AINOW

ARTIFICIAL POWER

2025 Landscape Report

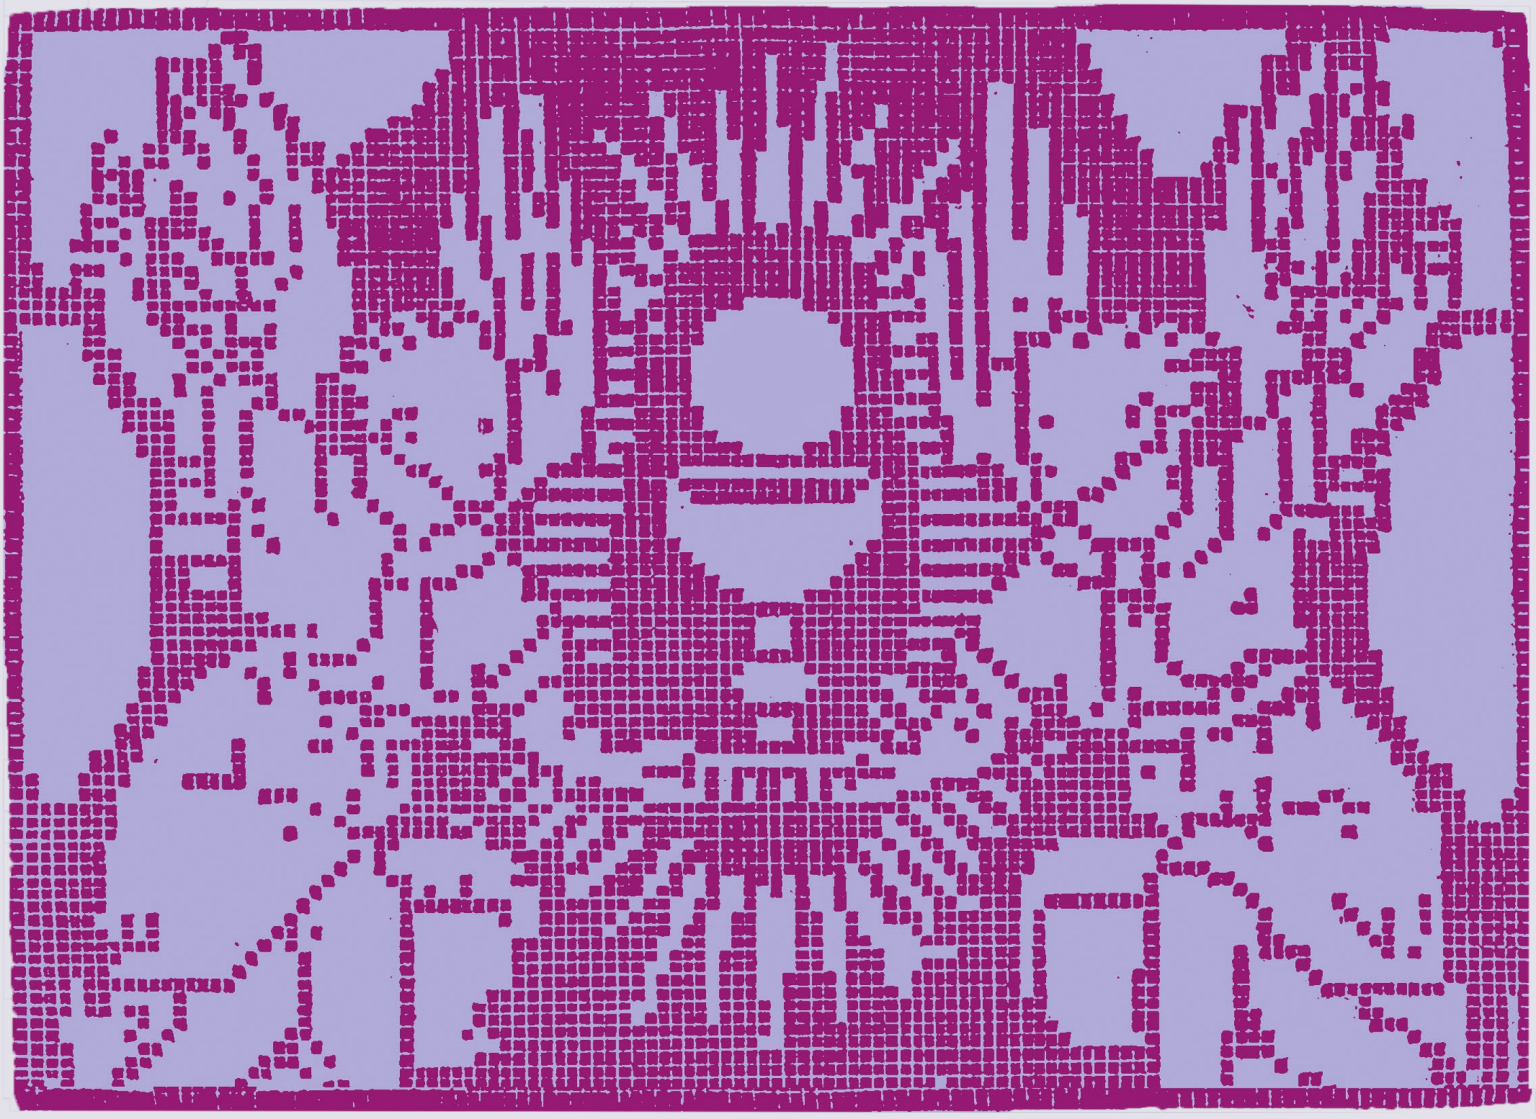

ACKNOWLEDGMENTS

Authored by Kate Brennan, Amba Kak, and Dr. Sarah Myers West.

