Secondary Stakeholders in AI: Fighting for, Brokering, and Navigating Agency

Leah Hope Ajmani
University of Minnesota
Minneapolis, Minnesota, USA
ajman004@umn.edu

Nuredin Ali Abdelkadir
University of Minnesota
Minneapolis, Minnesota, USA
ali00530@umn.edu

Stevie Chancellor
University of Minnesota
Minneapolis, Minnesota, USA
steviec@umn.edu

Abstract

As AI technologies become more human-facing, there have been numerous calls to adapt participatory approaches to AI development—spurring the idea of participatory AI. However, these calls often focus only on primary stakeholders, such as end-users, and not secondary stakeholders. This paper seeks to translate the ideals of participatory AI to a broader population of secondary AI stakeholders through semi-structured interviews. We theorize that meaningful participation involves three participatory ideals: (1) informedness, (2) consent, and (3) agency. We also explore how secondary stakeholders realize these ideals by traversing a complicated problem space. Like walking up the rungs of a ladder, these ideals build on one another. We introduce three stakeholder archetypes: the reluctant data contributor, the unsupported activist, and the well-intentioned practitioner, who must navigate systemic barriers to achieving agentic AI relationships. We envision an AI future where secondary stakeholders are able to meaningfully participate with the AI systems they influence and are influenced by.

CCS Concepts

• **Human-centered computing → Human computer interaction (HCI)**.

Keywords

artificial intelligence, participatory ai, stakeholders, agency

ACM Reference Format:
Leah Hope Ajmani, Nuredin Ali Abdelkadir, and Stevie Chancellor. 2025. Secondary Stakeholders in AI: Fighting for, Brokering, and Navigating Agency. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25), June 23–26, 2025, Athens, Greece.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3715275.3732071

1 Introduction

The recent proliferation of Artificial Intelligence (AI) has spurred calls to "turn" towards participatory approaches in AI development [8, 17, 22] called participatory AI (PAI). As its name suggests, PAI is the application of participatory design methods [42] to AI technology development, and is rooted in ethical, just, and community-oriented work [9, 35, 66].

A core question to PAI is about approach: *what does meaningful AI participation look like?* This question is especially important to answer as new AI advancements transform our basic notions of ethics [73], justice [1], and human-centeredness [13]. Prior PAI work relies on articulating current participatory practices [22] and their ability to moderate power imbalances [18]. In answering these questions, we can also understand the challenges and problems associated with current PAI – such as superficial participant engagement (i.e., participation-washing [59]). These works are crucial to understanding and evolving practices, especially as AI technologies are rapidly developing.

However, participation is about more than just approach. Participatory AI also raises person-centric questions: *who are we considering when engaging in participatory AI practices?* In traditional participatory design, there is often an end-user group or affected community in mind [57, 62, 64]. The subjects in these participatory approaches are often **primary stakeholders** – those who are clearly and directly affected by the system being developed. For example, participatory AI work about child welfare algorithms often engages with social workers, affected parents, and government officials (i.e., end users, affected communities, and regulators) [11, 58, 62]. This work has proven overwhelmingly successful at improving specific algorithms for specific stakeholder groups.

As AI fundamentally transforms who is allowed to be a stakeholder, there are unique concerns over participating with secondary stakeholders. **Secondary stakeholders** are those who are influenced by and have an influence on AI systems, but they have indirect relationships with these systems. Formally defined in organizational sciences [16], we articulate secondary stakeholders as those who (1) influence AI systems broadly but (2) lack a direct contract with an AI system. Previous work has theorized that, due to this indirectness, secondary stakeholders are difficult to reason about in the context of PAI [13, 38, 59]. However, if left unchecked, our current AI practices risk fueling the frustrations of secondary stakeholder groups.

Hollywood writers went on strike due to AI threats to their jobs [72]. Many people contribute to AI systems simply by being online [38]. Reddit moderators recently turned over 8,000 subreddits private in protest of Reddit Inc.'s new data relationships with AI companies [50]. In all of these incidents, the main stakeholder group is secondary; they have a murky and indirect relationship with the AI system. Yet, their frustrations are real and urgent. Our work relies on the fundamental premise: **an ethical AI future is one where everyone gets their due data rights.**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
FAccT '25, Athens, Greece
© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1482-5/25/06
https://doi.org/10.1145/3715275.3732071

To achieve this future, we argue for PAI work to consider secondary stakeholders with the same ideals we apply to primary stakeholders. As AI systems grow in scale, the number of people who are secondary stakeholders grows accordingly. The term stakeholder now includes the many examples mentioned above – from affected communities (in the context of welfare algorithms) to Hollywood writers, Reddit users, and moderators across the globe. Given the current frustrations that cut across these stakeholder groups and evidence that they are often exploited in AI contexts [38, 59], secondary stakeholders warrant increased attention from PAI practitioners. We need to actively translate PAI ideals from primary stakeholders to secondary ones. The first step here is to understand current barriers to achieving this translational effort.

In this paper, we empirically investigate the unique challenges of being a secondary stakeholder in an AI system. We start from theoretical foundations previously established at FAccT: due data rights in AI contexts require agency over one's data [1]. We expand this premise beyond data agency and describe meaningful AI participation as a variation of Arnstein's Ladder of Participation from civic engagement [2]. Meaningful participation is the step-by-step process of creating (1) informed, (2) consensual, and (3) agentic relationships with AI systems. With the ladder as our guiding lens, we present the results from semi-structured interviews with 12 secondary AI stakeholders about their experiences and frustrations with AI. We combine our interview findings into three stakeholder archetypes: the reluctant data contributor, the unsupported activist, and the well-intentioned practitioner. Each archetype has its own role to play in realizing a participatory future. For example, data contributors are often victims of non-consensual AI practices. Meanwhile, practitioners take on a brokerage role of trying to create meaningful participation. In sum, we contribute a general theoretical framework of meaningful participation, empirical evidence of how secondary stakeholders navigate this framework, and a call to action for deeper interrogations of AI stakeholder needs.

2 Related Work

2.1 Stakeholdership in AI

In 1984, Freeman introduced the term *stakeholder* in organizational sciences as a new way of framing ideal corporate strategy for firms [26]. Stakeholder theory asserts that the firm is a network of value-creation obligations for those who can affect or be affected by the realization of an organization's purpose (the wide definition) or those without whose support the organization would not exist (the narrow definition) [23]. Thinking about stakeholdership has spurred numerous taxonomies and frameworks to consider both inclusion and prioritization. For example, a common taxonomy considers the stakeholder's degree of closeness to the firm's impact; primary stakeholders are directly affected, while secondary stakeholders have less clear relationships with the firm [16]. In participatory design, the practitioner is tasked with identifying a group of "key" stakeholders whose concerns should be prioritized during the design process [27]. Recent HCI work has focused on using stakeholder prioritization as a mechanism to empower marginalized communities. For example, critical race theory for HCI argues that vulnerable or systemically marginalized groups must be considered as key stakeholders [48]. Though the concept originated in organizational sciences, stakeholders are a crucial concept in design.

In the development of AI/ML tools and systems, stakeholder theory has been heavily adapted to inform who should be considered during the design process. The idea here is that the AI system is the firm and has a set of value-creation obligations to others. Similar to organizational research, the term stakeholder carries moral weight; when we fail to include stakeholders in the design process, we are skirting a moral obligation to do right by them. This justice-oriented approach to stakeholder inclusion is common in many human-centered machine learning approaches [13, 44, 75]. In fact, many works focus on the material harm caused by flawed stakeholder reasoning in AI and ML systems. For example, Stapleton et al. [62] found that including parents in predictive child-welfare algorithm design fosters a richer understanding of the algorithm's potential shortfalls. Chancellor et al. [14] found that predictive mental health researchers have not coalesced on a conceptualization of stakeholders. This discord can cause dramatic differences in data handling procedures and limit human agency in predictive systems.

