Proceedings of the Eighth AAAI/ACM Conference on AI, Ethics, and Society (AIES 2025)

RelAItionship Building: Analyzing Recruitment Strategies for Participatory AI

Eugene Kim, Vaibhav Balloli, Berelian Karimian, Elizabeth Bondi-Kelly*, Benjamin Fish*

University of Michigan
genakim@umich.edu, ecbk@umich.edu

Abstract

Participatory AI, in which impacted community members and other stakeholders are involved in the design and development of AI systems, holds promise as a way to ensure AI is developed to meet their needs and reflect their values. However, the process of identifying, reaching out, and engaging with all relevant stakeholder groups, which we refer to as recruitment methodology, is still a practical challenge in AI projects striving to adopt participatory practices. In this paper, we investigate the challenges that researchers face when designing and executing recruitment methodology for Participatory AI projects, and the implications of current recruitment practice for Participatory AI. First, we describe the recruitment methodologies used in AI projects using a corpus of 37 projects to capture the diversity of practices in the field and perform an initial analysis on the documentation of recruitment practices, as well as specific strategies that researchers use to meet goals of equity and empowerment. To complement this analysis, we interview five AI researchers to learn about the outcomes of recruitment methodologies. We find that these outcomes are shaped by structural conditions of their work, researchers' own goals and expectations, and the relationships built from the recruitment methodology and subsequent collaboration. Based on these analyses, we provide recommendations for designing and executing relationship-forward recruitment methods, as well as reflexive recruitment documentation practices for Participatory AI researchers.

1 Introduction

Artificial intelligence for social impact (AISI) is being developed and deployed in many domains with promise to address pressing societal challenges, including healthcare (Kamineni et al. 2023) and conservation (Tuia et al. 2022). However, without involving stakeholders that are potentially affected by AI, there are risks to deploying AI broadly, such as encoding racism into AI systems (Benjamin 2019; Obermeyer et al. 2019) or supporting the centralization of power (Birhane et al. 2022b). Using participatory methods to carefully engage with stakeholders and impacted community members *while* developing AI systems realizes values of collective governance and socially beneficial AI (Bondi et al. 2021; Birhane et al. 2022a; Costanza-Chock 2020a). A body of work has been referring to the incorporation of participatory methods, drawn from theories of participatory design in computing and beyond, into AI development as "Participatory AI" (Kulynych et al. 2020).

This ethos of participation, while shared by many who study or aim to implement AISI and Participatory AI, can be difficult to achieve in practice (Lin et al. 2024). Barriers to implementing Participatory AI include a lack of clear standards and best practices (Delgado et al. 2023). Also, while Participatory AI idealizes the inclusion of marginalized and impacted communities, AI practitioners report difficulties in recruiting and engaging with certain impacted communities (Cooper and Zafiroglu 2024; Kallina and Singh 2024). Despite the crucial impact of recruitment on who participates (Vines et al. 2013), there is no systematic practical guidance on recruitment methods that AI researchers can use, especially for those who are new to participatory methods.

We aim to develop practical guidance on recruitment strategies by studying the current practice of AI researchers. Our main research questions in this work are: 1) What are the recruitment strategies used by AI researchers? 2) How do methods impact goals of Participatory AI, and 3) How can we improve outcomes of recruitment methods? First, we review a corpus of 37 AI projects to characterize and describe the current recruitment strategies used in AI, and note that recruitment methods are often not documented sufficiently for purposes of reproducibility and transparency. Then, we interview five AI researchers and practitioners to learn what methods tend to be more successful and why. By analyzing both our corpus data and the interviews, we learn that structural conditions, researchers' own goals and expectations for the project, and relationships between researchers and community groups are key factors that influence the outcomes and the kinds of impact accomplished by the AI system (Table 1). We then introduce recommendations to improve the outcomes of recruitment methodology towards achieving Participatory AI.

In the rest of the paper, we will provide an overview of related work in human-centered design and Participatory AI, discuss the methods we used for our review of existing practices of participation and our semi-structured interviews, the results of our analyses, and then introduce our recommendations for the design of recruitment strategies.

*These authors contributed equally.

Questions & Data	Takeaways
What are the recruitment strategies used in AI? **Data:** Corpus	• Decision users and data creators are often the ones recruited. • Researchers typically initiate recruitment. • Most recruitment is done via organizations and social media, websites, advertisements, and formal events. • Many details are not documented.
How do methods impact goals of Participatory AI? **Data:** Interviews	• Researchers' capacities to engage in participatory projects are limited by structural conditions. • Researchers' expectations and goals shape the methods they use and the relationships built with stakeholders. • Project outcomes depend on the relationships built by the researchers and participants.
How can we improve outcomes of recruitment methods? **Data:** Corpus, Interviews	• Relationship building is crucial when designing recruitment methodologies. • Building more processes and structural support, such as institutions that facilitate connections, can help improve recruitment methods. • Recruitment methodologies should be documented.

Table 1: Research Questions and Takeaways

2 Related Work

2.1 Participatory Methods in Computing

Participation as a research methodology in computing draws from the fields of participatory design (Spinuzzi 2005; Asaro 2000) and participatory action research (Costanza-Chock 2020a). Participatory approaches' focus on participant knowledge and deconstructing knowledge hierarchies (Frauenberger et al. 2015) lends itself well to design methodologies that aim to achieve socially just outcomes by elevating marginalized voices. Because of this, participatory approaches are key components to design methodologies in computing oriented towards social justice (Abebe et al. 2022; Irani et al. 2010; Bardzell and Bardzell 2011), and calls to apply principles from participatory research traditions to AI are increasing (Klein and D'Ignazio 2024; Costanza-Chock 2020b; Mohamed, Png, and Isaac 2020).

At the same time, participatory methods in AI development have been critiqued for "participation washing," in which participatory methods benefit the designers at the expense of the participants and manufacture consent for decisions made undemocratically (Sloane et al. 2022). In response to these critiques, scholars have worked to empirically characterize current Participatory AI practice (Cooper and Zafiroglu 2024; Groves et al. 2023; Kallina, Bohné, and Singh 2025), develop conceptual frameworks to evaluate existing participatory projects in AI according to the goals of participant empowerment and democratizing AI governance (Birhane et al. 2022a; Sloane et al. 2022; Berditchevskaia, Peach, and Malliaraki 2021; Feffer et al. 2023; Bondi et al. 2021; Corbett, Denton, and Erete 2023; Kallina, Bohné, and Singh 2025), and provide high level guidance and values to design participatory methods to achieve these goals (Berditchevskaia, Peach, and Malliaraki 2021; Feffer et al. 2023; Delgado et al. 2023; Gilman 2023; Bondi et al. 2021). The critical questions raised by previous work indicate that both practice and the social context of participatory methods are important in determining whether Participatory AI results in equitable, democratic AI development and systems.

2.2 Recruitment in Participatory AI

A key concern in Participatory AI is studying how technology designers shape the participatory process. Cooper and Zafiroglu (2024) examine the practices that designers, dubbed "participation brokers," use to shape who participates and how they participate in AI design. They note the degree of power that AI researchers have in determining these practices, and observe that AI researchers frequently form partnerships with community groups (Cooper and Zafiroglu 2024). In another vein, Lin et al. (2024) studies how AISI partnerships often fail to meet the needs of community partners due to issues such as funding constraints and mismatched expectations of AI's capabilities. They suggest that AI designers follow the lead of community members early in the project for AISI projects (Lin et al. 2024). Power and values-based recommendations for designing participatory methods are other key points (Costanza-Chock 2020a; Harrington, Erete, and Piper 2019).

Yet, there is still a need to provide practical guidance for participatory methods in AI. Kallina and Singh (2024) interview AI practitioners in corporate and organizational contexts, finding that 70% of respondents reported a lack of inclusivity in their samples (Kallina and Singh 2024). Respondents also note that certain stakeholder groups are difficult to identify and contact, in addition to having low incentives to participate. Other AI practitioners in corporate contexts similarly report needing clearer guidance as well as incentive structures to put participatory methods into practice (Delgado et al. 2023; Groves et al. 2023; Kallina, Bohné, and Singh 2025). As these works note, recruitment is key for achieving goals of Participatory AI; it configures who participates in the first place, and can impact the goals of sharing design control with impacted communities (Vines et al. 2013) and the values embedded in and performance of machine learning models (Diaz and Smith 2024). We aim to contribute to this body of work by clarifying the practice of participatory methods, particularly for recruitment.

