

The predatory fantasy of worker empowerment in AI marketing

JUSTINE ZHANG, University of Michigan, USA

SU LIN BLODGETT, Microsoft Research, Canada

NINA MARKL, University of Essex, UK

1 INTRODUCTION

A major recent advancement in the field of AI has been the packaging of AI models as products. Companies like OpenAI and Anthropic, ostensibly concerned with the engineering of superintelligent machines, now wrap fantasies of artificial general intelligence in sales pitches. For instance, the latest version of Claude is not just Anthropic’s “most intelligent model to date” [5]; it’s also “your thinking partner” that “solves your toughest challenges” while “amplifying your creative potential,” purchasable for yourself, your team, your enterprise, or your educational institution [6]. Just a few years earlier, the idea of AI loomed ominously over the labour market—humanlike robots would one day replace workers, threatening mass unemployment [12]. However, Claude and its competitors are now framed by their vendors as tools that *empower* workers, materializing a long-standing orientation towards technology in HCI and design circles expressed in keywords like “human-centred AI” and in Steve Jobs’ metaphor of computers as “bicycles for the mind” [33].

Critical computing scholars have repeatedly demonstrated how promises of technologically-enacted empowerment produce positive feelings about capitalist social orders while casting their structural injustices as private personal problems [25, 30, 34]. Adopting this sensibility, we examine tech companies’ messaging of AI as products for empowering workers. There is a stick to the carrot, of course—those who do not join the many organizations featured in positive testimonials (“GPT-4 deepens the conversation on Duolingo” [41]; “CodeGPT leverages Llama to boost developer productivity” [35]) will fall to the wayside in the AI-driven world that the AI industry is invested in bringing into being. AI, after all, is supposed to be for everyone, interpellated by the unspecified “you” recurring across companies’ websites. What does it say about you if you fail to avail yourself of the many learning resources, demonstrative “recipes” [38], and user-friendly interfaces these companies offer, in preparing for the future of work?

We suggest that AI companies have grafted a fantasy of AI-empowered work onto longer-running imaginaries that idealize certain types of workers. In particular, recent discourses about AI, that cast it as an endlessly-capable and rapidly advancing technology that “helps you do your best work” [4], are converging with another prevalent discourse that’s emerged over the last few decades. This discourse casts the ideal worker as possessing a diverse portfolio of skills, always self-improving, with an unshakeably positive disposition towards a labour market rendered dynamic and uncertain by inevitable technological progress [8, 55]. The figure of this so-called entrepreneurial worker indexes an unforgiving set of social conditions, as decades of sluggish economic growth have led to a persistently low demand for labour [9]. Capitalist economies have always come with ideologies for inducing workers to affirmatively opt into their own exploitation [14]. This reification of entrepreneurialism—as motivation for workers to press on in the face of unending precarity and seize whatever opportunities come along; as mechanism of responsabilization for their all-too-likely failures—is the latest iteration, tailored to the present moment.

Via critically analyzing the marketing performed by AI companies in the course of their project of political domination, we therefore argue that the AI industry exacts a steep ideological cost. Under the guise of empowerment, the fantasy of

Authors’ addresses: Justine Zhang, tisjune@umich.edu, University of Michigan, Ann Arbor, Michigan, USA; Su Lin Blodgett, sulin.blodgett@microsoft.com, Microsoft Research, Montréal, Quebec, Canada; Nina Markl, nina.markl@essex.ac.uk, University of Essex, Colchester, UK.

the AI-empowered worker entrenches workers' participation in systems that exploit them, and that tether them to the broader infrastructures of predation that constitute modern AI.

2 THE LATEST TRENDS IN AI MARKETING

In ongoing research, we've been performing a critical discourse analysis [20] of the messaging produced by AI companies to announce different versions of their large language models (LLMs). We've collected a corpus of webpages and technical reports, from OpenAI's release of GPT-1 in 2018 to Anthropic's release of Claude 4 in 2025. We find that companies' framings of what LLMs *are* have shifted, corresponding to an increasingly conspicuous fusion of academia and industry. Early models were presented as scientific objects (e.g., [18, 46]) in papers addressed to computer science researchers; announcements of the most recent models are primarily addressed to the people and firms who might potentially pay for the LLM products on offer. While companies have maintained an interest in artificial general intelligence (OpenAI: "our mission is to ensure that artificial general intelligence benefits all of humanity" [42]), they now market their models as the ultimate general-purpose tools.