An additional challenge in adapting stakeholder theory to AI/ML systems is the question of scale. In design, stakeholder inclusion is often proposed for small groups or community representatives. For example, Value-Sensitive Design (VSD) articulates indirect stakeholders as the large population of people who can be affected by a system without interacting with it [27]. Recent critiques highlight how this breadth of stakeholdership may neglect considerations of marginalization [21, 48]. Additionally, transitioning from individual needs to large stakeholder groups complicates engagement [20, 65]. When current AI systems are public-facing, is everyone a stakeholder? Engaging with AI stakeholders demands that we simultaneously consider contextual stakeholder positionalities while being able to abstract into shared ethical norms. One possible solution is to encourage AI practitioners to engage with more rigorous stakeholder analyses. A large thread of AI ethics work involves identifying relevant stakeholders and stakeholder needs for specific problems such as interpretability [63] and explainability [52]. Inspired by this work, our paper describes the specific ways that certain stakeholder groups realize participatory AI ideals.

2.2 Participation in AI

Distinct from "AI Stakeholders," a similar thread of work focuses on "AI Participation." While *stakeholder* originates from organizational sciences, *participant* originates from design [42]. Here, participants are the folks who we engage with during participatory design practices [34, 57]. Given these origins, work on participatory AI—a genus of participatory design focused on AI tools—often focuses on the *how* moreso than the *who*.

In AI ethics research, there is an overwhelmingly positive vision of *meaningfully participating with stakeholders*. Defining meaningful here has been a difficult research problem, dating back to civic engagement [2, 29]. In AI contexts, these frameworks inform how we give technological ownership to—in other words, participate with—a specific stakeholder group related to the project (e.g., end-users [7], community organizers [37, 64], domain experts [32, 36, 60]). Many of these creative approaches are successful at engaging with their intended participants.

In addition to the successes of meaningful participation, recent work has explored the harms of superficial participatory AI practices. Often dubbed as "participation washing," there is a class of stakeholder engagement approaches that do not constitute meaningful participation [59]. For example, Corbett et al. [18]'s review of PAI projects found that numerous papers were attempting to "fix" participant opinions rather than engage with them. Delgado et al. [22] describe a similar spectrum of consulting-based participation to ownership-based. To these forms of non-meaningful participation, recent work has turned its gaze toward practitioners themselves. Unpacking meaningful participation involves reflecting on the contextual power dynamics, such as the relationship between a researcher and their participants [9, 25, 71]. Indeed, there is a current seed in participatory AI that is starting to consider the stakeholders that have influence beyond the primary group of participants.

In this paper, we build on this seed by focusing on stakeholders who are not obvious participants, such as data contributors, activists, and practitioners. Inspired by work that calls for us to do right by these stakeholder groups [38, 43, 56, 59], despite their indirect relationships to AI, we consider the ideals of meaningful participation for secondary stakeholders. To do so, we must acknowledge that broadening who we consider complicates the participatory ideals established in prior work. What do ownership and agency look like when we start to scale stakeholder considerations?

3 Methods

Our research focuses on secondary stakeholders' experiences with AI participation. Therefore, we conducted semi-structured interviews with 12 individuals who are researchers, developers, activists, or data workers. We chose to interview people in these roles as previous work has heavily theorized how they influence AI systems [38, 51, 64]. We used an interpretivist methodology [24] to capture our participant's understandings of their lived experiences and frustrations. Therefore, we (1) ran semi-structured interviews with open-ended questions and member checks, (2) analyzed interviews with inductive open coding [41], and (3) communicated with participants as needed.

3.1 Participant Recruitment

We circulated our call for participation in relevant venues such as researcher Slack groups, advocacy mailing lists, and moderator forums. Interested individuals completed an intake survey to check they met our eligibility requirements: (1) were 18 years or older and (2) had participated as an AI researcher, developer, advocate, or data contributor over the past 24 months. We selected participants using purposeful stratified sampling techniques [49, 54]. Because we are interested in comparing and contrasting participation groups (researchers, developers, activists, and data workers), we wanted to ensure adequate representation and expertise in all groups. Our participants include international and US-based interviewees. Our study was approved by the university Institutional Review Board (IRB) #00020382. Given some of our participants are members of adjacent research communities and could easily be identified, we present our participant demographics in aggregate (Table 1). Table 2 details the participants' pseudonyms and their role descriptions. We chose to use pseudonyms instead of numeric identifiers to communicate the personal nature of our participants' stories. Company and institution names have been redacted unless the participant explicitly requested otherwise. Note that many participants embody more than one secondary stakeholder role. For example, Katie became an activist after her frustrating tenure as a data worker.

3.2 Interview Procedure

Our interview questions were designed to understand (1) the participant's experience with AI, (2) their frustrations with their current role in AI practices, and (2) their personal vision for a future of AI participation. Participants completed a written consent form before scheduling their interview. At the beginning of each interview, the interviewer verbally confirmed participant consent and reminded the interviewee they can decline any questions or redact any information after the interview. As the researchers may have personal beliefs and biases on AI, we followed Doyle [24]'s work on consistent member checking to make sure the participants' responses are interpreted correctly. For instance, by asking questions such as, *"So I'm hearing you say XYZ. Is that a fair interpretation?"* and encouraging participants to correct the interviewer's interpretations. We also validated our interpretations by communicating with participants after interviews as needed. The interviews were conducted virtually through Zoom and lasted approximately 1 hour. All participants were compensated $25 USD in gift cards for their participation.

We adopted an inductive approach [68] to our analysis inspired by interpretivist hermeneutic methodologies [24]—techniques to understand how people negotiate the meaning of a specific term (i.e., AI participant). Given how interview participants' relationships with AI are central to our research question, we employed a stratified version of open coding. We began with iterative open coding of each individual interview. Two researchers were involved in the analysis process. Next, the two coding authors transferred all of the codes into a virtual whiteboard coloring each code based on the participant's role (i.e. researcher, data worker, etc). This coloring allowed us to compare and contrast themes across groups. Finally, all authors contributed to the synthesis of major themes presented in this paper. Since synthesis in thematic analysis runs into the risk of diluting participant experiences into the author's narrative [10], we engaged in peer debriefing [39] among authors to shed light on our positionalities. We reached saturation after an initial round of interviews with three participants from each positionality (n=12). This sample size meets the requirements of interpretivist and constructivist methodologies [19, 24, 49, 61], particularly given our focus on individual positions about AI participation.

