

Reconfiguring Participatory Design to Resist AI Realism

Aakash Gautam
University of Pittsburgh
Pittsburgh, PA, USA
aakash@pitt.edu

ABSTRACT

The growing trend of artificial intelligence (AI) as a solution to social and technical problems reinforces AI Realism—the belief that AI is an inevitable and natural order. In response, this paper argues that participatory design (PD), with its focus on democratic values and processes, can play a role in questioning and resisting AI Realism. I examine three concerning aspects of AI Realism: the facade of democratization that lacks true empowerment, demands for human adaptability in contrast to AI systems’ inflexibility, and the obfuscation of essential human labor enabling the AI system. I propose resisting AI Realism by reconfiguring PD to continue engaging with value-centered visions, increasing its exploration of non-AI alternatives, and making the essential human labor underpinning AI systems visible. I position PD as a means to generate friction against AI Realism and open space for alternative futures centered on human needs and values.

CCS CONCEPTS

• **Human-centered computing** → **Interaction design theory, concepts and paradigms.**

KEYWORDS

agonism, adversarial design, politics, power, artificial intelligence

ACM Reference Format:

Aakash Gautam. 2024. Reconfiguring Participatory Design to Resist AI Realism. In *Participatory Design Conference 2024 (PDC ’24 Vol. 2)*, August 11–16, 2024, Sibul, Malaysia. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3661455.3669867>

1 INTRODUCTION

“It’s easier to imagine the end of the world than the end of capitalism”, writes Mark Fisher [21, pp. 1] to illuminate the change in global polity, which since the fall of the Soviet Union around 1991 has portrayed that there is no alternative to capitalism. Fisher calls this phenomenon “Capitalist Realism” [21]. In the era of capitalist realism, every aspect of our lives is governed by capitalist rules and values. At the same time, various mechanisms for shaping a collective imagination that does not see a possibility beyond capitalistic approaches are deployed, including conflating capitalism with democracy, emphasizing the power of political participation, and

contrasting the current order with tyrannical socialist governments that failed.

We have also entered the era of AI Realism, where problems that were previously attempted to be solved collectively through social and technical mechanisms are now proposed to be solved through corporate-controlled AI (c.f., [9, 11]). The “AI” I talk about here is not the statistical approach but rather the much-hyped larger ecological system that is being pushed by corporations. The growing use of generative AI systems in different workplaces [9], academic venues’ rules on the use of ChatGPT or other large language models [22], and the growing use of chatbots by existing services [26] all point to the fact that we have arrived at AI Realism. The shift to AI realism has been more rapid than that of capitalist realism though. The success of AlexNet in 2012 began a drive to push AI as *the* solution to a range of problems [10], and by 2022, with the release of ChatGPT, it had captured the public’s imagination as a panacea (e.g., [16, 20, 26]). In some cases, there have been concerted efforts to push AI as a solution, in others, we inadvertently accepted it.

More than the striking parallels, Capitalist and AI Realism are closely connected. AI is a tool that existed long before 2012, but *the* AI that is being pushed widely is more than a statistical tool. It is the tool as used by and defined by the capitalist; a technological mechanism that enables the concentration of resources and power, hidden behind the veil of neutral-seeming statistical formalism¹. The pervasive promotion of AI with claims of efficiency, productivity, and various other capital value creation — referred to as “AI snake oil” [37] — has distracted us from examining details of the value and the potential for harms that the system brings.

In the face of growing AI Realism, PD has a significant role to play. To illustrate this, I begin by questioning three aspects of AI that are promoted widely. First, the claims of “democratizing AI” lack the necessary distribution of power to the people affected by the AI system to be democratic. Second, the popular imagery of AI for personalization makes significant demands of human malleability and erodes opportunities for collective action. Third, the image of AI systems as efficiency-promoting obfuscates the labor required to enable them. Reflecting on these three aspects, I argue for reconfiguring PD methods to resist AI Realism. These reconfigurations include contesting narrow metrics to enable broader visions of the future, incorporating explorations of non-AI alternatives while identifying possibilities, and making visible all the human labor that enables AI systems. In doing so, I seek to (re)position PD as a force for setting agendas and upholding democratic participation, and, with it, centering human values in the era of AI Realism.



