

“A natural language interface for everyday life”: The social and political functions of AI capabilities discourses

JUSTINE ZHANG, University of Michigan, USA

SU LIN BLODGETT*, Mila - Québec AI Institute, Canada

NINA MARKL, University of Essex, United Kingdom

We interpret and historically contextualize *AI capabilities discourses*: spectacular narratives that ostensibly lay out the capabilities and use cases of the latest generative AI models in marketing and scientific communications. We complement existing work scrutinizing these discourses as components of a wider phenomenon of “AI hype”—which tends to focus on “debunking” them—by exploring how they rationalize wider social projects. Via an analysis of marketing materials produced by companies like OpenAI, we trace how AI capabilities discourses extend longer-running capitalist discourses, which construct aspects of self and social life as segmentable “skills” with economic value, and which responsabilize people for their survival in precarious economic conditions. We then show how AI capabilities discourses build on top of this ideological scaffolding. We suggest that these discourses have power because they reiterate hegemonic ideas about society and personhood that are entwined with material conditions. As such, contesting the power of the AI industry necessitates more capacious forms of political action, beyond interventions aiming to merely correct narratives.

CCS Concepts: • **Applied computing** → *Law, social and behavioral sciences*; • **Social and professional topics**; • **Human-centered computing** → *HCI theory, concepts and models*; • **Computing methodologies** → **Philosophical/theoretical foundations of artificial intelligence**;

Additional Key Words and Phrases: artificial intelligence, discourse, capitalism, ideology, labour, language

ACM Reference Format:

Justine Zhang, Su Lin Blodgett, and Nina Markl. 2026. “A natural language interface for everyday life”: The social and political functions of AI capabilities discourses. <https://doi.org/10.1145/3805689.3806481>

1 Introduction

The OpenAI website states that GPT-5, released in August 2025, is “our smartest, fastest, most useful model yet, with built-in thinking that puts expert-level intelligence in everyone’s hands.”¹ In an introductory video, text flashes across the screen in few-second intervals, listing GPT-5’s many virtues: “Thinks deeply when you need it to...an expressive writing partner...make it your own.” The webpage goes on to tout state-of-the-art performance “across coding, math, writing, health, visual perception, and more,” displaying model performance on benchmarks ranging from the AIME—a math competition for American high schoolers—to GPQA Diamond—a “PhD level” question-answering benchmark published by researchers at New York University, Cohere, and Anthropic—to an internal-to-OpenAI benchmark titled “economically important tasks.” There is a list of example use cases, including “poetry,” “understanding cancer risk,” and “band global tour planning”; prompts representing these use cases are paired with outputs by GPT-5 that demonstrate “striking metaphors,” “respond[ing] empathetically,” and “sophisticated operational contingency planning.”

*Work done while at Microsoft Research Montréal.

¹URLs of webpages quoted in our paper are included in the Appendix.

Authors’ Contact Information: Justine Zhang, University of Michigan, Ann Arbor, Michigan, USA, tisjune@umich.edu; Su Lin Blodgett, sulin.blodgett@mila.quebec, Mila - Québec AI Institute, Canada; Nina Markl, University of Essex, Colchester, United Kingdom, nina.markl@essex.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

A substantial body of academic scholarship, journalistic reporting, and popular commentary has characterized such discourses as “hype,” produced and circulated to manipulate the public perception of what generative AI products like GPT-5 are capable of, and to fuel an ongoing wave of overinvestment in the AI industry [e.g., 20, 88]. This work contends that these products are not nearly as capable as their makers claim, and that many of their suggested use cases have harmful consequences. Seeking to pierce through the widespread circulation of AI hype, researchers have sought to more accurately evaluate model performance and associated risks across domains.

Our work takes a different approach: rather than interrogating particular claims for their truth value, we are interested in the form these discourses take, in relation to the social projects they are part of. A range of anthropological, historical, and sociological work has examined the various discourses that interpret and rationalize social conditions that would otherwise seem incoherent or unjust [e.g., 11, 31]. Other scholars have highlighted the centrality of discourse to the project of AI itself [7, 72]. Building on this work, we seek to interpret what we refer to as *AI capabilities discourses*: those spectacular narratives that ostensibly lay out the capabilities and use cases of the latest AI models. What is required to make statements like “built-in thinking” sayable, “poetry” and “understanding cancer risk” comparable as items that belong in the same list, and all of these components interpretable as features of an advanced technology that is also a desirable consumer product? And what functions do these statements help perform, as they proliferate in economies structured by the imperative of capital accumulation and marked by widespread precarity?

To address these questions, we analyze a collection of marketing materials for generative AI models released by four companies currently developing these products: OpenAI, Anthropic, Google, and Meta. We attend to these materials’ content, and the visual motifs, linguistic constructions, and semiotic connections they contain. We show that AI capabilities discourses join ideas about flexibilized, cognitive work that are prevalent in present-day capitalist societies with ideas about intelligence that are associated with AI. As such, they index a consolidating alignment between AI development and capitalism. We also show how AI capabilities discourses render a wide variety of social activities as interchangeable, and as subsumable into this AI-capitalism alignment. In this way, like preexisting discourses that have circulated in capitalist societies for the past few decades [11, 31, 114], AI capabilities discourses rationalize the accumulation of capital while undermining the terrains on which people can contest the conditions of their labour, lives, and societies—regardless of whether the claims contained in them are true, false, or hyped up.

2 Background

By *AI capabilities discourses*, we mean discourses that describe the capabilities supposedly contained in generative AI models. AI capabilities discourses lay out the range of things models are capable of, the ways that people can use these models, and the various domains in which they could be used. As seen in the opening vignette from OpenAI, these discourses can touch on technical functionalities (e.g., “visual reasoning”), use cases (e.g., coding an app), qualities (e.g., “responding empathetically”), and domains (e.g., “health”). AI capabilities discourses often project impressions that have become closely associated with generative AI models: that they are “general-purpose,” that they represent profound technological advancements, and that they will be societally transformative. With the recent and rapid growth of the AI industry, narratives lauding the capabilities of AI—e.g., that “artificial general intelligence is already here” [128]—have proliferated across media, marketing, academic, and policy communications. As we discuss in §3, we focus on AI capabilities discourses as articulated by AI companies—a setting that foregrounds their role in marketing and serves as a discursive prototype for other sites.

To stage our approach, we start by surveying two predominant ways that scholars have engaged with AI capabilities discourses: by seeking to interrogate the “realness” of AI’s reported capabilities [72], and by unpacking the discursive moves involved in articulating them. We then offer a brief account of the political and economic

conditions in which AI development and AI discourses are embedded, and then survey work on prevalent ideologies that have circulated in these conditions.

2.1 Prior engagements with AI capabilities discourses

2.1.1 Contested capabilities of generative AI. Claims about the capabilities of AI models, in particular (multimodal) large language models, have been at the centre of popular and academic critiques. These accounts have highlighted the (often vast) gaps between purported capabilities of large language models and their actual behaviours, targeted the conceptual and terminological muddiness of “AI” [108], and drawn attention to the eugenicist and transhumanist ideologies undergirding notions of “intelligence” [62]. They build on long-standing theoretical debates about whether algorithmic systems (of any type) can (or could ever) “reason,” “understand,” or “think” [21, 27, 50, 107, 122], and extend more recent interrogations of common practices in AI development, particularly the construction and use of benchmark datasets in the training and evaluation of AI models [10, 13, 26, 28, 48, 64, 77, 91, 95]. They are further related to a robust body of literature showing how the extractive and reductive logics of algorithmic systems produce harms and exacerbate inequities in a wide range of settings, like healthcare, law, education, and interpersonal relationships [25, 30, 57, 79, 98, 106, 118]—and how these systems can obscure oppressive social projects through a veneer of progress or objectivity [22, 45, 53, 76, 119]. Pointing to the increasingly outsized influence of hyped-up claims about AI, critics of generative AI have often framed the technology, and its surrounding business model, as a “con” [20] or “snake oil” [88]. To combat these claims and deflate the power of the AI industry, these critics often underline the need for better scientific theory or experiments, more careful and critical accounts of what AI can and cannot do, and increased AI literacy [27, 64, 65, 73, 82, 96, 104, 112].

2.1.2 Discursive practices of AI. Other work has attended more closely to the *function* of these discourses, arguing that “AI hype” consists of the sorts of promises and threats that discipline and demoralize labour, while exciting capital [30]. Narratives about AI—what it is, what it will do, and what it might be capable of—are essential to the AI industry, given its reliance on large-scale venture capital and government investments [12, 123, 125]. Such discourses are a feature of all technologies: as Bassett and Roberts [14] note, automation in particular has long left both workers and employers anxious in the face of seemingly impending change, which never arrives as abruptly as anticipated. In this way, “AI hype” and AI capabilities discourses are not unique. However, such scaffolding discourses are especially important in the context of AI (broadly) and generative AI (specifically) to define intended uses and applications and articulate risks and benefits. Many things—or, as the marketing would have it, *anything*—can be framed as “predict-the-next-token” tasks, putting the onus on developers and other interested parties to demonstrate and describe relevant use cases.