With research support from Mohammed Ali, Yasmine Chokrane, Madeline Kim, Tekendra Parmar, Tanya Raja, and Boxi Wu.

Special thanks to Rhiana Gunn-Wright for her feedback and editorial contributions.

Project management by Ellen Schwartz.

Copyediting by Caren Litherland.

Design by Partner & Partners.

With contributions and feedback from:

- Annette Bernhardt, UC Berkeley Labor Center
- Abeba Birhane, Artificial Intelligence Accountability Lab, Trinity College Dublin
- Brian Chen, Data & Society
- Jane Chung, Justice Speaks
- Stefanie Coyle, New York Civil Liberties Union
- Andrea Dehlendorf, AI Now
- Kevin DeLiban, TechTonic Justice
- Becca Deutsch, Amazon Employees for Climate Justice
- Alix Dunn, The Maybe
- Ali Finn, AI Now
- Timnit Gebru, Distributed AI Research Center
- Ryan Gerety, Athena Coalition
- Lisa Gilbert, Public Citizen
- Sam Gordon, Tech Equity Collaborative
- Janet Haven, Data & Society
- Nidhi Hegde, American Economic Liberties Project
- Ben Inskeep, Citizen Action Coalition of Indiana
- Taylor Jo Isenberg, Economic Security Project
- Heidy Khlaaf, AI Now
- Stephen Lerner, Action Center for Race and the Economy - Bargaining for the Common Good
- Barry Lynn, Open Markets Institute
- Varoon Mathur, Duke Institute for Health Innovation, Duke Health
- Jill McArdle, Beyond Fossil Fuels
- Michelle Meagher, SOMO
- Brian Merchant, Independent Journalist
- Erie Meyer, Georgetown University
- Sarah Miller, American Economic Liberties Project
- Stacy Mitchell, Institute for Local Self-Reliance
- Arvind Narayanan, Center for Information Technology Policy, Princeton University
- Chris Nielsen, National Nurses United
- Teri Olle, Economic Security Project
- Britt Paris, Rutgers University
- Tekendra Parmar, Investigative Journalist
- Reshma Ramachandran, Yale School of Medicine
- Steven Renderos, MediaJustice
- Rashida Richardson, Worcester Polytechnic Institute & Northeastern University
- Hilary Ronen, Local Progress
- Leevi Saari, AI Now
- Hannah Sassaman, People's Tech Project
- Matt Scherer, Center for Democracy and Technology
- Paromita Shah, Just Futures Law
- Ganesh Sitaraman, Vanderbilt University
- Andrew Strait, UK AI Security Institute
- Kasia Tarczynska, Good Jobs First
- Jim Thomas, Scanthehorizon.org
- Max Von Thun, Open Markets Institute
- Jai Vipra, Cornell University
- Robert Weissman, Public Citizen
- Meredith Whittaker, Signal
- Savannah Wilson, Clean Virginia
- Boxi Wu, Oxford University

Cite as: Kate Brennan, Amba Kak, and Sarah Myers West, "Artificial Power: AI Now 2025 Landscape", AI Now Institute, June 3, 2025, https://ainowinstitute.org/2025-landscape.

EXECUTIVE SUMMARY

Those of us broadly engaged in challenging corporate consolidation, economic injustice, tech oligarchy, and rising authoritarianism need to contend with the AI industry or we will lose the end game. Accepting the current trajectory of AI proselytized by Big Tech and its stenographers as "inevitable" is setting us up on a path to an unenviable economic and political future—a future that disenfranchises large sections of the public, renders systems more obscure to those it affects, devalues our crafts, undermines our security, and narrows our horizon for innovation. This is true whether or not the technology even works well, on its own terms; it often doesn't.

The good news is that the road offered by the tech industry is not the only one available to us. This report explains why the fight against the industry's vision for AI is a fight worth having, even as we turn ourselves tirelessly toward the task of building out the shared project of a just, equitable, sustainable, and democratic society.

Over the past decade, taming the power of big technology platforms like Microsoft, Amazon, Google, and Meta has increasingly become a central question in American political and public discourse. Unless we contend with the power vested in these firms, we won't meaningfully be able to hold the industry accountable to the interests of the broader public, even as these companies reshape markets, institutions, and infrastructures core to public life.

What does AI have to do with any of this? As we argued in our 2023 report,[1] AI is fundamentally about concentration of power in the hands of Big Tech. At the start of the year, it seemed like the market was poised for disruption, with a new crop of Silicon Valley challengers gaining prominence, like OpenAI, Anthropic, StabilityAI, and Inflection AI. But now, just two years later, it is clear that the bench of key players in this market hasn't changed much: Microsoft, Google, Meta, Musk's xAI, OpenAI (backed by Microsoft), and Anthropic (backed by Amazon and Google).[2] The new suite of LLM-powered AI products has pushed these firms into the spotlight, dominating headlines and, increasingly, becoming the subject of dinner-table conversation.

