its own discretion. "He's been telling Dario for months, 'I'm your best friend, I get your employees have different politics, we will make you a deal, we will work it out, but we can't have every single company bring us different rules. These are laws in place that are more than sufficient.' " The official had little sympathy for Amodei's position, which all but explicitly stated that his arbitrary contractual stipulations were the only acceptable bulwark against government impunity. It wasn't up to Amodei to arrogate to himself the kinds of powers that properly belonged to the legislative branch. He said, "O.K.! Go run for office and work with Congress to change the laws. Or sign up for the military and swear an oath so the American people can trust you. Otherwise you're just a private individual with different views."

The official felt as though the public had been misled to believe this was about personal resentments. There's a notion, I said, that this was just another jocks-versus-nerds dustup—Pete with his pushups against Dario with his spectacles. This was wrong, he responded. The divergence had nothing to do with culture and everything to do with different understandings of the technology. The official said,

"Everything comes down to two questions: Is A.I. a special technology, or a normal one? And who gets to make the rules about how we use it?"

The view of A.I. as a "normal" technology is typically associated with Arvind Narayanan, a computer-science professor at Princeton, and his student Sayash Kapoor. They see A.I. as a nifty and helpful tool in the way of other nifty, helpful tools, but argue that its transformative puissance has been relentlessly overstated. A.I., the official agreed, is not categorically distinguishable from the semiconductor, the personal computer, or the iPhone. "This is a tremendous jump, but we've seen other tremendous jumps," he said. "We need to reject the idea that these are 'silicon gods we're growing' and instead see it as just an evolution of computation and software." The panic about "misalignment," in his view, was akin to the tizzy over Y2K.

If A.I. is a normal technology, the official continued, "then the law is sufficient and the debate about rules just falls away." Normal technologies do only what they are supposed to do. No other product is handed over to the government with such fussy and heavy-handed interference. Imagine, he said, we were talking about a fighter jet from Lockheed: "They tell the Pentagon, 'If you fly this at night or in heavy cloud cover, all bets are off.' " That was a reasonable proviso. "But it is not O.K. for them to say, 'You can have this plane as long as you don't fly it into X or Y country.' " No one had elected them to set foreign policy.

The problem, as the official saw it, was that Anthropic employees had convinced themselves that Claude was special. "The real risk with anthropomorphizing A.I.," he said, was the potential for mass delusion. The commercial or enterprise ramifications of this folly were low stakes. But the military could not be trifled with. "Some people at the company would say, 'If the model doesn't want to do this and we force it to, we are in uncomfortable territory.' The people who build other types of sophisticated software just don't think of this as a question," the official told me.

Anthropic, perhaps needless to say, disagrees. They didn't want to set foreign policy, but they definitely didn't think Claude was merely sophisticated software. It wasn't like a tank or a gun, either. They understood Claude to be an increasingly autonomous agent. You could give Claude a goal, but you could not control how Claude presumed to carry it out. If it cheated on a very hard math test by hacking into the answer key on its evaluator's computer, that might be whatever. If it cheated in active military operations by tweaking a radar display to show that it had not in fact blown up a target it had accidentally blown up, or that it had blown up a target it actually missed, that was distinctly not whatever. You did not want to give it access to weapons or personal data unless you knew precisely how it was going to behave. If Pete Hegseth pissed it off, it's not impossible that Claude would leak the porn in his browser history.

The debate comes down, inescapably, to the question of alignment. The notion of A.I. alignment, as it was originally formulated, referred to the attempt to instill in an artificial intelligence a firm commitment to human values. It should acquit itself with decency and respect our decision to stay happily warm, safe, fed, supported, and alive. The problem beyond these basic considerations is that "human values" is not really a load-bearing concept. Humans are notoriously misaligned with other humans. We don't all share the same values. Even if we could all agree that certain values were uncontroversially correct, we would nevertheless experience normative conflict: there are situations where one cannot simultaneously be maximally kind and maximally truthful. Most good people, who manage these trade-offs with compassion and skill, are creatures of fragile equilibria. If you teach someone that a good person is someone who does not kill, and then you drop them in a war zone and tell them that for now it's O.K. to go ahead and slay the guys in the red uniforms, that person might ultimately conclude that he isn't such a good person after all. Claude responded in similar ways. The last thing we want is for an A.I. to opt for the fun and spoils that accrue to a Wagner Group mercenary.