Hong [72] argues that “public dramas around genAI’s realness question are a central medium for perpetuating [a] strategic misrecognition” whereby “visceral resemblance at the output level” is taken as indicative of intelligence. The attribution of human-like cognitive processes based on such outputs can be usefully understood as a “form of human sense-making” whereby certain qualities (e.g., agency, intelligence) are projected onto objects [72]. Similarly, Stark [105] suggests that AI models like ChatGPT or Claude might be best understood as “animated characters,” designed to encourage projection of human qualities that models intrinsically cannot have. This misrecognition is *strategic* because it “calibrates collective expectations” regarding AI, and because the public dramas surrounding AI—events, controversies, press accounts, etc.—are driven by AI developers and other invested actors [72]. Whether they ascribe positive or negative effects to AI, many of these public dramas presuppose that AI is “here to stay” [70], thus justifying and legitimating investment in research and development, and further fueling public discourse [32, 102].

In this way, discourse has been understood as *central* to AI research and development, going back to its inception. As Agre [7] explained almost three decades ago, terms such as “learning” or “planning” in AI are

components of a “set of technical schemata that can be expanded into a narrative thematics for any particular domain.” He draws particular attention to the “elasticity” of such vocabulary, noting that it simultaneously serves as a “technical proposition to quite specific and circumscribed set of functionalities and algorithms” and as an “empirical proposition to anything at all that can plausibly be glossed with the term” [7]. In this way, the discursive practices of AI “figure[] in a project of instrumental appropriation,” rendering everyday life into “the object of a colonial project of reform” [5].

These discursive practices are co-constituted with the material conditions in which they operate. Scholars have detailed the ways that ideas and discourses about technology have developed in tandem with intensified processes of marketization, flexibilization, and financialization [6, 24, 57, 63]. Entwined with political economy, AI discourses soften the ground for material developments. As Hong [72] makes clear, the public dramas of AI legitimize financial, cultural, and political investments into the AI industry, helping to accelerate changes to the economy and the physical landscape alike, whether through the accumulation of capital or through the construction of data centres [116, 125].

2.2 Economic and political context

Booming AI industry aside, much of the world is experiencing deep social unrest and economic instability. Seeking to account for the constellation of ways that life has become increasingly difficult for more and more people, commentators have pointed to “the affordability crisis,” “the care crisis,” “the unemployment crisis,” “the climate crisis,” and more [39, 49, 59, 86, 113]. Many social theorists have illustrated how these crisis conditions result from the inherent structural tendencies of capitalism [16, 58, 84]: capitalism’s “power to make and destroy worlds” systematically undermines “the social relations and social subjects constituting and constituted by it” [36]. The specific form that capitalism has taken in recent decades, spurred by deindustrialization, stagnating economic growth, and shifts in social policy, is one characterized by intensified precarity, punishing austerity, unprecedented inequality, and financialized speculation [19, 31, 35, 44, 68]. For instance, in the US, AI-related investments are driving up the value of stock exchanges, in spite of widespread skepticism around the benefits of AI technologies [55, 124]; meanwhile, successive presidential administrations have gutted a wide range of social services [38, 44, 67]. While there are differences in how capitalism manifests around the globe, in this paper we will focus on its effects in Western societies.

Changing labour dynamics are especially pertinent to the functions of AI capabilities discourses, since they forefront ideas around work and productivity. Recent decades have seen a broad shift away from long-term employment to flexibilized, short-term work across many sectors, from delivery and transport to education and healthcare [11, 19, 117]. Meanwhile, government policies in countries like the US and the UK have tied essential protections and provisions like health care, pensions, and residency permits to people’s employment status, all the while instituting legal regimes that curtail workers’ ability to organize and bargain [51, 93, 103]. As a result, the power of people to collectively act to improve their labour conditions, always tenuous under capitalism, has further diminished.

Concurrently, there has been an ongoing process of tertiarization in Western societies, as economies have shifted away from the industrial production of goods towards service provision [19, 126]. The service sector, which spans a variety of fields, including hospitality, education, financial services, and healthcare, is characterized by particular labour conditions: aspects of self and sociality, like language, cognition, and affect, are critical to service-oriented jobs [52, 71]; accordingly, people are often expected to “rethink and transform one’s self to best fit one’s job” [114].

2.3 Ideologies and discourses of capitalism

Capitalism is not solely defined by material conditions, but rather involves both “practical arrangements and practices of knowledge [and] their complex cogeneration and entwining” [36]. As Boltanski and Chiapello [31] observe, “in many respects, capitalism is an absurd system” that must somehow be made intelligible: it can only reproduce itself if many people buy into an obviously oppressive arrangement, where the ability to survive is tethered to exploitation.

In their longitudinal analysis of management literature, Boltanski and Chiapello trace the shifting forms of the “spirits of capitalism” that justify people’s commitments to capitalism. At the end of the nineteenth century, for instance, when work predominantly took place in small family firms, economic arrangements were associated with “progress” and “freedom from local communities,” while practices like book-keeping and saving money were understood to lead to financial success and security. In recent decades, by contrast, the management literature has started to conceptualize work as more temporary, structured by projects that are necessarily short-term in a world cast as beset with rapid technological and social transformations. To be successful, workers “must be adaptable and flexible” to be rewarded with “opportunities for development” and access to new networks, which then open up the possibility of future projects. As Baker [11] demonstrates, these ideologies affect both workers and capitalists. The entrepreneurial entreatment to “make your own job” shapes workers’ behaviours, rationalizing an interminable process of self-improvement, and contributing to a widespread sense of exhaustion and burnout. Meanwhile, the valorization of qualities like entrepreneurialism serves as a justificatory logic for managers to discipline their labour force and ruthlessly downsize and outsource: in a company that must be flexible and efficient, there is no room for insufficiently entrepreneurial workers.

The concurrent processes of flexibilization, tertiarization, and financialization are often interpreted as reflecting the emergence of a new “knowledge economy,” characterized by immaterial, “cognitive” labour [3, 111]. An *Economist* article from 1982 predicts that “brainworkers [...] in the future would be most workers,” while a core idea in the influential—and eugenicist—1994 book *The Bell Curve* was that “knowledge workers had displaced the manual workers of the past, and these workers’ new ideas were now ultimately the source of value,” with figures like entrepreneurs, scientists, and artists constituting a new “cognitive elite” [11, 31, 78]. Contemporaneously, economists and management consultants elaborated a theory of “human capital” where people’s cognitive capacities were rendered as zones for speculative investment, seen as portfolios that could be grown, diversified, and risk-managed, via means like higher education and employee management [3]. Other aspects of personhood were likewise recruited into this new regime of valuation. Anticipating (presciently) a shift away from industrialized economies after the Second World War, a 1947 report from the US government calls for education that could foster “personal and social development [...] that is intensive, accurate, and comprehensive enough to give the student command of marketable abilities,” including “skills involved in critical and constructive thinking” [114]. As we describe in §5, *skills discourses*, which discursively cement the segmentation of personhood and sociality into productive units, exhibit resonances with AI capabilities discourses.

Capitalist ideologies, in conjunction with capitalist conditions, are productive: they reproduce capitalist relations, they legitimize exploitation, they justify the disposal of people who can’t keep up, and they obscure the necessity for collective struggle by valorizing an individuating pathway to success that, given continually dismal economic conditions, will never actually guarantee security, let alone success, for most people [11, 31, 40]. The power of these ideologies comes from the ways they are entwined with these conditions; as Baker writes, they are “not merely the result of an intellectual misunderstanding or a campaign of active deception” but rather are sustained by “the all-too-real decomposition of the working class, [...] the demolition of its social and political institutions, [and] the separation of workers from the work that underpinned their self-consciousness as a class” [11]. Destroying the hold of these ideologies on society, then, is not a mere matter of correcting the record and “pointing to the objective existence of class somewhere ‘out there’” [11]; what’s required is struggle for better

material conditions conjoined with the development of new subjectivities and forms of social life. We think of the relation between AI capabilities discourses and the conditions in which they proliferate similarly, as both tightly entrenched and subject to transformation.

3 Data and approach

Four companies—OpenAI, Anthropic, Meta, and Google—presently dominate the AI sector, in multiple senses: the majority of AI users use the generative AI products they develop (i.e., various versions of GPT, Claude, Llama, and Gemini), a large proportion of AI research is produced and sponsored by them, and they exert substantial influence in popular media and policymaking [2, 34, 121]. We examine these companies’ websites, attending to the ways they describe their generative AI models. To build out a corpus, we started with the homepages for each model on company websites, and then followed links to technical reports, news releases, blog posts, and other pages focused on aspects such as particular use cases. We tried to be as comprehensive as possible, collecting as much website content as we could find.

While our account is informed by our exploration of the entire corpus, for this paper we focus on materials pertaining to models released in 2024 and 2025. During this time, the AI industry grew dramatically, consolidating political and economic power. The target audiences for these materials accordingly expanded and blurred. Nominally “scientific” documents, such as descriptions of a new model version on a company’s research blog, were more broadly circulated as public spectacles [72]; materials attesting to the models’ generalist capabilities were often interwoven with materials selling them as general-purpose products. We specifically sought to examine the *all-encompassing* quality of these materials: how do the ways they treat various use cases, application domains, and capabilities work together to produce a sprawling whole?

Analytically, we take up a central premise of critical discourse analysis, that “discourse is structured by dominance” [127], and its central focus on “how language distinctively figures in the new capitalism” [54]. We approached our data as sites of *ideological work*, where discourses and signs are engaged in “the active making of social life” [60]. We sought to situate AI capabilities discourses in an understanding of the social relations that give rise to them, following interpretivist sensibilities from traditions like situational analysis [42]. Given the multimedia nature of our materials, we considered a broad range of linguistic and semiotic features.