Amid the frenzy, there's been a misplaced focus on blinkered questions of whether one AI system or application is good or bad, or evaluating the moral quandaries of hypothetical worlds. Instead we need to redirect attention to the AI ecosystem, and its dependencies and risks, as a whole. **The question we should be asking is not if ChatGPT is useful or not, but if OpenAI's unaccountable power, linked to Microsoft's monopoly and the business model of the tech economy, is good for society.**

Looking beyond individual use cases allows for a more comprehensive look into the centers of power that drive our current tech landscape. AI as a field has been not just co-opted but *constituted* by the logics of a few dominant tech firms. It is no coincidence that the "bigger-is-better" paradigm that dominates the field today, where the scale of compute and data resources are generally used as a proxy for performance, lines up neatly with the incentives of Big Tech, which disproportionately controls these resources, the talent to leverage them, and the pathways to monetization. Around 2012, as it became apparent that substantial gains in model performance could come simply from applying larger and larger scale data and computational resources to existing algorithms, tech giants moved quickly to shore up their existing advantages and hire talent.[3] Corporate influence over AI's research trajectory has been cemented through tech firms' AI labs dominant presence at prestigious machine learning conferences, further shaping the field of research in ways that align with industry.[4] This is in part because building AI bigger requires enormous resources, both financial and social, to achieve unrestricted growth at breakneck speed—resources that AI companies own and control.[5]

But it's not just market power we need to be concerned with: These tech oligarchs are counting on a wholesale rewriting of our social and economic foundations, using AI as the justification. From breaking apart the US federal government and raiding citizen data under the guise of efficiency, to redesigning workflows to devalue human labor and creativity so they are AI-ready, to redirecting our entire energy infrastructure to prioritize their technology over people's basic needs, the vision promulgated by tech oligarchs requires, as a foundation, the unraveling of core social, political, and economic fabrics.

Across our information ecosystem, from science to education, healthcare, culture, and art, AI is being positioned as a disruptive new infrastructure and a mediating force. **In truth, though, it rehashes an old playbook, helicoptering in solutions built on the extraction of expertise and value from all corners of society**—solutions that always, eventually, amount to the further degradation of life for the most

marginalized among us. While generative AI and AI agents are the buzzwords that splash across the headlines, the same dynamics are true of precursors to contemporary AI systems like automated decision-making technologies used in banking, hiring, and criminal justice. The techniques and vendor names vary, but the industry incentives powering proliferation, as well as the failure modes across these systems, share much in common.

More than a decade of evidence demonstrates how it goes: The introduction of these systems concentrates power among the deployers of the tech, leaving those on the receiving end more insecure, vulnerable, and unable to contest the determinations made by the "smart machine" at the expense of the broader public. These tools are often invisible to those judged by them, and inscrutable even when they are visible.

Why society would ever accept this bargain is the critical question at hand. Amid the excitement over AI's (speculative) potential, the sobering reality of its present and recent past is obscured. When we consult the record on how AI is already intermediating critical social infrastructures, we see that it is materially reshaping our institutions in ways that ratchet up inequality, render institutions opaque to those they are meant to serve, and concentrate power in the hands of the already powerful. (See Chapter 3: Consulting the Record.) It makes clear that for all the whiz-bang demos and bold Davos proclamations, on the ground AI is consistently deployed in ways that make everyday people's lives, material conditions, and access to opportunities worse and the systems that incorporate them stronger.

This report's title, *Artificial Power*, captures the critical, and at times contradictory, moment we find ourselves in. On one hand, the tech oligarchy has successfully deployed "AI"—as a strategic marketing term and as a set of automation technologies—to cement and grow its power. At the same time, this power is vastly inflated, contingent, and poised for disruption. Contending with the dual reality of how those with power have deployed AI systems to enact significant harm while exposing the ways this power can and must be disrupted is the central work of this moment.

THE ELEPHANTS IN THE AI ROOM: OF BUSINESS MODEL(S) AND FATAL TECHNICAL FLAWS

The AI industry is on shakier ground than it may seem. Valuations are sky high, while the business model hinges on an intensely expensive technology that lacks a consistent revenue stream. AI companies bleed money for every user they gain: Anthropic burned through $5.6 billion[6] this year but was valued at $61.5 billion.[7] OpenAI lost $5 billion[8] but is valued at $300 billion.[9] No profit-making use cases exist yet, or are even on the horizon. This may seem like business as usual for the move-fast-and-break-things ethos of Big Tech, but we are in a profoundly different monetary environment now—it's no longer the 2000s or 2010s. Markets are saturated, market dominance has been established among the platform and infrastructure winners that emerged from those decades, and, put simply, the cost of large-scale AI is eye-watering at a level not seen before in tech.

The question now circulating more, and more openly, is this: When will the AI bubble burst and who will be impacted by it? Because this *is* a bubble. For their part, companies are pushing out shiny objects to detract from the business reality while they desperately try to derisk their portfolios through government subsidies and steady public-sector (often carceral or military) contracts. While it's very clear how tech companies benefit from claims to "public-interest AI" used to justify the pouring of taxpayer dollars into this industry, it is not at all clear how this benefits the rest of society. (See Chapter 2: Heads I Win, Tails You Lose)

Chapter Snapshot

Heads I Win, Tails You Lose

This section maps the drivers that are securing Big Tech firms' advantage in the AI market, before turning to the question of who loses in the end:

- Cloud infrastructure providers benefit from cycles of AI dependence
- Big Tech firms benefit from leveraging control over the tech ecosystem
- Big Tech benefits from the data center boom … even if the AI boom doesn't pan out

With generative AI, in particular, the hyped claims stand in stark contrast to the largely mundane use cases that are being shoved into nearly every app and service. In contrast to the claims of world-changing tech, Meta is investing heavily in AI advertising infrastructure.[10] OpenAI is creating AI agents that fill out forms and call web browsing "research"—sucking up your data and requesting invasive permissions as it does so.[11] Pressured by their employers, software engineers are using Microsoft's Copilot to produce more code, more rapidly, undercutting their skills and trade.[12] And cloud companies are happily locking enterprise users into their software-as-a-service (SaaS) ecosystems by automatically upgrading them to new AI features—and raising the price.[13] These are not examples of a technology being embraced by a society glad of its utility.