2.3 Recruitment Methods

Existing recruitment methodologies seek to accomplish primary research goals, such as conducting surveys and interviews (Heerman et al. 2017), participant observation and other ethnographic methods (Negrin et al. 2022), co-design practices, and more. Other methodologies, drawn from research traditions such as participatory action research, also ensure diverse and representative sampling, encouraging equitable relationships with communities, and building organizational capacity (University of Michigan STPP program and Detroit Disability Power and the Detroit Justice Center; 2024; O'Brien et al. 2022; Beyond Chicago 2020). We refer to Table 2 for a brief overview of existing recruitment methods for a variety of purposes. Social scientists interviewing online communities, for example, highlight strategies such as participating in existing communities, using snowball sampling, and considering whether participants want to participate in the research projects (Hollingshead and Poole 2012). This suggests approaches to recruitment that are relational and dependent on the researcher's positionality.

Reflection on recruitment in the context of researchers' specific positionalities and power imbalances is also prevalent. For example, relationships and institutional context impact recruitment, the research questions that are pursued, and participant trust in researchers (Le Dantec and Fox 2015; Small and Calarco 2022; Gibson and Abrams 2003). In particular, Sum et al. (2025) examine how exploitative relationships between the university and local communities influence researchers' partnerships with community-based organizations. They ask researchers to build sustained relationships with community members, adopt individual tactics to support communities, and encourage institutions to revest in communities (Sum et al. 2025). Similarly, recruitment methods for non-"WEIRD" populations must also be contextual, flexible, and focus on needs and existing technological practices to be successful (De, Kanthawala, and Maddox 2025).

In short, much recruitment literature exists that AI researchers can draw on. Yet, little work has studied recruitment methods in AI or how those methods impact Participatory AI goals. We aim to help fill these gaps, towards improving recruitment methodology for Participatory AI.

3 Methods

We investigated recruitment practices in projects that use participatory methods during the design and development of AI by: 1) creating and analyzing a corpus of 37 projects in this space; 2) conducting institutional review board-approved semi-structured interviews with 5 AI researchers.

3.1 Corpus Creation and Analysis

Assembling the Corpus We conducted a literature review in order to build a corpus of AI projects that represented a diversity of modes of participation, as defined by Sloane et al. (2022). We drew case studies from existing reviews of Participatory AI, particularly (Birhane et al. 2022a; Delgado et al. 2023; Feffer et al. 2023; Berditchevskaia, Peach, and Malliaraki 2021; Kulynych et al. 2020), case studies of AI for Social Impact (Tambe, Fang, and Wilder 2022), recent publications at venues like FAccT, EAAMO, CHI, and AIES, and searches on Google Scholar for related terms like "Participatory AI." We adopted this approach in an effort to include projects using a wide variety of participatory methods not explicitly labeled as such, including projects where participation served a variety of goals (e.g. participation as labor vs. consultation (Sloane et al. 2022)) and impacts (e.g. participant empowerment (Delgado et al. 2023)). Appendix[1] Tables 10, 11, and 12 provide an overview of the countries represented by our sampled projects, the domains they were in, and the participation entry point (as defined by (Corbett, Denton, and Erete 2023)), respectively. To summarize, our corpus includes a wide variety of domains and both Global South and Global North contexts, with a majority of projects coming from Global North contexts. In addition, the majority of participation took place during "dataset development" and "deployment and monitoring" entry points. We determined saturation was reached and concluded the literature review when novel recruitment categories in AI projects were no longer identified.

Analyzing the corpus After compiling the corpus, we used inductive coding (Braun and Clarke 2006) with a goal to explore high-level themes and categories to characterize the recruitment methodologies used by AI researchers and practitioners. We found broad categories of recruitment methods, stakeholder types, and the actors who initiated the recruitment. We also coded each project by the level of empowerment afforded to the participants in the project (Delgado et al. 2023), to help us explore the question of how recruitment methods impact goals of Participatory AI. Two authors (EK, BK) coded the papers. Due to the small number of coders, we focused on building coding consensus via weekly meetings. Results are in Section 4.

3.2 ParticipAIte

We included the high-level categorizations from our qualitative coding into a preliminary database, which we call Par-

[1] https://realize-lab.github.io/participaite

	Diverse and representative samples	Equitable relationships with communities
Computing	Oversample minoritized groups, work with communities, remove barriers to participation, build participant trust, and use inclusive demographic qualities (Roscoe 2021)	Work with community groups in developing recruitment strategies (University of Michigan STPP program and Detroit Disability Power and the Detroit Justice Center; 2024)
Civics	Email and postal mail (Sharp and Anderson 2010)	Reach out to grassroots organizations, community leaders and local media (Nabatchi et al. 2012)
Social Sciences	Participate in existing communities to gain access, snowball sample, consider participants' perception of research (Hollingshead and Poole 2012)	Work with advocates and leaders of community groups (O'Brien et al. 2022)
Community Engagement	Develop culturally-specific engagement materials, work with community liaisons (Mancera et al. 2022)	Work with advocates and leaders of community groups, provide compensation (Beyond Chicago 2020)
Health and Medicine	Call and text using phone numbers from EHR (Langer et al. 2021)	Offer pay and accommodate meeting needs, and recruit through organizations, social/print media, clinics, snowball sampling (Leslie, Khayatzadeh-Mahani, and MacKean 2019)

Table 2: Non-exhaustive set of recruitment strategies in several fields

ticipAIte[2], and populated it with papers from our corpus. It includes information on recruitment methodology, including: a list of involved stakeholders, the recruitment strategy, who initiated recruitment, the specific call for participation, and lessons learned (see Appendix B).

3.3 Semi-Structured Interviews

For the semi-structured interviews, we aimed to gather more detail about the recruitment methodologies surfaced in our corpus, particularly from AI researchers with prior experience in recruiting participants for AI projects. Follow-up studies should also incorporate community group perspectives, which are essential for a more complete analysis of the empowerment of community groups.

We contacted ten researchers that were known to us to be researchers with Participatory AI experience via emails and direct messages, and interviewed five of them. Our selection was prioritized by our goal to converse with researchers who may have varied experiences in recruitment methodology, including the domain, format of participation, and goal of the participatory process. We acknowledge that the selection and response of participants were biased by our existing personal networks, interests, and reputations within certain research communities. See Table 3 for an overview of participant backgrounds.

Interviews lasted an hour, and were conducted via the video conferencing program Zoom. All interviews were recorded and transcribed with interviewee permission. We provided \$50 gift cards as compensation. The study was determined by the University of Michigan Institutional Review Board to be exempt from ongoing review (HUM00266324).

Interviews had two main components. For both components, participants had access to the ParticipAIte database, which contained an entry of one of their existing projects, as well as a project in our corpus. First, we asked about interviewees' previous and current experience with recruitment methodology, using database entries describing their recruitment methodology that we assembled from existing work about the projects. We aimed to elicit detailed narratives behind their experiences of developing recruitment methodology. The second main component was to discuss another recruitment case study in order to discuss documentation practices that would help with reproducibility. Finally, we asked for open ended, unstructured feedback on documentation about recruitment. See Appendix A for full interview questions.

As with the corpus, we performed an inductive coding on the interviews. Following thematic analysis (Braun and Clarke 2006), we developed higher level themes by analyzing data in context with the entire dataset, as well as descriptive themes. The results are in Sections 5 and 6.