4 The Meaningful Participation Ladder

We start by articulating a general framework of meaningful participation that applies to both primary and secondary stakeholders. One area that has deeply considered "meaningful" participation is civic engagement. Specifically, Sherry Arnstein [2] proposed the ladder of citizen participation, which models the degree to which civic processes redistribute power to citizens. Recent work in AI ethics has focused on translating this idea of the ladder of

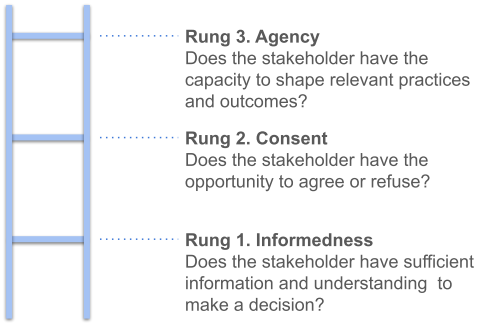


Figure 1: Overview of how we envision a ladder of meaningful participation. Each rung of the ladder is a participatory ideal that we argue must be realized for both primary and secondary AI stakeholders.

participation levels to build normative frameworks for participatory AI approaches [18, 22, 65]. We seek to marry this work with current AI ethics calls to create a general framework for considering meaningful participation. Given that our normative goal is to broadly do right by people, we translate Ajmani et al. [1]'s recent FAccT work on justice in AI to the ladder of participation. Specifically, we start from their position that *"agency is a contributor to justice and a product of consent policies."* In this section, we synthesize the work of Arnstein [2], Ajmani et al. [1], and their peers [8, 13, 17] to describe meaningful AI participation. Meaningful participation is the step-by-step process of building (1) informed, (2) consensual, and (3) agentic relationships with AI, where you cannot reach a higher rung without establishing the lower rungs first (Figure 1). In later sections, we use this ladder to describe each secondary stakeholders' role in AI systems and their complications in realizing these three participatory ideals.

Rung 1. Forms of Informedness: We define an informed stakeholder as an individual who has sufficient information and understanding at hand to make a decision about their participation [31, 47, 55]. Informedness of both primary and secondary stakeholders is a perennial challenge in AI research and development due to the scale of training data needed [15], the technical nature of AI development [74], and the fear of new technology [30]. Expertise gaps exacerbate peoples' ability to be informed – it is easier to make informed judgments about AI participation when you have relevant technical knowledge, which most people do not have. In settings with clear primary stakeholders, such as PAI research projects with human subjects, previous work has called for minimizing these expertise gaps between researchers and participants [8, 9, 74]. This work underscores how AI participation relies heavily on the informedness of stakeholders. Often, practitioners are responsible for informing the stakeholders they directly interact with [17]. However, when considering secondary stakeholders, there is minimal opportunity for interaction. In later sections, we describe how, though informedness is the most basic level of cultivating meaningful participation, it is extremely difficult to realize at the scale needed for secondary stakeholders.

Rung 2. Consent Builds on Informedness: Consent is about giving stakeholders the opportunity to agree to or refuse participation, free from coercive forces [28]. Consent scholars generally agree that uninformed consent is not true consent [5]; informedness and consent go hand-in-hand. For example, recent work evaluating common data consent procedures, like Terms of Service agreements, demonstrated that such broad agreements cannot feasibly create a consentful environment [1]. Often, these are our only mechanisms for soliciting consent from secondary stakeholder groups, such as data contributors. While there are repeated calls for legible [40], informed [70], and affirmative [33] consent procedures online, they have not come to fruition. Building on previous work on ethical tensions in AI/ML [15], consent in AI contexts must happen at a massive scale. Researchers and developers have techniques to solicit consent from primary stakeholders in PAI research projects [47], but it is unclear if these can properly scale. While consent for secondary stakeholders is a necessary step for meaningful participation, it is not an easy task. In this paper, we describe the tensions of getting consent-at-scale from well-intentioned AI practitioners and the frustrations of non-consensually engaging with AI.

Rung 3. Agency Beyond Consent: Even in ideal worlds with informed and consenting stakeholders, meaningful participation involves agency. In PAI, agency is about giving participants the capacity the shape the outcomes of their participation, such as resulting datasets [35], tools [60], and algorithmic recommendations [62]. Moreover, agency captures one's capacity to shape their own participation; agency involves allowing participants to voice their opinions freely. For an example of anti-agentic but informed and consensual PAI practices, we can look to Corbett et al. [18]'s examples of PAI projects where researchers tried to "fix" participant opinions rather than engage with them. On a broader scale that considers secondary stakeholders, Ajmani et al. [1] articulate data agency as one's capacity to shape action around their data, concluding that non-consensual data practices are anti-agentic and, more seriously, unjust.

In sum, meaningful participation involves designing for three participation ideals: informedness, consent, and agency. However, these are big and broad asks when engaging with secondary participants. Standard PAI techniques of achieving informedness and consent with primary stakeholders start to break down. Lending agency to AI stakeholders must happen in a scalable way. In the rest of this paper, we use this ladder as a framework to better understand the empirical experiences secondary stakeholders have with AI.

5 Secondary Stakeholder Archetypes

In this section, we present three archetypes of secondary AI stakeholders and how they traverse the meaningful participation ladder. The term "stakeholder" is often broadly and ambiguously defined in HCAI [14, 22]. We present the experiences of a specific class of "murky" AI stakeholders: those with complicated relationships to AI systems they affect. We find this aligns with the concept of secondary stakeholders from organizational sciences [16]: those (1) influenced by a firm's decisions but (2) outside formal contracts. However, we find that this definition is complicated by the ambiguity of AI. For instance, one participant, Fiona, had a contract as a paid moderator, but did not necessarily have a clear contract

with any AI system, as she did not realize her work was training one. We use *secondary stakeholder* to describe individuals who (1) influence AI systems broadly but (2) lack a direct contract with an AI system. While this working definition emerged from our analysis, participants occupied nuanced positions, elaborated below.

To start, we introduce the forms of AI participation our interviewees described. AI participation is split between human-centered activities that are grounded in traditional research methods [35, 62, 74] and data-centered activities. We find evidence that secondary stakeholders are delegitimized in their participation with AI systems.

To relegitimize their participation in AI systems, we frame stakeholders' contributions into three stakeholder archetypes: the reluctant data contributor, the unsupported activist, and the well-intentioned practitioner. For example, data contributors are often victims of non-consensual AI practices. Meanwhile, practitioners take on a brokerage role of trying to create meaningful participation. By conceptualizing each archetype along our ladder of meaningful participation, we identify where current gaps are in PAI practices. Each archetype has its own role to play and actions they take in realizing a participatory AI future.

5.1 Forms of Participation

Human-centered participation in AI research and development is a well-defined set of methods where "AI participants" are the people actively engaging in an activity (e.g., workshop, co-design, prototype evaluation [57]). These participatory efforts often engage with primary stakeholders, such as end-users or those directly affected by the AI system [36, 60, 62]. Our interviewees noted participating in AI-related development efforts, such as design workshops or research studies, as a clear way of defining "AI participants." Maddie, a non-profit industry researcher, initially described participation as research-based: *"When I hear the term participant, I'm immediately thinking about participants in our studies, the way that I'm a participant in your study."*

As participatory methods have evolved to be more collaborative between researchers and participants, our interviewees disagreed on whether practitioners—individuals who facilitate participatory AI workshops and activities—were participants themselves. This murky position of practitioners highlights their role as secondary stakeholders. For example, Sean is a researcher who uses participatory AI methods to build agriculture technology. He self-identified as an AI participant insofar as he participates with farmers (i.e., his research participants) to build a usable tool. In that way, his ideas and values bear on the final product. In contrast, Tyler is a researcher using participatory AI methods to build tools for online communities, such as Wikipedia. He noted that researchers are not inherently participants in the AI systems they are building. However, depending on the context, researchers can become members of the communities they study. For example, Tyler has become an avid Wikipedia editor on the side of his Wikipedia-based research. We find that human-centered participation has a clear set of primary stakeholders – the people who are engaging in the participatory design activity. Meanwhile, practitioners operate as secondary stakeholders; Tyler uses PAI methods as a researcher, but does not have a formal contract with any specific AI system. We find that this distance influences his perception of his own role as a participant, raising an open question of how researchers relate to "participatory AI" broadly.