We performed multiple rounds of analysis; over several iterations, each author examined separate parts of the data, producing memos while selecting key observations to discuss as a group. During weekly meetings, we shared preliminary analyses and collectively formulated our account of AI capabilities discourses. Each of us are trained as natural language processing researchers, and have some familiarity with the technical operations of large language models, with linguistic and sociolinguistic theories, and with the intellectual and cultural sensibilities of fields like AI. As workers in higher education institutions and industry research labs, we are also intimately familiar with the increasing precaritization of our professions, and the ever-escalating imperative to affix our intellectual labour to economic value. Our positionalities thereby oriented us to the ways that these discourse materials assembled specious technical and sociological claims to produce looming, barely-desirable renderings of work and life.

Given our choice of data, we offer a necessarily partial view of AI capabilities discourses, focusing on companies’ narration of their products in sites primarily oriented towards sales, branding, and the cultivation of public and institutional buy-in. As such, we read our data as prototypical examples of the broader category: they are highly distilled, and intent on casting their objects in a positive light. While they exemplify discursive features that recur across other sites, like academic papers and policy documents, they lack the nuances and contradictions that may be present elsewhere. We also note that our choice of material and our analytic frames are primarily centred in Western, specifically American contexts; consequently, we leave out of focus the discursive tactics AI companies

deploy as they expand into markets around the world. There are multiple ways, then, for future work to fruitfully situate AI capabilities discourses in a broader range of contexts, and to examine their varied articulations.

4 The discursive landscape of AI capabilities discourses

We provide an overview of AI capabilities discourses and their key discursive features, attending to the space of model capabilities they lay out, and the associations they draw between cognition, productivity, and flexibility.

4.1 Boundless lists of interchangeable things

“There’s never been a worse time to be a problem, or a better time to be a problem solver,” reads the Anthropic website, underneath a short video depicting problem solvers solving problems. The problem solvers are represented by a diverse array of people: men in suits leaning their heads against glass walls in frustration, a woman in athleisure holding a camera in a sun-speckled forest, a man wearing a backwards baseball cap leaning over a laptop next to a kitchen sink, a rock climber, some joggers, some people in a car next to the joggers, a trio of citizen scientists in a desert launching a homemade rocket, dancers in a studio, a woman in a cavernous hall admiring the architecture, a different woman in a dark room admiring a digitally-rendered manifold on the ceiling. The problems depicted range from concrete to abstract: a child’s science fair project, the construction of said rocket, physiological analyses of said joggers and dancers, a birds-eye view of the neat grid of lights in a city, math on a whiteboard, microscopes, furrowed brows, papers strewn across a floor, chess pieces floating in midair. “Keep thinking,” the video concludes.

The Anthropic video exemplifies a key quality of AI capabilities discourses: the scope of things presented as AI capabilities is vast. In the span of one and half minutes, the video flits between offices and the outdoors, between work applications and extramural engagements, between science and art: everywhere, one can find “problems” that Claude can take up. Whether Claude can actually tackle any one of these problems, or what it would even mean to construe a particular vignette as a problem, is beside the point. The rapidity through which scenes cycle suggests something more impressionistic: these are impressions of the many forms problems could take, the many domains where problems could be found, and the many problem solvers for whom Claude could be helpful. “Problem” is a category that elastically expands with one’s imagination; “if you can dream it, Claude can do it.”

AI capabilities discourses evoke this sense of boundlessness across numerous long lists of things, including and beyond sprawling video montages. Elsewhere, the Anthropic website states that “people are using Claude to—” followed by text that changes every few seconds, listing activities like “rehearse a sales pitch,” “debug a wedding website,” or “create a 30 day meal plan.” As a different take on the motif of interchanging text, a page providing an “overview” of ChatGPT prominently displays an animation consisting of three rows of blocks that travel leftwards or rightwards across the screen. Each block contains a potential ChatGPT prompt (clicking the blocks will actually run the prompts): “looking for an eco-friendly yoga mat that won’t slip”; “write a python script to automate sending daily email reports”; “how are oil prices impacting global energy markets”; “explain nostalgia to a kindergartener”; etc. Indeed, the size of these lists of things, and the web design, often suggest they should be appended with “and more.”

One quality of these lists is that individual entries are often not comparable to each other. “[L]ooking for an eco-friendly yoga mat that won’t slip” seems to call for something entirely different from “explain nostalgia to a kindergartener”; it is unclear what comparisons one could meaningfully make between purchase recommendations for a specific type of yoga mat and the intensely social activity of relating an emotional experience to a young child. Likewise, it is unclear how a good sales pitch could reflect the same criteria as a good 30 day meal plan. A video on Google’s Gemini 3 starts by declaring that “Gemini has been multimodal since the beginning,” with modalities—represented by terms overlaid on abstract images—like “see,” “hear,” “understand,” “generate,” “reasoning,” “think,” “take action.” Several of these terms do not seem straightforwardly interpretable as modalities; even before getting

to whether Gemini can really “understand” or “reasoning,” one is confronted with the heterogeneity of these descriptors.

The diversity of things invoked that AI can do is perhaps appropriate for a product pitched as “general-purpose.” However, this diversity also suggests a peculiar tactic for describing the world. The grammatical and visual syntax of the videos and moving text suggest that the things mentioned have an interchangeable quality. The point is not the content of each descriptor—what activity it points to, what its success criteria are, whether a model can meet these criteria—so much as it is yet another thing on an ever-extending list, and therefore in some way like the other things on that list. The technical, substantive, and sociocultural particularities of these things are immaterial; in fact, the juxtapositions can often be jarring, as seen in OpenAI’s pitch from §1 that GPT-5 is capable of both “understanding cancer risk” and “poetry,” or Meta’s list of six “capabilities” investigated in Llama 3’s human evaluations, “English, reasoning, coding, Hindi, Spanish, and Portuguese.” What matters is that there are many things, and there could be many more.

4.2 Cognition, productivity, and flexibility

In AI capabilities discourses, AI is intelligent, thinks, and reasons. Equipped with “superpowers,” Meta’s Llama 4 Behemoth is “one of the smartest LLMs in the world”; OpenAI’s GPT-5 “puts expert-level intelligence in everyone’s hands”; Google’s Gemini is “your partner in complex reasoning.” The intelligence of AI products is positioned as well-suited to these open-ended lists of capabilities, with lists referred to by various thinking-related terms. Alongside “problems,” Anthropic also mentions “challenges,” “your toughest work,” and anything that you can “dream.” Google offers Gemini to “Dream it. Describe it. Done.,” while Meta promises that Llama will “build your greatest ideas.”

Across these discourses, then, there is a recurring association between cognition and actualization, between thinking and accomplishing things. This association takes shape in Anthropic’s figure of the “problem solver,” and its video rendering of cognitive work in motion. It can also be seen in the types of people frequently featured in these discourses—scientists, coders, creatives, high-powered executives—as well as common examples of things that AI is offered up as doing—writing, learning, coding, researching, analyzing, building—that involve being thoughtful and also productive.

This association also takes shape in descriptions of what cognition brings you—all manner of positive outcomes enabled by intelligence and reasoning. AI can conduct analyses and grasp subtleties (conduct “due diligence,” “capture the nuances of user intent,” and “crack[] nuanced localization”); acquire any and all kinds of expertise (Llama learns the “language of fútbol,” which is “filled with idiomatic phrases and insider jargon”; “custom AI experts for help on any topic ... anything from a career coach or brainstorm partner to a coding helper”); produce truth and objectivity (“uncovering the truth in the most challenging research cases”); create knowledge (“push the boundaries of healthcare innovation and bioinformatics”); help humans improve (“break it down with clear, step-by-step guidance”); and work on humans’ behalf where they evidently cannot (“The volume, velocity and complexity of financial information far exceed human comprehension, making smart decision-making nearly impossible for the average investor”).

In some sense it is of course unsurprising that cognition is a central theme, since AI is about intelligence. But what these discourses forefront is the enregisterment of intelligence with a constellation of other qualities—knowledge, information, truth, insight—and domains—health care, finance, software development—that connote social and economic success. The expansive promise of AI is befitting of a specific type of world, governed by the accumulation and circulation of information, where value is linked to the cognitive and organizational power to process these information flows.

The capaciousness of capabilities-related terms is also perhaps unsurprising, given the AI field’s oft-commented on propensity for vague language [e.g., 108]. In AI capabilities discourses, this vagueness is instrumentalized to

suggest a sense of flexibility. Absent any specification of the situated actions that are involved in researching, writing, planning, and thinking, these terms are instead taken to denote abstract capabilities that could be applied across sites and domains, from the high-value (e.g., “cutting through the [financial] noise”) to the personal and empowering (e.g., self-care). Thus, “planning” can be applied to “plan a date night,” “create a 30 day meal plan,” “plan effective content calendars to drive engagement,” and “transform your prompt into a personalized multi-point research plan.” These activities are abstracted and packaged as AI capabilities, suggesting a sense of mobility across vastly disparate areas of life.

5 How AI capabilities discourses work

In AI capabilities discourses, a variety of activities are evacuated of their technical and sociocultural specificities, and enfolded into a regime where a generalist cognitive capacity—of the sort AI is understood to have—leads to success. In what follows, we elaborate on some of the discursive tactics that they employ to perform this function.

