

Participatory AI & Social Justice

Maria Luce Lupetti
Department of Architecture and
Design
Politecnico di Torino
Turin, Italy
maria.lupetti@polito.it

Cristina Zaga
Human-Centred Design
University of Twente
Enschede, NY, Netherlands
c.zaga@utwente.nl

Christina Harrington
Google Research
Atlanta, Georgia, USA
HCI Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
cnharrington@google.com

Laura Forlano
Northeastern University
Boston, Massachusetts, USA
l.forlano@northeastern.edu

Massimo Menichinelli
Elisava Research
Elisava, Barcelona School of Design
and Engineering (UVic-UCC)
Barcelona, Barcelona, Spain
mmenichinelli@elisava.net

Alessandro Bozzon
Knowledge and Intelligence Design.
Faculty of Industrial Design
Engineering
Delft University of Technology
Delft, Netherlands
a.bozzon@tudelft.nl

Q. Vera Liao
Computer Science and Engineering
University of Michigan, Ann Arbor
Ann Arbor, Michigan, USA
veraliao@umich.edu

Abstract

Artificial Intelligence (AI) and Machine Learning (ML) technologies are increasingly integrated into a variety of products and services. However, AI promises often fall short and unintended consequences multiply, causing a negative impact, especially on marginalised communities. In response, Participatory AI has emerged as a promising method to account for the consequences on people subject to AI and to help develop systems that are better aligned with societal values. Yet meaningful participation remains difficult to achieve. Participation tends to be short-term and consultative, or even a mask for hidden labour. Moreover, the aim of building knowledge that is situated and subjective often conflicts with the global scale and generalizability of AI, especially when it comes to foundational models. This panel unites experts to explore how to account for these issues, and more specifically, to collectively draw a picture of the limits of Participatory AI in terms of social justice.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; **Accessibility theory, concepts and paradigms**.

Keywords

Participatory AI, Social Justice, Participatory Design, Responsible AI, Marginalised Communities

ACM Reference Format:

Maria Luce Lupetti, Christina Harrington, Massimo Menichinelli, Cristina Zaga, Laura Forlano, Alessandro Bozzon, and Q. Vera Liao. 2026. Participatory AI & Social Justice. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3772363.3790077>

1 Introduction

As "Artificial Intelligence (AI)" (we use the term "AI" with quotation marks to denote its meaning in popular parlance, as suggested by [6]) gets pushed into an increasing number of products and services, unintended consequences multiply, exacerbating biases, inequalities, and injustice, especially for traditionally marginalized and 'otherized' communities [1, 10]. As a response, the academic and professional AI fields have recently become more and more committed to various forms of stakeholders and communities engagement, also encouraged by the recommendations and policies released over the last years by public bodies and institutions, such as the *Ethics guidelines for trustworthy AI* [11] in the European Union, the *Global AI Governance Action Plan* by Chinese government [14], the *Brasil AI Act* and more [2].

This growing interest in the mechanism of public participation in support of deliberation around AI developments, which has recently been framed as a participatory turn in AI design [5], is based on the belief that participation can enable AI systems to better reflect the values, preferences, and needs of users and other impacted stakeholders, or more broadly, empower stakeholders in shaping the design of AI systems [5].

Given its long tradition in participatory design work, the human-computer interaction (HCI) field, both in academic and professional



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3790077>

environments, has embraced this trend and started to produce knowledge and tools to support Participatory AI practices. On the one hand, resources are being created to better equip and support various stakeholders in understanding AI capabilities and what to do with these, as a way to enable collaborative development of AI applications. For instance, Yildirim and colleagues have developed a collection of slides illustrating AI capabilities [24] and 'Do-Reason-Know worksheets' [25] as a resource to support ideation around AI development. On the other hand, platforms and tools are being developed to support collective action to stir AI development and related practices, as in the case of Dynamo: a platform to support the Mechanical Turk community in forming publics around issues and then mobilizing [18]. In the six months that the platform was deployed on the web, the site was visited thousands of times, and 470 Turkers registered. Turkers have gathered on and used the platform for discussing twenty-two ideas for action, two of which have resulted in active campaigns: the *Guidelines For Academic Requesters*, a guide covering matters such as fair pay, how to respect Turkers' privacy, and how to respect Turkers' communities online; and the *Letter to Jeff Bezos*, a letter campaign intended to fight the media portrayals of Turkers as downtrodden, disempowered cogs in a machine and humanize them.