Data-centered participation deviates from the active type of participation found in participatory design methods [22]. Data-centered participation relies heavily on secondary stakeholder groups, such as data contributors and AI system developers. Yet, our interviewees cited data contribution and moderation as forms of AI participation. Building on previous literature that highlights how data work can be participation [38, 59], we find that individuals can participate in AI by contributing work that helps build AI training datasets.

Data-centered participation can include almost everyone who uses the internet. Our interviewees reflected on how they passively participate through data in their everyday activities. Scott develops AI tools for other developers, and he noted: *"I see participating in AI as contributing to the tech or contributing to the algorithms. For products like [AI tool], I do both."* Scott described building AI tools as "active" while contributing his code as training data is "passive," because he simply makes his code publicly available online. Many of our interviewees described similar situations; they produce data used for AI training data simply by engaging in adjacent activities, such as development, writing, and publishing online content. In terms of stakeholdership, this passive data contribution is secondary, as it exists outside of any formal contract. Yet, our interviewees described their own frustrations with feeling like everything they publicly create is unintentional training data. We elaborate on these frustrations in later sections.

Meanwhile, many of our interviewees similarly used the term "active" to describe their roles as intentional data workers. Data worker is the umbrella term used to describe participants who actively contribute, annotate, or moderate data that is known to be used for AI model training [38]. This includes volunteer work, such as contributing to Wikipedia, and paid positions, such as being an expert annotator[1] For example, employers often own the outcomes of their employees' work. As a positive example, Falen works as an expert data annotator for a mental-health-based AI company and self-identified as an AI participant. She was hired to train a model with her clinical mental health expertise. Falen positively noted: *"I'm helping train the machines so that they can help therapists do really important work."* However, Falen represents a rare case. Overwhelmingly, we find that data workers have little information, opportunities to consent, and agency.

In the following sections, we highlight how the scale of data-centered participation complicates the stakeholder considerations of human-centered participation. From data workers feeling exploited to practitioners feeling like their efforts are hollow, realizing participatory ideals is extremely frustrating for everyone.

5.2 The Reluctant Data Contributor: Navigating Informedness, Consent & Agency

In this section, we introduce **the reluctant data contributor**, who—if given sufficient information and the opportunity to consent to their current AI participation—would actively choose to

[1]We note that these roles (as data participants in an employment relationship vs. volunteering) exist in different legal contexts.

withdraw their data from AI training datasets. This archetype contrasts Falen, who is participating in AI willingly and positively. It is important to note that none of our participants described reluctant participation for human-centered efforts (like deliberate inclusion in research). Our interviewees felt reluctant about their data and data work being used to create training datasets for AI models.

5.2.1 Informedness. Our interviewees described how many contributors in data-centered participation are uninformed and, subsequently, would not agree to their data being used for AI training if they were informed. The most straightforward example of an informed but reluctant data contributor came from Sabrina, who described her complicated situation as a volunteer moderator. She moderates a platform that is a clear data source for LLMs. Given her other position as a researcher, Sabrina understood her data work was being used to train LLMs but was powerless to stop it. Sabrina noted her own unease in being an AI participant, using the term reluctance: *"It's something that I know, but I'm a reluctant participant in this."*

Sabrina's experience highlights how volunteer data workers, who are often crucial moderators for their online communities, are aware of how their data is used but are unwilling. At the same time, Sabrina described how many of her moderator colleagues are uninformed. *"I understand the basics whereas it's not something that [other moderators] are thinking about because it's just like if you know, you know."* Sabrina speaks to an important point – though many of our interviewees were informed, given their expertise, many other data workers are often uninformed.

Further complicating informedness, data-centered participation can be passive. Therefore, a myriad of internet users are uninformed data contributors. For example, Maddie described her realization that all of her publicly available content, such as her research blogs, is now training data:

> *"Most of that data is made by unknowing, perhaps unwilling participants. Who wrote the research blogs? Who wrote the books? Who wrote the Reddit posts? The Wikipedia articles? Their data is now training AI."*

Maddie empathized with this feeling of unwillingness embraced by Sabrina. Maddie continued to describe that, even though she is aware her data is probably being used for AI training, she had no say in the matter.

5.2.2 Consent. Both Maddie and Sabrina represent informed but non-consenting participants. While informedness describes a participant's level of understanding, consent describes whether they have actively agreed to their participation. In data-centered participation, our interviewees described a significant lack of opportunity for them to consent without leaving the platform altogether. Sabrina explicitly called out that her continued data work should not be seen as consent.

> *"It's like you can opt out of the community, or you can be relinquished to having your data used. Those are the only two options. There's no way to give back to your individual community, which means a lot to you."*

In this case, Sabrina's only recourse is to withdraw her data work entirely rather than simply withdraw from downstream AI uses. In other words, her engagement with her online community carries AI participation with it.

Our interviewees described how, even when they have consented, current data practices often take this agreement out of scope. For example, Fiona worked as a commercial content moderator and, therefore, had a contract where she consented to her data contribution being used by a company. However, she articulated that her consent was taken out of scope when she was forced to moderate gruesome content,

> *"When they did the screening and the interviews, they just told us that the job might be slightly challenging. They even showed us samples...but the content that we ended up moderating was directly from the war in [home country]. Everything was gruesome videos. It's the worst thing that one can imagine."*

Fiona felt informed at the time of consent but was then given no recourse to withdraw her consent in certain contexts. Contrast this to Falen's willing and consentful participation. Falen described how, given she annotates severe mental health content, her company offers the ability to refuse certain data annotation tasks: *"We can say this session is too hard to code...And they give us another one that's not as emotionally taxing."* Falen's experience offers a positive future where data contributors have multiple opportunities to agree to AI-training tasks and refuse tasks where contextually appropriate—turning *reluctant data contribution* into *willing data contribution.*

5.2.3 Agency. We conceptualize the reluctant data contributor as someone who has little to no agency over their data being currently used to train AI systems. Often, the only recourse these contributors have is dramatic non-use. This lack of agency is represented in Sabrina's feeling of voicelessness: *"You almost don't want to admit that your work is being exploited, but you don't have any sort of voice."*

Katie, a former data annotator on mTurk, described feeling like she had agency only after she reckoned with the task of annotating a politically charged dataset:

> *"I wish when I started this work that I had been more aware. I didn't realize at the time that what I was doing could be used for something that I personally didn't agree with. It was honestly a task tagging border crossing photos that stopped me, and I thought: wait a minute."*

In this situation, Katie did have the agency to choose whether to complete this annotation task or not. However, her lack of informedness clouded this choice, highlighting how agency builds on informedness. Since then, Katie has become a data rights activist to advocate for better practices around soliciting data annotation. Note that if the reluctant data contributor did have sufficient agency, much like Katie chose to stop her data annotation task, they would choose not to contribute their data.

5.3 The AI Activist: Fighting for Informedness, Consent & Agency

We interviewed multiple people with similar stories to Katie: they became activists after their own frustrating experiences of being AI-oriented workers. In this section, we introduce the archetype

of the **AI activist.** We describe this stakeholder as someone who fights for informedness, consent, and agency on behalf of both primary and secondary stakeholders. We find that activists fight for more meaningful participation by advocating for better general AI initiatives, such as AI literacy, and better specific practices from platforms, developers, and researchers.

5.3.1 Fighting for Informedness. We find that activists are currently fighting for informedness on two fronts: (1) public informedness on AI and (2) practitioner informedness on current practices. In terms of fighting for public informedness, Kayla described her advocacy work on AI literacy. She described AI literacy as a balancing act between informing the public, but also reducing the harm caused to those who do not understand the technical details of AI,

> *"I'm incredibly in favor of AI literacy... but with what other product do we ask people to have to be that literate? We don't expect ourselves to be fully literate about our kitchen appliances. We expect that regulators, government officials, and companies, are responsible for making sure the products we use are safe and fair and and ethical to some degree. To some extent, does this push towards AI literacy just displace that responsibility onto the individual."*

Kayla's fight is to consider *how* to inform participants and *who* is responsible when participants are left in the dark.

