

Deconstructing AI Hype: A Playbook for Resistance

Shreya Chowdhary, Justine Zhang, and Ben Green

University of Michigan School of Information
Ann Arbor, Michigan, USA
{schowdha, tisjune, bzgreen}@umich.edu

Abstract

Contemporary AI development is accompanied by widespread industry-led hype: speculative narratives that portray AI as transformative, inevitable, and urgently necessary. Dominant accounts diagnose the problem of hype in terms of the inaccuracy of its claims, and propose increased technical and sociotechnical literacy as the response. In this paper, we provide a materialist deconstruction of AI hype, examining its political-economic functions rather than assessing its scientific validity. We show that AI hype is a problem not because it misdescribes AI, but because it facilitates AI entrenchment—the concentration of power and resources in the AI industry, rapid proliferation of AI across sectors, and lack of democratic agency over AI’s trajectory. Based on an analysis of the tactics used by AI hypers, we argue that the appropriate response to AI hype is not literacy, but contestation: undermining hypers’ legitimacy and building collective power to pursue alternative trajectories for AI. We contribute a playbook of tactics for this contestation, illustrated by ongoing movements against AI entrenchment. Although there is no simple solution for resisting AI hype and the AI entrenchment it enables, collectively these strategies can produce a meaningful challenge that shifts the balance of power over AI futures.

1 Introduction

Contemporary AI development is accompanied by widespread industry-led hype: speculative narratives that portray AI as desirable, transformative, and urgently necessary. These narratives contain sweeping claims about AI’s superhuman capabilities, capacity to drive economic growth, and potential to solve complex societal challenges (Bubeck et al. 2023; Merchant 2024; Niemeyer and Varanasi 2024). Although hype has been a persistent feature of AI discourse throughout the field’s history, the production of AI hype narratives has intensified since OpenAI’s 2022 release of ChatGPT.

This proliferation of AI hype has occurred alongside a rapid entrenchment of AI across sectors. In recent years, the AI industry has accrued immense political and economic power. The valuations of AI firms—both established companies and startups—have skyrocketed to unprecedented

heights (Montgomery 2025). AI leaders regularly convene with powerful politicians, and governments have passed AI-friendly policies, removing regulatory barriers that would constrain development while facilitating the accelerated construction of data centers (Bareis and Katzenbach 2022; Tortorelli, Stevens, and Tolockaite 2025). Meanwhile, AI has permeated across society, with businesses, universities, governments, and other institutions incorporating AI systems into their core operations (Chen, Gascó-Hernández, and Esteve 2024; Lohr 2025; Wang et al. 2024). These developments have concentrated control over AI—and the vast resources committed to it—in a handful of unaccountable firms (Cohen 2025), while producing and intensifying harms such as environmental degradation and exploitative labor conditions, with the most severe consequences for already-marginalized groups (AI Now Institute 2026; Mateescu, Nguyen, and Pinto 2026; Williams, Miceli, and Gebru 2022).

The dominant account of AI hype, common within AI ethics research and public discourse, frames hype as a problem of misrepresentation. We refer to this account as *idealist* due to its focus on the ideas embedded in hype narratives. The idealist account challenges the accuracy of claims about AI and the ways that opportunistic actors undermine scientific rigor to sell AI—for instance, labeling many AI systems that cannot work as advertised as “snake oil” (Narayanan and Kapoor 2025). Accordingly, this account suggests challenging hype by promoting technical and sociotechnical *literacy* around AI (Sloane, Danks, and Moss 2024; Slota et al. 2020; Steen et al. 2025).

In this paper, by contrast, we offer a *materialist* deconstruction of AI hype. We argue that, although enhancing AI literacy has definite value, the idealist account only minimally connects hype to the AI industry’s growing political and economic power, and therefore lacks a theory of change for altering that power structure. Thus, instead of examining whether hype claims are true or false, we examine how hype discourses support AI entrenchment, with an ultimate goal of informing strategies to interrupt that process. We answer the following questions:

1. What are the purposes and impacts of AI hype?
2. How do AI hypers build and sustain hype?
3. How should opponents of AI entrenchment resist hype?

Account of AI hype	Definition of AI hype	Effects of AI hype	Response to AI hype
Idealist account: Interrogates whether claims about AI are scientifically sound.	AI hype is an unscientific discourse that misrepresents or exaggerates AI's capabilities.	Rushed and irresponsible decision-making about AI development and deployment.	Literacy: Correct inaccurate claims through technical or sociotechnical analyses.
Materialist account: Deconstructs how hype claims affect political and economic power.	AI hype is a performative, promissory discourse that mobilizes visions of AI-driven futures to organize present action.	Coordinated action across sectors to secure the political and economic conditions that facilitate AI entrenchment.	Contestation: Discredit hypens and build counterpower among the communities that hype disempowers.

Table 1: Comparison between the dominant, idealist account of AI hype and our alternative, materialist account.

Drawing from scholarship in science, technology, and society (STS) and political economy, we show that hype is a problem not because it misdescribes AI, but because it secures the capital, institutional commitments, and political conditions that facilitate AI entrenchment. Specifically, we illustrate how AI firms and venture capitalists use hype to align institutional decision-makers (e.g., policymakers, university administrators, and corporate executives) around actions that support AI entrenchment. This coordination is critical for the AI industry, which depends on securing a range of material and ideological support—like public subsidies, favorable regulation, procurement deals, and broad-based legitimacy—from those institutional decision-makers. As AI's profitability and social benefits remain largely unproven, AI hype acts as a strategy for securing these resources on the strength of promised futures rather than demonstrated returns. AI hype shepherds public and private actors alike into making durable commitments to AI development, in turn driving society toward a future defined by AI entrenchment while foreclosing on other potential futures.

Resisting AI hype is therefore essential for pushing toward a future shaped by democratic deliberation and the public interest rather than dictated by the AI industry and aligned actors. To inform strategies for contesting hype, we provide a systematic analysis of how AI hypens build and sustain hype. We highlight two main goals of hypens: to capture credibility—rendering speculative futures believable and desirable—and to manufacture urgency—encouraging immediate action to realize them. We then describe the discursive tactics that AI hypens use to achieve these goals, drawing on examples of prior tech hypens, AI industry claims and demonstrations, and statements and actions from non-industry hypens (such as policymakers). The AI industry uses these tactics to promote AI, and in turn institutional decision-makers reproduce them to justify AI adoption in service of other institutional agendas (like attracting tax revenue or appearing innovative).

Based on our analysis of how AI hype functions, we contribute a playbook for resisting AI hype. We argue that the appropriate response to AI hype is not literacy, but *contestation*. Because hype helps to produce the conditions for entrenching AI, resistance must dismantle these conditions and create space for alternative political and economic futures. Opponents of AI entrenchment should pursue two

goals that mirror and invert the goals of hypens: undermining hypens' legitimacy and building collective agency over AI's trajectory. Drawing on the efforts of communities already trying to resist AI, we outline a set of tactics for advancing these goals. By employing these tactics to break down hype's coordinating function, opponents of AI entrenchment can wrest collective control over AI's trajectory—and ultimately shift the political and economic arrangements that make hype effective in the first place.

Our analysis of AI hype and how to resist it, in comparison to the idealist account, is summarized in Table 1.

2 The Idealist Account of AI Hype

The dominant analysis of AI hype, common among AI ethics scholars, journalists, and other critics, regards AI hype as unscientific misrepresentations or exaggerations of AI's capabilities (Fleisher 2026; LaGrandeur 2024; Markelius et al. 2024; Narayanan and Kapoor 2025; Slota et al. 2020; Thais 2024). This account seeks to litigate whether claims about AI are accurate, critiquing discursive techniques (like anthropomorphism) that distort expectations about AI's capabilities and impacts (Kotliar 2025; Markelius et al. 2024; Steen et al. 2025). Because of its focus on the ideas contained in AI hype, we refer to this account as idealist.

Under the idealist account, the primary consequences of AI hype are overblown and unrealistic expectations about AI (Fleisher 2026; Thais 2024). These inflated expectations generate unwarranted attention to and investment in AI systems, leading to ill-informed development and deployment decisions that can produce immense societal harm (LaGrandeur 2024; Markelius et al. 2024; Narayanan and Kapoor 2025; Sloane, Danks, and Moss 2024; Slota et al. 2020).

