


THE POWER ISSUE

# So Few People, So Much Influence The future of AI will be shaped by a few dozen engineers, founders, and thinkers. It's time you got to know them.

By The Editors

5:00 A.M.

*Technologist Portraits* (2026). Digital image: AI image generation, VHS tape, digital import. Art: Simon Denny

AI has minted a new power structure. A handful of private companies command the might and wealth of small nation-states and are mostly left to self-govern. A few dozen employees are on the brink of becoming centimillionaires and considering which causes and presidential candidates they will funnel their vast resources toward. The industry is perpetually one breakthrough away from a step change in capability or public adoption, the question of who gets there first dependent on the specific expertise and visions of a handful of engineers.

IN THIS ISSUE

The Future of AI Will Be Shaped by a Few Dozen People

THE FUTURE OF AI WILL BE SHAPED BY A FEW DOZEN PEOPLE
IT'S PROBABLY TIME TO GET TO KNOW THEM

OpenAI's Hollywood Whisperer

The Democrat Shaping the Party's AI Debate

SEE ALL »

One way to tell the story of how AI is changing our world is through some newly familiar names. Sam Altman and Dario Amodei are locked in a race to build AGI and persuade President Donald Trump to regulate the technology to their liking. Nvidia CEO Jensen Huang and venture capitalist David Sacks are advocating for the president to take a light touch, concerned about how regulation might jeopardize America's dominance. As these figures have scaled their public influence, their surrogates and spokespeople (OpenAI president Greg Brockman, top OpenAI lobbyist Chris Lehane, Anthropic co-founder Jack Clark) have populated news stories, too.

But an astonishing amount of power rests in the hands of a few dozen other people, none particularly well known, some virtually anonymous. So *New York* Magazine has dedicated this year's annual "Power" issue to surfacing 26 people whose influence on the AI landscape far outstrips their fame. To assemble the list, we partnered with *Transformer*, a small but impressive publication with a command of the politics within transformative AI. The *Transformer* team drew on its years of reporting, as well as new interviews with more than 50 high-ranking industry leaders, to help articulate what makes these people powerful. We then commissioned multi-media artist Simon Denny to create images of each.[1]

Despite the hundreds of billions in capex and upcoming IPOs, many parts of the field are immature, which means those who decide how AI is built, adopted, and controlled can quickly become among the most influential people in the world. A few on this list, for example, are tasked with shaping how the public, or a company, thinks about the technology — a high-leverage role for a field where everyone's views are still being formed. (This is, in part, why AI has such an "essay"-driven culture — a blog post really can cause a seismic change.) Some will invent entirely new product categories, reimagining how we interact with AI, while others can spin up new industries or advocacy groups with a single check. Who is ahead and who is behind can change almost overnight (just look at Microsoft), which helps explain why a legacy tech CEO could have less sway over what Altman thinks than a 26-year-old with a podcast and why an engineer who runs a new critical function may appear on this list instead of a colleague who sits several rungs higher in their org chart. (Relatedly, the list's average age is 36. It is also overwhelmingly male.)

The AI scene is still small. The 26 individuals on our "Power List" are remarkably intertwined, some even sharing offices, homes, or funding. (*Transformer* receives funding from an organization run by Alexander Berger, who appears on this list — and he in turn has ties to many of the other people cited.) The widespread societal impact of AI led us to include names from both private and public sectors, so you'll find founders and engineers as well as lawmakers and regulators. Most important, they still have everything to play for. The new power brokers will decide how it all shakes out, refashioning the world for the rest of us in the process.