At the same time, Katie described much of her activism as fighting for more transparency about current practices. For example, she described the invisibility of data work in current AI practices,

> *"The tech companies not wanting to acknowledge that there is this workforce is very frustrating. Data workers put in essential work. These datasets have been built on the backs of data workers and to keep them invisible is disrespectful.*

In this case, informedness and transparency go hand-in-hand. Katie described her ideal world as one where data workers are informed about the tasks they are completing, and practitioners are informed about the labor behind data sources.

It is important to note that, especially in Katie's case, fighting for informedness requires multi-stakeholder support from researchers and developers. For example, both Katie and Fiona described their positive experiences of going to research conferences to tell their stories.

5.3.2 Fighting for Better Consent Procedures. Our interviewees described fighting for better consent procedures as a process of navigating the real-world limitations of current consent practices. For example, Katie described how the current consent mechanism on mTurk is impractical for data workers where "time is money:"

> *"For some tasks, the consent documents are very long. And I know that most workers skip them because time is money. They just hit the okay button and go."*

Therefore, Katie is fighting for concise consent documents that contain the necessary information about the task, including the intended data use. Kayla similarly described thinking about consent as context-specific rather than as a one-size-fits-all best model. Conceptually, Kayla described her fight as moving away from any singular consent model:

> *"Rather than thinking about a single consent model, what happens if we think in terms of coercion or coerciveness? Thinking about where people are giving up power because they have no other choice"*

Through her advocacy work, Kayla builds tools and frameworks for evaluating points of coercion in current data practices. To do so, Kayla advocates for conceptualizing coercion as an opposing force to consent. Inspired by feminist work, she repeated the mantra, *"consent with coercion is not consent."*

Katie and Kayla highlight how fighting for better consent procedures happens on two fronts. First, the pragmatic fight of figuring out appropriate consent procedures for a specific AI context – such as mTurk workers. Second, the conceptual fight of building more robust ways to evaluate current consent practices. Addressing both of these concerns is a crucial step toward agency, as agency can be communicated and afforded through consent.

5.3.3 Fighting for Agency. A lack of agency in AI relationships is a major motivating factor for AI activists. Kayla described how considering AI participation without considering agency can lead to non-meaningful participation,

> *"I think sometimes participation gets confused with meaningful ownership or control over technologies so we can have participants without agency."*

In other words, Kayla notes that using the term "participation" without creating space for ownership creates a problematic class of AI participants who do not have any say in how their participation affects the system. In an ideal world, this is considered unethical. Therefore, Kayla viewed advocating for agency and advocating for ethical AI as one in the same.

More concretely, Fiona is advocating for agency for her specific community of refugee moderators who were tasked with moderating violent content. Given the extreme consequences of Fiona's participation, her activism work involves community building and collective action. For example, she described creating a support group for her fellow moderators, a reparative measure from the isolation she experienced as a moderator.

> *"We were never given a chance to talk about them to anyone...We're having nightmares. We're having insomnia. We're having suicidal thoughts."*

Fiona views fighting for agency as fighting for support of past moderators and regulation for future moderators. Given that agency is one's capacity to shape action, Fiona argues that support and community are two crucial factors in lending agency to large stakeholder groups, such as commercial moderators.

Our activists described their own struggles with agency and resistance-based efforts. Positively, Kayla described an inspirational history of resistance-based design as a mechanism for groups to reclaim agency.

> *"There's a really fascinating counter-history of design, which is how people express and create tools that work for them at times when like a lot of their choice or ability to live freely has been foreclosed."*

These acts of reclaiming agency, power, and design are effortful. Kayla's conceptualization of agency as being effortful to gain was mirrored in Katie's experience as a data worker. Katie described the transition from data worker to data rights activist as a necessary step she needed to take to gain agency: *"It does take a lot of time and energy to advocate for data rights. You're going up, and you're speaking out."* While resistance-based efforts are necessary to create change, to what end are we putting AI activists into unsupported spaces and expecting them to speak up while they are in already burdened positions? We find that activists need multi-stakeholder support from researchers and designers to make meaningful change.

5.4 The Well-Intentioned Practitioner: Brokering Informedness, Consent & Agency

For this archetype, we focus on **well-intentioned practitioners**, AI researchers and developers who are truly attempting to engage in best practices for participatory AI. Although some practitioners, such as Tyler, did not view themselves as participants, they heavily engaged and participated with communities and recognized their own influence on AI development - effectively working as brokers of informedness, consent, and agency to other stakeholders. However, these well-intentioned practitioners also recognized the limits of brokering these values because of scale, which limited an ideal participatory future with AI.

5.4.1 Brokering Informedness. Our interviewees described successfully brokering informed AI relationships with primary stakeholder groups such as research participants and end-users. Practitioners often brokered informedness at the time of consent as an essential part of participatory AI – a required practice in human subjects research. Tyler described informing research participants as the process of: *"Giving people enough information to articulate their ideas and make the right decision."* Here, informedness is a necessary pillar in one's decision-making for research participants. Maddie noted that informedness in research settings is often risk-centered:

> *"We try to make [consent documents] very simple, but not in the way of removing important details. We try and make it as intuitive and readable as possible. So people understand the risks of what they're doing."*

We find that practitioners who use participatory AI methods are trying to create legible documents to achieve participant informedness during consent. However, given AI technology is relatively new and complex, our participants told us their participants are nervous because of the unknown outcomes of AI systems. Therefore, AI researchers and developers have to build participant relationships and understanding to supplement informedness with trust to move forward.

For example, Sean—a researcher who uses participatory AI methods—struggled to even introduce the idea of AI as a solution to his research participants. When working on AI tools for farmers, Sean mentioned:"*Trying to bring new ideas to those farmers was a major challenge because they fear that it will disrupt their way.*" To help manage this unease, Sean brought in agricultural extension workers as mediators between himself (the researcher) and the farmers. This relationship-building and brokering is feasible in well-defined research settings and served Sean well as he could better engage with the community of research participants.

However, as AI participation has expanded beyond well-defined research settings, informedness has become a widespread issue. Maddie often uses public datasets for training models, and described how she tries to "inform" a different stakeholder – data contributors. To reach this group, she described her research on blogs and social media, something she felt obligated to do:

> *"I feel it's my duty as a researcher to engage at least a little bit, so I write Medium articles about things people should know about large language models as a user."*

However, Maddie recognized that these broader efforts can feel "hollow"; she explained that there is currently no way to guarantee that these messages are reaching the people whose data she is using. We find that practitioners are well-equipped and often willing to broker informedness in human-centered participation, where they can converse with their research participants, but are unable to realize this ideal in data-centered participation.

5.4.2 Brokering Consent. We find that practitioners have similar difficulty in brokering consent – current practices don't scale well beyond human-centered participation. Many practitioners described their practices of creating consent documents, reminding AI participants they can withdraw at any time, and even reaffirming consent after engaging with participants. These are best practices in research settings.

However, similar to brokering informedness, standard consent practices break down with data-centered participation. Charles, a developer who uses a variety of AI-related methods, noted the difficulty of getting consent in a space with no policy guidance:

> *"A lot of AI projects do not require consent because it's this new space. People still don't know how to navigate it or how to regulate it, so a lot of people are building models without consent."*

Maddie also described her personal struggles with navigating consent in non-human subjects research. While she believes in individual informed consent for data usage, she had trouble operationalizing it for her research datasets because of their size:

> *"You cannot get individual consent for every single person. I mean, you could if platforms were better. They could have each person opting in or out."*

Maddie continued to speculate about how platforms could get consent for all data-centered contributions, referencing previous work that articulates consent-at-scale as a difficult and unsovled problem [4].