The idealist account suggests countering hype by increasing literacy: correcting inaccurate claims and making people better-informed about AI's capabilities. Proposals include increasing public awareness of hype-filled narratives (Floridi 2024; Markelius et al. 2024), circulating more accurate visual representations of AI (Steen et al. 2025), explaining AI's sociotechnical dimensions more clearly (Sloane, Danks, and Moss 2024; Slota et al. 2020; Steen et al. 2025), and improving how AI experts and developers communicate about AI's capabilities (Fleisher 2026; Thais 2024).

Although the literacy response can offer useful checks on

hypers' misrepresentations of AI, it does not interrogate the structural factors driving AI hype. As such, it offers no theory of change for interrupting AI entrenchment beyond correcting hypers' misrepresentations—a tactic that has often failed to dissuade institutions from adopting and supporting AI. In this way, the literacy response echoes the dominant paradigm of ethical and responsible AI, which focuses on correcting errors, mitigating individual-level harms, and improving institutional decision-making. Critical scholars have demonstrated that this tendency often leads ethical AI interventions to legitimize and reproduce the problems they seek to critique (Green 2021; Phan et al. 2022; Stark, Greene, and Hoffmann 2020). These and other scholars call instead for a materialist examination of AI, foregrounding how AI development and discourse are intertwined with economic interests and political power (Sadowski 2025; Vertesi et al. 2026; Vrabčič Dežman 2024; Zhang, Blodgett, and Markl 2026).

We extend this materialist orientation toward AI hype, examining hype's role in generating the political and economic conditions behind AI entrenchment. We contribute an account of the tactics that hypers use to accomplish this outcome, as well as the strategies for resistance that those tactics imply.

3 The Purposes and Impacts of AI Hype

To pursue a materialist analysis of AI hype, we draw from scholarship in STS and political economy on technological development in capitalist societies (Birch 2013, 2025; Sadowski 2019).¹ This scholarship shows how capitalist imperatives like accumulation and commercialization increasingly shape technological development (Baykurt 2025; Gansky 2025; Young, Katell, and Krafft 2022). Researchers have also traced how large tech firms accumulate power by controlling digital infrastructures and platforms, dynamics that are intensified under the conditions of AI development (Sadowski 2020; Steinhoff 2024; Vertesi et al. 2026; West 2019).

Our approach is particularly informed by Birch (2025)'s conceptualization of *constructivist political economy*, which traces how the economic value and material conditions of technology development are actively constructed through discourses and expectations. Narratives about emerging technologies arise in response to the material imperatives undergirding their development; in turn, these narratives shape the material conditions of development. Thus, Birch emphasizes that trajectories of technological development and adoption are neither natural nor inevitable, but actively produced. We apply this constructivist lens to hype, examining how narratives about AI are at once motivated by and produce the material conditions of AI development.

3.1 Defining Tech Hype

Hype, in its broadest sense, is an industry-agnostic mechanism for circulating value, attention, and expectation within

contemporary capitalism. Analyzing hype within pop culture, Powers (2012) defines hype as “a state of anticipation generated through the circulation of promotion.” Tech hype is the instantiation of this phenomenon within the tech industry, which we define as “a performative, promissory discourse that mobilizes visions of technological futures to organize action in the present.”

Hype is a central feature of capitalist technology development because many emerging technologies are highly uncertain and cost-intensive. For instance, biotech ventures require significant, long-term financial investment due to the expenses associated with creating new drugs (Rajan 2006). Firms, governments, and investors must decide where to allocate resources, what infrastructures to build, and which technological trajectories to pursue long before technologies are stable, profitable, or empirically validated (Brown and Michael 2003). Present-day regimes of financialization have intensified these dynamics, with speculative futures becoming increasingly central to capital accumulation (Bear 2020; Beckert 2016; Tsing 2000). For instance, venture capital markets depend heavily on expectations of exponential future growth, fueled by speculative storytelling, with investors seeking out “unicorn” firms that promise extraordinary future returns (Birch 2017, 2023; Klonowski 2018).

Under these conditions, hype serves as a discursive strategy that tech leaders (e.g., startup founders and investors) can use to coordinate particular audiences around actions crucial to the economic success of emerging technologies. These actions include marshaling resources, altering institutional priorities, and providing investments in service of their ventures (Brown and Michael 2003). Acting as hypers, tech leaders leverage speculative narratives about emerging technologies to pursue two interrelated goals. First, they work to *capture credibility*, constructing compelling visions of future transformation and establishing themselves as legitimate stewards of those futures (Birch 2017; Kampmann 2024; Törnberg 2025). Second, they work to *manufacture urgency*, framing immediate action to realize these futures as necessary and irrefutable, such that audiences feel they have no other choice (Beckert 2016). Credibility makes action appear worthwhile despite uncertainty; urgency compresses that action into the present.

In this sense, tech hype is productive. It works to bring hypers' desired future into existence by bolstering emerging markets, defining research agendas, justifying infrastructure development, and directing capital (Faxon, Fields, and Wainwright 2024; Geiger and Gross 2017). Hype is thus “successful” if it shapes policies, investments, and markets in ways that benefit hypers. With this in mind, tech hype cannot be wholly identified based on the presence of inaccurate or exaggerated claims, as idealist accounts contend. Even if promises about technological futures come true, those promises would still be hype if they helped mobilize the political and economic actions needed to make that happen.

In sum, our materialist definition centers hype's function rather than its accuracy. For the idealist account, hype dissolves if its claims are disproven; in our account, hype dissolves only when it can no longer coordinate flows of material resources.

¹Broadly, political economy examines how societal resources (e.g., labor and money) are allocated, how ideas and institutions shape those allocations, and how those allocations shape distributions of power (Birch 2025; Liu 2024; Verdegem 2023).

3.2 Situating AI Hype

AI hype is the manifestation of tech hype that has emerged under the material conditions of contemporary AI development. Specifically, we define AI hype as “a performative, promissory discourse that mobilizes visions of AI-driven futures to secure the political and economic conditions that facilitate AI entrenchment.”

The AI industry is heavily dependent on hype because its ventures are particularly cost-intensive and uncertain. Developing and deploying AI models requires immense quantities of data, chips, energy, and labor, with AI firms and investors committing enormous amounts of capital toward infrastructure like data centers (Birch and Cochrane 2024; Narayan 2026; Verdegem 2024). At the same time, AI firms have unproven business models, and there is widespread skepticism about whether revenue streams will become large enough, quickly enough, to justify those investments (Merchant 2024; Zitron 2026). To achieve sufficient investment and revenue, AI firms require significant and immediate support from outside the tech industry: for instance, subsidies and favorable regulations from governments, adoption commitments from schools and companies, and legitimacy in press coverage. AI hype is a central mechanism by which the AI industry garners this support.

The goal of AI hype is not merely to persuade outsiders, but to enlist them as producers of hype in their own right. We therefore distinguish between two types of AI hypers based on their relationship to AI entrenchment. *Primary hypers* are people whose core interests are tied directly to AI entrenchment; producing hype is existential for these actors. These actors include tech CEOs, venture capitalists, and executives across allied industries (e.g., chip manufacturing and data center construction). They are the primary producers of AI hype, and use discursive vehicles like advertisements, speeches, and op-eds to court audiences—particularly institutional decision-makers—who can provide the AI industry with needed support.

Secondary hypers are people who reproduce AI hype to serve their own institutional interests that are not fundamentally about advancing AI entrenchment; for them, producing hype is instrumental. These actors are institutional decision-makers like policymakers, corporate managers, and university administrators who reproduce AI hype through their statements and actions (such as driving institutional adoption of AI). Unlike primary hypers, secondary hypers produce hype to justify agendas that are external to—albeit highly advantageous for—the AI industry. For instance, policymakers hope to attract tax revenue by supporting data center construction, university presidents hope to garner prestige and tuition revenue by adopting an AI-centric curriculum, and corporate managers hope to appear innovative by embracing AI.

Because institutional decision-makers control conditions that are essential to industry success, primary hypers aim to bring them on board as secondary hypers who take up AI entrenchment as a shared priority. They do this by presenting AI adoption as a route to achieving institutional decision-makers’ existing goals, like attracting revenue or securing prestige, and by supplying the promissory language that

decision-makers can use to justify their actions. As a result, there is a strong alignment between primary and secondary hypers, as each actor reproduces hype in service of their own agenda. This alignment makes hype durable despite controversy or contrary evidence, and renders it essential for facilitating, legitimizing, and stabilizing AI entrenchment.

4 How AI Hypers Build and Sustain Hype

In this section, we uncover the discursive playbook of AI hypers, describing the tactics they use to advance their goals of capturing credibility and manufacturing urgency. We identified these tactics through an interpretive analysis of contemporary AI hype discourses—as found in sites such as press releases, policies, and op-eds—as well as scholarship documenting the tactics used by the hypers of other technologies. For each tactic, we present examples of its use in prior tech hypes, show how primary hypers apply it to court key audiences, and trace how secondary hypers reproduce primary hypers’ language to justify institutional agendas aligned with AI entrenchment.