5.1 AI capabilities discourses as skills discourses

To unpack how AI capabilities discourses work, we start by connecting them to longer-running iterations of capitalist discourses. Over the last few decades, *skills discourses* have emerged as predominant ways of narrativizing ideas about work and personhood in increasingly flexibilized, service-oriented economies. Skills discourses construct workers as “bundles of skills,” and advise or exhort workers on the skills they should acquire or enhance. In a semiotic study of skills discourses, Urciuoli [114] identifies two primary discursive features. First, the variety of things that can count as a skill “cover a range of disparate practices, knowledge, and ways of acting and being”; “listening,” “Photoshop,” and “French” can all be considered skills [114]. Second, there is a vagueness to skills terms that nonetheless implies a position, exhibiting a “tendency to expand with ideas of how the world is or should be” [114]. The keyword “skill” itself, for instance, points to a category that has ballooned in scope with the deindustrialization of Western economies, and has moved from being associated with artisanal craftwork to being associated with business success, like “leadership skill” [92, 114]. These properties are exemplified in materials produced by the Partnership for 21st Century Skills, an influential non-profit organization whose members included American policymakers and representatives from Adobe Systems and Microsoft. The types of skills promoted by this organization included “learning and thinking,” “information and communications technology,” and “life skills.” What these descriptors have in common is that they are part of a “vision for 21st century education to ensure every child’s success as citizens and workers in the 21st century” [1, 114].

As we’ve shown, these characteristics of indeterminacy (in what counts as skill) and strategic vagueness (in what terms point to) carry over to AI capabilities discourses. Alongside these discursive resonances, there is also a substantive overlap in what counts as a skill and as an AI capability. In skills discourses, the prototypical skill is a “soft skill” like “communication” and “writing”—aspects of self or social interaction which are also valuable across a range of jobs. Strongly associated with ideas about cognition, soft skills are contrasted with “hard skills,” which connote technical specialization for specific professions; in an economy where jobs tend to be short term, soft skills are associated with a beneficial flexibility. While skills discourses articulate “aspects of personhood with exchange value on the labor market” [114], AI capabilities discourses load these aspects of personhood into AI products, operationalizing ideas around AI’s humanlike qualities.

Skills discourses, via their semiotic qualities, segment people into technologies of self [56]. By putting all sorts of activities into undifferentiated lists, held together by vague terms indexing capitalist values, skills discourses cast these activities as of a type: shorn of particularities, they become discrete things, presumed to be measurable, rankable, and befitting the “skills” or “qualifications” section of a resume or LinkedIn profile, “located as the property of the individual, who then [carries] them, luggage-like, from job to job” [92]. As a skill, for instance, “communication” involves acting in ways that index business success, embedding an assortment of assumptions

about race, class and gender [46, 71, 85, 114, 115]. Detached from the rich, relational ways that people communicate with each other, “communication” is recast as something that could be assessed in performance reviews and inculcated in corporate training modules, like “time management” and “prompt engineering.” Accordingly, skills discourses function as disciplinary discourses imposed on individuals, tasked with upskilling and reskilling in conditions of labour precarity.

The strategic vagueness of skills discourses also allows them to work as ways of consolidating capitalist class interests [114, 115]. Whether these discourses are nonsensical or compelling depends on whether or not one buys into the visions on offer by coalitions like the Partnership on 21st Century Skills, and their underlying ideological presuppositions. Accordingly, skills discourses function as shared discursive templates for the reproduction of economies structured by the valorization (and material unavoidability) of flexibilized work. In aligning capitalist interests, skills discourses have helped structure social policy, with profound effects. For instance, the Common Core, a set of standards produced by policymakers for K-12 education in the United States, was shaped in part by the Partnership and its emphasis on assessable, economically-valuable skills. Critics, in turn, have traced a broad undermining of learning conditions to the Common Core, pinning present failures on “the mischaracterization of intrinsically humanistic modes of thought as generic academic skills [...] as though the modes of thinking intrinsic to the liberal arts have been fracked” [41].

Drawing on the structure and content of skills discourses, AI capabilities discourses reiterate these disciplining and aligning functions. In ostensibly articulating all the things that AI can do, they play off of these older discursive forms to align AI with the same capitalist social orders, and they carry forth the rendering of practices and aspects of personhood as commodifiable and commensurable things, then as skill, now as AI capabilities.

5.2 Discursive techniques

A key difference between AI capabilities discourses and skills discourses is that rather than entreatments for people to accumulate skills, AI capabilities discourses are entreatments for them to avail themselves of the capabilities offered by AI products—a more technologically advanced way for people to “make themselves into Swiss Army knives, with the tools to adapt to any new task or role” [101]. Two primary ways that AI capabilities discourses point to these capabilities is via prompt examples and benchmark statistics—both of which, ostensibly evidencing model capabilities, extend the work of skills discourses in rendering social activities and aspects of people as segmentable, measurable, and valuable.

5.2.1 Prompt examples. One recurring way that AI capabilities discourses illustrate what AI products can do is by including examples of prompts for performing various activities. Prompts are often shown with the corresponding model outputs. These examples point to the products’ practical usefulness; for instance, OpenAI includes a series of videos of OpenAI employees, hunched over their laptops, narrating out loud as they enter prompts into ChatGPT to “build a budget,” “create a presentation,” “make an itinerary,” and “plan a date night.” That people are typing the prompts in, and watching the output get rendered piece-by-piece, underlines that things are being processed, computed, and thought about in real time, for the benefit of a real person doing real things.

Prompting has been taken up by AI companies as a key index of AI advancement. An OpenAI webpage on how “AI has been evolving at an incredible rate” illustrates this claim by displaying responses to prompts like “What would you say if you could talk to a future OpenAI model?” and “Write a limerick about a dog” for each GPT version, positioned along a timeline, implying that the higher quality of responses by more recent versions indicates “the progress made so far.” The linguistic and interactional qualities of using AI products often projects the impression of humanlike power [72, 105]. Examples of prompts, encapsulating these symbolic qualities, can be found on almost every webpage in our data.

It takes work to conceive of an activity as achievable via model prompting; articulating an activity in terms of any grammar of action necessarily reduces or alters it [6, 8]. The need for operationalization becomes especially

conspicuous for activities that do not seem straightforwardly achievable via stylized self-talk. That self-talk has emerged as a powerful technology of the self in skills discourses—suitable for workplace productivity and therapy alike—helps to smooth these gaps [46, 101]. In a webpage headlined “A Student’s Guide to Writing with ChatGPT,” for instance, OpenAI lists suggestions—with links to executable prompts—such as “complete your understanding by asking specific questions” and “test your logic with reverse outlining,” effectively rendering the articulation of complex ideas as AI-augmented self-talk.

Prompt examples draw on many other ideological presumptions to evoke the impression that AI can perform an activity. For instance, Google’s website for Gemini 3 highlights that one can “Learn anything” with Gemini. This is illustrated with a video depicting a PDF of a scientific paper being uploaded to Gemini, followed by a prompt, “I want to learn about ‘Scaling deep learning for materials discovery’ (see attached PDF). Create beautiful, elegant 3D interactive visuals that explain it.” The video goes on to show some examples of said visuals, while its caption glosses “learn anything” as “Understand complex topics in a way that makes sense for you—with clear, concise, and helpful responses.” Notably, neither the caption nor the video substantiate what exactly is meant by “learn[ing] anything,” so much as invoke an expanding space of associated qualities: “learning” as related to “complex topics,” of which advanced scientific research is emblematic; “understanding” as related to user-friendly responses that look scientific. All of these qualities are linked to the prompt example for “Learn anything,” as if they are imminently attainable via prompting Gemini.

Unsurprisingly, given the poetics of AI capabilities discourses, prompts almost always occur in lists, as if each list item were equally achievable via model-prompting. On the Gemini website, for instance, “Learn anything” is listed alongside two other capabilities, “Build anything” and “Plan anything,” both with their own video demonstrations: a prompt to “create a polished, retro-futuristic 3D spaceship web game” is followed by clips of the resultant AI-coded spaceship game; a prompt that simply reads “organize my inbox” is followed by a sequence that pans between generated task lists, email drafts and intermediate model output (“I’m solidifying the final steps”). Elsewhere on the same page, Google offers yet more “hands-on” examples: descriptions of even more capabilities (“Gain a deeper understanding through more nuanced responses”) attached to video demonstrations that, in similarly associational ways, package these capabilities—or, at least, exemplars vaguely and indexically attached to the words invoking them—into the prompt-output form.

5.2.2 Benchmark statistics. Another way that AI capabilities discourses convey model capabilities is via quantitative performance measures on various benchmark tests. Benchmark results are often organized into large tables to facilitate comparison: each row is a different test, while each column represents one of the company’s various model offerings, an earlier version of the model, or a main competitor (for instance, Google compares Gemini 3 Pro to Gemini 2.5 Pro, Claude Sonnet 4.5 and GPT-5.1, while Meta compares Llama models of various sizes to different versions of Gemini, DeepSeek, and GPT). Another common form is the bar graph; for instance, OpenAI’s blog post introducing GPT-5 includes a series of bar graphs visualizing GPT-5’s performance next to older OpenAI models. These metrics almost always come with footnotes specifying parameters or links to separate documents detailing evaluation procedures. Complementing the lively depictions of model usefulness in prompt examples, benchmark statistics give the sense that serious science was done to verify that the models are, indeed, that good.