The discursive construction of LLMs as tools, in addition to artificial intelligences, embeds a remarkably instrumental view of language: what matters is not language's semiotic content or its social value per se, so much as its utility for performing tasks. Per OpenAI's paper describing GPT-2, "language provides a flexible way to specify tasks, inputs, and outputs all as a sequence of symbols" [47]; a key payoff of collecting a large training dataset is that it would contain "natural language demonstrations of tasks in as varied of domains and contexts as possible" [47], that a model could potentially take up. When GPT-2 was released, "tasks" had a relatively narrow, field-specific meaning: it pointed to the sorts of evaluations that, to natural language processing researchers, tested models for their language understanding abilities, like question answering, machine translation, and summarization. However, the paper also mentions more speculatively that "GPT-2 is able to write news articles about the discovery of talking unicorns" [47].

The scope of tasks that language can specify, and that LLMs can perform, has since ballooned; "just ask and ChatGPT can help with writing, learning, brainstorming, and more" [43]. Or, as rendered on Anthropic's website, "right now, people are asking Claude to"—followed by a span of text that changes every few seconds, listing such tasks as "review a job contract," "debug a wedding website," or "create a 30 day meal plan" [6]. Indeed, companies often format their descriptions of the apparently immense power of recent LLMs as lists of things. Alongside use cases, such lists could contain capabilities—per Meta, Llama 4's "class-leading capabilities" include "natively multimodal," "unparalleled long context," "expert image grounding," and "multilingual writing" [37]—and quantitative benchmarks—OpenAI's o4-mini model "outperforms its predecessors" on evaluation frameworks for coding, "agentic tool use," and multimodal reasoning, alongside high school math competitions [44]. These lists have an open-ended quality: they could be, and sometimes are, concluded with "and more." They are not meant to enumerate a closed ontology, as in earlier papers listing conventionalized NLP evaluations; rather, they suggest an unbounded landscape.

As numerous critics have shown, such claims are replete with logical faults [28], starting from the absurd premise of "general purpose" [48]. It is unclear how the generative paradigm LLMs use could appropriately perform the vast majority of functions on offer [10, 32, 52]. List entries are often clearly not of the same type (is "long context," i.e., the capacity to process large inputs, comparable to the ability to write in some unstated number of languages?). Nor could the large number of listed things plausibly be construed as verifiable or quantifiable; many commentators have emphasized that claims made by companies about their LLMs tend to be inflated in self-interested ways [26] or muddled by the contingencies of concept operationalization and measurement [31, 57].

3 AI-AUGMENTED SKILLS DISCOURSES

Rather than parsing AI marketing discourses as bad science, we interpret them in relation to what Urciuoli [55] terms “skills discourses,” that construct workers as “bundles of skills,” and that advise or exhort workers on the skills they should acquire or enhance. Skills, such as leadership, Photoshop, and French, “cover a range of disparate practices, knowledge, and ways of acting and being” [55]. Collections of skills are not coherent because they are of a type, but rather because they index certain social conditions with certain norms for interpreting them: “skills [are] aspects of personhood with exchange value on the labour market” [55], legible to managers and potential employers, and they are assumed to be comparable and quantifiable, even if specifically how is left open. There are many possible skills (as one career service website suggests, “each person has approximately 700 different skills in their repertoire” [55]); collectively, they triangulate a particular, post-Fordist world. Skills—often “soft,” or pertaining to communication—are highly valued in the information and service economies emblematic of the present [45, 51, 55].

Indeed, the peculiar discursive character of “skills” links to the social conditions in which skills discourses circulate. Since the deindustrialization of Western economies in the 1970s, the “skill” category has “expanded almost exponentially” [45], shifting from primarily technical abilities in particular crafts to vaguer, “softer,” personal qualities. In an increasingly precarious labour market, “workers had to become their own managers, eking out a living from serial opportunities” [51]. “Skills” enable such entrepreneurial workers to develop “a stock of qualifications” like a “portfolio” [14], thus maximizing their adaptability to any job that comes their way [8]. Recent additions to the skills pantheon that have cycled through with the arrival of the digital, social media, and AI eras, and amid the unceasingly dismal labour market, include coding, content creation, and prompt engineering.