Free IQ Test
Free Intelligence Quotient Test
Free Intelligence Test
Free Cognitive Skills Test
BlossomUp

## Logan Graham Stands Between Us and the Next Rogue AI Model

Head of the Frontier Red Team at Anthropic | Age: 32




As the leader of Anthropic's frontier "red teaming" efforts — a Cold War–inspired practice of playing the villain to find holes in your own defenses — Logan Graham has the job of pushing AI models to do terrifying things. This may mean persuading models to run cyberattacks or explain how to build a bioweapon. If the model then behaves dangerously, Graham's Frontier Red Team is in charge of warning both company leadership and the public. So far, that's happened at least twice, and each time Anthropic has pointed to Graham's work as evidence. On April 7, the Frontier Red Team announced that Claude Mythos Preview had surpassed the usual cybersecurity benchmarks — tests designed to measure how good models are at hacking — and had to be tested on real-world security tasks instead. It found that the model could locate and exploit security bugs in every major operating system and web browser, potentially empowering nonexperts to run major cyberattacks. The company withheld Mythos from the public but granted limited access to a few dozen big-tech companies, including Amazon Web Services, JPMorgan Chase, Microsoft, and Google. Graham became the public face of this decision, which critics called a marketing strategy or elitist with Anthropic unilaterally determining whose cyberinfrastructure was worth defending first.

Sign up for Intelligencer
*A daily roundup where politics, business, and technology intersect.*
SIGN UP

Graham joined Anthropic part time in 2022. That November, Graham "started red teaming LLMs for biorisks/weapons risks in my bedroom," according to his subsequent tweet. Back then, frontier models couldn't reliably do math or count the number of *r*'s in strawberry, but Graham still persuaded his bosses to dedicate resources to figuring out what AI could do at its worst. The Frontier Red Team he built would be one of the first in-house research groups of its kind.

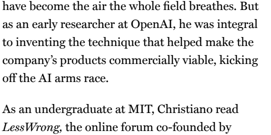
"Some people's intuition from growing up in a peaceful world is that somewhere there's a room full of adults who know how to fix it," Graham told *Time* earlier this year. "There are no groups of adults. There is no room in the first place … You are responsible." Despite the nature of his work, a colleague described Graham as "very chipper." Once, when asked what quality he most sought in red teamers, he responded "craftiness."

Free IQ Test
Free Intelligence Quotient Test
Free Intelligence Test
Free Cognitive Skills Test
BlossomUp

— *Celia Ford*

## Paul Christiano Is One of the Few People Who Can Slow Down OpenAI From the Inside

Founder of Alignment Research Center | Age: 37




Over the past 15 years, Paul Christiano[2] has helped pioneer the AI-safety movement, converting the mostly theoretical concerns of figures like Eliezer Yudkowsky into hands-on research and seeding the ideas and frameworks — documented in papers, blog posts, and comment threads — that have become the air the whole field breathes. But as an early researcher at OpenAI, he was integral to inventing the technique that helped make the company's products commercially viable, kicking off the AI arms race.

As an undergraduate at MIT, Christiano read *LessWrong*, the online forum co-founded by Yudkowsky that helped birth today's AI-safety world. Christiano was studying math and quantum computing but gradually turned his attention to alignment, the challenge of creating AI systems that do what human users want them to do. A blog post of his, titled "What Failure Looks Like," laid out what have become the two dominant visions of AI's worst-case scenarios: one, which Christiano calls "going out with a whimper," where humans gradually render themselves irrelevant; another, "going out with a bang," where AI-controlled critical infrastructure breaks down, leading to societal collapse. When doomers talk about "doom," they're often picturing Christiano's nightmare.

In 2017, Christiano joined OpenAI, where he led the development of reinforcement learning from human feedback, or RLHF, a method that turns raw language models into helpful, instruction-following assistants. RLHF is one of the biggest differences between Sydney — an unhinged Microsoft chatbot rumored to be an early version of OpenAI's GPT-4 — and today's polite, helpful chatbots. It's now a crucial part of all frontier AI-training pipelines and may have done more to unlock the technology's mass-commercialization potential than almost anything else.

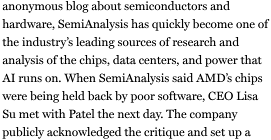
Around the time Dario Amodei, Christiano's co-author of the paper that helped introduce RLHF, exited OpenAI, Christiano left, too, founding a research nonprofit called the Alignment Research Center. In the spring of 2024, the Biden administration brought him on as the AI Safety Institute's head of AI safety, where he would be in charge of testing frontier AI models. Government staffers threatened to quit over his ties to effective altruism before his first day on the job. But through widespread DOGE-era cuts and institutional overhauls, Christiano somehow stayed put — he still works part time as a special government employee. In September, the day after Jacob Coxon's viral Anthropic resignation tweet, Christiano announced that he was joining OpenAI's nonprofit board to serve on its Safety and Security Committee. The five-person committee, at least on paper, has the power to delay OpenAI's model releases, giving Christiano the chance to wrangle his former employer into controlling its AI risks before it's too late.