1.1 Towards a critical understanding of participation in AI design

Alongside developing practical resources and tools for supporting stakeholders and community participation in AI design and related eco-socio-technical ecosystems, critical HCI scholars have started to report on the complexities of doing such work of participation around AI. Seminal in this regard is the work by Sloane and colleagues [22] who call for a cautious look at participatory AI practices and argue that there are very distinct ways in which participation unfolds around AI and machine learning (ML), more precisely, i.e., as work, as consultation, and as justice. Closer to participatory design ideals and democratic ambitions is the notion of *participation as justice*, where designers and technologists typically seek to engage in long-term partnerships with diverse stakeholders, with the ambition to address the structural issues that lead to algorithmic harm. While noble and close to participatory design traditions and ambitions, AI participation as justice is rare and difficult, due to various reasons, among which the extremely high complexity in the development pipeline, where multiple actors play diverse roles, and often lack a holistic view (as illustrated by [8]).

As a result, participation in most cases happens in specific moments of the AI development pipeline, often as a way to gather information about potential users and contexts of applications; what [22] call *participation as consultation*. In this mode of engagement, designers and technologists run episodic and short-term activities, such as workshops and hackathons, to identify context-specific needs together with stakeholders. This top-down approach is a way of "designing for" a particular group without really committing to their inclusion in the processes, with the result that 'systemic inequalities that can be hardcoded into consultation and representation protocols' [21, 22]. Both *participation as justice* and *participation as consultation*, with their benefits and limitations, fit

within the discourse surrounding participatory design methodologies and practices. As Sloane and colleagues [22] argue, however, the field of AI, and machine learning in particular, demands a massive amount of collective participation that often remains hidden. This hidden form of participation, addressed as *participation as work*, involves billions of people in the production of data that is used to train and evaluate ML models. Whether it is a paid form of participation, as in data labeling jobs (i.e., mTurk workers [17]), or unpaid, as in people's interaction with web features aimed at improving ML performance. In these forms of participation, data subjects, crowdworkers, and end users play a crucial role in ML (and AI more broadly) functioning, yet their work remains poorly compensated or, even, unpaid, and unacknowledged [22]. In short, AI comes with an inherent barrier for the public to genuinely contribute to its shaping, which Suresh and colleagues [23] refer to as the *participatory ceiling*.

A similar contribution, that helps us better grasp the different forms that participation can take and the related implications, is the work by Corbett and colleagues [4] who look back at Sherry Arnstein's *Ladder of Citizen Participation* [20] to underscore that participatory AI can ideally lead to various degrees of citizens' power, e.g., through forms of citizens' control, power delegation, and partnership, but also often lands itself as a form of tokenism, where design goals and ideas are already established prior the engagement, and there is no accountability mechanism in place to ensure that participation impacts on design decisions. Here, people are usually involved as informants, and participation is a form of consultation. Even worse, participation can be performative – 'designed to give the appearance of legitimacy without any real control' and thus actually resulting in non-participation.

In addition to moving a critique to mainstream participatory AI practices, this literature also shows the *complexity of defining what counts as participation*. If we look through the lens of participatory design tradition and scholarship, forms of 'participation as work', for instance, can hardly be considered as real participation as participants lack not only power, but also agency in the process.

Furthermore, issues of performativity and tokenism underscored by these works are known issues in participatory design literature, with authors arguing that in participatory practices, key design decisions are often predetermined, while the actual act of design occurs elsewhere and at different times—typically during technology development [12]. Furthermore, alternative epistemologies rooted in people's lived experiences are often devalued and subordinated to forms of 'scientific' or institutionalized knowledge. As a result, individuals who fall outside the normative assumptions of dominant design cultures are further marginalized [9]. The structural barriers that many marginalized communities face—such as those linked to geographic location, language and knowledge divides, financial constraints, and time limitations—remain largely unaddressed [7].

However, these challenges are even amplified in the context of participatory AI. The aspiration of participatory practices to promote inclusivity and enable public deliberation often conflicts with the existing configuration of the AI ecosystem, which is dominated by globally distributed yet centrally controlled infrastructures operated by technology corporations driven by capitalist imperatives [26]. Furthermore, the emphasis on localized, community-grounded knowledge that results from participatory engagements

sits uneasily alongside the design goals of foundation models, which prioritize generalizability and versatility across domains and geographies [23].