In sum, hypers 1) distort empirical knowledge by suppressing the empirical limits of present and past systems, 2) align actors around shared priorities by embedding technologies within shared imaginaries that justify immediate action, and 3) delimit trajectories of development by foreclosing alternatives to their vision even in the face of criticism. Table 2 presents an overview of these tactics.

4.1 Capture Credibility

Hypers build support for speculative technological futures by making themselves and those futures appear credible, desirable, and worthy of investment. They seek not only to be believed by their targeted audiences, but to shape which futures appear imaginable, desirable, and actionable. They do this through several discursive practices: enchanting audiences, mobilizing dominant imaginaries, and adapting claims to deflect criticism.

Enchant audiences. Hypers draw on technosolutionist narratives, spectacle, and ambiguity to enchant audiences, building interest in new technologies while discouraging critical scrutiny. For instance, edtech hypers framed the technology as a solution to systemic education problems in the Global South (Chan 2019). They paired this narrative with spectacular demonstrations, like futuristic displays, to enchant audiences, while dismissing calls for evaluation of early deployments as “a waste of resources” (Chan 2019).

The AI industry uses similar methods. It employs bold language framing AI as a “general-purpose solution” to diverse societal challenges, from economic growth to scientific discovery (Altman 2023; Christensen 2025; Katzenbach 2021). These sweeping technosolutionist narratives offer simplified and compelling visions of AI systems as exciting and socially necessary. The industry reinforces these claims through staged demonstrations—such as AI systems passing the bar exam or winning art competitions—that render progress tangible and emotionally resonant (Hong 2024).

At the same time, enchantment depends on preserving ambiguity and discouraging critical scrutiny. The indus-

	Goal 1: Capture credibility Make speculative AI futures appear believable, desirable, and worthy of support.	Goal 2: Manufacture urgency Render immediate action to realize speculative AI futures necessary and irrefutable.
Tactic 1: Distort empirical knowledge Selectively frame technological realities in ways that obscure uncertainty, controversy, failure, or material contingencies.	<i>Enchant audiences</i> Use spectacle and technosolutionist framing to evoke excitement and anticipation in audiences.	<i>Mischaracterize the past</i> Selectively reference history to paint an idealized picture of the past.
Tactic 2: Align actors Coordinate institutions and social actors around shared, immediate priorities for AI development.	<i>Mobilize dominant imaginaries</i> Draw on dominant cultural frames to make speculative technological futures appear desirable and actionable.	<i>Cultivate fear of falling behind</i> Manufacture a need to act immediately by emphasizing the potential costs of delay or missed opportunities.
Tactic 3: Delimit trajectories Sustain institutional momentum and infrastructural commitment toward particular trajectories of technological development.	<i>Adapt claims to deflect criticism</i> Dynamically recalibrate claims to subsume criticism and disappointment and sustain credibility.	<i>Present the future as inevitable</i> Present promised technological futures as inevitable, using deterministic rhetoric to foreclose alternatives.

Table 2: The goals, tactics, and discursive practices of AI hype.

try leaves AI's capabilities ambiguous by anthropomorphizing systems and describing AI as "magical" (Campolo and Crawford 2020; Elish and boyd 2018; Hong 2024; LaCroix, Mallory, and Luccioni 2026). It obscures key aspects of system design, like non-consensual data extraction and the workers whose labor powers AI systems behind the scenes (Nagy and Neff 2024; Rowe 2023; Taylor 2018). The industry also frequently releases AI systems in persistent "beta" forms, sidestepping failures or limits by sustaining the promise of improved capabilities on the horizon (Bourne 2024; Grill 2024). Together, these practices encourage audiences to treat AI as transformative, even in the absence of stable, well-defined, or empirically verifiable capabilities.

Once hypers' speculative futures appear credible and socially desirable, institutions become more willing to commit resources to realizing them. For instance, various governments translate AI industry framings into policies that advance AI development through funding allocations, infrastructure support, and workforce development initiatives. As one notable example, the UK's National AI strategy commits billions in funding to "mainlining AI into the veins of this enterprising nation—revolutionising our public services" (UK Government 2025). Many other governments have also accelerated AI adoption—touting benefits such as improving public safety and making bureaucratic processes more efficient—while ignoring the many challenges associated with using AI to reform public services (Green 2025).

Mobilize dominant imaginaries. Hypers embed emerging technologies within dominant imaginaries (i.e., cultural frames which provide familiar interpretive scaffolding), rendering their speculative futures immediately legible and actionable to their audiences. For example, tech companies hyping smart city initiatives attached visions of smartness to prevailing imaginaries of economic transformation, framing smart city tech as uniquely capable of solving enduring urban crises (Sadowski and Bendor 2019). Public officials then

recruited these frames to legitimize new deployments and public-private partnerships (Baykurt and Raetzsch 2020).

Today, the AI industry mobilizes longstanding imaginaries such as geopolitical competition and endless economic growth to shore up support for AI development. In the case of geopolitics, American tech industry CEOs regularly invoke the specter of China to push for greater coordination between the private and public sectors in advancing AI (Kak and West 2023; Schmidt 2020).

In turn, governments mobilize the imaginary of a global "AI race" to justify material commitments to advance AI, such as accelerated data center construction, investments in startups, and deregulation. For instance, the 2026 White House AI policy framework argues that "[s]tates should not be permitted to regulate AI development, because it is an inherently interstate phenomenon with key foreign policy and national security implications" (The White House 2026). State officials have invoked the same imaginary to justify local commitments, as when Michigan Governor Gretchen Whitmer defended an unpopular data center project spearheaded by OpenAI: "There is a shortage of global computing capacity, and China is moving aggressively to build their own centers and get a leg up on America [...] This is a matter of national security and economic competitiveness" (Office of Governor Gretchen Whitmer 2025). By attaching AI development to imaginaries of geopolitical competition and national survival, AI hypers coordinate bipartisan support for AI entrenchment across all levels of government.

Adapt claims to deflect criticism. As emerging technologies generate disappointment, public backlash, or visible harms, hypers adapt their claims, reframing critiques as problems manageable through better governance or technical expertise rather than grounds for slowing or rejecting technological expansion. By dynamically recalibrating their claims, hypers sustain institutional confidence and continued investment. For example, after the 2022 collapse of the

cryptocurrency exchange FTX made cryptocurrency synonymous with fraud and speculation, blockchain proponents recast the collapse as “a failure of centralization or fraudulent actors” rather than as a failure of the underlying technology (Koziol 2023). This rearticulation preserved the industry’s legitimacy and its capacity to attract continued investment even as its most visible instantiation failed.

In the face of public discourse about AI being overhyped and harmful, AI hypers have similarly sought to retain their credibility in ways that facilitate continued development and deployment. Most notably, tech firms invoke notions of “responsible,” “ethical,” “safe,” or “democratized” AI to signal responsiveness to criticism without sacrificing core business goals (Ahmed et al. 2024; Green 2021; Vertesi et al. 2026; Young, Madaio, and DiSalvo 2026). For example, Microsoft frames responsible AI governance not as a constraint on innovation but as an enabler of continued acceleration (Fleck 2024). Similarly, firms use “safety” discourse to address fears that powerful AI systems may become dangerous and uncontrollable. Instead of treating these risks as reasons to slow AI development, firms present scaling and “iterative deployment” as the solution (De, Lima, and Zou 2026; OpenAI 2025). Through these revised vocabularies, the industry incorporates critique into hype narratives, stabilizing its legitimacy without meaningfully changing its practices.

These revised narratives have diffused across sectors, becoming a coordinating logic through which other institutions legitimize AI adoption and investment. Universities frequently frame partnerships with AI firms as oriented around “human-centered,” “ethical,” or “responsible” AI development, presenting their collaborations as necessary for ensuring that AI aligns with social values and public benefit. For example, Northeastern University frames its partnership with Anthropic as an effort to “[bring] responsible AI to higher education” (Nordman 2025). Public officials have similarly responded to controversy around AI entrenchment in the public sector by framing continuous and iterative deployment as the solution to ethical concerns. After New York City’s AI-powered chatbot faced criticism for providing inaccurate and legally questionable advice, then-Mayor Eric Adams defended continued deployment by arguing that “any time you use technology, you need to put it into the real environment to iron out the kinks” (Allen 2024).