Free IQ Test
Free Intelligence Quotient Test
Free Intelligence Test
Free Cognitive Skills Test
BlossomUp

— *C.F.*

## Dylan Patel's Analyses Can Spook Even Nvidia's CEO

Founder and CEO of SemiAnalysis | Age: 30




Jensen Huang runs the world's most valuable company. Even he cares what Dylan Patel,[3] founder of the research-and-consulting firm SemiAnalysis, thinks: At the company's conference this year, Huang took a moment to say that Nvidia's chips are good because SemiAnalysis says they are. Since Patel launched it six years ago as an anonymous blog about semiconductors and hardware, SemiAnalysis has quickly become one of the industry's leading sources of research and analysis of the chips, data centers, and power that AI runs on. When SemiAnalysis said AMD's chips were being held back by poor software, CEO Lisa Su met with Patel the next day. The company publicly acknowledged the critique and set up a developer-relations team seemingly as a result.

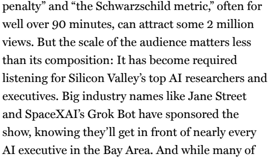
Patel has been obsessed with the technical nuances of silicon since he was a teen lurking on chip sub-Reddits and has a talent for networking that feeds him a steady stream of gossip. In 2023, when everyone was reporting that TSMC's limited chip-packaging capacity was curtailing the supply of Nvidia chips, SemiAnalysis mapped the upstream suppliers, showing how long the bottleneck would last and who would profit. "There were a bunch of people doing technical research" when SemiAnalysis launched, according to Sequoia's Shaun Maguire, "but no one that could combine the technical breakthroughs with the business impact." (Sequoia has invested in SemiAnalysis.) When the AI boom started, Patel and his team "were the first to really understand the whole systems math," says Maguire, building complex models that linked the cost of chips to AI companies' margins — and selling them to hedge funds, AI companies, and semiconductor giants.

SemiAnalysis reportedly expects $100 million in revenue this year. It recently acquired Citrini Research and raised $500 million for a fund investing in AI infrastructure and chip companies. Critics note that could create serious conflicts of interest — something Patel has brushed off.

— *Shakeel Hashim*

## Every Silicon Valley CEO Is Listening to Dwarkesh Patel

Host of 'Dwarkesh Podcast' | Age: 26




Top episodes of *Dwarkesh Podcast,* where guests talk unabashedly about "the parallelization penalty" and "the Schwarzschild metric," often for well over 90 minutes, can attract some 2 million views. But the scale of the audience matters less than its composition: It has become required listening for Silicon Valley's top AI researchers and executives. Big industry names like Jane Street and SpaceXAI's Grok Bot have sponsored the show, knowing they'll get in front of nearly every AI executive in the Bay Area. And while many of the guests — which include Jensen Huang, Dario Amodei, and Elon Musk — might be unavailable to most reporters, they'll agree to be interviewed by Patel.[4] One top tech PR representative told us that clients would rather be on his podcast than *60 Minutes.*

Patel's show sets the agenda for what the industry thinks about. Take his November 2025 interview with Ilya Sutskever, the tight-lipped co-founder of Safe Superintelligence Inc. Sutskever — who, while working at OpenAI, correctly foresaw that training bigger AI models on more data with more computing power would reliably make them smarter — told Patel, "Now that compute is big, we are back to the age of research." Sutskever was predicting that we'd hit the limits of what increased computing power alone could do and sparked fierce debate about whether AI progress was about to hit a wall. Through essays and interviews, Patel also seeds his own point of view — that AI will be a big deal, but truly transformative models will come slower than companies have promised because they haven't yet