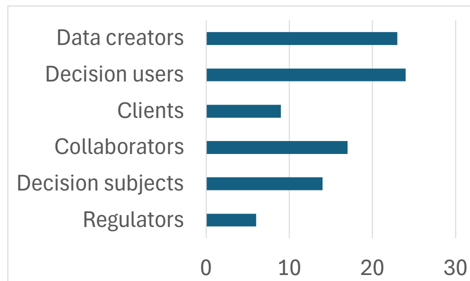

Figure 1: "Those recruited" papers (Appendix Table 6)

4 Recruitment Methodologies

Our corpus analysis was based on descriptive categories of recruitment methodologies, including: *Who is recruited*, *who recruits*, *recruitment strategy*, and *recruitment impacts on participant empowerment*. We also noted patterns in documentation. The Appendix lists the full corpus and analysis.

Who is recruited? In general, projects in our corpus tended to recruit either end users of an AI system ("Deci-

[2]https://realize-lab.github.io/participaite

ID	Role	Project experience discussed	Context	Academic Training
P1	Academia	AI for humanitarian purposes	GN	Computer Science
P2	Academia	AI for environmental and humanitarian purposes	GS & GN	Computer Science
P3	Academia	AI for humanitarian purposes	GS & GN	Computer Science
P4	nonprofit	AI and tech policy focused on marginalized groups' concerns	GN	Information Sciences
P5	Academia	AI for humanitarian purposes	GN	Computer Science

Table 3: Participants and descriptions (including Global South (GS) and/or Global North (GN) context focus)

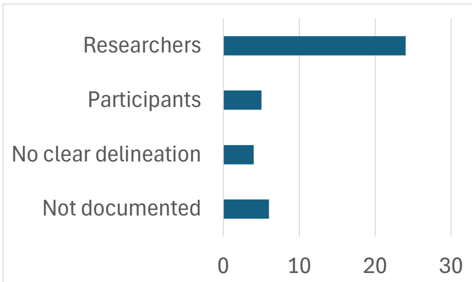

Figure 2: "Who recruits" papers (Appendix Table 5)

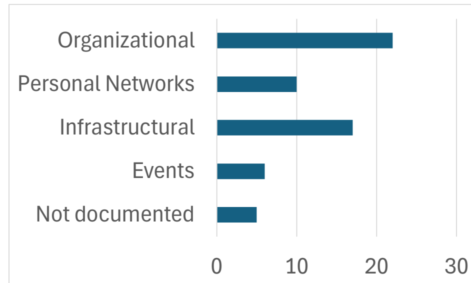

Figure 3: "Recruitment strategy" papers (Appendix Table 4)

sion users") or data labellers and workers ("Data creators") as participants. Results are presented in Fig.1 and Table 6 in the Appendix.

Works in the corpus also sometimes detailed the relevant stakeholder groups that they did not aim to recruit from (Lee et al. 2019; Yu et al. 2020; Newman et al. 2022; Cheng et al. 2021; Sendak et al. 2020), acknowledging how recruitment strategies determine who can or cannot participate. Some projects excluded stakeholders due to the challenges of engaging said stakeholders and their more indirect relationships to the system being developed, such as patients in a tool designed for healthcare workers (Sendak et al. 2020).

In our corpus, authors also noted that recruitment can shape barriers to participation. For example, enabling online participation and flexible scheduling were some strategies to ensure equitable participation and broader recruitment (Queerinai et al. 2023; Kuo et al. 2023; Lee et al. 2019). A lack of inclusion due to barriers could in turn lead to bias and generalization challenges (Awad et al. 2018; Calacci and Pentland 2022; Marian 2023; Bakker et al. 2022).

Who recruits? In addition to differences in who was recruited, there were also different stakeholders doing the recruitment (see Appendix Table 5 and Fig. 2). AI researchers or practitioners most frequently conducted recruitment in our corpus. Furthermore, some of our corpus provided descriptions of how the researchers or practitioners' previous experiences with the communities of their participants allowed access to research sites (Lee et al. 2019; Wilder et al. 2021; Eitzel et al. 2021). Our corpus also had a few rich descriptions of collaboration building between researchers and community partners (Delgado, Barocas, and Levy 2022; UNHCR Innovation Service 2019).

There were a few instances where the line between researcher and participant was blurred, as both parties engaged in the process of research (Queerinai et al. 2023; Nekoto et al. 2020). In our corpus, we also found several examples of participant-led research, where researchers were solicited by communities or developed the research problem in collaboration with communities (see Appendix for a list).

Recruitment Strategy Recruitment strategies were also quite varied, as shown in Fig. 3 and Table 4 in the Appendix. Recruitment strategies included working with organizations (organizational), drawing on personal networks, using social media or other mass media (infrastructural), and attending events like conferences (Table 4 details specific strategies in these categories). We found that most recruitment in our corpus was based on connecting to organizations, for example, by reaching out to community advocacy organizations or nonprofit organizations.

Empowerment Our corpus showcased several ways for recruitment methods to meet goals of equity and empowerment. For example, researchers ensured that their methods followed community norms via conversations with the community (Zhu et al. 2018), relevant domain experts (Kuo et al. 2023), and local institutions (Newman et al. 2022). Works in our corpus also aimed to build trust with marginalized communities by recruiting institutional actors that worked with these communities (Wilder et al. 2021; Kuo et al. 2023; Young, Magassa, and Friedman 2019; Frauenberger et al. 2015) and key organizers of communities (Calacci and Pentland 2022; Marian 2023).

While crowdsourcing has been noted as a form of participation that is considered labor, is more at the level of consulting or including (using Delgado et al. (2023)'s scale of participant empowerment), or offers a low amount of agency to participants (Nekoto et al. 2020), careful crowdsourcing can also empower community groups. A prominent example is Digital Workerism (Calacci 2022), where crowdsourcing was used to provide data for a worker-led campaign to investigate the impact of algorithmically determined wages on working conditions. In our corpus, researchers and practitioners used intentional strategies to reach marginalized communities during crowdsourcing, including decentralized and internet-based recruitment strategies to reach community members over large geographical divides (Queerinai et al. 2023; Nekoto et al. 2020), requiring few resources to recruit and participate (Li et al. 2023; Queerinai et al. 2023), and reaching out to strategic allies during the recruitment process (Nekoto et al. 2020; Irani and Silberman 2013). Offering some direct benefits to participants, such as easing existing workflow challenges (Kelling et al. 2012), or ensuring that the resulting data and knowledge production were of immediate benefit to participants (Li et al. 2023), were also important practices.

4.1 Documentation of Recruitment Strategies

So far, we have focused on those parts of recruitment methodologies that are documented in our corpus. However, projects varied in level of documentation. For example, Figs. 2 - 3 showed multiple projects lacking enough documentation for us to determine a categorization. Some projects described the specific actors and institutions they work with, while others only referred to the kind of institution. While omitting specific names to protect privacy is certainly an important consideration, we believe documentation of other aspects of recruitment are important for purposes like reproducibility and transparency. For instance, we noted several helpful examples where authors detailed reasons for excluding or including certain stakeholders, how researchers leveraged existing personal relationships during recruitment, or specific motivations for selecting certain recruitment strategies. These were not documented across the entire corpus. As such, we were motivated to interview AI researchers and learn about their experiences with recruitment methodology.

5 Results from Interviews

Our interviews focused on interviewees' previous experience with recruitment. To briefly summarize, each of the interviewees identified as researchers engaging in recruitment. They were primarily concerned with recruiting collaborators, mostly community organizations, other grassroots advocacy and activist groups, and members of specific communities, for co-design projects. To meet these recruitment goals, interviewees largely began by identifying partners with experience and existing relationships with communities the interviewee wished to engage with as a starting point, and reaching out most frequently through personal networks and cold calling. Once they established a relationship, e.g., with a nonprofit or government agency, they then often partnered with organizations for many core tasks, including: surveying communities, recommending policy, and introducing interviewees to other stakeholders. This aligns with our findings from the corpus, which showed the prevalence of organizational and personal network recruitment (Fig. 3), and slightly differs from formal social science recruitment methods, such as surveying or random sampling.

We now analyze and present our detailed findings in three connected themes. First, we explore how researchers' capacities to begin and engage in participatory projects are inherently limited by structural conditions. Second, we explore how researchers' expectations and goals shape the methods they use, and how methods can impact the relationships built with stakeholders in the domain. Finally, we argue that outcomes of participatory projects are fundamentally based on the relationships built by the researchers and participants.