This work is licensed under a [Creative Commons Attribution-NonCommercial International 4.0 License](https://creativecommons.org/licenses/by-nc/4.0/).

PDC ’24 Vol. 2, August 11–16, 2024, Sibul, Malaysia
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0654-7/24/08
<https://doi.org/10.1145/3661455.3669867>

¹Although statistical-AI and AI-as-capitalist-instrument are tied to some extent considering the corporate control over the data and other infrastructure [51].

2 LITERATURE REVIEW: AI AND THE CORE TENETS OF PD

Participatory design (PD) finds its roots in the 1970s with the workers' push for greater control and democratization amidst the introduction of new technology in the workplace [5]. Over the years, PD scholarship, within the workplace and beyond it, has highlighted tenets like carefully developed participation, mutual learning between workers and designers, and envisioning alternative technological futures [19]. These principles support creative cooperation in designing new sociotechnical systems, thus enabling potential space to challenge existing power dynamics [7, 35]. In some respect, the wheels of time have brought us back to the 1970s with the advent of AI in the workplace and beyond, with growing concerns of drastic changes in work practices and values [36].

AI systems bring a set of unique characteristics, some common with other large-scale technologies and a few unique to it. Among the unique ones, AI systems often lack transparency in how it was created and evaluated, in data collection, model development, fine-tuning, and impact evaluation [23, 39, 40, 42]. Second, AI system development is expensive and typically deployed on large scales, thus requiring infrastructure that only a few can afford [51]. Third, AI as a "modern technology" is justified for a fast-paced push, where rapid development and deployment evades auditing, evaluation, and even regulations. In this sense, the proliferation of AI embodies Silicon Valley's ethos of "move fast and break things" [47] or "ask for forgiveness, not permission" [15]. Given its pervasiveness, we are all being participated in its growth by providing it with data; at times, there seems like there is no alternative.

Efforts have been made to promote transparency and create possibilities for broader group engagement (e.g., [23, 34, 39, 40]). But these are technological solutions that assume AI is inevitable; the challenge of AI is a "social, cultural, political, and ethical one" [33, pp. 2] where the technical decisions intermingle with social responsibilities. The majority of participatory AI movements have been devoid of *empowered* engagement of broader stakeholders, including the possibilities to challenge existing power differences or opportunities for recourse and refusal [3, 6, 17, 52]. They fail to illuminate broader and alternative futures, thus further solidifying the position of existing AI as the inevitable "natural order". In such cases, Fisher contends, we "must reveal what is presented as necessary and inevitable to be a mere contingency just as it must make what was previously deemed to be impossible seem attainable" [21, pp. 17]. Revealing these alternative possibilities and enacting moves towards realizing them —i.e., resisting AI Realism—is where PD can play a significant role.

3 THE REAL BEHIND THE PROMOTED REALITY OF AI REALISM

Fisher, drawing from Lacan [29], argues for invoking the 'Real', that which threatens the promoted "reality". The Real is the "traumatic void that can only be glimpsed in the fractures and inconsistencies in the field of apparent reality" [21, pp. 18]. For Fisher the Reals that can be presented to challenge the realities in capitalist realism are ongoing and increasing environmental catastrophes challenging the growth fetish of capitalism, the burgeoning mental health crisis

Table 1: Contrast between the promoted reality and the Real in AI Realism

Promoted reality	Real
Democratized AI with transparency	Lack of data ownership and control
Personalized recommendations	Human malleability, machine intransigence
Automation and Efficiency	Human labor behind the infrastructure

that undermines the push for privatization of stress, and the proliferation of bureaucracy in capitalism despite the widespread claims of arriving at market-driven efficiency [21]. Drawing parallels, I share three Reals that could open avenues to question the reality promoted in AI Realism (Table 1).