Finally, the burden of enabling participatory practices that truly strive for inclusiveness and accountability, that center the voices of marginalized communities, and that aspire to open alternative—more just—pathways towards AI development, often is carried by individuals or groups of people who invest personal time and energy. Participatory AI stands as a form of activism that imposes a significant toll on the involved subjects (both researchers and participants), especially in terms of time, financial costs [3], and emotional labor [9]. Acknowledging these burdens not only raises questions of *what it takes for designers and researchers to be good Participatory AI practitioners*, but also *how participation, especially in its justice orientation, can be sustained over time*, even when project funding ends.

Thereafter, the aim of this panel is to bring issues of social justice to the central stage in the international Participatory AI discourse. More precisely, together with academic and industry experts, we will delve into questions of equitable access to participation, matters of fair access to information in these processes, as well as questions of effectiveness and sustainability of participatory AI both in corporate and civic contexts.

2 Tensions and questions

Critical AI and design scholarship warns us about the tensions and complexities of Participatory AI practices. In the panel, we will scaffold these with the ambition to set an agenda for Participatory AI committed to social justice. In particular, we will discuss:

- **How do we define participation?** *what are the criteria and conditions that qualify activities as 'participatory'?*
- *How do different disciplinary lenses affect whether we practice forms of participation?*
- **Can and should all Participatory AI work need to care for social justice?** *Is fighting for social justice compatible with AI research and development in the corporate environment?*
- **What are the key differences between grassroots and corporate forms of Participatory AI?** *Are these two incompatible? Can grassroots initiatives be systematically integrated into corporate innovation practices? If so, how?*
- **How do we prevent Participatory AI from ending up being merely performative,** *and what accountability mechanisms can we instantiate to make sure that participation and participants play a meaningful role?*
- **Is Participatory AI the best approach to account for social justice in AI systems development?** *What KPIs can we use to evaluate whether the approach achieves its ambitions and to what extent? What can be alternative approaches?*
- **Who is the Participatory AI designer/researcher?** *What are the necessary core competences in these practices? Where do we draw the line between the Participatory AI designer/researcher and participants?*
- **What role do artifacts play in Participatory AI** *Can traditional design skills be leveraged to support participation and dialogue?*

- **How do we ensure long-term sustainability of Participatory AI?** *What responsibilities do we bear as designers and researchers with regards to practicing participation bounded to the limits (temporal and economic) of funded research projects?*
- **How do we account for alternative epistemologies?** *How can we produce and share knowledge that is both scientifically valuable and societally useful? What language and dissemination channels should we use?*

3 Organizers and panelists

The panel will be organised by Maria Luce Lupetti, Christina Harrington, and Massimo Menichinelli. Harrington will also serve as moderator of the panel, which will be composed by five experts with experience from both industry and academia: Cristina Zaga, Laura Forlano, Alessandro Bozzon, and Vera Liao. Together, the organizers and panelists will bring a variety of perspectives and expertise in HCI, design and AI development, and social justice.

3.1 Organizers

3.1.1 Maria Luce Lupetti. is an Assistant Professor in Design at the Department of Architecture and Design at Politecnico di Torino. Her research is concerned with all matters of human entanglement with the artificial world, especially concerning complex technologies such as AI and robotics. She is co-director of POEL - Possible Entanglements Lab, a research group dedicated to investigating how people and technologies co-shape one another, and how design can help envision desirable configurations. She is the PI of *Participatory Design Justice for Ethical AI Transitions*, a three-year project funded by the Italian Ministry of Education and Research, under the FIS2 program. She also serves as Exhibit X section editor for *ACM Interactions Mag*, and sometimes as European Commission Expert (reviewer) on AI under the European Innovation Council (EIC) Accelerator Program.

3.1.2 Christina N. Harrington. is a research scientist at Google and affiliate faculty at Carnegie Mellon University. She addresses health and racial equity through human-computer interaction and design research approaches, most notably exploring community-based participatory methods as a way to support the needs of historically marginalized groups. Dr. Harrington has worked for over a decade to ideate, conceptualize and test technological interventions with communities that have historically been at the margins of mainstream design. Her work centers Black and Brown individuals, individuals with disabilities and impairments, and older adults in areas of community health advocacy, health information seeking, and community design of technology futures. Dr. Harrington has published at top HCI and design conferences and she was the recipient of the 2022 Skip Ellis Early Career Award from the Computing Research Association. Dr. Harrington holds a Ph.D. in Design from Georgia Tech, a Masters of Industrial Design from NC State University, and a B.S. in Electrical Engineering from Virginia Tech.