5.1 Structural Challenges to Recruitment

Before we begin to address the recruitment methodologies that are used by researchers, as well as their outcomes, we need to discuss what constraints and limitations researchers face.

Funding A first barrier is funding. Researchers noted: *"to support this kind of project is not as easy compared to say, defense and security,"* pointing to a clear lack of funding support for projects supporting community-based organizations (P2). Not only do the researchers need to get funding for themselves, they also need to provide funding for their collaborators (P2). Funding for the participating community-based organization is particularly important, as participants, staff, and the organization need to be paid for their time and any additional coordinating labor (P4). One interviewee also addressed the fact that continued funding is a significant problem to achieving participatory projects with a long-term goal and focus: *"there's not a ton of funding available for maintenance of systems"* (P3).

Time and effort A second barrier is the time commitment added on from deciding to pursue a participatory project. For example, researchers discussed how participation from organizations and communities required senior researchers to *"identify a problem at the right scope"* (P2) for a student to work on, as well as prototyping and creating pilot experiments to determine operational feasibility (P5). Due to the considerable time investments required from both the AI practitioners and the involved stakeholders, both parties would need to commit seriously to the length of the project (P2).

Logistics and coordination These challenges are compounded by logistical and legal concerns. P5, for example, was hampered by such issues when attempting to include a certain stakeholder group. Coordinating with communities and stakeholder groups can require considerable investment and project management, including overhead not necessary when working alone, such as additional meetings (P5), testing prototypes in the field (P5), and engaging in the domain (P2).

Ethical engagement and building trust The entrenched power dynamics between researchers and community groups is another barrier. One researcher discussed the difficulty in gaining the trust of certain groups: *"you're always working in the wake of whatever the university that you work at has done to that community"* (P4). Cognizance of structural harms and of ways that marginalized groups are often further marginalized by exploitative projects means that researchers need to evaluate whether the community groups and organizations that they are working with even have the capacity to participate in these projects, and if participation is putting an undue burden on the community groups (P4).

Lack of (relevant) expertise Some researchers are not trained in engaging with and recruiting community members or groups. One researcher stated: *"we just do not have any of the expertise needed to kind of navigate that world and do anything sensible"* (P1) when it comes to directly outreaching to community members. In fact, some interviewees were not involved with community engagement or recruitment despite being involved in Participatory AI projects, as teams of researchers had split into community-engagement and technical intervention subteams, with the technical researchers relatively separated from decisions about engaging with the community (P1, P5).

Furthermore, technical expertise can also serve as a structural challenge in recruitment. Several interviewees had explicit expectations to develop specific technical contributions as a result of a collaboration due to career incentives. This presents a very serious trade-off for interviewees: *"how much effort and time do you really want to spend on these AI for social good projects that may not fully leverage your technical expertise to work on?"* They suggested that, for researchers interested in creating social impact without engaging in the additional labor required to actively involve community organizations and groups, foregoing recruitment and working on open-source projects may be more efficient for achieving social impact (P2).

5.2 Impacts of Challenges and Goals on Recruitment

These challenges, coupled with researcher goals, ultimately incentivized researchers to find partners as a means of finding participants. This, in turn, shaped recruitment methods.

Personal connections and partners Working with organizations proved to be an effective strategy for researchers to engage in participatory projects, particularly with organizations that had both the capacity to work with researchers and were already embedded in the communities of interest. *This helped address challenges with logistics, engagement, expertise, and trust*, to name a few.

For example, P4 noted that organizations with a grassroots orientation enable better connection to membership and an ability to directly bring community members into the process of Participatory AI, while P1, P2, and P5 stressed the importance of having a community partner that already had connections to communities and the relevant expertise to work in these domains. P4 also found that going to the director of advocacy or programming at nonprofits was promising due to an indication that the organization had an increased capacity for collaboration. These existing connections to partners with capacity were significant in starting a collaboration.

At the same time, the structural barriers that made partnering with community groups more effective also led to challenges. P2 summarized the challenge as *"I think in general, it's not easy, because we have limited capacity, and these nonprofits also have their own agenda and priorities"* (P2). Furthermore, one interviewee reflected how success is often determined by an intuition surrounding the quality of the relationship: *"I do feel like a lot of recruitment happens in sort of...a ground-up way...a lot of it is vibes"* (P3). Ultimately, *"finding [any] collaborator is easier than finding a good collaborator"* (P2).

All of our interviewees reported heavily relying on personal networks in their recruitment processes; including academic and personal networks. In fact, our interviewees reported that relying on personal networks led to successful academic outcomes (P1-P5).

Mixed success from cold calls To find partners, the alternative to leaning on connections was cold calling potential partners. Cold calls, whether through email, phone, or direct messages, were reported to have mixed success rates - failures of collaborations starting from cold calls included no responses (P2, P4), or collaborations that were short-term or underwhelming (P2). P2 shared an anecdote where a close colleague emailed dozens of nonprofits in their local area, but were unable to successfully create a collaboration. While researchers admitted that they could not always know why a cold-call attempt failed, given that people are very busy and have other priorities, refusal to participate could indicate a deeper mistrust of researchers and the institutions they represent (P4), adding yet another difficulty to the process.

Establishing a relationship from cold calls is often easier when the researcher can lean on their existing reputation in the community (P2, P3). P4 said that *"If someone's not familiar with you, I think that some of this work comes across like weird marketing research ... like getting a cold call from an organization that wants to pay you..."* (P4). To address this challenge and to successfully get a response, interviewees suggested showing examples of previous existing projects (P2), emphasizing compensation (P4), and making a case that the community group has novel, important contributions to make (P4).

Challenges to achieve representativeness All of our interviewees engaged with goals of representativeness when they decided how and who to recruit. For instance, researchers were motivated to use cold calls to bring in partners they otherwise would not. P4 reflected: *"It would be a lot easier if we did only rely on our personal networks, because the [organizations] would be much more likely to participate"*, as they may already see the value in the project, but *"if you stay too close to home, then you're just talking to the same people with the same interests. So often, it's cold calls"* (P4).

One key challenge is that capacity limits interviewees' ability to achieve ideals of representativeness. Even though our interviewees aspired to include all relevant groups, this was not practically possible. One interviewee regretted not including one group of stakeholders, mentioning that even a small, biased sample would be better than having none (P5). The reasons cited by researchers for not including these specific stakeholder groups were a lack of capacity to navigate the logistical (e.g., institutional review board) obstacles to include these groups, noting that doing so would have added potentially years to the research study (P5).

Researchers also emphasized that representativeness is a priority throughout project design and development; questions of how to make participation *"successful and sustainable"* are not *"specific to recruitment"* (P3). The way people participate and to what degree they can influence the system, for instance, are similarly important. Recruitment is key for achieving this goal of representativeness throughout a project.

5.3 Relationship Impacts on Outcomes

The challenges in finding participants underscore significant trade-offs between pursuing and building relationships with community groups, and ensuring participatory projects can meet researchers' career needs, such as milestones and publishing requirements. Given that these kinds of tensions are largely structural, we found that the interviewees responded by trying to build longer-term relationships that were mutually beneficial to community partners and researchers.

Long-term partnerships Long-term partnerships were noted as having many benefits for interviewees and the partners they were working with. Ideally, when recruiting partners, researchers were looking for "long-term collaborations" (P2). Long-term relationships were noted as valued by community groups, who have made it clear to researchers that parachuting into community to produce research and leaving is not acceptable (P2). Other benefits included producing several research papers over a stand-alone application (P2) and establishing new governance relations between designers, governors, and community organizations (P4). Echoing concerns and critiques of exploitative projects, P4 described the frequently short-term nature of Participatory AI as something *"arguably broken about Participatory AI."*

Interviewees in our study discussed several more successful strategies for building longer-lasting, durable relationships between researchers and their collaborators (P1-P3). For example, P1 talked about a project in which the nonprofit center benefited, as the conducted research provided important administrative tasks and helped provide benefits to the community that the nonprofit served, despite the center choosing not to use the algorithm the researcher created in day-to-day usage. This relationship was mediated by a senior researcher's long-standing relationship with a collaborator in the domain who had years of experience collaborating with relevant nonprofits. In addition, P2 spoke to a model of relationship-building that was more scalable and accessible to junior researchers with fewer resources: senior researchers maintaining long-standing relationships with many partners and nonprofits to help scope projects for more junior researchers, and relying on existing long-term relationships between senior researchers and community partners. Thus, long-term relationships can be the basis of many separate projects, and provide continued impact for community partners.