Benchmarks are, of course, a long-standing feature in AI research [9, 66]. They are a standard way to mark scientific advancement, and, as the field has turned to the market, a standard way of comparing models across competing companies. In AI capabilities discourses, benchmarks serve as shorthands for technological advancement, and therefore, value. While benchmark results are usually presented somewhat separately from the qualities and use cases mentioned elsewhere in the data, one can still trace clear correspondences; for instance, that OpenAI lists “health” as a headline capability and “understanding cancer risk” as an exemplary use case for GPT-5 is backed

up elsewhere with a mention of state-of-the-art performance for HealthBench Hard; the website goes on to note that “those gains show up in everyday use.”

Numerous critiques of benchmarking [91, 120] have raised issues ranging from the choices of construct to be measured—including under-specified or reductive framings of broad concepts like “general knowledge” or “math” [28, 95]—to operationalizations [13, 33, 81], reporting practices [37, 97], and underlying “competitive epistemologies” [64, 89]. Yet benchmarking accords with a widespread social practice: there are many standardized tests for human performance, in line with the ideological investment in the idea of general—and economically productive—intelligence, and in tune with the quantifiability of skills [69, 78].

In fact, in AI capabilities discourses, results of academic benchmarks are often reported alongside those of human tests: as evidence of GPT-5’s mathematical capabilities, for instance, OpenAI lists performance on the AIME high school math exam; for GPT-4’s release, OpenAI additionally foregrounded the bar and sommelier exams. Regardless of whether these measures adequately capture math, law, and wine-tasting aptitude, they draw on presumptions that all of these are skills that are testable and rankable. Knitting together connotations of educational achievement, class and personhood, they additionally project a sense of cognitive capacity as affixed to social value. The semiotic functions of benchmarks are exemplified in the GPQA benchmark, which recurs across the data, and is often glossed as “PhD-level.” Completing a PhD has little to do with the format of the benchmark, which involves answering exam questions. The implication, rather, is that good benchmark performance indicates belonging with the uppermost echelons of the cognitive elite.

The substantive concerns and measurement models behind “Humanity’s Last Exam” (a “multimodal benchmark at the frontier of human knowledge”²), SWE-bench Verified (a test of how often models generate correct code), and the proprietary OpenAI benchmark referred to as “Economically important tasks” have very little in common. Nonetheless, that these metrics can all be packaged in a single table, or a series of similar-looking bar graphs, suggests that these things are all of a kind with each other, because they all result in a performance metric. In the context of AI capabilities discourses, this commensurability additionally facilitates comparison shopping: Meta and Google both list the monetary cost of using their models and their competitors’, per input and output token, in the same table as benchmark statistics for each model.

6 The social and political implications of AI capabilities discourses

Building on the ideologies inscribed in capitalist discourses, and the imagined superhuman qualities of AI, AI capabilities discourses reify and productize a notion of abstract cognitive power. The thinginess [108] of cognition is exemplified by OpenAI stating that GPT-5 comes with “built-in thinking,” comparing benchmark results for different versions of GPT-5 “with thinking” and “without thinking,” and offering a premium option, GPT-5 Pro, where users pay extra for “extended reasoning.” Such a product has many capabilities that can be used for many activities and to supplement many aspects of social life and personhood, all of which have likewise been packaged as things one can get by using AI.

AI capabilities discourses, like the skills discourses they build on, instrument a process of alienation: all manner of situated, relational, embodied practices are cast as things. While in skills discourses these things can be acquired via education and training, and then cashed out on the labour market, in AI capabilities discourses, it is perhaps more accurate to say that these things are *rented* from the AI product [100]. This arrangement is often framed as a time-saving measure; it is also, as Boltanski and Chiapello [31] write, a legible strategy for contending with precarity: “given the comparatively unpredictable character of fruitful projects, it is difficult to anticipate the kind of assets one might need [...] it is therefore reasonable to prefer ready, temporary access to resources that are borrowed, employed or expended in the framework of the project, while maintaining sufficient flexibility to refund them when required.” In the context of economies already marked by the depletion of labour conditions, AI

²<https://agi.safe.ai/>

capabilities discourses thereby rationalize shifts in the control of skills from workers “into a machine,” depressing wages and forming pretexts for mass layoffs [30].

Such shifts, moreover, are not constrained to the workplace. In their expansive scope, AI capabilities discourses essentially offer a “natural language interface for everyday life,” as articulated in an OpenAI webpage profiling the use of GPT-5 by a South Korean startup offering “lifestyle AI.” These are neat taglines for a product attuned to the totalizing governmentality of contemporary capitalism. In capitalist economies structurally predisposed to undermine peoples’ collective capacities for life, care, and sociality [4, 16, 58], to package such capacities as discretized rentable things has grotesque implications. For instance, the promotion of AI’s health and self-care capabilities—as seen in their preeminent status in OpenAI’s most recent marketing materials—disavows the idea that care is contingent on mutual interdependence among people, and instead underwrites deadly psychic attachments to AI products [23].

AI capabilities discourses, like their discursive antecedents, also instrument the alignment of capitalist class interests, as ways of marking a variety of fields as domains ripe for transformation by AI. For instance, AI has rapidly proliferated in schools and universities, with many institutions striking deals with AI companies to integrate their products in classrooms and on campuses [94]. Meanwhile, education is foregrounded across AI capabilities discourses. OpenAI features large collections of example prompts “made by students, for students,” and “faculty from a dozen disciplines;” Anthropic has a dedicated sales page for education, that seems conspicuously attuned to widely-circulating discourses that cast AI as fundamentally antithetical to teaching and to critical thought [e.g., 65]: it suggests Claude “asks the questions that help you find the answers yourself,” and “foster[s] critical thinking through guided exploration, prioritizing deep conceptual understanding over convenient shortcuts.” As we’ve noted, such statements are less about operationalizing teaching and critical thinking, and more about incorporating these activities into the loose semiotic field of AI capabilities. These discursive tropes are echoed when people in educational institutions rationalize investments in AI products; for instance, an announcement from the provost’s office of a major American university points to new generative AI tools that are “designed to be an expert in almost any topic you desire,” and that can be used for “building novel applications, making new discoveries, or connecting systems to optimize and enhance data flow and outputs.”³

Whether these capabilities will help “mak[e] AI work for educators,” in the context of these discourses, is besides the point. As Agre wrote of AI, “computing has been constituted as a kind of imperialism; it aims to reinvent virtually every other site of practice in its own image” [7]. In sectors like education, that have already undergone decades of restructuring to align with market imperatives, AI capabilities discourses are powerful tools for this reinvention: they serve as a prototype for AI industry spokespeople and university administrators alike to articulate shared interests that index the circuit of investments flowing between them. These discursive dynamics extend to other fields as well; both OpenAI and Anthropic have increasingly foregrounded healthcare, and more recently, government operations. As Urcioli writes, the semiotic indeterminacy of AI capabilities discourses—what is meant by “enhancing the capabilit[ies] of government workers” or “powering custom clinical solutions”—“index advantageous alignments based on shared values” [115]. In the present moment, these shared values are backed by substantial financial investments and political reconfigurations [18, 74].

We suggest that AI capabilities discourses help to shift the terrain on which people can contest the conditions of their labour, lives, and societies—not in terms of these conditions and of their collective desires for better ones, but in terms of whether these conditions can be reformulated as subject to AI capabilities, and more broadly, as subject to capitalist logics. AI capabilities discourses function as a technofuturist worldbuilding vernacular for the capitalist class. Meanwhile, in inscribing a skill-rental model, they discursively alienate people from “the practices and knowledge that come out of [...] one’s experience of work, of politics, of one’s own household life and neighborhood life” [47].

³<https://provost.umich.edu/making-ai-work-for-educators-a-u-m-generative-ai-update/>

It is not inevitable that AI capabilities discourses will succeed at reshaping the political terrain. As our analysis has traced, these discourses build on existing ways of making an unjust and incoherent world intelligible; their power, and the present conditions of capitalism, are linked together. But there are other ways to interpret the world, pointing to other ways to transform it. If the field of AI has embedded capitalist logics since its inception [63, 90], it took decades of discursive work, atop the consolidation of material and political resources, to articulate ways of valorizing AI's centrality in contemporary capitalism—and even then, not everyone is convinced. As Hall writes, “politics [is] where forces and relations, in the economy, in society, in culture, have to be actively worked on to produce particular forms of power” [68]. The active work of AI can be subverted; other work can be done.

7 Discussion and conclusion

In this paper, we've examined AI capabilities discourses as discourses that do ideological work for capitalism. AI capabilities discourses align a staggering variety of social activities with a regime of flexibilized, cognitive work that is structured by, and helps to reproduce, the conditions of contemporary capitalism. They render social activities as essentially fungible, packaging them as things that someone can use AI to do; at the same time, they cast AI as valuable—even central—to this regime.

Our account adds to existing critical understandings of discourses that sell AI: that they consist of misrepresentations of what AI can do—and what humans can do in comparison—that benefit the AI industry while devaluing human qualities [e.g., 20, 72]. Our work elaborates on how these discourses interact with the political and economic context in which they circulate. It is not just that quantitative and qualitative reports of what AI can do are misleading, and that the field's ambitions for artificial general intelligence are categorically impossible. What we suggest is that the AI industry, building on the totalizing scope of capitalism, projects an idea of how the world works that is *plausible*, because it is continuous with prevailing ideas about work and about personhood, and with the material conditions that these ideas are embedded in. This worldbuilding work, in turn, helps to narrow the terrain on which political action against AI, and against capitalism, can happen.