3.1.3 Massimo Menichinelli. is an Associate Professor at Elisava, Barcelona School of Design and Engineering (UVic-UCC) and director of Elisava Research, its research department. Doctor of Arts in New Media (Aalto University – Aalto Media Lab) and Master of Science in Industrial Design (Politecnico di Milano – School of

Design), since 2005, he has worked and researched on the role of design in developing ecosystems and infrastructures that support collaborative, distributed and open processes. He researches how to visualize and co-design design processes, how to measure the impact of design activities, how to develop and support social and digital innovation initiatives, how to develop and research platforms and how to promote the collaboration and creation of ecosystems between university and industry. He has participated in and coordinated several European research projects and given classes, lectures and workshops in several countries globally, including Italy, Spain, Finland, Germany, UK, Mexico, Colombia, South Korea, Singapore, China, Australia, USA. He has published several scientific articles and books on the topics of Open Design, Digital Fabrication, Makers, Fab Labs and Social Innovation. His books have been published in English, Spanish, Italian, French and Japanese.

3.2 Panelists

3.2.1 Cristina Zaga. is a tenured Assistant Professor of the Human-Centered Design group, the lead researcher at DesignLab, and the director of the JEDAI Network at the University of Twente (The Netherlands). Cristina holds a cum laude Ph.D. in Computer Science. Her research focuses on critical AI, design, and HCI, specifically resisting, reclaiming, and reimagining the roles of AI and robots in work and care. She pioneers transdisciplinary approaches that incorporate new materialism and agonistic pluralism as philosophies of action and social justice, and futuring as a design lens. She is an editor for scholarly journals (ACM THRI associate editor) and serves as an associate editor and co-organizer for the CHI, HRI, and DRS conferences. She spearheads initiatives like the Social Justice and AI network (JEDAI), which aims to counter-resist, reclaim, and reimagine AI in collaboration with communities. She has received numerous awards for her work, including the NWO Science Prize for DEI initiatives (2022), the Dutch Higher Education Award (2022), and the Google Women Techmaker Award and scholarship (2018).

3.2.2 Laura Forlano. is a Fulbright award-winning and National Science Foundation funded scholar, is a disabled writer, social scientist and design researcher. She is a Professor in the departments of Art + Design and Communication Studies in the College of Arts, Media, and Design and Senior Fellow at The Burnes Center for Social Change at Northeastern University. She is the author of *Cyborg* (with Danya Glabau, MIT Press 2024) and an editor of three books: *Bauhaus Futures* (MIT Press 2019), *digitalSTS* (Princeton University Press 2019) and *From Social Butterfly to Engaged Citizen* (MIT Press 2011). She received her Ph.D. in communications from Columbia University.

3.2.3 Alessandro Bozzon. is Professor of Human-Centered Artificial Intelligence at TU Delft, where he leads the Knowledge & Intelligence Design group. His research explores user modeling, recommender systems, crowdsourcing, and machine learning, aiming to design AI that is trustworthy, explainable, and aligned with human values. He has published extensively in leading HCI journals and conferences. He collaborates across academia and industry to develop intelligent technologies that address societal challenges while enhancing human creativity, decision-making, and well-being.

3.2.4 Q. Vera Liao. is an Associate Professor at University of Michigan CSE. Her research is focused on examining and mitigating risks of emerging technologies. Her most recent work focuses on transparency of AI technologies (e.g. explainability, evaluation, uncertainty communication) and its intersection with important aspects of human experiences with AI such as trust, reliance and control. Prior to joining U-M in 2025, she worked at Microsoft Research, where she was part of the FATE (Fairness, Accountability, Transparency, and Ethics of AI) group, and IBM T.J. Watson Research Center, where her research contributed to IBM products such as AI Explainability 360, Uncertainty Quantification 360, and Watson Assistant.

3.3 Audience engagement approach

The panel discussion will address the tensions and questions listed in section 2. Panelists will be prompted to address the various topics not only through directly asking questions, but also by showing key frameworks and seminal theories addressing matters of social justice in participatory AI literature. These include, but are not limited to, the *Ladder of Citizen Participation* [4], the concept of *Participatory ceiling* [23], the critique of participation as *tokenistic* [16], and *performative practice* [4, 22].