Relationship challenges and deployment Yet, long-term partnerships are difficult to achieve. One interviewee recalled a collaboration that broke down right before the deployment of an algorithm, despite the CEO of the nonprofit being very enthusiastic about the project in the initial stages. When reflecting on what could have been improved, the interviewee suggested that the researchers and the nonprofit needed an "explicit collaboration" from the beginning to avoid failure: *"Maybe just having them attend way more meetings, even on the [technical] side, would help them feel even more like part of this."* Unfortunately, in this case, the algorithm did not get to be deployed. Three of our interviewees (P1, P2, P5) reported that the algorithmic interventions created for nonprofit partners were not still currently deployed after the researchers "handed off" the project to the organizations. This happened when the project would end, and the researchers would no longer be involved in the project, often moving to other universities (P1 P2, P5), or when the nonprofit would fundamentally alter the way they collect data, rendering the algorithm obsolete (P1, P2). The challenges maintaining relationships between researchers and community-based organizations during the development of the algorithm may have contributed to limited algorithmic impacts.

Outcomes beyond deployment Sometimes, benefits to nonprofits or community partners existed beyond algorithmic impacts. Other benefits mentioned in our interviews included receiving funding (P2, P4), gaining legitimacy in the eyes of governments and other decision-makers (P3), establishing data collection practices initiated by researchers (P1), or even developing AI literacy (P2). This indicates, then, that truly impactful participation might not be measured in terms of algorithm deployment alone, but also holistically with respect to other benefits.

6 Discussion

Given the challenges we found to successfully recruiting for Participatory AI, there is a clear need for improved recruitment methodology. Design recommendations for Participatory AI from previous work can and should apply to recruitment methods, such as recommendations to examine and challenge power (D'ignazio and Klein 2023), to prioritize marginalized communities (Costanza-Chock 2020a), and to ensure research meaningfully benefits the involved communities (Erete, Rankin, and Thomas 2021). The interviewees (see Section 5) wanted to achieve these kinds of goals, but the challenges they faced were significant. We now discuss possible approaches to improving recruitment strategies.

6.1 Building Community

Our findings highlighted the highly relational aspect of recruitment. From researchers' reputations to building trust to achieving representativeness, recruitment methodologies often succeeded or failed based on the relationships between researchers and communities.

We encourage researchers to expand conceptions of what is a legible recruitment strategy and who is able to use these strategies. For instance, in addition to finding partners, researchers can establish themselves within the networks of communities that they aim to work with by helping organize or attending workshops that focus on prominent issues in local communities, or regularly volunteering with or organizing alongside community groups. This echoes findings from studies for researchers to build personal relationships with the communities they work with (Sum et al. 2025; Harrington, Erete, and Piper 2019). Forefronting trust and building strong relationships in recruitment also helps tackle a core question – who is actually participating in Participatory AI – by encouraging participation from marginalized groups.

There are existing mechanisms for researchers to join community-building efforts. Interviewees mentioned building spaces for community groups and organizations to create relationships with AI practitioners, such as at the FAccT conference, where *"if you make space for more and more and more community-based organizations and advocacy groups to present there, you do build real relationships"* (P4). Another interviewee (P2) mentioned creating academic spaces where a principal investigator can maintain many relationships with community organizations to support PhD projects, and summer programs where a senior academic matches undergraduate and masters students to community groups. P3 discussed creating a formal network of community organizations and advocacy groups with similar concerns: *"the project is more well known and global ... so sometimes people find us."* Maintaining such a network included holding regular events for interested groups and extending regular invitations to engage with the researchers, such as relevant research colloquiums and symposiums, resulting in long-term connections and collaborations over time (P3, P4). Although labor-intensive, building spaces that can help connect expertise, potential funding, and need is a promising approach to growing personal and professional networks without an individual researcher or community needing to do it all themselves.

Once researchers are able to establish themselves within a network, it is far more likely that community organizations will approach a researcher with a proposal. One interviewee said, after participating in a project with an organization that was organizing similar community groups, *"this relationship unlocked all these relationships to local community-based organizations"* (P4). Describing how these relationships eventually led to a collaboration, P4 described how *"a contact was looking for some help from us . . . [and] as that conversation evolved . . . we realized there was a lot of overlap."* Once an interviewee had one successful collaboration, they were able to focus on recruiting similar organizations with similar structural needs that could be addressed by the algorithmic intervention that they had created (P3), showing a reputational benefit for similar communities.

Besides benefits to the researcher, earlier engagement with communities can help all participants get involved at earlier stages of projects and better engage with other stakeholders, potentially building a sense of shared ownership and algorithmic interventions with a clear and sustained social impact. Moreover, technical interventions made with engagement from the relevant communities early on are key to anticipating needed changes in the design phase, reflecting recommendations to "go out in the field" (Perrault et al. 2020). Earlier engagement also makes it easier to clearly agree on the potential benefits at the beginning of the project, using the findings in this work and others including Lin et al. (2024); Ismail et al. (2023); Erete et al. (2025) as guidance to frame such conversations. When researchers are in community with the broader, local community, there are possibilities for directly recruiting community members themselves in a grassroots manner, rather than recruiting through formal institutions that represent them.

While building community is promising to further Participatory AI goals, we acknowledge that the challenges detailed in Section 5 remain. Individually, a researcher may need to adjust expectations and narratives to make the work *"legible as a research contribution to a specific community"* (P3). Without changes to the institutions that support researchers and communities alike, fully overcoming many of these challenges, from funding to coordination to sustaining long-term relationships, will remain difficult. We believe that sharing challenges and successes will help improve recruitment methodology.

6.2 Reflexive Documentation

During our analysis of the corpus, we saw significant disparities in the detail of documentation of recruitment strategies. When we interviewed researchers, we learned many more details about recruitment than from the researcher's published papers. These stories included the details and difficulties of finding funding, navigating specific social networks, engaging with specific actors to build relationships with community partners, and working with a set of institutions, such as universities and conferences, to help researchers build the relationships to recruit. As such, we recommend better documenting recruitment practices. Our recommendation to document these "backstage" aspects (Bødker and Kyng 2018) draws from a tradition of researcher reflexivity to improve transparency and create better recruitment methodologies in the field. Another benefit of introducing this reflexivity is to make the invisible work of recruitment and establishing relationships more visible to the research community, which can initiate change in how this labor is valued by institutions (Le Dantec and Fox 2015). Reflexivity is also a tool that can allow researchers to more clearly understand the complexities of power dynamics that impact their research projects, including the hierarchies, interests, and worldviews driving the decisions in research (Miceli et al. 2021). We now attend to three factors of reflexive documentation for participatory projects: social and personal networks, institutional support, and learning from refusal. These three factors are introduced as a starting point for reflexive documentation.

Personal networks In our corpus of Participatory AI projects, how recruitment leveraged personal networks was not always very clear, yet our interviewees emphasized the importance of personal networks to their recruitment methodologies. We suggest that researchers record which relationships were key in building community partnerships, how these relationships were initially formed (as raised by P2), and how they resolved any conflicts in the relationships. Detailed narratives of recruitment are given in Le Dantec and Fox (2015); Delgado, Barocas, and Levy (2022), and Tandon et al. (2022), but even brief descriptions would help provide additional context for others recruiting.