3.1 Participation Without Power Is Not Democratic

In many Participatory AI efforts, participation focuses on inclusion for design input rather than a critical mechanism to support people in setting an agenda. For instance, on reviewing 80 participatory AI projects, Delgado et al. [17] found that only three projects involved stakeholders throughout the design process, and only ten involved them in setting the scope and purpose of AI. The effect of this is that those efforts presuppose AI as a given with its claims of benefits unchecked; participation is constrained to small-scale inputs on improving the development and deployment of AI.

From this perspective, participatory AI is deployed as a mechanism to co-opt human labor toward acquiescence [1, 3]. This weakened participation stems from perverse incentive structures benefitting corporate AI — they benefit from expanding AI adoption with the facade of people's participation. There have been exception cases of community ownership and governance of data practices and AI systems, including deciding whether to use AI (e.g., [12]). However, constraints such as limited time and resources are often cited as barriers, leaving open questions of *how* communities can be empowered in governing AI systems. Indeed, the opacity and rapid evolution of AI systems raise questions on what even constitutes meaningful participation.

Bratteteig and Verne [6], discussing if AI makes PD obsolete, argue for participation in (1) envisioning ideas with knowledge of what AI can and cannot do, (2) concretizing design ideas by engaging with values and visions, and (3) evaluating the impact of the decision. While these are important steps, they are not sufficiently political. Knowing the capabilities, which is a motivation for explainable AI, does not challenge the techno-deterministic stance that it promotes; participation is co-opted to attend to technical solutions conveniently solved through AI rather than engaging in institutional or social possibilities. Similarly, discussing higher-level values and visions only has the risk of arriving at consensus on systems that in reality may necessitate surveillance and other large-scale data collection mechanisms that the participants may not be aware of. Transparency and co-designing processes involved in developing and deploying AI systems needs to be added to the

process of making collective decisions about AI. Further, for participants to evaluate the impact of AI systems especially in the design phase is challenging since effects are unknown until development and, often, seen after long-time deployment [6]. Thus, participation should involve collectively establishing and bringing into practice accountability structures before the development or deployment of AI technologies.

3.2 Personalization Demands Human Malleability and Erodes Mutual Learning

Personalization is one of the most visible aspects of AI. Marketing rhetoric pushes personalization as a universally desired feature, with one McKinsey report declaring, “Consumers don’t just want personalization, they demand it” [2]. Despite widespread concerns about privacy and lack of awareness regarding the types of data being collected [31], AI systems continue to collect data for personalization, which indicates users’ limited agency in challenging this juggernaut drive towards personalization.

Personalization surfaces two tensions with PD. First, personalization in AI systems demands human malleability. This is not a new phenomenon: people accommodate technology by changing their behavior and practices around the technology [32, 46]. But with AI, we are required to accommodate it both while using it and, since AI is trained in use, in training it. Gyldenkerne et al. [25], for example, in studying AI use in electronic health records in Denmark, highlight how an AI system in providing personalized insights imposed a new process of data collection even though the clinician’s existing processes were sufficient for their purposes. It not only changed the way clinicians recorded patient data but since it was rife with ambiguity, it made additional demands on the clinicians to attend to the data processes, moving them away from the primary purpose of their work. Moreover, human flexibility demanded in training the AI reifies AI’s assumptions about us and our practices, leaving little room for us to change those. Going back to the case of the clinicians, despite the mismatch in their practices and the new AI-based processes, they did not have room for recourse. Instead, they had to devise additional workarounds [25]. This challenge of ensuring users’ control over their tools and their environment is a classic problem that PD has grappled with [8], but now with a much-pervasive technological system.

Second, personalization creates its own context, influencing the user to think about aligning to the demands of the system rather than finding opportunities to make common ground with others. It raises challenges for surfacing shared experiences and engaging in collective action. We are forced to work on pushing back on the limitations and impositions of our own spaces rather than drawing commonalities among a larger mass to push against the system. This echoes the Marxist notion of ideology — that AI systems are *the* thing — and acquiescing the alienation that it brings. A significant exercise of abstraction is necessary to arrive at shared experiences and opportunities for mutual learning; these skills are not easily accessible to all. Personalization seeks to make participation individualistic. For PD, which relies on forming publics — of people forming attachments to shared matters of concern [14] — this poses a challenge.