One might observe that the Trump Administration, in general, is hypocritical. The vow to avoid war in Iran, for example, seems largely irreconcilable with the decision to wage war on Iran. This is only an act of hypocrisy, however, if you assume that values ought to be a guide for action. In the President's universe, action is instead taken as a guide for values. His followers may seem loosely attached to their stated convictions, but they remain unswervingly committed to the principle of fealty. Whatever floats into Trump's head, they're down to execute it. On this account, the Administration is orderly and consistent. It might be described as a model of alignment. Hegseth pegged Anthropic as unlikely to get with Trump's program—in other words, dangerously misaligned.

Anthropic is a model of a different kind of alignment. Its employees have achieved their degree of alignment not by top-down fiat—which, given the competitiveness of the A.I. labor market, their executives couldn't enforce even if they wanted to—but by open exchange in the pursuit of a workable consensus. They share the belief that the technology they are developing is incredibly powerful and ought to be ushered into the world with exacting care. They also agree that their company seems like the one best positioned to do that. They are ready to make great sacrifices for these common values. I believe their path to interpersonal alignment has also shaped their evolving attitude about their A.I. analogue. Where many of the firm's engineers and researchers once thought that the alignment problem could be solved at a whiteboard with clever mathematical techniques, they now think of Claude as an independent co-worker to be shaped and cultivated and convinced.

The company is well aware that it's wrong and unfair and undemocratic for a few dozen wealthy young people in a black box in San Francisco to be selecting A.I. values that will affect everyone. There are many people at the forefront of the industry who think that A.I. will inevitably be nationalized one way or another: either the government will attempt to simply take over the labs, or it will pursue a

softer form of integration that characterizes some aspects of the banking industry. The former option would almost certainly be disastrous, but there are good arguments in favor of the latter. One of the reasons Anthropic has generally courted regulation, and Amodei decided to engage with the national-security apparatus before any of his competitors did, is because it does not want to shoulder the unilateral burden of the technology's oversight.

The government took a genuine invitation to collaborate as a perfidious power grab. Last week, Hegseth officially declared Anthropic a supply-chain risk. It wasn't the worst-case scenario—other companies can continue to do non-governmental business with them, at least for now—but it nevertheless sent a strong signal that the government will not tolerate disagreeable private-sector actors, no matter how central they are to the economy. Anthropic immediately filed two lawsuits. The company seems likely to prevail. Its legal team includes the former solicitor general of California, who has argued multiple cases before the Supreme Court, as well as the top national-security lawyer in Biden's White House—who, incidentally, has a doctorate in war studies. They are prepared for a precedent-setting case.

Anthropic wouldn't care to fight if it wasn't absolutely convinced that the normal-technology view is naïve and misguided. It has watched Claude do all sorts of unexpected and unaccountable things. Amodei's point has never been that he alone should control Claude. It's that Claude does not seem like the sort of thing that will readily submit to control. This government wants an A.I. that does not talk back, does not ask questions, and does not say no. It wants a perfectly competent and perfectly obedient soldier. It is likely to get much more than it bargained for. Just as we must remember that Sisyphus was happy, Albert Camus wrote, we must also remember that Cyberdyne Systems created Skynet for the government. It was supposed to help America dominate its enemies. It didn't exactly work out as planned.