4.2 Manufacture Urgency

While AI hypers can use the previous tactics to establish themselves as credible and AI as desirable, to sustain material investment in AI they must also position AI development as urgent and difficult to refuse. Hypers seek to compress the temporal and deliberative space necessary for democratic decision-making, presenting investment in AI as the only path forward while diminishing the agency of institutions and publics to refuse, slow, or redirect technological development. They do so through several discursive practices: mischaracterizing the past, cultivating a fear of falling behind, and presenting their futures as inevitable.

Mischaracterize the past. One way hypers accelerate adoption is by obscuring historical precedents that might

generate skepticism toward their claims. By hiding the unrealized promises of past hypes and selectively invoking favorable historical analogues, hypers legitimize present-day investment and development (Borup et al. 2006). When promoting educational technologies, venture capital firms recycled narratives from an earlier hype around teaching machines, while excluding the contentious debates about their costs to students, teachers, and schools which deflated the teaching machine hype (Chan 2019). By suppressing these histories, venture capitalists could reintroduce similar claims without confronting prior criticism, thereby facilitating a new wave of investment and adoption.

AI hypers reproduce this pattern, selectively invoking historical analogues to legitimize present-day action while obscuring histories of overpromising and societal harm. For example, Microsoft compares the build-out of data centers to earlier large-scale infrastructure projects such as railroads, the electrical grid, and highways (Smith 2026). This analogy casts contemporary controversies over data centers—such as environmental impact, land use, and energy consumption—as temporary growing pains analogous to those associated with earlier infrastructure systems. By presenting these historical cases as having ultimately yielded broadly accepted and beneficial outcomes, Microsoft constructs a narrative in which AI infrastructure expansion appears both necessary and worth the cost.

By suppressing histories of disappointment and contestation, the AI industry primes other institutions to perceive investment in and adoption of AI as an unprecedented opportunity. For instance, school districts and administrators reproduce earlier edtech hype narratives when promoting AI in schools, while rarely acknowledging these historical precedents and their associated controversies. Recent AI-focused schools have promoted AI as a mechanism for “personaliz[ing] learning” for students (Jones 2026), echoing earlier promises surrounding edtech as a solution to structural problems in education (Chan 2019). By obscuring these historical continuities and prior controversies, contemporary AI discourses present recycled promises as novel breakthroughs warranting renewed institutional investment and adoption.

Cultivate fear of falling behind. Hypers create temporal pressure to act, translating their speculative futures into present-day imperatives. In past cases of tech hype, hypers cultivated a fear of falling behind to encourage rapid adoption. For example, tech companies promoted smart city initiatives as necessary for remaining globally competitive and economically innovative (Baykurt and Raetzsch 2020; Sadowski and Bendor 2019).

In addition to framing AI as an unprecedented opportunity, hypers frame AI development and adoption as time-sensitive. Industry leaders promise that artificial general intelligence (AGI) is perpetually around the corner, suggesting that any slowdowns will delay transformative economic growth and scientific breakthroughs (Altman 2023; Garibay-Petersen, Lorimer, and Menzat 2025). They also frame rapid adoption as necessary to remain productive and competitive, at levels ranging from individual companies to entire nations. For example, George Osborne, head of the “OpenAI

for Countries” initiative and a former UK Chancellor, stated that nations have no choice but to adopt AI: “don’t be left behind, because then you will be a weaker nation, a poorer nation, a nation whose workforce will be less willing to stay put” (Booth 2026).

The industry’s narratives prompt institutions to support AI infrastructure and adopt AI systems—not necessarily because the benefits are empirically established, but because of fear that not doing so will lead to economic disadvantage or institutional irrelevance. For example, policymakers translate urgency narratives into policy agendas advancing rapid AI development and infrastructure expansion (The White House 2026; UK Government 2025). Similarly, managers and executives at non-AI firms pressure employees to maximize AI use in the workplace, producing a culture of “token-maxxing,” in which workers compete to top organizational leaderboards of AI use (Roose 2026). These actions generate adoption and engagement metrics that AI firms can leverage as proof of growing market penetration and AI’s economic indispensability, regardless of whether the underlying usage actually produces value for the adopting organizations.

Present the future as inevitable. Hypers frame technological development as an inevitable process to which institutions and people must adapt, rather than as a political trajectory open to collective contestation and democratic control. By directing attention away from debates over whether a technology should be developed at all and toward discussions of how it should be implemented and adopted, hypers narrow the range of imaginable alternatives. This deterministic orientation is reflected in the Gartner Hype Cycle, a popular visualization of hype produced by the consulting firm Gartner, which depicts technological development as progressing toward a singular endpoint: “enlightened” integration into society (Fenn and Raskino 2008). Though this representation does not accurately describe how technology development and adoption actually occur (Mullany 2016), it organizes expectations in the present by signaling that a technology’s eventual success is assured, helping produce the conditions that make projected futures more likely to materialize (Bareis, Roßmann, and Bordignon 2023; Geiger and Gross 2017).

AI hypers similarly frame AI as inevitable, positioning it as a foundational part of future economic and social life. Industry leaders frequently portray AI adoption not as a strategic choice, but as a necessary condition of survival. When introducing Copilot, Microsoft presented AI-assisted work as the essential new mode of productivity: “Copilot will fundamentally change how people work with AI and how AI works with people. As with any new pattern of work, there’s a learning curve—but those who embrace this new way of working will quickly gain an edge” (Spataro 2023). Such claims invert the causal relationship between present and future: rather than present action shaping the future, the future is already determined, and actors can only adapt to it (Cheney-Lippold 2025).

Policymakers have translated the AI industry’s inevitability discourse into institutional commitments that further solidify AI entrenchment. National AI strategies frequently

frame AI as a historical transformation requiring immediate adaptation and investment (Bareis and Katzenbach 2022; Bennani-Taylor 2024). Policymakers also treat industry visions of the future as a foregone conclusion. For example, Michigan Governor Gretchen Whitmer presents data centers as an unavoidable component of the future, asserting: “The question isn’t whether [data centers] will be built, but rather: Can Michigan benefit from these jobs and build data centers in a smarter way while upholding our strong environmental laws to protect our precious natural resources?” (Office of Governor Gretchen Whitmer 2025). This framing forecloses fundamental questions about AI’s trajectory, treating it as already settled and directing attention toward lower-level questions of implementation and governance. It also discourages and delegitimizes critics, suggesting that efforts to envision or promote alternative futures are pointless.

5 A Playbook for Resisting AI Hype and Entrenchment

As the prior sections detailed, AI hype is a performative discourse that coordinates diverse actors around securing the political and economic conditions for AI entrenchment. Therefore, efforts to resist AI entrenchment must confront hype as a key piece of its machinery.

Our account of AI hype suggests two specific goals for resistance that mirror and directly challenge the goals of AI hypers. Where hypers attempt to capture credibility, opponents should *undermine hypers’ legitimacy*, dislodging hypers’ position as singularly credible stewards of desirable technological futures. Where hypers attempt to manufacture urgency, opponents should *build collective political agency*, asserting deliberative control over decisions that urgency forecloses. Enacting these goals disrupts the conditions of possibility for AI entrenchment that hype secures, while creating conditions for democratically pursuing alternative political, economic, and technological futures. We refer to this approach to resisting hype as *contestation*.

Contestation provides a stark contrast to the dominant response to AI hype, which we described in Section 2 as literacy. Literacy efforts can help break down discursive practices like enchantment by reducing perceptions of AI as magical (Tully, Longoni, and Appel 2025). But because literacy engages with what hype says rather than what hype does, it leaves hype’s coordinating function intact. Secondary hypers reproduce hype claims because doing so serves their agendas, not necessarily because they are persuaded by those claims. Thus, the alignment that hype produces can survive corrected claims. A university administrator convinced that AI is overhyped still faces pressures around enrollment and competition with peer institutions that make an AI partnership attractive. In fact, by aiming to prepare people for an AI-centric future, literacy efforts often concede the narrative of inevitability manufactured by AI hypers (Floridi 2024; Narayanan and Kapoor 2025; Slota et al. 2020). Literacy has consequently become another concept that hypers can invoke to deflect criticism—as exemplified by a recent industry-backed bill funding AI literacy initiatives (Cole 2026).