3.3 Claims of Efficient AI Hide Away the Hidden Labor

Proclamations of AI’s potential to drive efficiency and unprecedented scale abound. Hidden behind the veneer of magical-seeming machine intelligence lies the indispensable critical infrastructure of human labor engaged in essential tasks of collecting, cleaning, and labeling data. This is not accounted for in the claims of AI’s efficiency! For instance, the ImageNet dataset which was the foundation for the AlexNet model in 2012, required upwards of 25,000 people (Amazon Mechanical Turks) to build the dataset, with more than 45% of them making less than USD\$5 a week [30]. This labor is rarely counted.

The labor is typically sourced out to the Global South where worker conditions and compensations are, at best, questionable (e.g., [24, 50, 52]). The indispensable human labor is obfuscated behind claims of technological progress of AI. Paradoxically, the very AI that relies on the labor of many is employed to justify efficiency drives, leading to layoffs elsewhere [49]. A prominent example of this is the recent Duolingo layoff, where workers at most precarious positions were laid off citing AI-enabled efficiency².

There are two facets to the problem. First, the promotion of values around efficiency and scale brought by AI normalizes the practice of hiding and often exploiting large human labor [24]. The growing reliance on labor from the Global South enabled through complex sub-contracting practices has created mechanisms to evade accountability. Those who use and are affected by AI are placed at a distance from those who build the infrastructure for that AI, enabling very little space for collective action and solidarity [52].

Second, “efficiency” is presented as the only metric that should guide our endeavors. It has broader societal implications. Examining the reality of what kind of labor is hidden and what kinds of work continue to be prized, we notice that the push for efficiency is top-down. The acceptance of efficiency requires acquiescing to unilateral decisions. It leaves no room to explore alternative visions (e.g., promoting adaptability or interconnectedness) that we could collectively design [27], even when efficiency may conflict with our collective well-being [43]. These two facets — hiding the labor and leveraging the labored artifact to squeeze greater efficiency — enable those with the power to create and deploy AI tools to concentrate wealth and resources.

4 DISCUSSION: PD IN THE ERA OF AI REALISM

PD, with its emphasis on forming attachments to matters of concern [14], is uniquely positioned to question the relentless push for AI and create space to determine *if and how* AI should be developed and deployed.

AI is not a modular unit; it is ecological and does not function in isolation. It is not monolithic either. Each AI system taps into and brings with it a myriad of standards, processes, and values, each of which in turn influences our actions and values. However, AI Realism pushes an instrumental and utopic view of AI, focused solely on outcomes (e.g., efficiency and growth) while obfuscating

²<https://www.washingtonpost.com/technology/2024/01/10/duolingo-ai-layoffs/>

the complex processes and tradeoffs involved. In contrast, PD engages with processes that emphasize forming publics for collective envisionment, surfacing diverse values, and negotiating decisions over the technologies that impact them. More generally, there is an inherent tension: PD acknowledges diverse human conditions and seeks to support pluralism and individual agency in collectively realizing a shared vision of the future, whereas AI Realism seeks to — or rather requires — imposing AI as *the* solution, removing it from the nuanced contextual differences and uncertainties inherent in our human condition. This oppositional stance of PD can be constructive, allowing it to steer the relentless drive of AI Realism toward values that matter.

Research on agonistic pluralism in PD (e.g., [4, 18]) provide guidance on creating the necessary friction to realize positive change. Scholars argue for embracing an adversarial stance rather than avoiding conflict. It is needed in the face of AI Realism. Drawing from it, I propose three initial reconfiguration of participatory methods to create space for resisting AI Realism. First, contesting dominant metrics to support engagements with broader visions and possibilities for the future. Second, including possibilities of developing non-AI solutions as alternatives to distill needs and situations that truly require AI solutions. Third, surfacing and engaging with the necessary human labor that is required to make the AI system function before making decisions on incorporating AI. Embracing these would allow participatory approaches to move beyond questioning whether AI works or how to make it work, towards surfacing systemic issues, assumptions, and impositions that underlie AI Realism.