In addition to the invited panelists and organizers, the audience will also be actively engaged in the panel discussion. We will structure the debate as an "*open chair panel*": a dynamic, interactive discussion format in which members of the audience (one at a time) are invited to play an active role in the conversation. The "open chair" concept allows attendees to step up and contribute actively with their thoughts, not only asking questions, but also sharing perspectives on the topic at hand. As the format fosters a collaborative, inclusive and engaging environment where diverse viewpoints are welcomed, we see the method fitting particularly well with the topic at hand: *broadening participation to the AI discourse and valuing alternative viewpoints*.

3.4 Expected outcomes and contribution to the HCI community

The panel discussion aligns strongly with a growing body of academic and practice-based research that acknowledges the value and momentum of participatory practices in the AI ecosystem. In the last five years, in fact, the topic has been increasingly addressed in academic papers and there has been a growing engagement by the AI and HCI (and beyond) community through dedicated initiatives, such as academic workshops including the workshop on *Emerging Practices in Participatory AI Design in Public Sector Innovation* [19], the establishment of a *Special interest group on Participatory AI* [13], and the strengthening of *PAIRS, a Participatory AI Research & Practice Symposium* [15], that for two years now runs concurrently to the Global AI Impact Summit.

In contrast to these initiatives, our panel will distinctively serve as a platform to gather diverse perspectives on the methodological space of Participatory AI and its limitations and challenges in terms of social justice implications. We will delve into practical as well as theoretical questions of practicing Participatory AI, with a particular attention to the role and impact of individuals and institutions on the effects of these practices.

At the end of the workshop, we aim to provide the audience with a richer understanding of the conceptual and practical ways there can be to reimagine participatory approaches to address systemic inequities embedded in data practices, algorithmic design, and governance. Our ambition is to collectively imagine how participatory design practices around AI can be revised to sustain processes that acknowledge hidden labor, redistribute power, and enable shared ownership over AI technologies.

Finally, we aim to make available to the wider public the insights and perspectives shared during the panel. On the one hand, we aim to distill key insights that emerged from the panel into a *short argumentative article* to be published in an HCI public platform, e.g., *ACM Interactions Mag* or similar. On the other hand, we will invite the broader HCI community (through relevant SigCHI mailing lists) to contribute to answering our questions 2 after the panel. Ultimately, we will use the answers to create an open-access online repository of perspectives, ideas, and recommendations for practicing Participatory AI and Social Justice.

Acknowledgments

The panel is partially funded by: 1) the Italian Ministry of University and Research, FIS 2 funding program (Project reference PARJAI – CUD E53C24003890001); 2) JEDAI: Social Justice in AI Project 20003432 ENW Diversity Initiative award; 3) the Dutch Research Council's (NWO) Perspectief project TACIT (Inclusive Technologies for Access and Social Participation) under grant number P23.018; and 4) the European Erasmus+ projects "Open Access Contents on Design for Equality, Diversity, and Inclusion for Higher Education Programmes" (project acronym: EDIDesK; action type: KA220-HED (cooperation Partnerships in Higher Education; project reference: 2023-1-IT02-KA220-HED-000153774) and "RenAIssance - Revolutionizing Education and Nurturing Artificial Intelligence Skills for Sustainable Advancement in cultural and creative industries" (project acronym: RenAIssance; action type: Alliances for Sectoral Cooperation on Skills (implementing the 'Blueprint'); project reference: 101246744). The content of this panel does not reflect the official opinion of the European Union. Responsibility for the information and views expressed in the publication lies entirely with the author(s).