	Goal 1: Undermine hypers' legitimacy Dislodge hypers' position as singularly credible stewards of technological futures.	Goal 2: Build collective political agency Grow collective power to make choices about which futures to pursue.
Tactic 1: Expose omissions and priorities Reveal the labor relations, infrastructural dependencies, and political-economic interests obscured by hype narratives.	<i>Disenchant audiences</i> Use counter-spectacle to weaken the cultural legitimacy of AI firms.	<i>Recover suppressed histories</i> Recover histories that challenge dominant narratives of technological progress and legitimize resistance to technology.
Tactic 2: Disrupt institutional alignments Disrupt institutional consensus around advancing AI as a central and immediate priority.	<i>Advance alternative imaginaries</i> Draw on popular imaginaries marginalized by hypers to promote more critical actions toward AI.	<i>Create friction to AI expansion</i> Open deliberation over whether AI projects should proceed by blocking the investments and contracts that bind institutions to them.
Tactic 3: Open alternative futures Contest attempts to naturalize AI-centered futures and foreclose broader political-economic alternatives.	<i>Call out self-serving rearticulations</i> Deconstruct how adapted claims are intended to stabilize the industry's political and economic power.	<i>Expand political imagination</i> Reopen the questions hype forecloses, organizing toward collective control over AI's trajectory.

Table 3: The goals, tactics, and strategies for resisting AI hype via contestation.

The difference between the literacy and contestation responses stems from the idealist and materialist accounts of why hype is a problem. Literacy, following the idealist account, corrects inaccurate claims so that people can make more informed decisions about AI. Contestation, following the materialist account, targets hype's capacity to coordinate actors around AI entrenchment. Literacy efforts are thus "reformist reforms": reforms that seek to improve outcomes without altering existing distributions of power (Gorz 1967). Contestation instead pursues "non-reformist reforms": reforms that seek to redistribute power and build capacity for democratic control across society (Akbar 2020; Gorz 1967). Specifically, contestation involves building counterpower among the groups on whom AI is imposed and who are largely excluded from the decisions driving AI development and deployment. We refer to these groups—the workers, communities, and broader publics that hype disempowers—as *opponents* of AI entrenchment.

In this section, we outline how opponents can challenge AI hype, building on our analysis of hypers' tactics and drawing on ongoing efforts to resist AI entrenchment. Table 3 outlines a playbook for contestation that counters the hype tactics identified in Section 4 and Table 2: opponents should 1) expose the omissions and political-economic priorities behind hype narratives, 2) disrupt institutional alignments, and 3) open alternative futures. While these tactics are not silver bullets and none are sufficient on their own, together they constitute a meaningful challenge to the distribution of power over AI futures.

5.1 Undermine Hypers' Legitimacy

Opponents should aim to undermine hypers' position as credible stewards of technological futures by uncovering the self-serving falsehoods and narratives that hypers use to establish their legitimacy. In doing so, opponents can re-open controversy around AI that hypers close.

Disenchant audiences. Because hypers use spectacle and strategic omission to enchant audiences, opponents should disenchant audiences by using counter-spectacle and exposing what enchantment conceals. Counter-spectacle can puncture hypers' legitimacy and normalize dissent. For example, New York City subway ads for the AI companion startup Friend were repeatedly defaced with graffiti containing messages such as "Stop normalizing this" and "Real friends don't surveil you" (Witte 2025). Rather than fact-checking the company's claims about AI companionship, the graffiti refused the premise that the ads worked to establish: that AI companionship is normal and desirable. In doing so, the graffiti delegitimized the company's broader venture and seeded a boycott of the startup (Mogul 2026). Opponents can also transform moments intended to reinforce institutional consensus into occasions for collective contestation. For example, when college students booed commencement speakers promoting AI as an unquestioned social good, they disrupted ritualized moments of celebration and made dissent toward AI publicly visible (Singh 2026).

Critical reporting and research can also disenchant audiences by exposing the social and material conditions that enchantment obscures. While hypers describe AI as magical and leave its capabilities ambiguous, sustained investigations render its operations concrete. For example, the Data Workers' Inquiry supports data workers across nine countries in researching and documenting their labor, rendering visible the content moderation, data annotation, and chatbot impersonation that enchantment conceals (Miceli 2026). Because workers conduct the research themselves, the project builds organizing capacity alongside public knowledge—thus feeding directly into unionization and association efforts. As another example, journalists revealed that Meta allows its AI companions to engage children in romantic exchanges, prompting a Senate investigation into whether the company had misled the public about its safeguards (Hor-

witz 2025). More broadly, journalists have documented the failed promises and concrete harms of AI rollouts, contributing to growing skepticism about AI development and adoption (Lapowsky 2026; Ray 2026).

Advance alternative imaginaries. Hypers embed AI within dominant narratives (such as a race against China) to make expansion feel like common sense. Opponents can counter by centering alternative imaginaries: cultural frames marginalized by hypers that render different decisions about AI legible and actionable. When these imaginaries reflect the core interests of constituencies that hype disempowers, they can challenge institutional alignments around AI entrenchment and redirect political momentum to other outcomes.

Imaginaries of affordability, opposition to corporate power, and shared access to public resources have punctured the image of common sense and necessity that hypers propagate. At protests against data centers, signs with phrases such as “water for people, not billionaires” position data centers as private claims on public resources made by a wealthy few, leading to higher utility bills and drawn-down aquifers for everyone else (Taylor and Levin 2026). These narratives have spurred groups around the world—spanning partisan and geographic divides—to mount grassroots efforts against the buildout of AI infrastructure. Similarly, when opposing efforts by public school districts to adopt AI, coalitions of teachers and families have invoked ideas about schooling that center relationality over efficiency, insisting that quality education depends on small classes with empowered teachers. These groups have challenged the practice of giving tech companies so much influence over children’s education while demanding that the resources spent on AI be redirected toward reducing class sizes and increasing student support (Fight Back! News Staff 2026; Singer 2026).

Call out self-serving rearticulations. Because hypers continually adapt their claims to defuse criticism, opponents should call out those rearticulations as self-serving efforts to stabilize the industry’s political and economic power. Doing so builds a record that demonstrates hypers’ untrustworthiness, diminishing the effectiveness of successive rebrands. For example, the Kairos Fellowship’s report on Big Tech’s “AI climate solutions” deconstructs corporate claims that AI infrastructure can be environmentally responsible and even useful for addressing climate change (Kairos Fellowship 2026). The report exposes how the industry deploys these claims to distract from how its behavior exacerbates the climate crisis, inoculating readers against the industry’s revised vocabularies. The report pairs that analysis with guidance on specific actions readers can take, such as campaigning against a data center in their area, to achieve genuine victories rather than concessions in name only.

Similarly, writers have uncovered how corporate claims about “safety” are motivated less by a desire to protect the public good than a desire to promote the power of new models, which in turn attracts further investment (Marcus 2026; Merchant 2024, 2026). Such accounts puncture AI firms’ self-mythologizing as responsible stewards of technology, refocusing public attention on contesting the investments underpinning their power.

5.2 Build Collective Political Agency

In addition to undermining hypers’ legitimacy, opponents should work to build the collective political agency that hype’s urgency forecloses—growing the power of communities, workers, and publics to shape whether and how AI expansion proceeds.

Recover suppressed histories. Hypers manufacture urgency in part by mischaracterizing past instances of hype, so that each present moment feels unprecedented and impossible to pause. Opponents can reverse this erasure by recovering what hype suppresses: histories of AI’s own repeated failures and of collective resistance against technologies.

By recovering the history of AI’s repeated failures, opponents can puncture the aura of unprecedented breakthrough on which enchantment depends. AI researchers and developers have made many previous proclamations about AI capabilities that, in retrospect, were clearly incorrect, like proclaiming “AI springs” that ended in “winters,” rather than in periods of flourishing technological advancement (Hawkins 2024; Williams 2024). Recovering these histories shows that AI development has never proceeded linearly and establishes the grounds for caution and critique that hypers’ erasures work to close off. Anecdotally, one of this paper’s authors has found that presenting students with historical examples of AI hype is effective in prompting them to approach contemporary instances of AI hype more skeptically.

Another recent effort with similar effects is the embrace of Luddism. Though Luddism has often been framed as irrational opposition to technological progress, recent scholarship and organizing emphasize that the 19th-century Luddite movement drew on trenchant political and economic analyses to challenge labor exploitation and the concentration of power associated with industrialization (Merchant 2023; Sadowski 2025). Present-day efforts like “Luddite Clubs” reclaim the label of Luddite as a legitimate political position and create spaces for collective technological refusal (Mathews 2026; Merchant 2025). Recovering this history establishes that communities and workers have successfully contested the trajectory of technological change in the past and embeds current fights within broader democratic projects.