4.1 Beyond Dominant Metrics: Engaging in Value-Centered Visions

The incessant focus on scale and efficiency — a root element of modern-day capitalism — finds fuel in AI systems. These metrics fail to capture the complex values that we care about. Yet, in the rapid pace of incorporation, we rarely have an opportunity to question what values AI systems enforce and whether we want them.

There is a need to design alternative visions collectively. I echo Huybrechts et al. [27, pp. 7], who posit that “designing ‘visions’ can turn the tension between addressing the focus on the big issues and the close attention to the particularity of relations into a dynamic dialogue that can repoliticise design.” However, given the dominant focus on scale and efficiency, there is a need to first contest those metrics and discuss alternatives such as commoning and post-growth models (e.g., [44, 48]).

As we engage in designing for alternative visions, it is equally important, given the global pervasiveness of AI systems, to engage both with the local priorities and the global realities [5], acknowledging and attending to the potential impacts, tradeoffs, and tensions at play across these levels [13, 24]. Moreover, AI is not a monolithic system. Deconstructing AI into its constitutive parts is necessary to understand the vision it embodies. To this end, we may, for example, bring together the precarious workers early in the AI development pipeline and the higher-paid, slightly empowered workers — the recent layoffs suggest no worker is truly in control — in the design process of alternative visions and possibilities.

4.2 Plural Possibilities: Exploring Non-AI Alternatives

Moving past the dominant metrics is an important first step. A subsequent step along the line involves ensuring that PD methods make space for participants to enact their visions through non-AI alternatives.

Specifically engaging with non-AI alternatives is important for the very act of co-designing AI systems, even when they are critical of other existing AI systems, enforces the belief that AI is necessary or inevitable. Exploring non-AI alternatives challenges the assumptions that AI systems are better or necessary. Further, the perception of AI as a widely-encompassing solution has made it easier to brush aside organizational and/or social problems. Engaging in non-AI alternative solutions may surface issues that cannot be fixed by technology.

While exploring non-AI alternatives may not necessarily lead to the refusal of AI, it will afford space to critically reflect and distill aspects of the participants’ needs and identify situations where AI is truly needed. In doing so, participants can deepen their understanding of the situation, and if AI is deployed, critically understand the role the AI system plays in the solution. This can help them gain some power over the AI systems.

Moreover, exploring non-AI alternatives opens possibilities to reflect on what would be lost when a situation is reformulated to make it amenable to AI ecosystem. Scholars warn about “algorithmic governmentality,” where only outcomes that can be measured and realized through algorithms will be pursued [38, 41]. It is a reductive move, taking us away from uncertainties and politics to perfectly objective outcomes that have no room for deliberation and recourse [41]. Exploring non-AI alternatives enables space to engage with uncertainties and the inherent messiness in our collective visions and values, potentially leading to plural possibilities of being.

4.3 Delegation Inversion: Making Visible the Human Labor

Earlier, I presented two stages where human labor was essential to make the AI systems function: in creating AI systems such as with data labeling work, and our efforts to accommodate AI despite its rigidity. Both these essential labor are hidden from the popular narrative of AI. PD methods should surface and engage with it, particularly making it central in the decision-making process when participants are exploring AI-based possibilities. This echoes scholars who have long argued for making work visible [24, 45].

For PD to resist AI Realism, we have to surface and engage with the realities of the hidden labor at the time of design envisionment. To enable this, I build on Latour, who, using a door as an exemplary technology, expounds on the value of “delegation”, i.e., the work the technology does that humans no longer have to do [28]. I contend that for pervasive systems like AI, PD methods can be configured to engage participants in “delegation inversion” — cataloging the additional work and sacrifice humans have to make for the technology to function. For instance, participants could create a “labor accountability card” to record information such as the number of data labelers, total labeling hours, and the working conditions of those contributing to the system supporting their envisioned AI

solution. Similarly, to enhance shared experiences regarding human adaptability, participants could note places of friction in using existing AI systems and where they needed to be flexible.