Institutional support As researchers navigate their own capacities for engaging in Participatory AI, guidance on finding funding is key. One interviewee in our study strongly suggested documenting how researchers recruited funding in Participatory AI projects (P2). Beyond sharing strategies for identifying such funding sources, it would be useful to share details about the structures of institutions that support Participatory AI development. As an example of this, one of our interviewees spoke to how replicating successful institutions in other universities that enable Participatory AI development could help their own university support similar endeavors to improve recruitment outcomes (P1).

Refusal and non-response Researchers tended to document the successful recruitment strategies, but often had stories of unsuccessful recruitment strategies that were not documented. Recruitment non-response and the failure of strategies to produce successful partnerships are also part of the recruitment process, and should be documented. Narratives of these failures and how they were overcome can provide starting points for other researchers designing recruitment methodology. As P5 put it, *"finding out what went wrong is almost more useful than finding out what went right, because there are certain things that...everyone would do differently in each of the papers."*

Documenting the instances of refusal to participate can provide light on which communities are underrepresented in Participatory AI and why. AI designers have called for supporting refusal in design (Barocas et al. 2020), and P4 noted that *"[academic and nonprofit spaces] can be... almost violent"* for certain people, highlighting the importance of refusal as a mode of participation. Using data collected about refusal, Participatory AI can create methods to incorporate refusal as an active mechanism of governance and accountability, potentially building on models of informed refusal in bioethics (Benjamin 2016).

6.3 Towards Representativeness

Finally, representativeness and meaningful participation are key goals in Participatory AI, yet our interviews highlighted the difficulty achieving these goals due to logistical challenges and challenges with cold calling potential partners, for instance. Right now, finding partners and communities is done largely ad-hoc. Even when done well, it is hard to learn what worked well, let alone to reproduce the methodology in new projects and new domains. In part, this is a problem of documentation, but it is also because researchers may not have the same access or context other researchers had. We believe that there would be value in combining relationship building with a more systematic recruitment methodology. Ideally, a systematic recruitment methodology could be utilized by many researchers to improve representativeness, even in new contexts, and could allow repeated evaluation of methodology across many projects. A focus on recruitment methodology can highlight the key role that affected communities should play in developing AI.

7 Limitations and Future Directions

The selection of the corpus and interviewees reflect our positionalities as researchers based in a research institution in the Global North, as well as our personal networks and publication goals. In addition to interviewing AI researchers, community groups that have partnered with AI researchers could provide more insights on these challenges, particularly from their perspectives and needs. Existing recruitment practices of other kinds of practitioners of Participatory AI, such as community groups that organize and gather data about AI systems, are also under-represented and deserve further study. Together, these directions can help advance the goal of improving Participatory AI and AI for Social Impact practice.

Acknowledgments

We would like to thank our interviewees for taking the time to converse and share their perspectives with us. We would also like to thank the reviewers, as well as members of the Computational Social Science group, Realize Lab, the Political Economy and Algorithms group, Nathan Kim, Shreya Chowdhary, and others at the University of Michigan, for feedback at crucial stages of the paper.

References

Abebe, V.; Amaryan, G.; Beshai, M.; Ilene; Gurgen, A. E.; Ho, W.; Hylton, N. R.; Kim, D.; Lee, C.; Lewandowski, C.; and others. 2022. Anti-Racist HCI: notes on an emerging critical technical practice. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–12.

Asaro, P. M. 2000. Transforming society by transforming technology: the science and politics of participatory design. *Accounting, Management and Information Technologies*, 10(4): 257–290. Publisher: Elsevier.

Awad, E.; Dsouza, S.; Kim, R.; Schulz, J.; Henrich, J.; Shariff, A.; Bonnefon, J.-F.; and Rahwan, I. 2018. The moral machine experiment. *Nature*, 563(7729): 59–64. Publisher: Nature Publishing Group.

Bakker, M.; Chadwick, M.; Sheahan, H.; Tessler, M.; Campbell-Gillingham, L.; Balaguer, J.; McAleese, N.; Glaese, A.; Aslanides, J.; Botvinick, M.; and others. 2022. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35: 38176–38189.

Bardzell, S.; and Bardzell, J. 2011. Towards a feminist HCI methodology: social science, feminism, and HCI. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 675–684.

Barocas, S.; Biega, A. J.; Fish, B.; Niklas, J.; and Stark, L. 2020. When not to design, build, or deploy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 695–695.

Benjamin, R. 2016. Informed refusal: Toward a justice-based bioethics. *Science, Technology, & Human Values*, 41(6): 967–990. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Benjamin, R. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity. ISBN 978-1-5095-2639-0.

Berditchevskaia, A.; Peach, K.; and Malliaraki, E. 2021. Participatory AI for humanitarian innovation. *London: Nesta*.

Beyond Chicago. 2020. Why Am I Always Being Researched?

Birhane, A.; Isaac, W.; Prabhakaran, V.; Diaz, M.; Elish, M. C.; Gabriel, I.; and Mohamed, S. 2022a. 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*, EAAMO '22, 1–8. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9477-2.

Birhane, A.; Kalluri, P.; Card, D.; Agnew, W.; Dotan, R.; and Bao, M. 2022b. The values encoded in machine learning research. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, 173–184.

Bondi, E.; Fang, F.; Hamilton, M.; Kar, D.; Dmello, D.; Choi, J.; Hannaford, R.; Iyer, A.; Joppa, L.; Tambe, M.; and others. 2018. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32. Issue: 1.

Bondi, E.; Xu, L.; Acosta-Navas, D.; and Killian, J. A. 2021. Envisioning Communities: A Participatory Approach Towards AI for Social Good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 425–436. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-8473-5.

Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2): 77–101. Publisher: Taylor & Francis.

Bødker, S.; and Kyng, M. 2018. Participatory design that matters—Facing the big issues. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(1): 1–31. Publisher: ACM New York, NY, USA.

Calacci, D. 2022. Organizing in the end of employment: information sharing, data stewardship, and digital workerism. In *Proceedings of the 1st Annual Meeting of the Symposium on Human-Computer Interaction for Work*, 1–9.

Calacci, D.; and Pentland, A. 2022. Bargaining with the black-box: Designing and deploying worker-centric tools to audit algorithmic management. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–24. Publisher: ACM New York, NY, USA.

Cheng, H.-F.; Stapleton, L.; Wang, R.; Bullock, P.; Chouldechova, A.; Wu, Z. S. S.; and Zhu, H. 2021. Soliciting stakeholders' fairness notions in child maltreatment predictive systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–17.

Cooper, N.; and Zafiroglu, A. 2024. From Fitting Participation to Forging Relationships: The Art of Participatory ML. *arXiv preprint arXiv:2403.06431*.

Corbett, E.; Denton, E.; and Erete, S. 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.

Costanza-Chock, S. 2020a. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.

Costanza-Chock, S. 2020b. Introduction: #TravelingWhileTrans, Design Justice, and Escape from the Matrix of Domination. In *Design Justice*. 1 edition.

De, A.; Kanthawala, S.; and Maddox, J. 2025. Who Gets Heard? Calling Out the" Hard-to-Reach" Myth for Non-WEIRD Populations' Recruitment and Involvement in Research. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 855–867.

Delgado, F.; Barocas, S.; and Levy, K. 2022. An uncommon task: Participatory design in legal AI. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1): 1–23. Publisher: ACM New York, NY, USA.

Delgado, F.; Yang, S.; Madaio, M.; and Yang, Q. 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.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

Diaz, M.; and Smith, A. D. 2024. What Makes An Expert? Reviewing How ML Researchers Define" Expert". In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 358–370.

D'ignazio, C.; and Klein, L. F. 2023. *Data feminism*. MIT press.

Eitzel, M.; Solera, J.; Hove, E. M.; Wilson, K.; Ndlovu, A. M.; Ndlovu, D.; Changarara, A.; Ndlovu, A.; Neves, K.; Chirindira, A.; and others. 2021. Assessing the potential of participatory modeling for decolonial restoration of an agro-pastoral system in rural Zimbabwe. *Citizen Science: Theory and Practice*, 6(1).