Sabrina—in part due to her position as both a researcher and a moderator—believes this is a solved problem, and very clearly placed the onus on platforms and platform design. She described the problematic lack of incentive structures for platforms and companies to design more responsible consent mechanisms.

> *"We know that it's possible for people to consent and to have this. But [companies] are not doing it, and they're not going to. They don't have to."*

Sabrina is referencing the lack of laws around obtaining meaningful data consent. Sabrina brought up that there have been promising movements for better consent procedures, such as the AI Bill of

Rights [67], but there is little incentive for platforms to realize these calls into design. Moreover, there is no punishment for using problematic consent procedures, such as Terms of Service agreements.

5.4.3 Brokering Agency. Finally, we found practitioners heavily consider the agency of the AI participants they engage with. Our interviewees described agentic practices—practices that afford agency to participants—in their own disciplines. For example, when Sean was researching and developing AI tools for farmers, he consulted farmers at every major step of the process.

> *"We went to the field and collected the needs of the farmers before we eventually came back. We came up with our model and then gave it to the farmers to get feedback from their usage. Based on that feedback, we improved the model and eventually deployed it."*

Feedback solicitation and iteration are just one mechanism for giving participants agency in the AI development process. Anthony, an AI tool developer, described soliciting feedback on a larger scale with user surveys. However, agentic practices take time. Even in well-defined and well-resourced environments, such as academic research and AI development, our interviewees wanted increased support for brokering meaningful relationships. For example, Spencer, who develops AI tools for the public sector, described the time-consuming process of building rapport and, thereby, more agentic relationships with his participants:

> *"We are not delivering these public services, right? We are not on the ground. We don't understand every possible constraint. A lot of time goes into understanding, having conversations with people who know the right answers, and making sure that the assumptions we make are valid."*

While this step of building understanding is crucial to ethical participatory AI practices, it is time-consuming. Therefore, Spencer proposed that researchers and developers get more support by having longer project timelines and increased resources so they can do right by the people who would be affected by their system. We find that practitioners are engaging in creative practices—such as rich feedback solicitation—to broker agency, but doing so can run counter to the fast-pace of current AI development.

6 Discussion

6.1 Recommendations for Meaningful Participation

We return to the central question of PAI work posed earlier and one central to FAccT [17, 22, 65]: *how do we meaningfully participate with stakeholders?* Many of our interviewees described actionable ways to improve AI participation. For example, Maddie and Sabrina suggested default opt-in consent mechanisms. Spencer called for longer project timelines to incorporate participatory methods. In terms of the meaningful participation ladder, practitioners could broker agency with activists by inviting them to research venues, fostering a shared agenda. At the same time, we find that many of these recommendations rely on powerful stakeholders, such as technology companies, making them beyond the scope of any single archetype presented in our work.

Participating with secondary stakeholders has been heavily theorized in prior work [38, 59], but is challenging due to scale and ambiguity about who to include and how to include them. Specifically, we align with Sloane et al. [59]'s claim that all participation is work, and failure to change our PAI practice to include all stakeholders can lead to exploitative practices. In fact, to ignore AI labor is to ignore deep ethical and societal injustice [38]. Our empirical work and results show that the secondary stakeholders we interviewed are frustrated by current AI systems and their own anti-agentic relationships with them. Further, our participants highlighted that, without deep consideration of secondary stakeholders in PAI practice, we risk furthering exploitative practices of these groups rather than meaningful ones.

6.2 Using our Archetypes

In addition to providing empirical evidence of risks with PAI and secondary stakeholders [38, 59], we contribute several archetypes to the PAI space that focus attention and reasoning. We anticipate that many readers of this paper are one of our archetype—FAccT researchers may see themselves as well-intentioned practitioners. Like Maddie and Tyler, they are equipped to broker informedness, consent, and agency in human-centered participatory settings but feel limited in using the same strategies in data-centered settings, to the extent that Maddie's work feels futile. Using our archetypes may enable these practitioners to direct their efforts toward secondary stakeholders in a more nuanced and strategic way. For example, by applying the first rung of our ladder (i.e., informedness) to our most disempowered archetype (i.e., the reluctant data contributor), well-informed practitioners might consider how their brokerage could best inform data contributors to AI systems.

Additionally, future research could use our archetypes to identify contextual user needs, akin to design personas. These contextual descriptions are crucial when operationalizing ideals like informedness, consent, and agency. For example, Katie mentioned that being a data contributor meant often ignoring documentation, because *"time is money"*. Our archetype articulates a stakeholder need in achieving informedness—to be informed in a way that is not labor-intensive; research with data contributors thus should consider the value of conciseness alongside legibility [40], transparency [3], and literacy [6]. Similarly, our activist archetype requested support for multi-stakeholder action, highlighting a gap in current infrastructure to realize their positive visions. Like personas, our archetypes are both a descriptive synthesis of our participants' experiences and also can be generative for new design and technical PAI futures.

6.3 Recognizing Systemic Barriers

It is important to note that not all of the onus is on researchers and practitioners to improve AI relationships with secondary stakeholders. While practitioners should engage with the above recommendations, our findings highlight the limits of achieving participatory AI. Our interviewees echoed the current calls in AI ethics literature for increased AI regulation and literacy campaigns [45, 69]; platform redesign for meaningful data consent [1, 33]; and social responsibility of tech companies [12]. Practitioners hold privilege but face systemic barriers, such as tech companies failing to design better consent mechanisms. Our findings stress these real constraints.

Making participatory ideals a norm requires dismantling systemic barriers through grand interventions, such as AI regulation. Without this, we face an ethical glass ceiling—able to make incremental progress yet blocked from transformative change.

We find that, on one hand, each archetype has its own role to play in realizing a participatory future—even recognizing this as their "duty." On the other hand, navigating these roles is overwhelmingly frustrating, laborious, and littered with systemic barriers. The issue is not just that participatory ideals are lacking, but that they are structurally unachievable in the status quo.

7 Limitations

We used non-random stratified sampling to recruit twelve participants across our roles of interest (researchers, developers, activists, and data workers). This sample size is limited and not comprehensive. For example, we had no participants who identified as genderqueer or gender non-conforming, despite recent efforts to create gender-diverse participatory AI environments [53]. As with most interpretivist work, our methods are designed to create findings that prioritize transferability rather than generalizability [61]. We find common ideals across many of our interviewees despite their different roles, identities, and experiences, indicating that our results transfer to various contexts. Following design traditions, we treat archetypes like personas—composite representations that guide decision-making rather than exhaustive classifications [46]. AI encompasses far more than three archetypes, and articulating all possibilities would require a series of papers. Future work could explore the concept of AI Participation through more scalable techniques such as large-scale surveying, discourse analysis of online conversations, or quantitative methods of participant impact on datasets.

8 Conclusion

In sum, we empirically investigate the experience of being a secondary stakeholder in the current age of AI. We start from the position that participatory AI efforts have laid important groundwork for building more ethical AI systems. However, there is a current gap between applying PAI methods to primary stakeholders and applying PAI ideals to secondary stakeholders. We seek to close this gap by first establishing three general ideals of meaningful participation: informedness, consent, and agency. We use these ideals as a lens to analyze 12 semi-structured interviews with secondary AI stakeholders. We find that there are important archetypes to consider, each with its own orientation toward and struggle with meaningful participation. For example, AI practitioners are attempting to broker more meaningful AI relationships but struggle to do so at the scale required to inform data contributors. We discuss how practitioners can use our archetypes as a roadmap for furthering both the empirical and theoretical study of AI participation.