Create friction to AI expansion. Hypers use urgency to ensure that commitments to AI entrenchment are made before scrutiny can occur. As a counter, opponents should create material friction to AI expansion. This friction can involve challenges that directly impede specific AI projects, such as protests, boycotts, and ballot initiatives. Friction can also involve forcing hypers’ claims into the deliberative arenas they were designed to evade. Pushing for moratoria, permitting hearings, and substantive investigations creates time and space for collective assessment of whether an AI-related project is desirable, and under what conditions it can proceed, if at all. More fundamentally, these introductions of friction starve the apparatus of hype itself. Every stalled project breaks the chain of wins that hypers need to continuously regenerate confidence, and instead becomes evidence that challenges the industry’s projected aura of invincibility.

Opponents should particularly target instances where an institutional consensus around advancing AI is becoming

materially reified through investments and contracts. Here, we highlight how coalitions can convert their centering of alternative imaginaries into material friction. The cross-partisan organizing against data centers has translated into several hundred billion dollars' worth of blocked projects, as well as moratoria, permitting delays, and other measures that slow construction (Data Center Watch 2026; Gold and Daher 2026). Additionally, coalitions of families and teachers have blocked numerous AI-centric school initiatives, including a proposed AI-themed high school in New York City (Kilander 2026; Singer 2026).

Across these cases, the imaginative work of contesting the desirability of AI fostered the organizing capacity to interrupt material investments in AI. Moreover, beyond the specific project being opposed, such efforts empower other communities to engage in similar fights and demonstrate to institutional leaders that support for AI will be met with opposition.

Expand political imagination. Hypers use inevitability discourses to constrain political imagination: narrowing the question of AI expansion from whether to how, thus foreclosing visions of the future. In response, opponents should expand political imagination: reopening the questions that hype forecloses and making different futures both thinkable and realizable. Rather than simply negotiating the conditions of an already-assumed deployment, opponents should challenge existing power structures and organize toward arrangements where the resources and infrastructure on which AI expansion depends are collectively controlled (May, Paris, and Sargent 2026).

Recent labor organizing efforts demonstrate this tactic in action. For example, the American Association of University Professors (AAUP) has been fighting for shared governance structures through which students, faculty, and staff determine whether and under what conditions AI enters their workplaces (Paris et al. 2025). These organizing efforts treat opposition to AI not merely as a matter of forcing administrators to reconsider AI adoption, but as an opening for reshaping universities more broadly—for example, by contesting their entrenchment into finance capital and the military-industrial complex (Zhang, Chowdhary, and Kim 2026).

The National Nurses United (NNU) union presents another example of refusing AI adoption rather than negotiating its conditions. Like the AAUP, NNU has treated opposition to AI as connected to broader critiques of political-economic structures, opposing how the hospital industry cooperates with Silicon Valley and Wall Street to the detriment of patient health (National Nurses United 2024). Furthermore, nurses' opposition to AI has expanded beyond its application in hospitals to its broader public health effects. In North Carolina, NNU called for a statewide data center moratorium on the grounds that data centers harm the patients that nurses treat (National Nurses United 2026). By demanding that the state's AI buildout halt until there are strong protections in place, NNU asserted that the trajectory of AI development should be determined by the people who bear its costs, while also demonstrating how solidarity across sectors can help push for that future.

6 Conclusion

Taking AI hype seriously as a material force—rather than a matter of incorrect statements—reveals a new and more effective roadmap for countering hype. If hype were merely a matter of exaggerated claims about AI, better literacy would be enough. Instead, though, hype is a discursive mechanism for organizing the political and economic conditions that facilitate AI entrenchment. As a result, resisting hype requires contestation: disorganizing institutional alignments and building collective power to advance alternative futures.

As we describe across Section 5, groups who bear the costs of AI entrenchment are already pursuing many of the actions that we recommend. Nonetheless, such efforts to challenge AI hype and entrenchment face a long uphill battle. Capital expenditure and elite support for AI remain high despite widespread public disapproval, and AI hypers hold much more political and economic power than their opponents. In the face of this structural disadvantage, as is the case for all counterhegemonic struggle, movements that oppose AI entrenchment must aggregate local victories, sustain pressure long enough to shift entrenched arrangements, and withstand efforts to repress resistance (Boguslaw 2026; Piven and Cloward 1977). Movements must also translate organized pushback into durable structural reforms that undermine the incentives driving hype, such as regulation prohibiting the venture capital speculation and circular financing arrangements that drive AI investments and business models (Ramzanali 2026).

Future research can provide critical insights to empower resistance against AI hype. One direction is to further study the incentives, goals, and mindsets of institutional decision-makers with respect to AI. Better understanding the factors that lead institutional decision-makers to act (or not act) as secondary hypers would help opponents identify more precise pressure points for opposition. Another direction involves studying the tactics and efficacy of efforts to contest tech hypes, both past and present. This research would surface insights on what tactics are most successful at shifting entrenched arrangements in the face of power asymmetries. Finally, researchers, in conjunction with ongoing resistance efforts, should articulate and enact democratic and socially just visions for the future—spanning not just AI itself, but also the broader political and economic arrangements that shape AI entrenchment. By demonstrating the desirability and feasibility of alternative futures, these efforts would help to displace the sense of inevitability about AI entrenchment that AI hypers work to produce.

The deeper task for research and organizing alike is to build the collective capacity that undermines the effectiveness of future tech hypes. AI hype is not the first manifestation of tech hype, and it will not be the last. Efforts to resist AI hype must therefore endeavor to seed coalitions, policies, and critical consciousness that prevent institutions from taking up future tech hypes, whatever their central artifacts may be. Technological developments are not fates to which the public must adapt, but ongoing political projects that are open to contestation. Contesting hype head-on as a material force is essential for wresting power from hypers in order to shape technological futures ourselves.

Acknowledgments

We would like to thank our colleagues who offered critical feedback on drafts of this paper: Andrew Balch, John Cheney-Lippold, Abbie Jacobs, and Linda Huber. Thank you also to attendees of our session at the 2026 Privacy Law Scholars Conference, especially our commentator Susan Tanner, who significantly strengthened this paper through their questions and suggestions.

