


08-21-2025 | TECH

AI is already shaping the future. So why do so few of us get to decide what that future will be?

‘The point is not to slow AI. It is to decide, in public, which AI we want and on what terms.’

SHARE 



[Animation: RDVector/Adobe Stock]



BY **ENRIQUE DANS**

LISTEN TO THIS ARTICLE

[More info](#)



0:00 / 7:35

In Silicon Valley boardrooms, a small group of executives is quietly making decisions that will shape the lives of billions. And most of us won't know what those decisions are until it's too late to change them.

In July, the White House published “**America's AI Action Plan**,” a 28-page document that reads like an industrial policy for a new arms race.

Buried in Pillar I is a line that tells you exactly where U.S. policy is headed: Revise the NIST **AI** Risk Management Framework to eliminate references to misinformation, diversity, equity, inclusion, and climate change.

When governments start crossing out those words by design, it's fair to ask who is setting the terms of our technological future—and for whose benefit.

This is more than rhetoric. The same plan boasts of rolling back the prior administration's AI order, loosening oversight, and fast-tracking infrastructure and energy for data centers. It recasts **artificial intelligence** primarily as a geopolitical race to “win,” not as a societal system to govern. It's a perspective less about stewardship and more about deal-making, a style of governance that treats public policy like a term sheet. That framing matters: When the policy goal is speed and dominance, accountability becomes a “nice-to-have.”



Enrique Dans

Subscribe to Enrique Dans on Medium to understand and reflect on the effects of technology on people, companies and society. An article and a handful of carefully curated links per day.

[SIGN UP](#)

THE EUROPEAN PATH

Europe has chosen a completely different sequence: Set guardrails first, then scale. The EU AI Act entered into force in August 2024 and phases in obligations through 2026, with enforcement designed around risk. Imperfect? Sure. But the message is unambiguous: Democratic institutions—not just corporate PR—should define acceptable uses, disclosures, and liabilities before the technology is everywhere.

Meanwhile, the center of gravity in AI sits with a handful of firms that control compute, models, and distribution. Consider compute—the accelerator-based computing capacity (GPU/TPU time) required to train and run modern AI—as well as models and distribution. **Analysts still peg Nvidia’s share of AI accelerators at around 90%**, and hyperscalers lock up capacity years in advance. That scarcity shapes who can experiment, who can’t, and who pays whom for access.

When the head of state approaches technology policy like an investment banker, those negotiations aren’t about public interest; they’re about maximizing a deal, often for the state’s coffers, and sometimes for political capital.

Opacity compounds the problem. **OpenAI’s own GPT-4 technical report** declines to disclose the training data, model size, or compute used, explicitly citing competition and safety. Whatever you think of that rationale, it means society is being asked to accept consequential systems while remaining largely blind to what went into them. “Trust us” is not governance.

CONCENTRATED POWER

If you want a small but vivid example of how private choices ripple into public life, look at what happened when OpenAI released a flirty voice called Sky that many **thought sounded like Scarlett Johansson**. After public backlash, **the company paused the voice**. A cultural boundary was drawn not by a regulator or a court, but by a product team, a crisis comms cycle, and a corporate decision. That’s a lot of power for a very small group of people.

Power also shows up on the utility bill. Google’s latest environmental reporting **links a 48% increase in greenhouse emissions since 2019 to data-center growth for AI**, and documents **6.1 billion gallons of water used in 2023 for cooling**—numbers that will rise as we scale. **Mistral’s**

life cycle analysis goes further, estimating per-prompt energy and water use for its models. Every “ask the model” has a footprint; multiply by billions and you can’t pretend it’s free, no matter how much of a climate change denialist you may be.

So yes, the United States is “winning the race”—to concentrate decisions that affect expression, employment, education, and the environment in a tiny circle of boardrooms. The result is a democratic deficit. The public is reduced to spectators, reacting to *faits accomplis* instead of setting the rules.

THE ALTERNATIVE

What would it look like to flip the script? Start by treating AI as infrastructure that requires public capacity, not just private CapEx. The [National AI Research Resource](#) pilot reflects the right instinct: Give researchers and startups shared access to compute, data, and tools so that inquiry isn’t gated by hyperscaler contracts. Make it permanent, well-funded, and independent, because open science dies when access is controlled by NDA.

Second, attach conditions to public money and public procurement. If agencies and schools are going to buy AI, they should demand basic disclosures: which data were used for training; what guardrails govern outputs; which independent tests the model has passed; and an energy-and-water ledger tied to time and place, not annual averages. If a vendor can’t meet those bars, they don’t get the contract. That’s not “anti-innovation.” It’s market discipline aligned with public values.

Third, separate layers to curb lock-in. Cloud providers shouldn’t be able to mandate that their model must use their chips to run their services as the default. Interoperability and data portability aren’t romantic ideals; they are how you keep a sector competitive when three firms control the stack.

Fourth, transparency must mean more than model cards written by the vendor. For systems above a certain scale, we should require auditable disclosures to qualified third parties—on training data provenance, evaluation suites, and post-deployment performance. If that sounds onerous, that’s because consequence at scale is onerous. We’ve learned this in every other critical infrastructure.

Finally, align the environmental story with reality. Water and energy disclosures must be real-time, facility-specific, and verified. “Water positive by 2030” doesn’t help a town whose aquifer is being drained this decade. If companies want to be first to ship frontier models, they should also be first to implement 24/7 carbon-free energy procurement and hard water budgets tied to local hydrology.

A DEEPER DANGER

There’s a deeper danger when national technology strategy is run like a business portfolio: Efficiency and revenue become the primary metrics, overshadowing the harder-to-quantify needs of citizens. In the private sector, sacrificing ethics, transparency, or long-term stability for a profitable deal can be chalked up to shareholder value. In government, that same trade-off erodes democracy itself, concentrating decisions in even fewer hands and normalizing a “profit-first” lens on matters that should be about rights, safeguards, and public trust.

The point is not to slow AI. It is to decide, in public, which AI we want and on what terms. The U.S. is capable of both ambition and restraint; we did it with aviation, with medicine, with finance. AI should be no different. If we leave the big choices to a few firms and a few political appointees, we’ll get a future built *for* us, not *by* us. And the price of rewriting it later will be higher than anyone is admitting today.



Enrique Dans

Subscribe to Enrique Dans on Medium to understand and reflect on the effects of technology on people, companies and society. An article and a handful of carefully curated links per day.

SIGN UP

*The early-rate deadline for Fast Company's **Most Innovative Companies Awards** is Friday, September 5, at 11:59 p.m. PT. **Apply today.***



Sign up for our weekly tech digest.

Email *



SIGN UP

This site is protected by reCAPTCHA and the Google [Privacy Policy](#) and [Terms of Service](#) apply.

[Privacy Policy](#)

ABOUT THE AUTHOR

Enrique Dans has been teaching Innovation at IE Business School since 1990, hacking education as Senior Advisor for Digital Transformation at IE University, and now hacking it even more at Turing Dream. He writes every day [on Medium](#). [More](#)

EXPLORE TOPICS

ARTIFICIAL INTELLIGENCE