These efforts of surfacing and engaging with the labor will help illuminate the cost of developing and deploying AI solutions, thus seeding informed decision-making on whether it is worth it. It can potentially create space for enacting alternative visions that do not demand such cost.

These three reconfigurations of PD methods — contesting narrow metrics while envisioning possibilities, incorporating explorations of non-AI alternatives, and illuminating the often-hidden labor enabling AI — aim to position PD as a powerful mechanism to resist AI Realism. They surface societal costs incurred by AI Realism, while simultaneously empowering people to shape or refuge technologies based on the values they care about.

However, this is a conversation at an early stage of PD's evolution in the era of AI Realism. There is critical work ahead for the PD community, both in research and practice. We, as PD researchers and practitioners, must continue to challenge the dominant drive of AI solutionist approaches and ensure that the public is empowered to enact the future that *they* desire. It is, indeed, a call for us to push back on AI Realism toward democratic, caring, and just futures shaped through collective decision-making. The stakes could not be higher.

5 CONCLUSION

The fundamental question of AI Realism is about making decisions on the distribution of resources and power. It is asking, “Who should have the decision-making power on whether, what, how, when, where, and why AI systems should shape our collective future?” As it stands, the power differences are highly skewed towards those who already control AI, and challenging it — i.e., resisting AI Realism — is where I see PD play a significant role. PD has always been political, but in the face of AI Realism, we need a vigorous renewal of the political. I argue for repositioning PD as a mechanism to create friction in the incessant focus on outcome that AI Realism pushes by engaging with the processes enabling AI and the values inherited by adopting it. Resisting the relentless push for AI adoption is an attempt to prioritize human values and justice over the corporatization and centralization of resources and power. As such, this work attends to the theme of PDC 2024 in arguing to extend the notion of participation beyond the specific methods towards being a driver for setting agendas in the face of AI Realism and establishing initiatives to reflect and facilitate alternative futures. This calls for an alliance of researchers and practitioners who see PD as a democratic force to look outward and leverage PD methods to address root issues in the era of AI Realism, rather than inward at refining our methods.

ACKNOWLEDGMENTS

I am grateful to Aditya Nayak and Deborah Tatar for their feedback.