9 Statements

Ethical Considerations Statement. This paper is human-subjects work and was approved by the University IRB #00020382. In the spirit of meaningful participation, we heavily considered how to broker informedness, consent, and agency among all of our interviewees. In addition to written consent forms before the interview, we verbally explained the project in colloquial terms and reaffirmed consent. At both the beginning and end of each interview we reminded interviewees that they could withdraw their consent at any time or redact any information. Furthermore, we engaged with member checking strategies including communicating with participants after their interviews to confirm our own interpretations.

Positionality Statement. All of the authors of this work identify as well-intentioned practitioners. We all have experience with both human-subjects research and using data from public sites, such as social media. Therefore, it is important to note that the researchers of this work have reaped the benefits of the broad consent procedures and anti-agentic practices we critique in this paper. In keeping with member-research techniques, we believe our own struggles with realizing participatory ideals allowed us to appreciate our participants' experiences rather than overshadow them.

Adverse Impact Statement. Misinterpretations of our work could lead to a dilution of meaningful participation in AI systems. We are not suggesting that ameliorating the frustrations of data contributors, activists, and practitioners will solve all structural problems with AI systems. Therefore, adversarial actors could use our paper as a "moral magic" rather than a provocation for richer stakeholder consideration.

Acknowledgments

We are grateful to all of our participants for sharing their stories. Thank you to Nick Vincent and our reviewers for their feedback. This work was supported by the National Science Foundation (NSF) under Award #2332841.

References

[1] Leah Ajmani, Logan Stapleton, Mo Houtti, and Stevie Chancellor. 2024. Data agency theory: A precise theory of justice for AI applications. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. ACM, New York, NY, USA, 631–641.

[2] Sherry R Arnstein. 1969. A ladder of citizen participation. *J. Am. Inst. Plann.* 35, 4 (July 1969), 216–224.

[3] Solon Barocas and Helen Nissenbaum. 2009. On notice: The trouble with notice and consent. In *Proceedings of the engaging data forum: The first international forum on the application and management of personal electronic information*. papers.ssrn.com.

[4] Solon Barocas and H Nissenbaum. 2014. Big data's end run around anonymity and consent. *Book of Anonymity* (June 2014).

[5] Melanie A Beres. 2007. 'spontaneous' sexual consent: An analysis of sexual consent literature. *Fem. Psychol.* 17, 1 (Feb. 2007), 93–108.

[6] Maalvika Bhat and Duri Long. 2024. Designing interactive explainable AI tools for algorithmic literacy and transparency. In *Designing Interactive Systems Conference*. ACM, New York, NY, USA.

[7] Christian Bird, Denae Ford, Thomas Zimmermann, Nicole Forsgren, Eirini Kalliamvakou, Travis Lowdermilk, and Idan Gazit. 2022. Taking Flight with Copilot: Early insights and opportunities of AI-powered pair-programming tools. *ACM Queue* 20, 6 (Dec. 2022), 35–57.

[8] 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 *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, New York, NY, USA.

[9] Elizabeth Bondi, Lily Xu, Diana Acosta-Navas, and Jackson A Killian. 2021. Envisioning communities: A participatory approach towards AI for social good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, New York, NY, USA.

[10] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3, 2 (Jan. 2006), 77–101.

[11] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward algorithmic accountability

in public services: A qualitative study of affected community perspectives on algorithmic decision-making in child welfare services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA.

[12] Mark Anthony Camilleri. 2024. Artificial intelligence governance: Ethical considerations and implications for social responsibility. *Expert Syst.* 41, 7 (July 2024), e13406.

[13] Stevie Chancellor. 2023. Toward practices for human-centered machine learning. *Commun. ACM* 66, 3 (March 2023), 78–85.

[14] Stevie Chancellor, Eric P S Baumer, and Munmun De Choudhury. 2019. Who is the "human" in human-centered Machine Learning: The case of predicting mental health from Social Media. *Proc. ACM Hum. Comput. Interact.* 3, CSCW (Nov. 2019), 1–32.

[15] Stevie Chancellor, Michael L Birnbaum, Eric D Caine, Vincent M B Silenzio, and Munmun De Choudhury. 2019. A taxonomy of ethical tensions in inferring mental health states from social media. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. ACM, New York, NY, USA, 79–88.

[16] Max E Clarkson. 1995. A stakeholder framework for analyzing and evaluating corporate social performance. *Acad. Manage. Rev.* 20, 1 (Jan. 1995), 92–117.

[17] Ned Cooper and Alexandra Zafiroglu. 2024. From fitting participation to forging relationships: The art of participatory ML. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Vol. 19. ACM, New York, NY, USA, 1–9.

[18] Eric Corbett, Emily Denton, and Sheena Erete. 2023. Power and Public Participation in AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, New York, NY, USA, 1–13.

[19] Natalie Cunningham and Teresa Carmichael. 2017. Sampling, interviewing and coding: Lessons from a constructivist grounded theory study. In *European conference on research methodology for business and management studies*. books.google.com, 78–85.

[20] Christopher A Le Dantec and Carl DiSalvo. 2013. Infrastructuring and the formation of publics in participatory design. *Soc. Stud. Sci.* 43, 2 (April 2013), 241–264.

[21] Janet Davis, Lisa P Nathan, and Others. 2015. Value sensitive design: Applications, adaptations, and critiques. *Handbook of ethics, values, and technological design: Sources, theory, values and application domains* (2015), 11–40.

[22] 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 *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23, Article 37)*. ACM, New York, NY, USA, 1–23.

[23] Sergiy D Dmytriyev, R Edward Freeman, and Jacob Hörisch. 2021. The relationship between stakeholder theory and corporate social responsibility: Differences, similarities, and implications for social issues in management. *J. Manag. Stud.* 58, 6 (Sept. 2021), 1441–1470.

[24] Susanna Doyle. 2007. Member checking with older women: a framework for negotiating meaning. *Health Care Women Int.* 28, 10 (Nov. 2007), 888–908.

[25] Christopher Frauenberger, Judith Good, Geraldine Fitzpatrick, and Ole Sejer Iversen. 2015. In pursuit of rigour and accountability in participatory design. *Int. J. Hum. Comput. Stud.* 74 (Feb. 2015), 93–106.

[26] R Edward Freeman. 2010. *Strategic Management: A Stakeholder Approach*. Cambridge University Press, Cambridge, England.

[27] Batya Friedman. 1996. Value-sensitive design. *Interactions* 3, 6 (Dec. 1996), 16–23.

[28] Jaclyn Friedman and Jessica Valenti. 2019. *Yes Means Yes!: Visions of female sexual power and a world without rape*. Seal Press, Seattle, WA.

[29] Archon Fung and Erik Olin Wright. 2001. Deepening Democracy: Innovations in Empowered Participatory Governance. *Polit. Soc.* 29, 1 (March 2001), 5–41.

[30] Sucheta Ghoshal, Rishma Mendhekar, and Amy Bruckman. 2020. Toward a Grassroots Culture of Technology Practice. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1 (May 2020), 1–28.

[31] Gail E Henderson. 2011. Is informed consent broken? *Am. J. Med. Sci.* 342, 4 (Oct. 2011), 267–272.

[32] Naja Holten Møller, Irina Shklovski, and Thomas T Hildebrandt. 2020. Shifting concepts of value: Designing algorithmic decision-support systems for public services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. ACM, New York, NY, USA.