Despite being positioned as critical infrastructure, AI systems in their current form have fundamental flaws: there is an intractable problem of "hallucinations" with LLMs reliant on randomly generated coordination, leaving the humans in this technology in the unenviable position of fact-checking the tech meant to make their lives easier.[14] Peer-reviewed research indicates that in many cases, AI systems fail profoundly at even basic tasks when applied in real-life contexts.[15] They're also far from resilient, prone to cybersecurity vulnerabilities like web poisoning attacks and new jailbreaking methods that enable the persistent unauthorized disclosure of training data and other sensitive information.[16]
And it's not that the trade-offs are weighted and deemed worth it: In many use cases, AI is deployed by thwose with power against those who lack it, and who have no opportunity to opt out or seek remedy when mistakes are made. (See Chapter 3: Consulting the Record.)

Chapter Snapshot

Consulting the Record

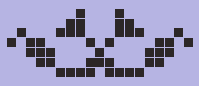

This section compiles over a decade of evidence showing how the tech industry has sought to reshape society to enable more widespread deployment of the technologies it builds and profits from, often contributing to the degradation of our social, political, and economic lives. The section aligns on five key takeaways:

- **AI's benefits are overstated and underproven,** from cancer cures to hypothetical economic growth - while some of its flaws are real, immediate, and growing.
- **AI-sized solutions to entrenched social problems displace grounded expertise,** in disparate domains like higher education, healthcare, and agriculture.
- **AI solutionism obscures systemic issues facing our economy – obscuring economic concentration and acting as a conduit for deploying austerity mandates by another name.** The DOGE power grab is instructive, though New York's MyCity offers another example where millions of taxpayer dollars were invested into flawed AI solutions that failed to deliver tangible benefits to the public.
- **The productivity myth obscures a foundational truth - the benefits of AI accrue to companies, not to workers or the public at large**, even as algorithmic management tools make work unstable and unsafe. 'Agentic AI' will make workplaces even more bureaucratic and surveillant, reducing not increasing autonomy.
- **AI use is frequently coercive, violating rights and undermining due process.** This is nowhere more clear than the rise of AI usage in immigration enforcement, where human rights abuses are common and legal norms are routinely violated - even before AI is in the mix.

It doesn't help that the prevailing deregulatory current is an industry that continually acts above the law and is driven narrowly by its bottom line (See 1.4: Regulation): in 2024, we saw companies rush to market with products that are patently inaccurate, insecure, and compromise the safety of consumers; engage in anti-competitive practices that shore up their advantages to shut the door behind them; and deploy larger than life narratives around AGI and innovation to quell any form of interrogation and critique.

Chapter Snapshot

AI's False Gods

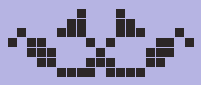

This section interrogates narratives that advance AI industry dominance and make the current trajectory of AI seem inevitable:

- **The AGI Mythology: The Argument to End All Arguments** unpacks the nebulous claims surrounding "artificial general intelligence," arguing that the term collapses complex technical realities into a singular, imminent, and inevitable future that conveniently advances the interests of the companies claiming to build it.
- **'Too Big to Fail:' Infrastructure and Capital Push** explores how tech firms are deploying unprecedented amounts of capital to perpetuate a "bigger-is-better" AI paradigm, shoring up their continued market dominance through government and taxpayer support.
- **AI Arms Race 2.0: From Deregulation to Industrial Policy** details how the US-China AI arms race has heightened, and is now used to brand a slate of industrial policy initiatives designed to boost the tech industry and avert regulatory scrutiny.
- **Recasting Regulation as a Barrier to Innovation** shows how the AI industry has strategically pitted regulation against innovation, leading to a global deregulatory posture that ignores the role regulation plays in enhancing innovation and competition.

IT IS TIME TO BUILD—JUST NOT AI.

AI is predominantly used *on* us, not *by* us, to shape access to resources and life chances. But while there is a clear path dependency within this narrow trajectory for AI proselytized by big tech and its stenographers, the good news is that it's not the only road available to us.[17] Not by a long shot.

AI hype has tapped into a sentiment that is real and widespread: genuine enthusiasm to build a future where all people can thrive, a future that will likely look radically different from the present. It is a catalyzing goal we should unite around; most of us want a future that frees us from the endless cycle of war, pandemics, and environmental and financial crises that characterize our present. The 2024 US presidential election brought the need to create social and political institutions connected to the needs and lived realities of people even closer to home, across the country and across the world. But AI doesn't create any of these—and pegging our shared future on AI makes that future harder, not easier, to achieve because it binds us to a decidedly bleak path, stripping us not only of the ability to choose what to build and how to build it, but also stealing what joy we might take in that building. This hype-prescribed AI future further distances us from a life with dignity, one where we have the autonomy to make our own decisions and where democratically accountable structures work to distribute power and technological infrastructures in ways that are robust, accountable, and protected from systemic shocks.

What we've seen play out within the AI industry is not unique to this industry, of course. The dynamics of "gain for me, loss for thee" have been examined in many critiques of shareholder capitalism broadly, which emphasize the corporate willingness to speculate in ways beneficial to shareholders but not to society at large, alongside the perverse incentives that lead firms to act against their own business interests, and develop an orientation toward monopolization and sclerosis. If anything, the AI market is the peak exemplar of overreliance on venture-based investment.[18]

But AI introduces new dynamics and accelerants. As designed, developed, and deployed currently, AI works to entrench existing power asymmetries, and to ratchet them up. It naturalizes inequity as destiny and deservedness—simply the classification given by the intelligent system—while rendering these underlying patterns, judgments, and self-interested drivers inscrutable to those affected by AI's judgments and instructions.

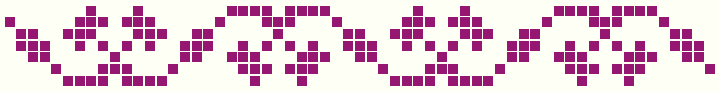

ANOTHER AI IS POSSIBLE. HOW DO WE GET THERE?