REFERENCES

- [1] Alex Ahmed. 2020. We Will Not Be Pacified Through Participation. <https://assets.pubpub.org/zy0iuc4f/61d44af5-cf5b-4c40-8b0e-5ea160a5e956.pdf>. *Tech Otherwise* 15, 10 (2020).
- [2] Nidhi Arora, Daniel Ensslen, Lars Fiedler, Wei Wei Liu, Kelsey Robinson, Eli Stein, and Gustavo Schüler. 2021. The value of getting personalization right—or wrong—is multiplying. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personalization-right-or-wrong-is-multiplying>.
- [3] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Jason Gabriel, and Shakir Mohamed. 2022. Power to the people? opportunities and challenges for participatory AI. *Equity and Access in Algorithms, Mechanisms, and Optimization* (2022), 1–8. <https://doi.org/10.1145/3551624.3555290>
- [4] Erling Björgvinsson, Pelle Ehn, and Per-Anders Hillgren. 2012. Agonistic participatory design: working with marginalised social movements. *CoDesign* 8, 2-3 (2012), 127–144.
- [5] Susanne Bødker and Morten Kyng. 2018. Participatory design that matters—Facing the big issues. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 1 (2018), 4. <https://doi.org/10.1145/3152421>
- [6] Tone Bratteteig and Guri Verne. 2018. Does AI make PD obsolete? exploring challenges from artificial intelligence to participatory design. In *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial-Volume 2*. 1–5. <https://doi.org/10.1145/3210604.3210646>
- [7] Tone Bratteteig and Ina Wagner. 2012. Disentangling power and decision-making in participatory design. In *Proceedings of the 12th Participatory Design Conference: Research Papers-Volume 1*. 41–50. <https://doi.org/10.1145/2347635.2347642>
- [8] Tone Bratteteig and Ina Wagner. 2016. What is a participatory design result?. In *Proceedings of the 14th Participatory Design Conference: Full papers-Volume 1*. 141–150. <https://doi.org/10.1145/2940299.2940316>
- [9] Davide Castelvetti. 2022. Are ChatGPT and AlphaCode going to replace programmers? *Nature* (2022). <https://doi.org/10.1038/d41586-022-04383-z>
- [10] Brian Christian. 2020. *The alignment problem: Machine learning and human values*. WW Norton & Company.
- [11] James Clayton. 2024. CES 2024: AI pillows and toothbrushes - is it all getting a bit silly? <https://www.bbc.com/news/technology-67959240>.
- [12] Donavyn Cofey. 2021. Maori are trying to save their language from Big Tech. *Wired UK (April 2021)* (2021).
- [13] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- [14] Christopher A Le Dantec and Carl DiSalvo. 2013. Infrastructuring and the formation of publics in participatory design. *Social Studies of Science* 43, 2 (2013), 241–264. <https://doi.org/10.1177/0306312712471581>
- [15] Jamie Davies. 2019. Silicon Valley's 'ask for forgiveness, not permission' attitude is wearing thin. <https://www.telecoms.com/mobile-devices/silicon-valley-s-ask-for-forgiveness-not-permission-attitude-is-wearing-thin>.
- [16] Luigi De Angelis, Francesco Baglivo, Guglielmo Arzilli, Gaetano Pierpaolo Privitera, Paolo Ferragina, Alberto Eugenio Tozzi, and Caterina Rizzo. 2023. ChatGPT and the rise of large language models: the new AI-driven infodemic threat in public health. *Frontiers in Public Health* 11 (2023), 1166120.
- [17] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23. <https://doi.org/10.1145/3617694.3623261>
- [18] Carl DiSalvo. 2015. *Adversarial design*. Design Thinking, Design Theory.
- [19] Pelle Ehn. 2008. Participation in design things. In *Proceedings of the tenth anniversary conference on participatory design 2008*. Indiana University, 92–101. <https://doi.org/10.5555/1795234.1795248>
- [20] Andy Extnance. 2023. ChatGPT has entered the classroom: how LLMs could transform education. *Nature* 623, 7987 (2023), 474–477. <https://doi.org/10.1038/d41586-023-03507-3>
- [21] Mark Fisher. 2022. *Capitalist realism: Is there no alternative?* John Hunt Publishing.
- [22] Association for Computing Machinery. 2023. ACM Policy on Authorship. <https://www.acm.org/publications/policies/new-acm-policy-on-authorship>.
- [23] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92. <https://doi.org/10.1145/3458723>
- [24] Mary L Gray and Siddharth Suri. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- [25] Christopher H. Gyldenkerne, Gustav From, Troels Mønsted, and Jesper Simonsen. 2020. PD and the challenge of AI in health-care. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 26–29. <https://doi.org/10.1145/3384772.3385138>
- [26] Erin Hale. 2023. ChatGPT is giving therapy. A mental health revolution may be next. <https://www.aljazeera.com/economy/2023/4/27/could-your-next-therapist-be-ai-tech-raises-hopes-concerns>.
- [27] Liesbeth Huybrechts, Maurizio Teli, Mela Zuljevic, and Mela Bettega. 2020. Visions that change. Articulating the politics of participatory design. *CoDesign* 16, 1 (2020), 3–16.