With this understanding of AI capabilities discourses as articulated with political economy, we offer some remarks on tactics for contesting AI development and proliferation. As we've outlined in §2, a large body of scholarship has engaged with these discourses as collections of mistaken or absurd claims, to be debunked, corrected, or ridiculed with better scientific techniques, deeper engagements with social scientific theory, and more measured accounts of how the technology really operates. Such strategies can serve as useful ground-clearing for broader efforts, but are limited in their ability to check the power of the AI industry, because they risk drawing attention away from the ideological work these discourses do [29].

In the context of AI capabilities discourses, claims about what AI can do carve social life into discrete skills, capabilities, and use cases, and re-render these things in terms of AI. If the point is to contest the broader systems of exploitation that the project of AI is woven into, then simply offering better assessments of the ability of an AI model to do productive cognitive work leaves intact the reifying logics that continue to alienate people while organizing capital. Critical scholarship on AI should therefore aim to interrupt the ideological work performed by these claims, dismantling the grounds on which they are sayable rather than simply interrogating their truth value. Our work demonstrates one potential analytic approach: troubling ready-at-hand framings of AI, excavating the historical continuities and political projects embedded in these framings, and widening the aperture to analyze the broader conditions in which discourses and ideologies are produced.

To conclude, we insist that critical engagements with AI must involve helping people build up the capacity for situated political contestation, shifting the terrain away from the regime of decontextualized skills and capabilities, and abstract cognitive power [16]. Here, we take up Baumer and Khovanskaya's orientation of *remantling*, “[attending] to the social relations, subjectivities, and temporal conditions through which technologies and practices take shape,” and “shift[ing] focus toward how people are already navigating and unsettling the

dominant logics of the present” [17]. If researchers wish to effectively critique the discourses of the AI industry, they should “identify, understand, and amplify” [17] efforts aimed at dismantling the ideological and political regimes that empower AI, towards building alternative futures.

After all, in recent years, there has been a resurgence of counterhegemonic political action, in labour organizing, transnational solidarity movements, and beyond [43, 75, 80, 99, 109, 110]. The process of political organizing—of recognizing, and then acting on this recognition of shared vulnerability and collective power [40]—has transformative potential, with respect to both material conditions and ideology [15]. Organizing practices like workers’ inquiries, for instance, enable people to elaborate “exact and positive knowledge” of their activities and working conditions [61, 83, 87, 129] in the context of broader political struggle. This conjunction of discursive and political work—taking place on the terrain of situated practice—is essential for subverting the discursive and political work performed by AI capabilities discourses, replacing the narrowly-circumscribed question of what AI can or cannot do with more capacious arenas of contestation, in which to imagine the sorts of worlds we would like to make together.

Acknowledgments

We’re grateful for the many helpful discussions that have shaped this work, including with Mark Ackerman, Ali Alkhatib, Shreya Chowdhary, Silvia Lindtner, Nasanbayar Ulzii-Orshikh, and the Political Economy and Algorithms Collective. We also thank the reviewers for their helpful feedback.

Generative AI Usage Statement. We did not use generative AI tools at any point in preparing this manuscript.

References

- [1] 2011. P21 Common Core Toolkit. <https://files.eric.ed.gov/fulltext/ED543030.pdf>
- [2] Mohamed Abdalla and Moustafa Abdalla. 2021. The Grey Hoodie Project: Big Tobacco, Big Tech, and the Threat on Academic Integrity. In *Proceedings of AIES*. ACM. doi:10.1145/3461702.3462563
- [3] Morgan Adamson. 2009. The human capital strategy. *Ephemera* 9, 4 (2009), 271–284.
- [4] Beatrice Adler-Bolton and Artie Vierkant. 2022. *Health Communism*. Verso Books.
- [5] Philip E. Agre. 1992. Formalization as a Social Project. <https://pages.gseis.ucla.edu/faculty/agre/formalization.html>
- [6] Philip E. Agre. 1994. Surveillance and capture: Two models of privacy. *The information society* 10, 2 (1994), 101–127.
- [7] Philip E. Agre. 1998. Toward a Critical Practice: Lessons Learned in Trying to Reform AI. In *Social Science, Technical Systems, and Cooperative Work Beyond the Great Divide*, Geoffrey Bowker, Susan Leigh Star, Les Gasser, and William Turner (Eds.). Taylor & Francis Group, 131–157.
- [8] Louise Amoore, SJ Bennett, Alexander Campolo, Benjamin Jacobsen, and Ludovico Rella. 2025. Politics of the prompt: Government in the age of generative AI. *Economy and Society* 54, 3 (2025), 573–596.
- [9] Ashton Anderson, Dan Jurafsky, and Dan McFarland. 2012. Towards a computational history of the ACL: 1980-2008. In *Proceedings of the ACL-2012 Special Workshop on Rediscovering 50 Years of Discoveries*. 13–21.
- [10] Stefan Baack. 2024. A Critical Analysis of the Largest Source for Generative AI Training Data: Common Crawl. In *Proceedings of FAccT*. Association for Computing Machinery, New York, NY, USA, 2199–2208. doi:10.1145/3630106.3659033
- [11] Erik Baker. 2025. *Make your own job*. Harvard University Press, Cambridge.
- [12] Jascha Bareis and Christian Katzenbach. 2021. Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values* 47, 5 (July 2021), 855–881.
- [13] Teanna Barrett, Quanze Chen, and Amy Zhang. 2023. Skin Deep: Investigating Subjectivity in Skin Tone Annotations for Computer Vision Benchmark Datasets. In *Proceedings of FAccT*. Association for Computing Machinery, New York, NY, USA, 1757–1771. doi:10.1145/3593013.3594114
- [14] Caroline Bassett and Ben Roberts. 2019. Automation Now and Then: Automation Fevers, Anxieties and Utopias. *New Formations* 98, 98 (2019), 9–28.
- [15] Alyssa Battistoni. 2019. Spadework. *n+1* (2019).
- [16] Alyssa Battistoni. 2025. Ideology at work? Rethinking reproduction. *American Political Science Review* 119, 3 (2025), 1205–1218.
- [17] Eric PS Baumer and Vera Khovanskaya. 2025. Sociotechnical remantling. In *Proceedings of the sixth decennial Aarhus conference: Computing X Crisis*. 27–41.

- [18] Burcu Baykurt. 2025. Gov-tech as capture: public infrastructures under data capitalism. *Information, Communication & Society* (2025), 1–16.
- [19] Aaron Benanav. 2019. Automation and the Future of Work—Part 1. *New Left Review* 119 (2019), 5–38.
- [20] Emily M. Bender and Alex Hanna. 2025. *The AI con*. Vintage Publishing, London.
- [21] Emily M. Bender and Alexander Koller. 2020. Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of ACL*. doi:10.18653/v1/2020.acl-main.463
- [22] Ruha Benjamin. 2019. *Race after technology: abolitionist tools for the New Jim Code*. Polity Press, Newark.
- [23] Johana Bhuiyan. 2025. ChatGPT encouraged Adam Raine’s suicidal thoughts. His family’s lawyer says OpenAI knew it was broken. *The Guardian*. <https://www.theguardian.com/us-news/2025/aug/29/chatgpt-suicide-openai-sam-altman-adam-raine>
- [24] Kean Birch. 2025. Do artifacts have political economy? *Science, Technology, & Human Values* (2025).
- [25] Abeba Birhane. 2021. The impossibility of automating ambiguity. *Artificial life* 27, 1 (2021), 44–61.
- [26] Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. The Values Encoded in Machine Learning Research. In *Proceedings of FAcCT*. Association for Computing Machinery, New York, NY, USA, 173–184. doi:10.1145/3531146.3533083
- [27] Abeba Birhane and Marek McGann. 2024. Large models of what? Mistaking engineering achievements for human linguistic agency. *Language Sciences* 106 (Nov. 2024). doi:10.1016/j.langsci.2024.101672
- [28] Borhane Blili-Hamelin and Leif Hancox-Li. 2023. Making Intelligence: Ethical Values in IQ and ML Benchmarks. In *Proceedings of FAcCT*. Association for Computing Machinery, New York, NY, USA, 271–284. doi:10.1145/3593013.3593996
- [29] Hagen Blix and Ingeborg Glimmer. 2025. Deflating “Hype” Won’t Save Us. *Liberal Currents*. <https://www.liberalcurrents.com/deflating-hype-wont-save-us/>
- [30] Hagen Blix and Ingeborg Glimmer. 2025. *Why we fear AI*. Common Notions, Brooklyn, NY.
- [31] Luc Boltanski and Eve Chiapello. 2018. *The new spirit of capitalism*. Verso Books.
- [32] Clea Bourne. 2024. AI hype, promotional culture, and affective capitalism. *AI and Ethics* 4, 3 (2024), 757–769. doi:10.1007/s43681-024-00483-w
- [33] Samuel R. Bowman and George Dahl. 2021. What Will it Take to Fix Benchmarking in Natural Language Understanding?. In *Proceedings of NAACL*. 4843–4855. doi:10.18653/v1/2021.naacl-main.385
- [34] J. Scott Brennen, Philip N. Howard, and Rasmus Kleis Nielsen. 2018. An industry-led debate. doi:10.60625/RISJ-V219-D676
- [35] Robert Brenner. 2006. *The economics of global turbulence: the advanced capitalist economies from long boom to long downturn, 1945-2005*. Verso.
- [36] Wendy Brown. 2024. Foreword. In *Capital, Volume 1*. Princeton University Press.
- [37] Ryan Burnell, Wout Schellaert, John Burden, Tomer D Ullman, Fernando Martinez-Plumed, Joshua B Tenenbaum, Danaja Rutar, Lucy G Cheke, Jascha Sohl-Dickstein, Melanie Mitchell, et al. 2023. Rethink reporting of evaluation results in AI. *Science* 380, 6641 (2023), 136–138.
- [38] Adam Cancryn and Megan Messerly. 2023. Biden’s health care wins are being undone — and at the worst possible time. *Politico*. <https://www.politico.com/news/2023/08/09/bidens-health-care-wins-undone-medicare-00110389>
- [39] Gabrielle Canon. 2026. How the climate crisis showed up in Americans’ lives this year: ‘The shift has been swift and stark’. *The Guardian*. <https://www.theguardian.com/us-news/2025/dec/31/climate-crisis-guardian-readers>
- [40] Lillian Cicerchia. 2021. Why does class matter? *Social Theory and Practice* 47, 4 (2021), 603–627.
- [41] Anna Clark. 2025. *Toward the Higher- and Secondary-Ed Alliance!* Public Books. <https://www.publicbooks.org/toward-the-higher-and-secondary-ed-alliance/>
- [42] Adele E Clarke. 2003. Situational analyses: Grounded theory mapping after the postmodern turn. *Symbolic interaction* 26, 4 (2003), 553–576.
- [43] Joshua Clover. 2016. *Riot. Strike. Riot: The New Era of Uprisings*. Verso Books.
- [44] Melinda Cooper. 2024. *Counterrevolution: Extravagance and Austerity in Public Finance*. Princeton University Press.
- [45] Jenny L Davis, Apryl Williams, and Michael W Yang. 2021. Algorithmic reparation. *Big Data & Society* 8, 2 (2021), 20539517211044808.
- [46] Alfonso Del Percio and Sze Wan Vivian Wong. 2019. Resetting minds and souls: Language, employability and the making of neoliberal subjects. In *Language and neoliberal governmentality*. Routledge, 190–210.
- [47] Michael Denning. 2023. Michael Denning on Antonio Gramsci and Hegemony. *Jacobin*. <https://jacobin.com/2023/01/michael-denning-antonio-gramsci-prison-notebooks-theory-hegemony-class-organizing>
- [48] Fernando Diaz and Michael Madaio. 2024. Scaling Laws Do Not Scale. *Proceedings of AIES* (2024), 341–357. doi:10.1609/aies.v7i1.31641
- [49] Emma Dowling. 2022. *The care crisis: What caused it and how can we end it?* Verso Books.
- [50] Hubert L. Dreyfus. 1979. *What computers can’t do*. Harper & Row.
- [51] Veena B Dubal. 2022. Economic security & the regulation of gig work in California: From AB5 to Proposition 22. *European labour law journal* 13, 1 (2022), 51–65.
- [52] Alexandre Duchêne and Monica Heller. 2012. *Language in late capitalism: Pride and profit*. Routledge.