References

- [1] Haleh Afshar. 2013. The politics of fear: what does it mean to those who are otherized and feared? *Ethnic and Racial Studies* 36, 1 (2013), 9–27.
- [2] Daniela Atanasovska and Lejla Robeli. 2025. Brazil's AI Act: A New Era of AI Regulation. <https://gdprlocal.com/brazils-ai-act-a-new-era-of-ai-regulation/>
- [3] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–8.
- [4] Eric Corbett, Remi Denton, and Sheena Erete. 2023. Power and public participation in ai. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–13.
- [5] 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.
- [6] Journal Editors. 2025. For Authors. <https://dukeupress.edu/critical-ai#ForAuthors>
- [7] Salma Elsayed-Ali, Elizabeth Bonsignore, and Joel Chan. 2023. Exploring challenges to inclusion in participatory design from the perspectives of global north practitioners. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–25.
- [8] Cristina Gonzalez-Gonzalo, Eric F Thee, Caroline CW Klaver, Aaron Y Lee, Reinier O Schlingemann, Adnan Tufail, Frank Verbraak, and Clara I Sánchez. 2022. Trustworthy AI: Closing the gap between development and integration of AI systems in ophthalmic practice. *Progress in retinal and eye research* 90 (2022), 101034.
- [9] Christina Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–25.
- [10] Christina N Harrington, Aashaka Desai, Aaleyah Lewis, Sanika Moharana, Anne Spencer Ross, and Jennifer Mankoff. 2023. Working at the intersection of race, disability and accessibility. In *proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*. 1–18.
- [11] AI Hleg. 2019. Ethics guidelines for trustworthy AI. *B-1049 Brussels* (2019).
- [12] Thomas Lodato and Carl DiSalvo. 2018. Institutional constraints: the forms and limits of participatory design in the public realm. In *Proceedings of the 15th Participatory Design Conference: Full Papers-Volume 1*. 1–12.
- [13] Kyle Montague, Shaun Lawson, Pam Briggs, Valentina Nisi, Nuno Jardim Nunes, and Hugo Nicolau. 2025. SIG on Participatory AI: Identifying Research Gaps and Building Community Infrastructure for Citizen-Centred AI. In *Proceedings of the 16th Biannual Conference of the Italian SIGCHI Chapter (CHIItaly '25)*. Association for Computing Machinery, New York, NY, USA, Article 125, 2 pages. doi:10.1145/3750069.3769313
- [14] China's Ministry of Foreign Affairs. 2025. Global AI Governance Action Plan. https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202507/t20250729_11679232.html
- [15] PAIRS. 2025. PAIRS: Participatory AI Research & Practice Symposium. <https://www.pairs.site/>
- [16] Organizers Of Queerinaï, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, et al. 2023. Queer in AI: A case study in community-led participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1882–1895.
- [17] Joel Ross, Andrew Zaldivar, Lilly Irani, and Bill Tomlinson. 2009. Who are the turkers? worker demographics in amazon mechanical turk. *Department of Informatics, University of California, Irvine, USA, Tech. Rep* 49 (2009).
- [18] Niloufar Salehi, Lilly C Irani, Michael S Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, and Clickhappier. 2015. We are dynamo: Overcoming stalling and friction in collective action for crowd workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1621–1630.
- [19] Devansh Saxena, Zoe Kahn, Erina Seh-Young Moon, Lauren Marietta Chambers, Corey Jackson, Min Kyung Lee, Motahhare Eslami, Shion Guha, Sheena Erete, Lilly Irani, et al. 2025. Emerging Practices in Participatory AI Design in Public Sector Innovation. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–7.
- [20] Arnstein Sherry. 1969. A Ladder of Citizen's Participation. *Journal of American Institute of Planners* 35, 4 (1969), 216–224.
- [21] Mona Sloane. 2019. On the need for mapping design inequalities. *Design Issues* 35, 4 (2019), 3–11.
- [22] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation is not a design fix for machine learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–6.
- [23] Harini Suresh, Emily Tseng, Meg Young, Mary Gray, Emma Pierson, and Karen Levy. 2024. Participation in the age of foundation models. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1609–1621.
- [24] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, et al. 2023. Creating design resources to scaffold the ideation of AI concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2326–2346.
- [25] Nur Yildirim, Susanna Zlotnikov, Deniz Sayar, Jeremy M Kahn, Leigh A Bukowski, Sher Shah Amin, Kathryn A Riman, Billie S Davis, John S Minturn, Andrew J King, et al. 2024. Sketching ai concepts with capabilities and examples: ai innovation in the intensive care unit. In *Proceedings of the 2024 CHI conference on human factors in computing systems*. 1–18.
- [26] Meg Young, Upol Ehsan, Ranjit Singh, Emmet Tafesse, Michele Gilman, Christina Harrington, and Jacob Metcalf. 2024. Participation versus scale: Tensions in the practical demands on participatory AI. *First Monday* (2024).