- [28] Jim Johnson. 1988. Mixing humans and nonhumans together: The sociology of a door-closer. *Social problems* 35, 3 (1988), 298–310.
- [29] Jacques Lacan. 1988. *The Seminar of Jacques Lacan*. WW Norton & Company.
- [30] Fei-Fei Li. 2010. ImageNet: Crowdsourcing, benchmarking & other cool things. https://www.image-net.org/static_files/papers/ImageNet_2010.pdf. CMU VASC Seminar (2010).
- [31] Mengqi Liao and S Shyam Sundar. 2021. How should AI systems talk to users when collecting their personal information? Effects of role framing and self-referencing on human-AI interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14. <https://doi.org/10.1145/3411764.3445415>
- [32] Q Vera Liao, Muhammed Mas-ud Hussain, Praveen Chandar, Matthew Davis, Yasaman Khazaeni, Marco Patricio Crasso, Dakuo Wang, Michael Muller, N Sadat Shami, and Werner Geyer. 2018. All Work and No Play? Conversations with a Question-and-Answer Chatbot in the Wild. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3173574.3173577>
- [33] Daria Loi, Thomas Lodato, Christine T Wolf, Raphael Arar, and Jeanette Blomberg. 2018. PD manifesto for AI futures. In *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial-Volume 2*. 1–4. <https://doi.org/10.1145/3210604.3210614>
- [34] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229. <https://doi.org/10.1145/3287560.3287596>
- [35] Michael J Muller and Allison Druin. 2012. Participatory design: the third space in human-computer interaction. In *Human Computer Interaction Handbook*. CRC Press, 1125–1153.
- [36] Mark Muro, Robert Maxim, and Jacob Whiton. 2019. *Automation and Artificial Intelligence: How machines are affecting people and places*. <https://www.brookings.edu/research/automation-and-artificial-intelligence-how-machines-affect-people-and-places/>.
- [37] Arvind Narayanan. 2019. How to recognize AI snake oil. <https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf>. *Arthur Miller Lecture on Science and Ethics* (2019).
- [38] Aditya Nayak. 2021. Boredom, Time-Perception and Algorithmic Governmentality. *Journal of AI Humanities* 7 (2021). <https://doi.org/10.46397/JAIH.7.1>
- [39] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data cards: Purposeful and transparent dataset documentation for responsible ai. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1776–1826. <https://doi.org/10.1145/3531146.3533231>
- [40] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 33–44. <https://doi.org/10.1145/3351095.3372873>
- [41] Antoinette Rouvroy, Thomas Berns, and Liz Carey-Libbrecht. 2013. Algorithmic governmentality and prospects of emancipation. *Réseaux* 177, 1 (2013), 163–196.
- [42] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15. <https://doi.org/10.1145/3411764.3445518>
- [43] Barry Schwartz. 2020. Why efficiency is dangerous and slowing down makes life better. <https://psyche.co/ideas/why-efficiency-is-dangerous-and-slown-down-makes-life-better>.
- [44] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2023. Post-growth Human-Computer Interaction. *ACM Transactions on Computer-Human Interaction* 31, 1 (2023), 1–37. <https://doi.org/10.1145/3624981>
- [45] Lucy Suchman. 1995. Making work visible. *Commun. ACM* 38, 9 (1995), 56–64. <https://doi.org/10.1145/223248.223263>
- [46] Lucy Suchman. 2007. *Human-machine reconfigurations: Plans and situated actions*. Cambridge University Press.
- [47] Jonathan Taplin. 2017. *Move Fast and Break Things*. Little, Brown and Company.
- [48] Maurizio Teli, Marcus Foth, Mariacristina Sciannamblo, Irina Anastasiu, and Peter Lyle. 2020. Tales of institutioning and commoning: participatory design processes with a strategic and tactical perspective. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 1*. 159–171. <https://doi.org/10.1145/3385010.3385020>
- [49] Catherine Thorbecke. 2024. The tech sector is pouring billions of dollars into AI. But it keeps laying off humans. <https://www.cnn.com/2024/01/13/tech/tech-layoffs-ai-investment/index.html>.
- [50] Ding Wang, Shantanu Prabhat, and Nithya Sambasivan. 2022. Whose AI Dream? In search of the aspiration in data annotation.. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–16. <https://doi.org/10.1145/3491102.3502121>
- [51] Meredith Whittaker. 2021. The steep cost of capture. *Interactions* 28, 6 (2021), 50–55. <https://doi.org/10.1145/3488666>
- [52] Adrienne Williams, Milagros Miceli, and Timnit Gebru. 2022. The exploited labor behind artificial intelligence. *Noema Magazine* 13. <https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence/>

³Bruno Latour published the article under a pseudonym. Please see the footnote on page 304 of the article.