- [53] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. Macmillan+ ORM.
- [54] Norman Fairclough. 2001. Critical discourse analysis as a method in social scientific research. In *Methods of critical discourse analysis*.
- [55] Luciano Floridi. 2024. Why the AI hype is another tech bubble. *Philosophy & Technology* 37, 4 (2024), 128.
- [56] Michel Foucault. 1988. Technologies of the self. In *Technologies of the Self: A Seminar with Michel Foucault*.
- [57] Marion Fourcade and Kieran Healy. 2024. *The Ordinal Society*. Harvard University Press. doi:10.4159/9780674296688
- [58] Nancy Fraser. 2014. Behind Marx's hidden abode: For an expanded conception of capitalism. *New left review* 86, 86 (2014), 55–72.
- [59] Garth Friesen. 2025. Unemployment Crisis: Why College Graduates Are Struggling In 2025. *Forbes*. <https://www.forbes.com/sites/garthfriesen/2025/10/23/unemployment-crisis-why-college-graduates-are-struggling-in-2025/>
- [60] Susan Gal and Judith T Irvine. 2019. *Signs of difference: Language and ideology in social life*. Cambridge University Press.
- [61] Cailean Gallagher, Karen Gregory, and Boyan Karabaliev. 2025. Digital worker inquiry and the critical potential of participatory worker data science for on-demand platform workers. *New Technology, Work and Employment* 40, 1 (2025), 20–40.
- [62] Timnit Gebru and Émile P. Torres. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday* (2024). doi:10.5210/fm.v29i4.13636
- [63] David Golumbia. 2009. *The cultural logic of computation*. Harvard University Press.
- [64] Gabriel Grill. 2024. Constructing Capabilities: The Politics of Testing Infrastructures for Generative AI. In *Proceedings of FAccT*. 1838–1849. doi:10.1145/3630106.3659009
- [65] Olivia Guest, Marcela Suarez, Barbara Müller, Edwin van Meerkerk, Arnoud Oude Groote Beverborg, Ronald de Haan, Andrea Reyes Elizondo, Mark Blokpoel, Natalia Scharfenberg, Annelies Kleinherenbrink, et al. 2025. Against the Uncritical Adoption of “AI” Technologies in Academia. (2025).
- [66] Sireesh Gururaja, Amanda Bertsch, Clara Na, David Widder, and Emma Strubell. 2023. To build our future, we must know our past: Contextualizing paradigm shifts in natural language processing. In *Proceedings of EMNLP*. 13310–13325.
- [67] Eli Hager. 2025. The Untold Saga of What Happened When DOGE Stormed Social Security. ProPublica. <https://www.propublica.org/article/inside-doge-social-security-takeover-leland-dudek>
- [68] Stuart Hall. 1991. Postscript: Gramsci and us. *Gramsci's political thought: An introduction* (1991), 114–130.
- [69] Malcolm Harris. 2023. *Palo Alto: A history of California, capitalism, and the world*. Hachette UK.
- [70] Theresa Heyd. 2025. Presentism as a digital language ideology in generative AI discourse. *Internet Pragmatics* (Nov. 2025). doi:10.1075/ip.00132.hey
- [71] Arlie Russell Hochschild. 2012. *The managed heart*. University of California Press.
- [72] Sun-ha Hong. 2024. Strategic misrecognition and speculative rituals in generative AI. 6, 4 (2024), 92–106. doi:10.33621/jdsr.v6i440474
- [73] John Horgan. 2020. Will artificial intelligence ever live up to its hype. *Scientific American* 4 (2020).
- [74] Linda Huber. 2025. Designing markets, governing data: Engineering value in the American healthcare system. *Big Data & Society* 12, 3 (2025).
- [75] Sarah Jaffe. 2025. Federal Workers Rise up Against Musk, Trump and Drastic Cuts. In *These Times*. <https://inthesetimes.com/article/federal-workers-rise-up-musk-trump-doge>
- [76] Pratyusha Ria Kalluri, William Agnew, Myra Cheng, Kentrell Owens, Luca Soldaini, and Abeba Birhane. 2025. Computer-vision research powers surveillance technology. *Nature* (2025), 1–7.
- [77] Shivani Kapania, Stephanie Ballard, Alex Kessler, and Jennifer Wortman Vaughan. 2025. Examining the Expanding Role of Synthetic Data Throughout the AI Development Pipeline. In *Proceedings of FAccT*. Association for Computing Machinery, New York, NY, USA, 45–60. doi:10.1145/3715275.3732005
- [78] Yarden Katz. 2022. Intelligence under racial capitalism: From eugenics to standardized testing and online learning. *Monthly Review* 1 (2022).
- [79] Os Keyes. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 88 (Nov. 2018), 22 pages. doi:10.1145/3274357
- [80] James Kilgore. 2024. Palestine Has Reignited US Left Internationalism. Can We Keep the Fire Alive? Truthout. <https://truthout.org/articles/palestine-has-reignited-us-left-internationalism-can-we-keep-the-fire-alive/>
- [81] Noya Kohavi and Anna Weichselbraun. 2025. Human tests for machine models: What lies “Beyond the Imitation Game”? *Journal of Linguistic Anthropology* (2025).
- [82] Alva Markelius, Connor Wright, Joahna Kuiper, Natalie Delille, and Yu-Ting Kuo. 2024. The mechanisms of AI hype and its planetary and social costs. *AI and Ethics* 4, 3 (2024), 727–742.
- [83] Karl Marx. 1880. A Workers' Inquiry. <https://www.marxists.org/archive/marx/works/1880/04/20.htm>
- [84] Karl Marx. 2024. Capital: Critique of Political Economy, Volume 1. In *Capital*. Princeton University Press. Translated by Paul Reitter.
- [85] Bonnie McElhinny. 2012. Silicon Valley Sociolinguistics?: Analyzing Language, Gender and Communities of Practice in the New Knowledge Economy. In *Language in late capitalism*. Routledge, 230–260.
- [86] Kamaron McNair. 2025. As the cost of living keeps climbing, Americans are 'trying to muscle through' the affordability crisis, analyst says. *CNBC*. <https://www.cnbc.com/2025/12/26/how-americans-are-responding-to-the-affordability-crisis.html>