We suggest that the discourses produced by AI companies can be read as riffs on skills discourses, appending to “skills” “capabilities” and “tasks” that collectively evoke a multivariate information-age aptitude. If being a good worker involves the accumulation of skills, why not use an AI helper that can increase your skill portfolio, thus “expand[ing] what you can do” [6]? Indeed, many of the example use cases listed across company websites—develop a sales pitch, debug some code, plot some data—point to standard competencies of a worker that could plausibly be employed at one of the many companies offering product testimonials. Just as skills discourses blur the lines between self and work [51, 55], such use cases are interleaved with others that are more oriented towards self-management—create a morning routine, plan a “mental health day” [43], help to pick an outfit.

Having AI handle all of these things means that one is capitalizing on the latest technological advances in “frontier intelligence” [6] to “find new insights,” “increase productivity” [43] and “build your greatest ideas” [36]—all prototypical qualities of the entrepreneurial worker-self. Indeed, LLMs are perhaps ideal instruments for a contemporary, highly contextualized mode of self-improvement, characterized by “textual practices that allow individuals to reflect on themselves and to monitor or adjust their life projects and ambitions” [17], and inculcated through professional training schemes and care modalities like cognitive behavioural therapy [51, 58]. For instance, to “give Claude a try,” Anthropic suggests such inputs as “How can I build a meaningful side project in life’s margins?” and “After 15 years of real work, how do I write as myself, not another ‘hustle culture’ clone?” [6]. If “language is the medium through which a neoliberal rationality is circulated” [17], then perhaps LLM-as-worker-empowering-product represents neoliberalism’s pinnacle.

4 RISKS AND MITIGATIONS

AI marketing discourses evince a brutal indifference to the myriad things that people do. It is unclear whether LLMs can “extract patterns in these PDFs” [6], “give me ideas for what to do with my kids’ art” [43], and “write a beginner’s

guide to gardening” [6] in ways that are meaningfully accountable to the particular social relations that structure such activities [10]. Moreover, as with skills discourses, to list these activities in a boundless, undifferentiated collection annihilates the sociocultural distinctions between them [55]. Stripped of substance, they are rendered as comparable, interpretable with respect to their exchange value, and doable with the help of whichever AI product is being sold. As such, in evoking a spectacular fantasy of “doing your best work,” these discourses “enact a widespread devaluation of the human qualities that the technology purports to reproduce” [29], projecting a regime where workers do not own but rather rent a staggering diversity of means of production from AI companies [49].

“Empowered” by AI, the entrepreneurial worker reproduced in such discourses cannot escape their domination [15]. AI companies may promise workers enhanced powers to make work better, but possible improvements are limited, absent a reversal of long-term economic trends, more redistributive social programs, and less exploitative workplaces. The dire need for political action is instead laundered as a matter of personal responsibility, “neatly resolv[ing] the contradiction between the imperative to work under capitalism and the scarcity of work created by its advanced corporate form” [8]. AI will help you become your best self in work and life; if that doesn’t happen, you’re not using it correctly. The fantasy of AI thus instruments the structural disempowerment of workers, subscribing them to an endless cycle of exhaustion [13] while foreclosing their recognition of their collective vulnerability [8, 15].

Enjoined to make a neocolonial infrastructure of land and labour expropriation [39, 53, 56]—packaged as a sleek user interface—a central part of their work lives, workers are cast as the ultimate unethical consumer under capitalism. AI companies do nod to the various costs of building and deploying AI systems at scale; companies—especially in increasingly lengthy technical reports—speak of “mitigations” [40], “responsibility” [23], and “risks” [7]. But what sorts of risks—itsself a loaded construction [49]—must be addressed, and how, is refracted through these companies’ interests: a heavy focus on AI safety (a well-funded “epistemic community” [3] with quasi-millennarian concerns [22]), an occasional (depoliticized [11]) interest in bias, scant discussion of climate, labour and data (perhaps because, per a former Meta executive, asking permission for training data collection would “kill the AI industry” [50]), and a conspicuous absence of references to deployment in genocidal violence and surveillance [19, 24]. Prospective LLM users are to engage with such matters in the grammar of conspicuous consumption: if you’re put off by OpenAI’s many newsworthy scandals, why not try Anthropic instead? Refusal [21]—to use such systems, to invest in such political formations—is, of course, not a valid consumer choice.