Erete, S.; Corbett, E.; Smith-Walker, N.; Cunningham, J. L.; Gatz, E.; Park, T.; Perry, T.; Wilcox, L.; and Denton, R. 2025. Towards Equitable Community-Industry Collaborations: Understanding the Experiences of Nonprofits' Collaborations with Tech Companies. *Proceedings of the ACM on Human-Computer Interaction*, 9(2): 1–31. Publisher: ACM New York, NY, USA.

Erete, S.; Rankin, Y. A.; and Thomas, J. O. 2021. I can't breathe: Reflections from Black women in CSCW and HCI. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3): 1–23. Publisher: ACM New York, NY, USA.

Feffer, M.; Skirpan, M.; Lipton, Z.; and Heidari, H. 2023. From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 38–48. Montr\'{e}al QC Canada: ACM. ISBN 9798400702310.

Finn, A.; Jones, P.-L.; Mahelona, K.; Duncan, S.; and Leoni, G. 2022. Developing a Part-Of-Speech tagger for te reo Māori. In *Proceedings of the Fifth Workshop on the Use of Computational Methods in the Study of Endangered Languages*, 93–98.

Flanigan, B.; Gölz, P.; Gupta, A.; Hennig, B.; and Procaccia, A. D. 2021. Fair algorithms for selecting citizens' assemblies. *Nature*, 596(7873): 548–552.

Frauenberger, C.; Good, J.; Fitzpatrick, G.; and Iversen, O. S. 2015. In pursuit of rigour and accountability in participatory design. *International journal of human-computer studies*, 74: 93–106. Publisher: Elsevier.

Freedman, R.; Borg, J. S.; Sinnott-Armstrong, W.; Dickerson, J. P.; and Conitzer, V. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, 283: 103261. Publisher: Elsevier.

Gibson, P.; and Abrams, L. 2003. Racial difference in engaging, recruiting, and interviewing African American women in qualitative research. *Qualitative Social Work*, 2(4): 457–476. Publisher: Sage Publications.

Gilman, M. E. 2023. Democratizing AI: Principles for Meaningful Public Participation. *Data & Society*.

Groves, L.; Peppin, A.; Strait, A.; and Brennan, J. 2023. Going public: the role of public participation approaches in commercial AI labs. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1162–1173.

Halfaker, A.; and Geiger, R. S. 2020. Ores: Lowering barriers with participatory machine learning in wikipedia. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2): 1–37. Publisher: ACM New York, NY, USA.

Harrington, C.; Erete, S.; and Piper, A. M. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on human-computer interaction*, 3(CSCW): 1–25. Publisher: ACM New York, NY, USA.

Heerman, W. J.; Jackson, N.; Roumie, C. L.; Harris, P. A.; Rosenbloom, S. T.; Pulley, J.; Wilkins, C. H.; Williams, N. A.; Crenshaw, D.; Leak, C.; and others. 2017. Recruitment methods for survey research: findings from the mid-south clinical data research network. *Contemporary clinical trials*, 62: 50–55. Publisher: Elsevier.

Hollingshead, A.; and Poole, M. S. 2012. *Research methods for studying groups and teams: A guide to approaches, tools, and technologies*. Routledge.

Imran, M.; Alam, F.; Qazi, U.; Peterson, S.; and Ofli, F. 2020. Rapid damage assessment using social media images by combining human and machine intelligence. *arXiv preprint arXiv:2004.06675*.

Irani, L.; Vertesi, J.; Dourish, P.; Philip, K.; and Grinter, R. E. 2010. Postcolonial computing: a lens on design and development. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 1311–1320.

Irani, L. C.; and Silberman, M. S. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 611–620.

Ismail, A.; Thakkar, D.; Madhiwalla, N.; and Kumar, N. 2023. Public health calls for/with AI: an ethnographic perspective. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2): 1–26. Publisher: ACM New York, NY, USA.

Kallina, E.; Bohné, T.; and Singh, J. 2025. Stakeholder Participation for Responsible AI Development: Disconnects Between Guidance and Current Practice. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 1060–1079.

Kallina, E.; and Singh, J. 2024. Stakeholder Involvement for Responsible AI Development: A Process Framework. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–14. San Luis Potosi Mexico: ACM. ISBN 9798400712227.

Kamineni, M.; Ötleş, E.; Oh, J.; Rao, K.; Young, V. B.; Li, B. Y.; West, L. R.; Hooper, D. C.; Shenoy, E. S.; Guttag, J. G.; and others. 2023. Prospective evaluation of data-driven models to predict daily risk of Clostridioides difficile infection at 2 large academic health centers. *Infection Control & Hospital Epidemiology*, 44(7): 1163–1166. Publisher: Cambridge University Press.

Katell, M.; Young, M.; Dailey, D.; Herman, B.; Guetler, V.; Tam, A.; Bintz, C.; Raz, D.; and Krafft, P. 2020. Toward situated interventions for algorithmic equity: lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 45–55.

Kelling, S.; Gerbracht, J.; Fink, D.; Lagoze, C.; Wong, W.-K.; Yu, J. Y.; Damoulas, T.; and Gomes, C. 2012. ebird: A human/computer learning network for biodiversity conservation and research. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, 2229–2236. Issue: 2.

Klein, L.; and D'Ignazio, C. 2024. Data Feminism for AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, 100–112. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704505.

Kulynych, B.; Madras, D.; Milli, S.; Raji, I. D.; Zhou, A.; and Zemel, R. 2020. Participatory Approaches to Machine Learning. Published: International Conference on Machine Learning Workshop.

Kuo, T.-S.; Shen, H.; Geum, J.; Jones, N.; Hong, J. I.; Zhu, H.; and Holstein, K. 2023. Understanding frontline workers' and unhoused individuals' perspectives on ai used in homeless services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–17.

Langer, S. L.; Castro, F. G.; Chen, A. C.-C.; Davis, K. C.; Joseph, R. P.; Kim, W. S.; Larkey, L.; Lee, R. E.; Petrov, M. E.; Reifsnider, E.; Youngstedt, S. D.; and Shaibi, G. Q. 2021. Recruitment and Retention of Underrepresented and Vulnerable Populations to Research. *Public health nursing (Boston, Mass.)*, 38(6): 1102–1115.

Le Dantec, C. A.; and Fox, S. 2015. Strangers at the gate: Gaining access, building rapport, and co-constructing community-based research. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, 1348–1358.

Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; and others. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on human-computer interaction*, 3(CSCW): 1–35. Publisher: ACM New York, NY, USA.

Leslie, M.; Khayatzadeh-Mahani, A.; and MacKean, G. 2019. Recruitment of caregivers into health services research: lessons from a user-centred design study. *Research involvement and engagement*, 5: 1–9. Publisher: Springer.

Li, R.; Kingsley, S.; Fan, C.; Sinha, P.; Wai, N.; Lee, J.; Shen, H.; Eslami, M.; and Hong, J. 2023. Participation and Division of Labor in User-Driven Algorithm Audits: How Do Everyday Users Work together to Surface Algorithmic Harms? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19.

Liao, Q. V.; and Muller, M. 2019. Enabling value sensitive AI systems through participatory design fictions. *arXiv preprint arXiv:1912.07381*.

Lin, H.; Karusala, N.; Okolo, C. T.; D'Ignazio, C.; and Gajos, K. Z. 2024. "Come to us first": Centering Community Organizations in Artificial Intelligence for Social Good Partnerships. ArXiv:2409.06814 [cs].

Mancera, B. M.; Sy, A.; Williams, C. D.; and Hargreaves, M. K. 2022. Utilizing a social-ecological health promotion framework to engage diverse populations for recruitment in the All of Us research program. *Journal of Community Engagement and Scholarship*, 13(2).

Marian, A. 2023. Algorithmic transparency and accountability through crowdsourcing: A study of the nyc school admission lottery. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 434–443.

Miceli, M.; Yang, T.; Naudts, L.; Schuessler, M.; Serbanescu, D.; and Hanna, A. 2021. Documenting computer vision datasets: an invitation to reflexive data practices. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 161–172.

Mohamed, S.; Png, M.-T.; and Isaac, W. 2020. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33: 659–684. Publisher: Springer.