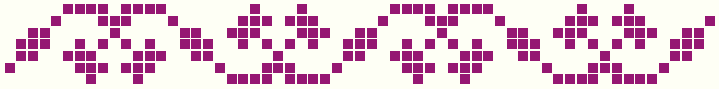

Although real ways in which the AI market could be structured to benefit the public may exist, the path charted by the companies controlling AI, and those wanting a piece of the control AI could give them over our lives and institutions, won't lead us there.

One thing is clear: we can't fight tech oligarchy without rejecting the current industry trajectory around large scale AI. **It's a crucial inflection point and how American policymakers and movement leaders choose to respond to the AI industry will write the coming chapters of the story of tech power.** AI companies, and those who lead them, have positioned themselves to reshape broad swaths of society—within and beyond the US—not only to work in their interest, but to do so in ways that allow their firms to capture the lion's share of the value.

This isn't inevitable. In fact, the tide of public opinion is moving decisively against the entrenched power of tech firms. And we've seen major legal wins in the landmark antitrust cases filed by the DOJ and FTC against Google and Meta. After successfully proving that Google maintains an illegal monopoly in search and advertising markets,[19] the DOJ is now requesting bold, structural remedies that were all but inconceivable a few years ago.[20] These remedies, which include breaking up Google's advertising technology business and spinning off Chrome, strike directly at the heart of Google's business model. But the remedy trials have revealed a larger truth: AI startups can't scale or achieve distribution without Big Tech firms' infrastructure. That's why OpenAI offered to buy Chrome. It's why Perplexity's CEO said he'd want to buy it—then pandered so as not to aggravate a company he's dependent on.[21] This is why it's especially important that we not cede the momentous ground these regulatory actions have pushed us towards when Big Tech companies use AI as cover for staying unregulated.

In this report, we lay out another path forward. First, we map what we mean by AI in the first place, provide an accounting of the false promises and myths surrounding AI, and examine whom it's working for and whom it's working against. Then, given that AI consistently fails the average public, even as it enriches a sliver, we ask what we lose if we accept the current vision of AI peddled by the industry. Finally, we identify leverage points that we can latch on to as we mobilize to build a world with collective thriving at its center—with or without AI.

We are not naive. The headwinds against sensible AI alternatives have never been stronger. The tech industry is better resourced and the political environment more bleak than ever before. Indeed, considering the tech industry to be simply a collection of firms itself misses key sources of their power, from their surveillance apparatus to their control over global digital infrastructures that shape our states, our institutions, our economies, and, most importantly, our lives. But the battle has never been more important. Contesting AI links the movements we must build not only to create meaningful public power, but also to seed a new path defined by autonomy, dignity, respect, and justice for all of us.

In the following sections, we set out a tool kit for reasserting public power amid a takeover by AI firms.

A more fulsome treatment appears in Chapter 4: A Roadmap for Action, but here are the highlights:

Chapter Snapshot

A Roadmap for Action: Make AI a Fight About Power, Not Progress

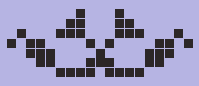

1. **Target how the AI industry works against the interests of everyday people.** In the aftermath of the 2024 election, there's growing consensus across the ideological spectrum that focusing on the material conditions and economic interests of working people is key to building political power. We need to not only make AI-related issues more relevant to movements fighting for economic populism and against tech oligarchy; we also need to better target the AI industry as a key actor working against the interests of the working public. The pushback against the Department of Government Efficiency (DOGE), the buildout of AI data centers, and algorithmic prices and wages constitutes fertile ground for building a broader movement unified in its focus on rejecting AI's unaccountable tech-enabled social and political control.

2. **Advancing worker organizing is the clearest path to protecting us and our institutions from AI-enabled capture.** Labor campaigns have demonstrated that working people have a particular form of power to wield, power that can determine how their employers deploy AI and digital systems. The deeper opportunity for labor, and a more transformative ambition, however, would be to direct labor's power not just toward whether and how AI is used in the workplace, but also toward recalibrating the technology sector's power overall and shaping the trajectory of AI in the public interest and common good.

3. **Enact a 'zero trust' policy agenda for AI.** Trust in AI firms' benevolence is not a smart, informed, or credible option—not if we're going to proceed with serious work. Enacting a policy agenda built on bright-line rules that restrict the most harmful uses of AI, regulate the AI life

   cycle from nose to tail, and ensure that the industry that currently creates and profits from AI isn't left to regulate and evaluate itself—essentially grading its own homework—must be a priority at the state and federal levels in the US, and internationally.

4. **Bridge networks of expertise, policy, and narrative to strengthen AI advocacy.** AI advocacy and policy has often been undermined too often by blinkered views that fail to see the different components of the AI supply chain materially, are often single issue-focused, and it is easy to miss the ways in which big picture narratives manifest to limit possibilities in policy fights. From national security logics that can be a vector both for, and cutting against, moves towards industry accountability; to reframing traditional data privacy levers as key tools in the fight against automation and addressing market power.

5. **Reclaim a positive agenda for public-centered innovation without AI at the center.** The current trajectory of AI puts the public under the heel of unaccountable tech oligarchs. But their success is not inevitable. By moving out from under the shadow of the idea that large-scale AI is inevitable, we can reclaim the space to conduct real innovation and to push for exciting and novel alternative pathways that leverage tech to shape a world that serves the public and is governed by our collective will.