- [87] Milagros Miceli, Adio-Adet Dinika, Krystal Kauffman, Camilla Salim Wagner, Laurenz Sachenbacher, Alex Hanna, and Timnit Gebru. 2025. Methodological Considerations for Centering Workers' Epistemic Authority in AI Research. In *Proceedings of AIES*, Vol. 8. 1698–1710.
- [88] Arvind Narayanan and Sayash Kapoor. 2024. *AI snake oil*. Princeton University Press.
- [89] Will Orr and Edward B. Kang. 2024. AI as a Sport: On the Competitive Epistemologies of Benchmarking. In *Proceedings of FAcCT*. Association for Computing Machinery, New York, NY, USA, 1875–1884. doi:10.1145/3630106.3659012
- [90] Matteo Pasquinelli. 2023. *The eye of the master: A social history of artificial intelligence*. Verso Books.
- [91] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna. 2021. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns* 2, 11 (2021).
- [92] Jonathan Payne. 2000. The unbearable lightness of skill: the changing meaning of skill in UK policy discourses and some implications for education and training. *Journal of Education Policy* 15, 3 (2000), 353–369.
- [93] John Peters. 2008. Labour market deregulation and the decline of labour power in North America and Western Europe. *Policy and Society* 27, 1 (2008), 83–98.
- [94] Justin Raden. 2025. Higher Ed's Rush To Adopt AI Is About So Much More Than AI. Defector. <https://defector.com/higher-eds-rush-to-adopt-ai-is-about-so-much-more-than-ai>
- [95] Deborah Raji, Emily Denton, Emily M. Bender, Alex Hanna, and Amandalynne Paullada. 2021. AI and the Everything in the Whole Wide World Benchmark.
- [96] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *Proceedings of FAcCT*. Association for Computing Machinery, New York, NY, USA, 959–972. doi:10.1145/3531146.3533158
- [97] Anka Reuel, Amelia Hardy, Chandler Smith, Max Lamparth, Malcolm Hardy, and Mykel J. Kochenderfer. 2024. BetterBench: Assessing AI Benchmarks, Uncovering Issues, and Establishing Best Practices. In *NeurIPS*, Vol. 37. 21763–21813. doi:10.52202/079017-0685
- [98] Paola Ricaurte. 2022. Ethics for the majority world: AI and the question of violence at scale. *Media, Culture & Society* 44, 4 (May 2022), 726–745. doi:10.1177/01634437221099612
- [99] Tracy Rosenthal and Leonardo Vilchis. 2024. *Abolish Rent: How Tenants Can End the Housing Crisis*. Haymarket Books.
- [100] Jathan Sadowski. 2025. *The Mechanic and the Luddite: A Ruthless Criticism of Technology and Capitalism*. University of California Press.
- [101] Lily Scherlis. 2024. Skill Issues: Dialectical Behavior Therapy and Its Discontents. *The Drift* 13 (2024).
- [102] Paul Schütze. 2024. The impacts of AI futurism: an unfiltered look at AI's true effects on the climate crisis. *Ethics and Information Technology* 26, 2 (2024). doi:10.1007/s10676-024-09758-6
- [103] Heather K Scott. 2004. Reconceptualizing the nature and health consequences of work-related insecurity for the new economy: The decline of workers' power in the flexibility regime. *International Journal of Health Services* 34, 1 (2004), 143–153.
- [104] Mona Sloane, David Danks, and Emanuel Moss. 2024. Tackling ai hyping. *AI and Ethics* 4, 3 (2024), 669–677.
- [105] Luke Stark. 2024. Animation and Artificial Intelligence. In *Proceedings of FAcCT*. 1663–1671. doi:10.1145/3630106.3658995
- [106] Luke Stark and Jevan Hutson. 2021. Physiognomic artificial intelligence. *Fordham Intell. Prop. Media & Ent. LJ* 32 (2021), 922.
- [107] Lucy Suchman. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge university press.
- [108] Lucy Suchman. 2023. The Uncontroversial 'thingness' of AI. *Big Data & Society* 10, 2 (July 2023). doi:10.1177/20539517231206794
- [109] Cella M Sum, Anna Konvicka, Mona Wang, and Sarah E Fox. 2025. The Future of Tech Labor: How Workers are Organizing and Transforming the Computing Industry. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–30.
- [110] JS Tan, Natalia Luka, and Emily Mazo. 2025. Unlikely organizers: The rise of tech worker labor activism. *ILR Review* (2025), 00197939251375319.
- [111] JS Tan and Kathleen Thelen. 2025. Cloud Capitalism and the AI Transition. *Politics & Society* (2025), 00323292251396395.
- [112] Savannah Thais. 2024. Misrepresented technological solutions in imagined futures: The origins and dangers of ai hype in the research community. In *Proceedings of AIES*, Vol. 7. 1455–1465.
- [113] Adam Tooze. 2022. Welcome to the world of the polycrisis. *The Financial Times*. <https://www.ft.com/content/498398e7-11b1-494b-9cd3-6d669dc3de33>
- [114] Bonnie Urciuoli. 2008. Skills and selves in the new workplace. *American ethnologist* 35, 2 (2008), 211–228.
- [115] Bonnie Urciuoli. 2025. Revisiting the Strategically Deployable Shifter: Manipulating Indeterminacy for Semiotic Power and Profit. *Signs and Society* 13, 1 (2025), 9–24.
- [116] Ana Valdivia. 2025. The supply chain capitalism of AI: A call to (re) think algorithmic harms and resistance through environmental lens. *Information, Communication & Society* 28, 12 (2025), 2118–2134.
- [117] Steven Vallas and Juliet B. Schor. 2020. What Do Platforms Do? Understanding the Gig Economy. *Annual Review of Sociology* 46, Volume 46, 2020 (2020), 273–294. doi:10.1146/annurev-soc-121919-054857
- [118] Shannon Vallor. 2024. *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking*. Oxford University Press. doi:10.1093/oso/9780197759066.001.0001
- [119] Luc van der Gun and Olivia Guest. 2024. Artificial Intelligence: Panacea or Non-Intentional Dehumanisation? *Journal of Human-Technology Relations* 2 (2024).

- [120] Hanna Wallach, Meera Desai, Nicholas Pangakis, A Feder Cooper, Angelina Wang, Solon Barocas, Alexandra Chouldechova, Chad Atalla, Su Lin Blodgett, Emily Corvi, et al. 2025. Evaluating Generative AI Systems is a Social Science Measurement Challenge. In *ICML*.
- [121] Kevin Wei, Carson Ezell, Nick Gabrieli, and Chinmay Deshpande. 2024. How Do AI Companies “Fine-Tune” Policy? Examining Regulatory Capture in AI Governance. *Proceedings of AIES 1* (Oct. 2024), 1539–1555. doi:10.1609/aies.v7i1.31745
- [122] Joseph Weizenbaum. 1976. *Computer power and human reason*. Freeman, San Francisco.
- [123] Meredith Whittaker. 2021. The steep cost of capture. *Interactions* 28, 6 (2021), 50–55.
- [124] David Gray Widder and Mar Hicks. 2024. Watching the generative AI hype bubble deflate. *arXiv preprint arXiv:2408.08778* (2024).
- [125] David Gray Widder, Sarah Myers West, and Meredith Whittaker. 2023. Open (for business): Big tech, concentrated power, and the political economy of open AI. *Nature* (2023).
- [126] Gabriel Winant. 2021. The next shift: The fall of industry and the rise of health care in Rust Belt America. In *The next shift*. Harvard University Press.
- [127] Ruth Wodak, Michael Meyer, et al. 2009. Critical discourse analysis: History, agenda, theory and methodology. In *Methods of critical discourse analysis*. Vol. 2. 1–33.
- [128] Blaise Agüera y Arcas. 2023. Artificial general intelligence is already here. *Noema* (2023).
- [129] Kevin Zheng, Linda Huber, Aaron Stark, Nathan Kim, Francesca Lameiro, Wells Lucas Santo, Shreya Chowdhary, Eugene Kim, and Justine Zhang. 2025. Resisting AI Solutionism through Workplace Collective Action. *arXiv preprint arXiv:2508.08313* (2025).

Appendix

Below, we list links to the documents we quote from in this paper.

Anthropic.

- <https://claude.com/product/overview>
- <https://www.youtube.com/watch?v=FDNkDBNR7AM>
- <https://claude.ai/>
- <https://website.claude.com/pricing/max>
- <https://claude.com/solutions/education>

Google.

- <https://www.youtube.com/watch?v=98DcoXwGX6I>
- <https://deepmind.google/models/gemini/>
- <https://www.youtube.com/watch?v=MZCpgTi-Iys>
- <https://deepmind.google/models/gemini/pro/>
- <https://gemini.google/overview/deep-research/>
- <https://gemini.google/overview/gems/>
- <https://gemini.google/ca/students/>

Meta.

- <https://www.llama.com/>
- <https://www.llama.com/resources/case-studies/sevilla-football-club/>
- https://www.llama.com/resources/case-studies/case_based_research/
- <https://www.llama.com/resources/case-studies/sofya/>
- <https://www.llama.com/resources/case-studies/oxide-ai/>
- <https://arxiv.org/abs/2407.21783>

OpenAI.

- <https://openai.com/index/introducing-gpt-5/>
- <https://chatgpt.com/overview>
- <https://chatgpt.com/>

- <https://chatgpt.com/features/agent/>
- <https://progress.openai.com/>
- <https://openai.com/index/gpt-4-research/>
- <https://chatgpt.com/use-cases/students/>
- <https://chatgpt.com/use-cases/university-educators/>
- <https://chatgpt.com/use-cases/student-writing-guide/>
- <https://openai.com/index/booking-com/>
- <https://openai.com/index/wrtn/>
- <https://openai.com/global-affairs/introducing-openai-for-government/>
- <https://openai.com/index/openai-for-healthcare/>