5 CONCLUSION

In selling a product that projects a particular imaginary of entrepreneurial subjecthood, AI companies are extending a long-running assault on the working class [16, 27], anointed decades ago in Margaret Thatcher’s proclamation that “there is no such thing as society.” In the world that these companies are trying to make, workers will continue to be exploited and exhausted, or discarded as surplus, and this will continue to be made out to be workers’ own faults.

Do workers really subscribe to this predatory fantasy of AI-empowered work? In our own jobs, it can often seem as if overwhelming numbers of people—colleagues, managers, students—are using AI products. But perhaps this broad uptake is the result of coercion as well as enchantment, enacted via AI-by-default product designs and managerial directives. And, in the past few years, there have been numerous examples that suggest other, countervailing imaginaries. At many workplaces, including those figured in marketing discourses as sites for AI-enabled transformation, workers have been recognizing their interdependence and fashioning this recognition into counterhegemonic political action [1, 2, 54].

REFERENCES

- [1] 2025. No Tech for Apartheid. <https://www.notechforapartheid.com/>
- [2] 2025. #NoTechForICE. <https://notechforice.com/>
- [3] Shazeda Ahmed, Klaudia Jazwińska, Archana Ahlawat, Amy Winecoff, and Mona Wang. 2024. Field-building and the epistemic culture of AI safety. *First Monday* (2024).
- [4] Anthropic. 2025. Claude. <https://claude.ai/>
- [5] Anthropic. 2025. Claude Opus 4. <https://www.anthropic.com/claude/opus>
- [6] Anthropic. 2025. Meet Claude. <https://www.anthropic.com/claude>
- [7] Anthropic. 2025. System Card: Claude Opus 4 & Claude Sonnet 4. <https://www-cdn.anthropic.com/6be99a52cb68eb70eb9572b4cafad13df32ed995.pdf>
- [8] Erik Baker. 2025. *Make Your Own Job: How the Entrepreneurial Work Ethic Exhausted America*. Harvard University Press.
- [9] Aaron Benanav. 2019. Automation and the Future of Work—1. *New Left Review* 119 (2019), 5–38.
- [10] Emily M Bender and Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics*. 5185–5198.
- [11] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050* (2020).
- [12] Erik Brynjolfsson and Andrew McAfee. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & company.
- [13] Ajay Singh Chaudhary. 2024. *The exhausted of the Earth: politics in a burning world*. Watkins Media Limited.
- [14] Eve Chiapello and Luc Boltanski. 2018. *The new spirit of capitalism*. Verso Books.
- [15] Lillian Cicerchia. 2021. Why does class matter? *Social Theory and Practice* 47, 4 (2021), 603–627.
- [16] Melinda Cooper. 2024. *Counterrevolution: Extravagance and Austerity in Public Finance*. Princeton University Press.
- [17] 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.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*. 4171–4186.
- [19] Benj Edwards. 2024. Claude AI to process secret government data through new Palantir deal. *Ars Technica* (2024). <https://arstechnica.com/ai/2024/11/safe-ai-champ-anthropic-teams-up-with-defense-giant-palantir-in-new-deal/>
- [20] Norman Fairclough. 2001. Critical discourse analysis as a method in social scientific research. *Methods of critical discourse analysis* 5, 11 (2001), 121–138.
- [21] Patricia Garcia, Tonia Sutherland, Marika Cifor, Anita Say Chan, Lauren Klein, Catherine D’Ignazio, and Niloufar Salehi. 2020. No: Critical refusal as feminist data practice. In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing*. 199–202.
- [22] Timnit Gebru and Émile P Torres. 2024. The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday* (2024).
- [23] Google. 2025. Gemini. <https://deepmind.google/models/gemini/>
- [24] Anonymous Google and Amazon workers. 2021. We are Google and Amazon workers. We condemn Project Nimbus. *The Guardian* (2021). <https://www.theguardian.com/commentisfree/2021/oct/12/google-amazon-workers-condemn-project-nimbus-israeli-military-contract>
- [25] Daniel Greene. 2021. *The Promise of Access: Technology, Inequality, and the Political Economy of Hope*. MIT Press.
- [26] Gabriel Grill. 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.
- [27] Stuart Hall. 1991. Postscript: Gramsci and us. *Gramsci’s political thought: An introduction* (1991), 114–130.
- [28] Alex Hanna and Emily Bender. 2025. The AI Con: How to Fight Big Tech’s Hype and Create the Future We Want. (2025).
- [29] Sun-ha Hong. 2024. Strategic misrecognition and speculative rituals in generative AI. *Journal of Digital Social Research* 6, 4 (2024), 92–106.
- [30] Lilly Irani. 2019. Chasing innovation: Making entrepreneurial citizens in modern India. In *Chasing Innovation*. Princeton University Press.
- [31] Abigail Z Jacobs and Hanna Wallach. 2021. Measurement and fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 375–385.
- [32] Karen Spärck Jones. 2004. Language modelling’s generative model: Is it rational. *Computer Laboratory, University of Cambridge, Cambridge, UK* (2004).
- [33] Julian Krainin and Michael R Lawrence. 1990. *Memory & Imagination: New Pathways to the Library of Congress*. Encyclopedia Britannica Educational Corporation.
- [34] Cindy Lin and Silvia Margot Lindtner. 2021. Techniques of use: Confronting value systems of productivity, progress, and usefulness in computing and design. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–16.
- [35] Meta. 2025. Community Stories. <https://www.llama.com/community-stories/>
- [36] Meta. 2025. Llama. <https://www.llama.com/>
- [37] Meta. 2025. Llama 4. <https://www.llama.com/models/llama-4/>