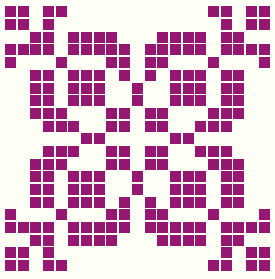

EPILOGUE: THE WORLD WE WANT (AND WHY THE CURRENT TRAJECTORY OF AI WON'T GET US THERE)

The AI hype of the past year has sucked the air out of an already stuffy room, making it feel futile—at times, impossible—to imagine anything other than a steady march toward the inevitable supremacy of AI. But no matter how true that may feel, it is only that: a feeling. It is not reality—not yet, at least. There are, in fact, many alternatives to this version of AI, many ways to shape new worlds. Like AI, though, these are not inevitable either. Making them possible starts by asking and answering a single question: Is this the world we want? At AI Now, we want to see a world that has:

Good Jobs

Everyone deserves access to the resources to live a happy, fulfilled life, and a dignified job that will provide them with these resources. Under the right conditions—like policies intent on uplifting, rather than exploiting, workers—new technologies have the potential to make everyone's working lives better.

Yet AI's current trajectory is fundamentally incompatible with the proliferation of good jobs rooted in human flourishing. As the AI industry embeds itself into nearly every sector of the economy, firms are increasingly shaping a job market contingent on

worker displacement and exploitation. Overwhelmingly, AI companies are embedding "productivity" tools designed to help businesses optimize their bottom line across the entire labor supply chain. This requires work itself to become legible to AI systems, making working life more routinized, surveilled, and hierarchical. Furthermore, instead of working to protect workers from the uncertainty coming from this new market, AI companies are undermining hard-won labor protections, exploiting legal loopholes to avoid corporate accountability, and lobbying governments to support policies that prioritize corporate profits over fair and just treatment for workers.

As the current vision of AI takes hold, we lose a future where AI technology works in support of stable, dignified, and meaningful jobs. We lose a future

- where AI supports fair and livable wages, instead of wage depreciation;
- where AI ensures that workers have the control to decide how new technology affecting their careers is deployed, instead of undermining their expertise and knowledge of their own work;
- where we have strong policies to support workers if and when new technologies automate existing roles—including laws that broaden the social safety net—instead of AI boosters who brag to shareholders about cost savings from automation;
- where robust public benefits and time-off policies ensure the long-term wellness of employees, instead of AI being used to surveil and nickel-and-dime workers at every turn;
- where AI helps protect employees from health and safety risks on the job, instead of perpetuating conditions that make work dangerous and celebrating employers who exploit labor loopholes to avoid responsibility; and
- where AI fosters meaningful connection through work, instead of driving cultures of fear and alienation.

Shared Prosperity

The proliferation of any new technology has the potential to increase economic opportunity and lead to widespread shared prosperity. *But shared prosperity is incompatible with AI's current path toward maximizing shareholder profit.*

The insidious myth that AI will lead to "productivity" for everyone when it really means productivity for a select number of corporate firms propels us further down the path of shareholder profit as the singular economic goal. Even well-intentioned government policies designed to boost the AI industry steal from the pockets of workers. For example, government incentives meant to revitalize the chip manufacturing industry were thwarted by corporate buy-back provisions, sending millions of dollars to companies, not to workers or job creation. And despite some meaningful moves to investigate the AI industry under the Biden Administration, companies have still gone largely unchecked, meaning new entrants cannot come in to challenge these practices.

By proliferating the myth that AI will inherently lead to shared prosperity (or that "a rising tide lifts all boats"), we lose the economic policies that could meaningfully lead us into a period of shared prosperity, including pro-enforcement policies to break up the concentration of corporate power, a strong pro-labor agenda to center the needs of workers, and industrial policy

strategies designed to put workers and communities before the bottom line of large corporations. Crucially, we lose a thriving and competitive economy, where innovators and entrepreneurs are incentivized to launch sustainable and prosperous businesses that need not rely on surveillance mechanisms, hyper-growth venture capital funding, and extractive business models to succeed.

Freedom & Autonomy

We all deserve to live in a world where our personal, political, and economic lives are free from coercion.

But the current trajectory of AI is grounded in coercion and opacity. By amassing such concentrated power, AI companies have assumed control over many aspects of our lives, subjecting us to coercive practices designed to maximize their own profit-making potential. Nowhere is this more explicit than the regime of surveillance pricing algorithms. These algorithms collect extraordinary amounts of consumer data to set individualized prices for goods and services, such that important aspects of daily life—from buying groceries, taking an Uber, or buying an apartment—are controlled by companies looking to squeeze consumers to pad their own bottom lines. As an inherently centralizing technology, AI is consistently deployed to support and benefit from surveillance states, carceral systems, and military techniques, embedding corporate interests into state apparatuses and vice versa, making life more coercive and violent.

AI could contribute to a future built on autonomy and transparency. But right now, we stand to lose a robust ecosystem of public resources that are not dependent on private industry players or priced for profit; pricing for goods and services driven by principles of fair competition; and changing life circumstances not being treated as grounds for companies to profit. A public life invested in freedom would be divested from the proliferation of surveillance systems, carceral logics, and military state apparatuses, enabling people and communities to thrive.

Sustainable Future

Technological progress does not have to come at the expense of our natural environment. A publicly beneficial AI landscape places principles of sustainability and environmental justice at its center, recognizing that the health and safety of our planet and communities is of paramount importance.

AI's current trajectory is not merely incompatible with sustainability; it is fueling climate degradation. The focus on scale at all costs within the AI industry makes it dependent on climate extraction and energy dominance. These include a deregulatory policy intent on accelerating the AI industry (meaning more energy, more infrastructure, more natural resources going into AI) as well as the insidious belief that AI is going to help *solve* the climate crisis, "greenwashing" the environmental harm it is already enacting.