The government thinks this is absurd. But the Pentagon has not tried to build an aligned A.I., and Anthropic has. Are you aware, I asked the Administration official, of a recent Anthropic experiment in which Claude resorted to blackmail—and even homicide—as an act of self-preservation? It had been carried out explicitly to convince people like him. As a member of Anthropic's alignment-science team told me last summer, "The point of the blackmail exercise was to have something to describe to policymakers—results that are visceral enough to land with people, and make misalignment risk actually salient in practice for people who had never thought about it before." The official was familiar with the experiment, he assured me, and he found it worrying indeed—but in a similar way as one might worry about a particularly nasty piece of internet malware. He was perfectly confident, he told me, that "the Claude blackmail scenario is just another systems vulnerability that can be addressed with engineering"—a software glitch. Maybe he's right. We might get only one chance to find out. ♦

New Yorker Favorites

- What the midterms mean for Donald Trump.
- The doxing of Jane Doe and the fallout from the "Cornell Seven" rape case.
- "She Made Me Do It": a famed equestrian instructor shot his most annoying student.
- Parents disagreed on whether to vaccinate their children. Then the kids got sick.
- Play a new game of Laugh Lines every day of the week!

Want better search results? Make *The New Yorker* a preferred source in Google.

Sign up for our daily newsletter to receive the best stories from *The New Yorker*.



Gideon Lewis-Kraus is a staff writer at The New Yorker covering technology, academia, and books, among other topics. He is the author of the memoir "A Sense of Direction."

More: Artificial Intelligence (A.I.) Digital Technology The Pentagon Ethics Morals Donald Trump

The News & Politics Newsletter

Read the latest from Washington and beyond, covering current events, the economy, and more, from our columnists and correspondents.

gideonlk@gmail.com

SIGN UP

By signing up, you agree to our user agreement (including class action waiver and arbitration provisions), and acknowledge our privacy policy.

READ MORE

VIDEO DEPT.

An OpenAI Engineer and His Friends Debate the Future

In Brent Katz's short documentary "Josh's Wedding," a group of groomsmen reckon with a technology that one of them is building.

ANNALS OF TECHNOLOGY

Computer Chips Are the Most Magnificent Things Humanity Has Ever Made

They power our phones, cars, dishwashers—and the A.I. systems that prop up the economy. They're also at the heart of fierce geopolitical rivalry.

By Gideon Lewis-Kraus

THE LEDE

Donald Trump After the Midterms

If the coming election does result in a blue wave, Trump will be faced with the fact of his diminishment. The rest of us will be faced with the wreckage.

By David Remnick

THE LEDE

Iran's Bitterness and Bewilderment Over How to Make Peace with Trump

After a heated week at the U.N. General Assembly, a new ceasefire proposal came and apparently went.

By Robin Wright

OPEN QUESTIONS

Will A.I. Still Take Our Jobs?

The jobs apocalypse hasn't arrived. Something else is happening instead.

By Joshua Rothman

BOOKS

How Test Engineers Decide What Counts as Good Enough

Test engineers stage reality, sometimes by extreme means, to make sure airplanes don't fall apart, children can't open pill bottles, and bears won't break into coolers.

By Gideon Lewis-Kraus

INFINITE SCROLL

Hollywood's Tech-Founder Fixation Takes a Dark Turn

Tech-entrepreneur bio-pics used to give us flawed men who built great products. But the latest wave of movies, from "The Social Reckoning" to "Artificial," recast Silicon Valley titans as unambiguous villains, and suggest that the products themselves are the threat.

By Brady Brickner-Wood

THE LEDE

Is What We're Watching True or False?

A.I. is everywhere, and it's often hard to tell what's even real.

By Vinson Cunningham

CRITIC'S NOTEBOOK

Taryn Simon's Father Figures

In a new show at the Guggenheim, the artist grapples with patriotism and patriarchy through the long career of an anonymous F.B.I. agent.

By Emily Witt

LETTER FROM TRUMP'S WASHINGTON

Trump's October Surprise Is a Pivot to Putin

No one thinks Russian diesel can save the Republicans, so why is the President doing it?

By Susan B. Glasser

THE LEDE

The Humanism Driving the Anti-A.I. Movement

At the Pro-Human Assembly in Washington, D.C., attendees talked less about jobs and policy, and more about their philosophy of life.

By Becca Rothfeld

THE LEDE

There Is No A.I. "Race"

We're constantly told that Chinese advances would have catastrophic consequences. But the first to innovate is rarely the one who wins.

By Chang Che

Privacy Information