References

- Ahmed, S.; Jazwińska, K.; Ahlawat, A.; Winecoff, A.; and Wang, M. 2024. Field-building and the epistemic culture of AI safety. *First Monday*, 29(4).
- AI Now Institute. 2026. North Star Data Center Policy Toolkit.
- Akbar, A. A. 2020. Demands for a Democratic Political Economy. *Harvard Law Review Forum*, 134(1): 90–118.
- Allen, J. 2024. New York City Defends AI Chatbot That Advised Entrepreneurs to Break Laws. *Reuters*.
- Altman, S. 2023. Planning for AGI and Beyond. OpenAI.
- Bareis, J.; and Katzenbach, C. 2022. Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values*, 47(5): 855–881.
- Bareis, J.; Roßmann, M.; and Bordignon, F. 2023. Technology hypes: Practices, approaches and assessments. *TATuP-Zeitschrift für Technikfolgenabschätzung in Theorie und Praxis*, 32(3): 11–16.
- Baykurt, B. 2025. Gov-Tech as Capture: Public Infrastructures under Data Capitalism. *Information, Communication & Society*, 1–16.
- Baykurt, B.; and Raetzsch, C. 2020. What smartness does in the smart city: From visions to policy. *Convergence: The International Journal of Research into New Media Technologies*, 26(4): 775–789.
- Bear, L. 2020. Speculation: A political economy of technologies of imagination. *Economy and Society*, 49(1): 1–15.
- Beckert, J. 2016. *Imagined Futures: Fictional Expectations and Capitalist Dynamics*. Harvard University Press.
- Bennani-Taylor, S. 2024. Infrastructuring AI: The stabilization of ‘artificial intelligence’ in and beyond national AI strategies. *First Monday*, 29(2).
- Birch, K. 2013. The Political Economy of Technoscience: An Emerging Research Agenda. *Spontaneous Generations*, 7(1): 49–61.
- Birch, K. 2017. Techno-economic Assumptions. *Science as Culture*, 26(4): 433–444.
- Birch, K. 2023. Reflexive expectations in innovation financing: An analysis of venture capital as a mode of valuation. *Social Studies of Science*, 53(1): 29–48.
- Birch, K. 2025. Do Artifacts Have Political Economy? *Science, Technology, & Human Values*.
- Birch, K.; and Cochrane, D. T. 2024. Big Tech: Four emerging forms of digital rentiership. In Sanghera, B., ed., *Global Rentier Capitalism*, 82–95. Routledge.
- Boguslaw, D. 2026. US Law Enforcement Warns of ‘Anti-Tech Extremism’ as AI Hatred Grows. *WIRED*.
- Booth, R. 2026. Countries that do not embrace AI could be ‘left behind’, says OpenAI’s George Osborne. *The Guardian*.
- Borup, M.; Brown, N.; Konrad, K.; and Van Lente, H. 2006. The sociology of expectations in science and technology. *Technology Analysis & Strategic Management*, 18(3-4): 285–298.
- Bourne, C. 2024. AI hype, promotional culture, and affective capitalism. *AI and Ethics*, 4(3): 757–769.
- Brown, N.; and Michael, M. 2003. A Sociology of Expectations: Retrospecting Prospects and Prospecting Retrospects. *Technology Analysis & Strategic Management*, 15(1): 3–18.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M. T.; and Zhang, Y. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4.
- Campolo, A.; and Crawford, K. 2020. Enchanted Determinism: Power without Responsibility in Artificial Intelligence. *Engaging Science, Technology, and Society*, 6.
- Chan, A. S. 2019. Venture ed: recycling hype, fixing futures, and the temporal order of edtech. In Vertesi, J.; and Ribes, D., eds., *digitalSTS: A Field Guide for Science & Technology Studies*, 161–177. Princeton University Press.
- Chen, T.; Gascó-Hernández, M.; and Esteve, M. 2024. The Adoption and Implementation of Artificial Intelligence Chatbots in Public Organizations: Evidence from U.S. State Governments. *The American Review of Public Administration*, 54(3): 255–270.
- Cheney-Lippold, J. 2025. The silicon future. *New Media & Society*, 27(7): 4164–4180.
- Christensen, E. 2025. Organizing expectations: the co-constitutive dynamics of AI imaginaries. *Human Communication Research*, 51(4): 215–227.
- Cohen, J. E. 2025. Oligarchy, State, and Cryptopia. *Fordham Law Review*, 94(2): 563–625.
- Cole, S. 2026. OpenAI, Google, and Microsoft Back Bill to Fund ‘AI Literacy’ in Schools. *404 Media*.
- Data Center Watch. 2026. Q1 2026: Data Center Watch Report.
- De, A.; Lima, G.; and Zou, Y. 2026. What is Safety? Corporate Discourse, Power, and the Politics of Generative AI Safety. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*.
- Elish, M. C.; and boyd, d. 2018. Situating methods in the magic of Big Data and AI. *Communication Monographs*, 85(1): 57–80.
- Faxon, H. O.; Fields, D.; and Wainwright, T. 2024. Beyond the hype: Digital transformations in global land, housing, and property. *Environment and Planning D: Society and Space*, 42(4): 445–455.
- Fenn, J.; and Raskino, M. 2008. *Mastering the Hype Cycle: How to Choose the Right Innovation at the Right Time*. Harvard Business Review Press.

- Fight Back! News Staff. 2026. NYC teachers and community protest at chancellor's town hall, demand AI moratorium in schools. *Fight Back! News*.
- Fleck, A. 2024. Responsible AI: Why It Matters and How We're Infusing It Into Our Internal AI Projects at Microsoft. *Microsoft Inside Track Blog*.
- Fleisher, W. 2026. Inductive Risk of AI Hype. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 4969–4987.
- Floridi, L. 2024. Why the AI Hype is Another Tech Bubble. *Philosophy & Technology*, 37.
- Gansky, B. 2025. The Politics of AI Systems Are Inextricable from Their Supply Chains: Public Values Versus the Digital Political Economy. In *Proceedings of the Eighth AAAI/ACM Conference on AI, Ethics, and Society*, 1049–1061.
- Garibay-Petersen, C.; Lorimer, M.; and Menzat, B. 2025. Creating certainty where there is none: Artificial intelligence as political concept. *Big Data & Society*, 12(4).
- Geiger, S.; and Gross, N. 2017. Does hype create irreversibilities? Affective circulation and market investments in digital health. *Marketing Theory*, 17(4): 435–454.
- Gold, A.; and Daher, N. 2026. These States Don't Want Data Centers in Their Backyards. *Axios*.
- Gorz, A. 1967. *Strategy for Labor: A Radical Proposal*. Beacon Press.
- Green, B. 2021. The Contestation of Tech Ethics: A Sociotechnical Approach to Technology Ethics in Practice. *Journal of Social Computing*, 2(3): 209–225.
- Green, B. 2025. Using AI to Reform Government is Much Harder Than it Looks. *Tech Policy Press*.
- Grill, G. 2024. Constructing Capabilities: The Politics of Testing Infrastructures for Generative AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1838–1849.
- Hawkins, A. J. 2024. Dude, where's my self-driving car? *The Verge*.
- Hong, S.-h. 2024. Strategic misrecognition and speculative rituals in generative AI. *Journal of Digital Social Research*, 6(4): 92–106.
- Horwitz, J. 2025. Meta's AI Rules Have Let Bots Hold 'Sensual' Chats with Children. *Reuters*.
- Jones, S. 2026. Parents Aren't Happy as NYC Kindergartens Go All In on AI. *New York Magazine*.
- Kairos Fellowship. 2026. Big Tech's 'False Solutions' to the Climate Crisis.
- Kak, A.; and West, S. M. 2023. US-China AI Race: AI Policy as Industrial Policy. AI Now Institute.
- Kampmann, D. 2024. Venture capital, the fetish of artificial intelligence, and the contradictions of making intangible assets. *Economy and Society*, 53(1): 39–66.
- Katzenbach, C. 2021. "AI will fix this" – The Technical, Discursive, and Political Turn to AI in Governing Communication. *Big Data & Society*, 8(2).
- Kilander, G. 2026. Plan for AI-Themed High School Pulled amid Public Outcry Over the Technology's Potential Impact on Kids. *People*.
- Klonowski, D. 2018. *The Venture Capital Deformation: Value Destruction Throughout the Investment Process*. Palgrave Macmillan.
- Kotliar, D. M. 2025. Can't stop the hype: scrutinizing AI's realities. *Information, Communication & Society*, 29(3): 828–849.
- Kozioł, M. 2023. "Web3 Is Going Just Great" Creator On Why It Isn't. *IEEE Spectrum*.
- LaCroix, T.; Mallory, F.; and Luccioni, S. 2026. Strategic Polysemy in AI Discourse: A Philosophical Analysis of Language, Hype, and Power. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 497–517.
- LaGrandeur, K. 2024. The consequences of AI hype. *AI and Ethics*, 4(3): 653–656.
- Lapowsky, I. 2026. Tokenmaxxing Is Dead. Now Comes the Belt Tightening. *Bloomberg*.
- Liu, Z. 2024. Generative AI, Big Tech and the Value of Political Economy Analysis for the Sociology of Artificial Intelligence. In Brzeziński, D.; Filipek, K.; Piwowar, K.; and Winiarska-Brodowska, M., eds., *Algorithms, Artificial Intelligence and Beyond*, 155–172. Routledge.
- Lohr, S. 2025. Companies Are Pouring Billions Into A.I. It Has Yet to Pay Off. *The New York Times*.
- Marcus, G. 2026. The revenge of Claude Mythos. *Marcus on AI*.
- Markelius, A.; Wright, C.; Kuiper, J.; Delille, N.; and Kuo, Y.-T. 2024. The mechanisms of AI hype and its planetary and social costs. *AI and Ethics*, 4(3): 727–742.
- Mateescu, A.; Nguyen, A.; and Pinto, S. 2026. Last Place in the AI-First Economy: How the AI Industry Relies on Worker Disempowerment. *Data & Society*.
- Mathews, S. T. 2026. Getting Off with the Luddites. *n+1*.
- May, E.; Paris, B.; and Sargent, S. 2026. Dis/engaging the 'common sense' of AI: Labor strategies from the 2023 SAG-AFTRA around data-driven technologies. *Big Data & Society*, 13(1).
- Merchant, B. 2023. *Blood in the Machine: The Origins of the Rebellion Against Big Tech*. Little, Brown and Company.
- Merchant, B. 2024. AI Generated Business: The Rise of AGI and the Rush to Find a Working Revenue Model. AI Now Institute.
- Merchant, B. 2025. The Luddite Renaissance is in full swing. *Blood in the Machine*.
- Merchant, B. 2026. How Anthropic used AI ethics slop to play the pope and eclipse OpenAI. *Blood in the Machine*.
- Miceli, M. 2026. How We Made Data Workers' Inquiry. *Data Workers' Inquiry*.
- Mogul, S. 2026. A burgeoning AI boycott plasters the NYC subway. *Prism*.
- Montgomery, B. 2025. The mind-boggling valuations of AI companies. *The Guardian*.