Nabatchi, T.; Gastil, J.; Weiksner, G. M.; and Leighninger, M. 2012. *Democracy in motion: Evaluating the practice and impact of deliberative civic engagement*. Oxford University Press.

Negrin, K. A.; Slaughter, S. E.; Dahlke, S.; and Olson, J. 2022. Successful recruitment to qualitative research: A critical reflection. *International Journal of Qualitative Methods*, 21: 16094069221119576. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Nekoto, W.; Marivate, V.; Matsila, T.; Fasubaa, T.; Kolawole, T.; Fagbohungbe, T.; Akinola, S. O.; Muhammad, S. H.; Kabongo, S.; Osei, S.; and others. 2020. Participatory research for low-resourced machine translation: A case study in african languages. *arXiv preprint arXiv:2010.02353*.

Newman, N.; Immorlica, N.; Leyton-Brown, K.; Lucier, B.; Quinn, J.; and Ssekibuule, R. 2022. Kudu: An Electronic Agricultural Marketplace in Uganda.

Obermeyer, Z.; Powers, B.; Vogeli, C.; and Mullainathan, S. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464): 447–453. Publisher: American Association for the Advancement of Science.

O'Brien, J. E.; Brewer, K. B.; Jones, L. M.; Corkhum, J.; and Rizo, C. F. 2022. Rigor and respect: recruitment strategies for engaging vulnerable populations in research. *Journal of interpersonal violence*, 37(17-18): NP17052–NP17072. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Perrault, A.; Fang, F.; Sinha, A.; and Tambe, M. 2020. Artificial intelligence for social impact: Learning and planning in the data-to-deployment pipeline. *AI Magazine*, 41(4): 3–16.

Puussaar, A.; Montague, K.; Peacock, S.; Nappey, T.; Anderson, R.; Jonczyk, J.; Wright, P.; and James, P. 2022. SenseMyStreet: Sensor Commissioning Toolkit for Communities. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–26. Publisher: ACM New York, NY, USA.

Queerinai, O. O.; Ovalle, A.; Subramonian, A.; Singh, A.; Voelcker, C.; Sutherland, D. J.; Locatelli, D.; Breznik, E.; Klubicka, F.; Yuan, H.; J, H.; Zhang, H.; Shriram, J.; Lehman, K.; Soldaini, L.; Sap, M.; Deisenroth, M. P.; Pacheco, M. L.; Ryskina, M.; Mundt, M.; Agarwal, M.; Mclean, N.; Xu, P.; Pranav, A.; Korpan, R.; Ray, R.; Mathew, S.; Arora, S.; John, S.; Anand, T.; Agrawal, V.; Agnew, W.; Long, Y.; Wang, Z. J.; Talat, Z.; Ghosh, A.; Dennler, N.; Noseworthy, M.; Jha, S.; Baylor, E.; Joshi, A.; Bilenko, N. Y.; Mcnamara, A.; Gontijo-Lopes, R.; Markham, A.; Dong, E.; Kay, J.; Saraswat, M.; Vytla, N.; and Stark, L. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *2023 ACM Conference on Fairness, Accountability, and Transparency*, 1882–1895. Chicago IL USA: ACM. ISBN 9798400701924.

Roscoe, R. 2021. Designing for diversity: Inclusive Sampling. *Ergodesign & HCI*, 9.

Sendak, M.; Elish, M. C.; Gao, M.; Futoma, J.; Ratliff, W.; Nichols, M.; Bedoya, A.; Balu, S.; and O'Brien, C. 2020. " The human body is a black box" supporting clinical decision-making with deep learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 99–109.

Sharp, A.; and Anderson, K. 2010. Online citizen panels as an advance in research and consultation: A Review of pilot results. *Commonwealth Journal of Local Governance*, (6): 33–54.

Shen, H.; Wang, L.; Deng, W. H.; Brusse, C.; Velgersdijk, R.; and Zhu, H. 2022. The model card authoring toolkit: Toward community-centered, deliberation-driven AI design. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 440–451.

Shi, Z. R.; Yuan, Y.; Lo, K.; Lizarondo, L.; and Fang, F. 2020. Improving efficiency of volunteer-based food rescue operations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 13369–13375. Issue: 08.

Sinders, C. 2020. Feminist data set. *Clinic for Open Source Arts. https://carolinesinders. com/wp-content/uploads/2020/05/Feminist-Data-Set-Final-Draft-2020-0517. pdf.*

Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 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.

Small, M. L.; and Calarco, J. M. 2022. *Qualitative literacy: A guide to evaluating ethnographic and interview research*. Univ of California Press.

Spinuzzi, C. 2005. The Methodology of Participatory Design.

Sum, C. M.; Zhi, J.; Cook, A. N.; Cooper, P. J.; Lozano, A.; Johnson, T.; Perez, J.; Ghani, R.; Skirpan, M.; Eslami, M.; and others. 2025. " You're in a Ferrari. I'm Waiting for the Bus": Confronting Tensions in Community-University Partnerships. *Proceedings of the ACM on Human-Computer Interaction*, 9(2): 1–28. Publisher: ACM New York, NY, USA.

Suresh, H.; Movva, R.; Dogan, A. L.; Bhargava, R.; Cruxên, I.; Cuba, Á. M.; Taurino, G.; So, W.; and D'Ignazio, C. 2022. Towards intersectional feminist and participatory ML: A case study in supporting Feminicide Counterdata Collection. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 667–678.

Swanson, A.; Kosmala, M.; Lintott, C.; Simpson, R.; Smith, A.; and Packer, C. 2015. Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Scientific Data*, 2(1): 150026. Publisher: Nature Publishing Group.

Tambe, M.; Fang, F.; and Wilder, B., eds. 2022. *AI for Social Impact*.

Tandon, U.; Khovanskaya, V.; Arcilla, E.; Hussein, M. H.; Zschiesche, P.; and Irani, L. 2022. Hostile ecologies: Navigating the barriers to community-led innovation. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–26. Publisher: ACM New York, NY, USA.

Tuia, D.; Kellenberger, B.; Beery, S.; Costelloe, B. R.; Zuffi, S.; Risse, B.; Mathis, A.; Mathis, M. W.; Van Langevelde, F.; Burghardt, T.; and others. 2022. Perspectives in machine learning for wildlife conservation. *Nature communications*, 13(1): 792. Publisher: Nature Publishing Group UK London.

UNHCR Innovation Service. 2019. A Goat Story. https://medium.com/unhcr-innovation-service/a-goat-story-3ed6bdd2b237.

University of Michigan STPP program and Detroit Disability Power and the Detroit Justice Center;. 2024. Community Partnerships Playbook.

Vines, J.; Clarke, R.; Wright, P.; McCarthy, J.; and Olivier, P. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 429–438. Paris France: ACM. ISBN 978-1-4503-1899-0.

Wilder, B.; Onasch-Vera, L.; Diguiseppi, G.; Petering, R.; Hill, C.; Yadav, A.; Rice, E.; and Tambe, M. 2021. Clinical trial of an AI-augmented intervention for HIV prevention in youth experiencing homelessness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 14948–14956. Issue: 17.

Wood, C.; Sullivan, B.; Iliff, M.; Fink, D.; and Kelling, S. 2011. eBird: engaging birders in science and conservation. *PLoS biology*, 9(12): e1001220. Publisher: Public Library of Science San Francisco, USA.

Young, M.; Magassa, L.; and Friedman, B. 2019. Toward inclusive tech policy design: a method for underrepresented voices to strengthen tech policy documents. *Ethics and Information Technology*, 21: 89–103. Publisher: Springer.

Yu, B.; Yuan, Y.; Terveen, L.; Wu, Z. S.; Forlizzi, J.; and Zhu, H. 2020. Keeping Designers in the Loop: Communicating Inherent Algorithmic Trade-offs Across Multiple Objectives. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, 1245–1257. Eindhoven Netherlands: ACM. ISBN 978-1-4503-6974-9.

Zhu, H.; Yu, B.; Halfaker, A.; and Terveen, L. 2018. Value-Sensitive Algorithm Design: Method, Case Study, and Lessons. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW): 1–23.