- [38] Meta. 2025. Llama Cookbook. <https://www.llama.com/resources/cookbook/>
- [39] Milagros Miceli and Julian Posada. 2022. The data-production dispositif. *Proceedings of the ACM on human-computer interaction* 6, CSCW2 (2022), 1–37.
- [40] OpenAI. 2023. GPT-4 System Card. <https://cdn.openai.com/papers/gpt-4-system-card.pdf>
- [41] OpenAI. 2024. Duolingo. <https://openai.com/index/duolingo/>
- [42] OpenAI. 2025. About. <https://openai.com/about/>
- [43] OpenAI. 2025. ChatGPT. <https://openai.com/chatgpt/overview/>
- [44] OpenAI. 2025. Introducing OpenAI o3 and o4-mini. <https://openai.com/index/introducing-o3-and-o4-mini/>
- [45] 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.
- [46] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [47] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [48] Inioluwa Deborah Raji, Emily Denton, Emily M Bender, Alex Hanna, and Amandalynne Paullada. [n. d.]. AI and the Everything in the Whole Wide World Benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- [49] Jathan Sadowski. 2025. *The Mechanic and the Luddite: A Ruthless Criticism of Technology and Capitalism*. Univ of California Press.
- [50] Mia Sato. 2025. Nick Clegg says asking artists for use permission would ‘kill’ the AI industry. *The Verge* (2025). <https://www.theverge.com/news/674366/nick-clegg-uk-ai-artists-policy-letter>
- [51] Lily Scherlis. 2024. Skill Issues: Dialectical Behavior Therapy and Its Discontents. *The Drift* 13 (2024).
- [52] Chirag Shah and Emily M Bender. 2022. Situating search. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*. 221–232.
- [53] Jasmina Tacheva and Srividya Ramasubramanian. 2023. AI Empire: Unraveling the interlocking systems of oppression in generative AI’s global order. *Big Data & Society* 10, 2 (2023), 20539517231219241.
- [54] Ben Tarnoff. 2020. The Making of the Tech Worker Movement. *Logic* (2020). <https://logicmag.io/the-making-of-the-tech-worker-movement/full-text/>
- [55] Bonnie Urciuoli. 2008. Skills and selves in the new workplace. *American ethnologist* 35, 2 (2008), 211–228.
- [56] Ana Valdivia. 2024. The supply chain capitalism of AI: a call to (re) think algorithmic harms and resistance through environmental lens. *Information, Communication & Society* (2024), 1–17.
- [57] 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*.
- [58] Hannah Zeavin. 2021. *The distance cure: A history of teletherapy*. MIT press.