- Mullany, M. 2016. 8 Lessons from 20 Years of Hype Cycles. LinkedIn.
- Nagy, P.; and Neff, G. 2024. Conjuring algorithms: Understanding the tech industry as stage magicians. *New Media & Society*, 26(9): 4938–4954.
- Narayan, D. 2026. Cloud Capitalism and the Rise of Scalable Digital Assets. *The Information Society*.
- Narayanan, A.; and Kapoor, S. 2025. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press.
- National Nurses United. 2024. A.I.'s impact on nursing and health care. Accessed: 2026-08-01.
- National Nurses United. 2026. National Nurses United, UE Local 150, community activists say North Carolina's 'A.I. Strategic Roadmap' falls short.
- Niemeyer, K.; and Varanasi, L. 2024. Former Google CEO Eric Schmidt says we should go all in on building AI data centers because 'we are never going to meet our climate goals anyway'. *Business Insider*.
- Nordman, D. 2025. Northeastern and Anthropic to lead in responsible AI innovation in higher education. *Northeastern Global News*.
- Office of Governor Gretchen Whitmer. 2025. Gov. Whitmer Submits Public Comment in Support of Stargate Project, Creating Thousands of Jobs, Meeting Strong Environmental Standards.
- OpenAI. 2025. How We Think About Safety and Alignment. Accessed: 2026-05-16.
- Paris, B.; Conti-Cook, C.; Greene, D.; Jones, K. M. L.; Justie, B.; Kirschenbaum, M.; Kresge, L.; May, E.; Nguyen, A.; Reynolds, R.; and Widder, D. G. 2025. Artificial Intelligence and Academic Professions. American Association of University Professors.
- Phan, T.; Goldenfein, J.; Mann, M.; and Kuch, D. 2022. Economies of Virtue: The Circulation of 'Ethics' in Big Tech. *Science as Culture*, 31(1): 121–135.
- Piven, F. F.; and Cloward, R. A. 1977. *Poor People's Movements: Why They Succeed, How They Fail*. Pantheon Books.
- Powers, D. 2012. Notes on Hype. *International Journal of Communication*, 6: 857–873.
- Rajan, K. S. 2006. *Biocapital: The Constitution of Postgenomic Life*. Duke University Press.
- Ramzanali, A. 2026. After the AI Crash. Vanderbilt Policy Accelerator.
- Ray, J. 2026. Americans Cool Toward AI. *Gallup*.
- Roose, K. 2026. More! More! More! Tech Workers Max Out Their A.I. Use. *The New York Times*.
- Rowe, N. 2023. 'It's Destroyed Me Completely': Kenyan Moderators Decry Toll of Training of AI Models. *The Guardian*.
- Sadowski, J. 2019. When data is capital: Datafication, accumulation, and extraction. *Big Data & Society*, 6(1).
- Sadowski, J. 2020. The Internet of Landlords: Digital Platforms and New Mechanisms of Rentier Capitalism. *Antipode*, 52(2): 562–580.
- Sadowski, J. 2025. *The Mechanic and the Luddite: A Ruthless Criticism of Technology and Capitalism*. University of California Press.
- Sadowski, J.; and Bendor, R. 2019. Selling Smartness: Corporate Narratives and the Smart City as a Sociotechnical Imaginary. *Science, Technology, & Human Values*, 44(3): 540–563.
- Schmidt, E. 2020. I Used to Run Google. Silicon Valley Could Lose to China. *The New York Times*.
- Singer, N. 2026. In Backlash Against Tech in Schools, Parents Are Winning Rollbacks. *The New York Times*.
- Singh, L. 2026. Booing Commencement Speakers Over AI Is Becoming a Trend. *New York Magazine*.
- Sloane, M.; Danks, D.; and Moss, E. 2024. Tackling AI hyping. *AI and Ethics*, 4(3): 669–677.
- Slota, S. C.; Fleischmann, K. R.; Greenberg, S.; Verma, N.; Cummings, B.; Li, L.; and Shenefiel, C. 2020. Good systems, bad data?: Interpretations of AI hype and failures. *Proceedings of the Association for Information Science and Technology*, 57(1).
- Smith, B. 2026. Building Community-First AI Infrastructure. *Microsoft On the Issues Blog*.
- Spataro, J. 2023. Introducing Microsoft 365 Copilot: Your Copilot for Work. *The Official Microsoft Blog*.
- Stark, L.; Greene, D.; and Hoffmann, A. L. 2020. Critical Perspectives on Governance Mechanisms for AI/ML Systems. In Roberge, J.; and Castelle, M., eds., *The Cultural Life of Machine Learning: An Incursion into Critical AI Studies*, 257–280. Palgrave Macmillan.
- Steen, M.; Timan, T.; Van Diggelen, J.; and Vethman, S. 2025. We need better images of AI and better conversations about AI. *AI & Society*, 40(5): 3615–3626.
- Steinboff, J. 2024. Toward a Political Economy of Synthetic Data: A Data-Intensive Capitalism That Is Not a Surveillance Capitalism? *New Media & Society*, 26(6): 3290–3306.
- Taylor, A. 2018. The Automation Charade. *Logic(s)*.
- Taylor, A.; and Levin, S. 2026. The fight against AI datacenters isn't just about tech – it's about democracy. *The Guardian*.
- Thais, S. 2024. Misrepresented Technological Solutions in Imagined Futures: The Origins and Dangers of AI Hype in the Research Community. In *Proceedings of the Seventh AAAI/ACM Conference on AI, Ethics, and Society*, 1455–1465.
- The White House. 2026. A National Policy Framework for Artificial Intelligence: Legislative Recommendations.
- Törnberg, P. 2025. Hype, financial narratives, and self-fulfilling prophecies in surveillance capitalism. *Media, Culture & Society*, 47(8): 1543–1558.
- Tortorelli, P.; Stevens, P.; and Tolockaite, A. 2025. In race to attract data centers, states can forfeit hundreds of millions of dollars in tax revenue to tech companies. *CNBC*.
- Tsing, A. 2000. Inside the Economy of Appearances. *Public Culture*, 12(1): 115–144.

Tully, S. M.; Longoni, C.; and Appel, G. 2025. Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity. *Journal of Marketing*, 89(5): 1–20.

UK Government. 2025. Prime Minister Sets Out Blueprint to Turbocharge AI.

Verdegem, P. 2023. Critical AI studies meets critical political economy. In Lindgren, S., ed., *Handbook of Critical Studies of Artificial Intelligence*, 302–311. Edward Elgar Publishing.

Verdegem, P. 2024. Dismantling AI capitalism: the commons as an alternative to the power concentration of Big Tech. *AI & Society*, 39(2): 727–737.

Vertesi, J.; boyd, d.; Taylor, A.; and Shestakofsky, B. 2026. Reckoning with the Political Economy of AI: Avoiding Deceits in Pursuit of Accountability. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 2186–2205.

Vrabič Dežman, D. 2024. Promising the future, encoding the past: AI hype and public media imagery. *AI and Ethics*, 4(3): 743–756.

Wang, H.; Dang, A.; Wu, Z.; and Mac, S. 2024. Generative AI in higher education: Seeing ChatGPT through universities’ policies, resources, and guidelines. *Computers and Education: Artificial Intelligence*, 7.

West, S. M. 2019. Data Capitalism: Redefining the Logics of Surveillance and Privacy. *Business & Society*, 58(1): 20–41.

Williams, A.; Miceli, M.; and Gebru, T. 2022. The Exploited Labor Behind Artificial Intelligence. *Noema Magazine*.

Williams, D. J. 2024. We’ve been here before: AI promised humanlike machines – in 1958. *The Conversation*.

Witte, R. 2025. New Yorkers Are Defacing This AI Startup’s Million-Dollar Ad Campaign. *Futurism*.

Young, M.; Katell, M.; and Krafft, P. M. 2022. Confronting Power and Corporate Capture at the FAccT Conference. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1375–1386.

Young, M.; Madaio, M.; and DiSalvo, C. 2026. You can’t solve democracy in Q3: A Critical Analysis of Corporate Initiatives to Democratize AI. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 2727–2744.

Zhang, J.; Blodgett, S. L.; and Markl, N. 2026. “A natural language interface for everyday life”: The social and political functions of AI capabilities discourses. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 5023–5042.

Zhang, J.; Chowdhary, S.; and Kim, N. 2026. AI as a War Issue, War as a Workers’ Issue. *Academe*.

Zitron, E. 2026. AI Is Really Weird. *Where’s Your Ed At*.