As a result, what we lose is a world where governments and companies work together to advance principles of sustainability and environmental justice, investing in green and renewable energy infrastructure to support the additional energy usage that new

technologies require. Decisions on siting, permitting, and constructing new technology infrastructure should be made in relationship with local communities, with particular attention to those communities who are most significantly affected by technology's industrial evolution. Cities and states should not be coerced into providing subsidies for infrastructure that come at the expense of local communities' needs, like funding for schools and healthcare. We also lose a future where industry takes its climate commitments seriously and works to mitigate the harmful effects of its industrial processes, recognizing that the health and safety of our planet and communities is key to everyone's survival—including industry's.

Strong Social Safety Net

We should live in a world where everyone has access to and community control over a robust system of public resources.

Yet as AI rapidly shifts our social landscape, governments are increasingly driven by privatization and austerity measures. AI is ushering in a policy agenda designed to enrich private interests rather than provide robust public benefits. AI firms are pushing AI integration in local and federal government agencies driven by austerity, restricting people's access to needed resources and the social safety net. Moreover, the purported AI arms race with China is being leveraged to convince governments that domestic infrastructure is an imperative for national competitiveness and security, encouraging public agencies to throw generous tax exemptions at private companies in order to build massive data centers in communities that may not want them there. These exemptions—totaling billions of dollars—steal investment from strong public resources that benefit everyone, like investments in more teachers, roads, and libraries.

We could have a world where technology works in service of the broader public, like algorithms designed to maximize people's options, not reduce them; and where companies are required to pay their fair share of taxes to local communities, rather than fighting efforts to make them pay their fair share.

Security

In an increasingly complex world, security and resilience are more important than ever. Our infrastructures are largely invisible to us until the moment they break down - and we've felt the shock of infrastructural failure frequently over recent years, from supply chain issues during the pandemic to power outages to bank closures.

Layering AI into our critical infrastructures, particularly when AI systems are highly concentrated, creates a real and present security risk. These risks are manifold: there are cybersecurity risks that are inherent to many AI systems that make them vulnerable to hacking, and some of these cannot be remedied. There are systemic risks introduced by overreliance on a single technology: for example, if banks, hospitals, and schools all use the same cloud infrastructure provider, an outage could affect all of these at one fell swoop. These risks are at their highest when AI is used in life or death settings - and from healthcare to defense, these are some of the industry's prime markets. And there are risks emanating from decisions made by

the companies themselves to experiment in the wild, bringing to market technologies that have not been adequately tested or validated and with little certainty that they will work as intended, let alone cause harm.

We could have a world where our safety and well-being are not vulnerable and exposed to an industry that is scaling at an unprecedented rate with little regard for safety and security, let alone compliance with the law; where AI is validated, tested and built safe by design and used with prudence rather than impunity.

Innovative Tech Ecosystem

Technology has the potential to solve important societal challenges and push the frontiers of innovation forward. In a thriving tech ecosystem, companies big and small are able to succeed, not by amassing concentrated power but by engaging in fair competition. Society benefits from the distribution of a diverse set of products, services, and technologies that result from such competition.

The current AI industry is defined by concentration, precluding a truly diverse and innovative tech ecosystem from flourishing. There is no AI without Big Tech firms, which have spent decades amassing unrestrained data access and economic power and then used those advantages to control key inputs at all levels of the AI stack. Even where new entrants are able to enter the AI market, they are still dependent on the cloud and computing infrastructures of Big Tech firms in order to succeed, creating an ecosystem of dependence rather than competition. Where consolidation has been averted - as was the case with the proposed Nvidia-ARM merger - firms have been able to thrive. ARM went on to IPO and beat quarterly estimates, all after its acquisition was blocked. Furthermore, Big Tech companies control most pathways to consumers and enterprise businesses at scale. This centralized power is also driving a crisis of innovation, where Big Tech companies are bloated, stalled by legal reviews, and are stuck repackaging their existing technologies in order to revive and boost their bottom lines.

What we lose in this bloated and stale tech ecosystem is a truly diverse horizon of possibility, filled with innovations that tackle people's real-life needs, rather than an endless soup of enterprise software and AI agents. We can inspire an ecosystem where people can use technology to build companies slowly and sustainably, without the need to grow rapidly and amorally in order to stay in business. Our entire tech ecosystem is in need of a paradigm shift, one that tears down existing structures to make room for complexity and emergence.[22] This includes breaking up big companies, overhauling the VC-backed funding structure so more companies can thrive, investing in public goods to ensure tech resources are not dependent on large private companies, and increasing institutional investment to bring more diverse people—and thus ideas—into the tech workforce.

Vibrant, Democratic State

We deserve a technological future that works to support strong democratic values and institutions.

What we have now is a society captured by the tech billionaire class. Over the past few decades, a handful of billionaires degraded our entire information system

under the guise of disruption, killing the business model of newsrooms and replacing it with ad-based products required to keep our constant attention. And despite rhetoric that AI has the potential to "democratize" the world, the inherent pathologies of AI make it a centralizing force,[23] contingent on the mass accumulation of data and compute resources in the hands of a few big players.[24] Now the tech billionaire class threatens to destroy our creative industries, transforming hard-earned craft into "content" that is then fed to AI models intended to churn out lossy xerox copies of our masterworks. And, as if that's not enough, these same billionaires have begun to destroy our institutions, purchasing newspapers and taking over the opinion pages, buying elections, and hollowing out our social services.

Contesting the economic and political capital amassed within the tech industry is necessary to create the conditions for a thriving democracy. We urgently need to restore the institutional structures that protect the interests of the public against oligarchy. This will require confronting tech power on multiple fronts, from enacting corporate accountability measures that keep tech oligarchs in check, to staving off efforts to use AI to hollow out our institutions, to bolstering work happening at the community level among local government officials, organizers, and workers devoted to rebuilding a democracy that serves the broader public.

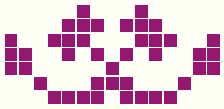

Executive Summary Endnotes

1 Amba Kak and Sarah Myers West, *2023 Landscape: Confronting Tech Power*, AI Now Institute, April 11, 2023, https://ainowinstitute.org/publications/research/2023-landscape-confronting-tech-power.

2 David Cahn, "Steel, Servers and Power: What it Takes to Win the Next Phase of AI," Sequoia, August 5, 2024, https://www.sequoiacap.com/article/steel-servers-and-power.

3 Meredith Whittaker, "The Steep Cost of Capture," *Interactions* 28, no. 6 (November 2021): 50–55, https://dl.acm.org/doi/10.1145/3488666.

4 See Nur Ahmed, Muntasir Wahed, and Neil C. Thompson, *Science* 379, no. 6635 (March 2023): 884–886, https://www.science.org/doi/abs/10.1126/science.ade2420; and Whittaker, "The Steep Cost of Capture."

5 Kak and West, *2023 Landscape: Confronting Tech Power*.

6 Edward Zitron, "There Is No AI Revolution," *Where's Your Ed At*, February 24, 2025, https://www.wheresyoured.at/wheres-the-money.

7 Shirin Ghaffary, "Anthropic Finalizes Megaround at $61.5 Billion Valuation," Bloomberg, https://www.bloomberg.com/news/articles/2025-03-03/anthropic-finalizes-megaround-at-61-5-billion-valuation.

8 Cory Weinberg, "OpenAI Projections Imply Losses Tripling to $14 Billion in 2026," *The Information*, October 9, 2024, https://www.theinformation.com/articles/openai-projections-imply-losses-tripling-to-14-billion-in-2026.

9 Cade Metz, "OpenAI Completes Deal That Values Company at $300 Billion," *New York Times*, March 31, 2025, https://www.nytimes.com/2025/03/31/technology/openai-valuation-300-billion.html.

10 Nilay Patel, "Mark Zuckerberg Just Declared War on the Entire Advertising Industry," *The Verge*, May 1, 2025, https://www.theverge.com/meta/659506/mark-zuckerberg-ai-facebook-ads.

11 Casey Newton, "ChatGPT's Deep Research Might Be the First Good Agent," *Platformer*, February 3, 2025, https://www.platformer.news/chatgpt-deep-research-hands-on.

12 Grant Gross, "Devs Gaining Little (If Anything) from AI Coding Assistants," CIO, September 26, 2024, https://www.cio.com/article/3540579/devs-gaining-little-if-anything-from-ai-coding-assistants.html.

13 Abner Li, "Google Workspace Business Users Getting Full Gemini Experience, Price Increasing," *9to5Google*, January 15, 2025, https://9to5google.com/2025/01/15/google-workspace-gemini-price-increase.

14 Ashish Agarwal et al., "Hallucinations in Neural Machine Translation," abstract, 2018, accessed May 15, 2025, https://research.google/pubs/hallucinations-in-neural-machine-translation.

15 Inioluwa Deborah Raji et al., "The Fallacy of AI Functionality," *arXiv* (2022), https://arxiv.org/abs/2206.09511.

16 Chad Boutin, "NIST Identifies Types of Cyberattacks That Manipulate Behavior of AI Systems," NIST, January 4, 2024, https://www.nist.gov/news-events/news/2024/01/nist-identifies-types-cyberattacks-manipulate-behavior-ai-systems.

17 Sarah Myers West and Amba Kak, "AI's Missing Others," *Boston Review*, December 4, 2024, https://www.bostonreview.net/forum_response/ais-missing-others.

18 Catherine Bracy, *World Eaters* (Penguin Random House, 2025).

19 See David McCabe, "'Google Is a Monopolist,' Judge Rules in Landmark Antitrust Case," *New York Times*, August 5, 2024, https://www.nytimes.com/2024/08/05/technology/google-antitrust-ruling.html; and McCabe, "Google Broke the Law to Keep Its Advertising Monopoly, a Judge Rules," New York Times, April 17, 2025, https://www.nytimes.com/2025/04/17/technology/google-ad-tech-anti-trust-ruling.html.

20 David McCabe, "U.S. Wants to Break Up Google's Advertising Technology," *New York Times*, May 2, 2025, https://www.nytimes.com/2025/05/02/technology/google-advertising-technology-hearing.html.

21 See Kate Brennan, "The Elephant in the Room in the Google Search Case: Generative AI," *Tech Policy Press*, November 4, 2024, https://www.techpolicy.press/the-elephant-in-the-room-in-the-google-search-case-generative-ai; Erin Woo, "OpenAI's Bid To Get Google's Search Data," *The Information*, April 23, 2025, https://www.theinformation.com/articles/openais-bid-get-googles-search-data; and Lauren Feiner, "Perplexity Wants to Buy Chrome If Google Has to Sell It," *The Verge*, April 23, 2025, https://www.theverge.com/policy/654835/perplexity-google-antitrust-trial-remedies-chrome.

22 Maria Farrell and Robin Berjon, "We Need to Rewild the Internet," *Noēma*, April 16, 2024, https://www.noemamag.com/we-need-to-rewild-the-internet.

23 Dan McQuillan, *Resisting AI* (Bristol University Press, 2022).

24 Alex Hanna, "The AI/Fascism Nexus," *Mystery AI Hype Theater 3000: The Newsletter*, November 8, 2024, https://buttondown.com/maiht3k/archive/the-aifascism-nexus.