[33] Jane Im, Jill Dimond, Melody Berton, Una Lee, Katherine Mustelier, Mark S Ackerman, and Eric Gilbert. 2021. Yes: Affirmative consent as a theoretical framework for understanding and imagining social platforms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21, Article 403)*. ACM, New York, NY, USA, 1–18.

[34] F Kensing and J Blomberg. 1998. Participatory design: Issues and concerns. *Comput. Support. Coop. Work* (1998).

[35] Tzu-Sheng Kuo, Aaron Lee Halfaker, Zirui Cheng, Jiwoo Kim, Meng-Hsin Wu, Tongshuang Wu, Kenneth Holstein, and Haiyi Zhu. 2024. Wikibench: Community-driven data curation for AI evaluation on Wikipedia. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, Vol. 35. ACM, New York, NY, USA, 1–24.

[36] Min Hun Lee, Daniel P Siewiorek, Asim Smailagic, Alexandre Bernardino, and Sergi Bermúdez i Badia. 2020. Co-design and evaluation of an intelligent decision support system for stroke rehabilitation assessment. *Proc. ACM Hum. Comput. Interact.* 4, CSCW2 (Oct. 2020), 1–27.

[37] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D Procaccia. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proc. ACM Hum. Comput. Interact.* 3, CSCW (Nov. 2019), 1–35.

[38] Hanlin Li, Nicholas Vincent, Stevie Chancellor, and Brent Hecht. 2023. The Dimensions of Data Labor: A Road Map for Researchers, Activists, and Policymakers to Empower Data Producers. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. ACM, New York, NY, USA, 1151–1161.

[39] Yvonna S Lincoln and Egon G Guba. 1985. *Naturalistic Inquiry*. SAGE.

[40] Ewa Luger, Stuart Moran, and Tom Rodden. 2013. Consent for all: revealing the hidden complexity of terms and conditions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2687–2696.

[41] Cormac McGrath, Per J Palmgren, and Matilda Liljedahl. 2019. Twelve tips for conducting qualitative research interviews. *Med. Teach.* 41, 9 (Sept. 2019), 1002–1006.

[42] Michael J Muller and Sarah Kuhn. 1993. Participatory design. *Commun. ACM* 36, 6 (1993), 24–28.

[43] Naja Holten Møller, Claus Bossen, Kathleen H Pine, Trine Rask Nielsen, and Gina Neff. 2020. Who does the work of data? *Interactions* 27, 3 (April 2020), 52–55.

[44] Terrence Neumann, Maria De-Arteaga, and Sina Fazelpour. 2022. Justice in misinformation detection systems: An analysis of algorithms, stakeholders, and potential harms. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, New York, NY, USA, 1504–1515.

[45] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence* 2, 100041 (2021), 100041.

[46] Lene Nielsen. 2019. *Personas - user focused design* (2 ed.). Springer, London, England.

[47] Office for Human Research Protections (OHRP). 2010. Informed Consent Checklist (1998). https://www.hhs.gov/ohrp/regulations-and-policy/guidance/checklists/index.html. Accessed: 2024-7-3.

[48] Ihudiya Finda Ogbonnaya-Ogburu, Angela D R Smith, Alexandra To, and Kentaro Toyama. 2020. Critical Race Theory for HCI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, New York, NY, USA, 1–16.

[49] Anthony J Onwuegbuzie and Nancy L Leech. 2007. A call for qualitative power analyses. *Qual. Quant.* 41, 1 (Feb. 2007), 105–121.

[50] Kari Paul. 2023. How social media's biggest user protest rocked Reddit. *The Guardian* (Dec. 2023).

[51] Jennifer Pierre, Roderic Crooks, Morgan Currie, Britt Paris, and Irene Pasquetto. 2021. Getting Ourselves Together: Data-centered participatory design research & epistemic burden. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama Japan) *(CHI '21, Article 406)*. ACM, New York, NY, USA, 1–11.

[52] Alun Preece, Dan Harborne, Dave Braines, Richard Tomsett, and Supriyo Chakraborty. 2018. Stakeholders in Explainable AI. *arXiv [cs.AI]* (Sept. 2018).

[53] Organizers Of Queerinai, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, J Hetvi, Huan Zhang, Jaidev Shriram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, Martin Mundt, Milind Agarwal, Nyx Mclean, Pan Xu, A Pranav, Raj Korpan, Ruchira Ray, Sarah Mathew, Sarthak Arora, St John, Tanvi Anand, Vishakha Agrawal, William Agnew, Yanan Long, Zijie J Wang, Zeerak Talat, Avijit Ghosh, Nathaniel Dennler, Michael Noseworthy, Sharvani Jha, Emi Baylor, Aditya Joshi, Natalia Y Bilenko, Andrew Mcnamara, Raphael Gontijo-Lopes, Alex Markham, Evyn Dong, Jackie Kay, Manu Saraswat, Nikhil Vytla, and Luke Stark. 2023. Queer In AI: A Case Study in Community-Led Participatory AI. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. ACM, New York, NY, USA, 1882–1895.

[54] Oliver C Robinson. 2014. Sampling in Interview-Based Qualitative Research: A Theoretical and Practical Guide. *Qual. Res. Psychol.* 11, 1 (Jan. 2014), 25–41.

[55] Janet Salmons. 2017. Getting to Yes: Informed Consent in Qualitative Social Media Research. In *The Ethics of Online Research*, Woodfield Kandy (Ed.). Advances in Research Ethics and Integrity, Vol. 2. Emerald Publishing Limited, 109–134.

[56] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21, Article 39)*. ACM, New York, NY, USA, 1–15.

[57] Elizabeth B-N Sanders. 2002. From user-centered to participatory design approaches. In *Design and the social sciences*. CRC Press, 18–25.

[58] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A human-centered review of algorithms used within the U.S. child welfare system. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–15.

[59] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation is not a design fix for machine learning. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '22, Article 1)*. ACM, New York, NY, USA, 1–6.

[60] C Estelle Smith, Bowen Yu, Anjali Srivastava, Aaron Halfaker, Loren Terveen, and Haiyi Zhu. 2020. Keeping community in the loop: Understanding Wikipedia stakeholder values for machine learning-based systems. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, New York, NY, USA, 1–14.

[61] Robert Soden, Austin Toombs, and Michaelanne Thomas. 2024. Evaluating interpretive research in HCI. *Interactions* 31, 1 (Jan. 2024), 38–42.

[62] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldechova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA.

[63] Harini Suresh, Steven R Gomez, Kevin K Nam, and Arvind Satyanarayan. 2021. Beyond expertise and roles: A framework to characterize the stakeholders of interpretable machine learning and their needs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21, Article 74)*. ACM, New York, NY, USA, 1–16.

[64] Harini Suresh, Rajiv Movva, Amelia Lee Dogan, Rahul Bhargava, Isadora Cruxen, Angeles Martinez Cuba, Guilia Taurino, Wonyoung So, and Catherine D'Ignazio. 2022. Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Feminicide Counterdata Collection. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, New York, NY, USA, 667–678.

[65] Harini Suresh, Emily Tseng, Meg Young, Mary Gray, Emma Pierson, and Karen Levy. 2024. Participation in the age of foundation models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA.

[66] Ningjing Tang, Jiayin Zhi, Tzu-Sheng Kuo, Calla Kainaroi, Jeremy J Northup, Kenneth Holstein, Haiyi Zhu, Hoda Heidari, and Hong Shen. 2024. AI failure cards: Understanding and supporting grassroots efforts to mitigate AI failures in homeless services. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 713–732.

[67] The White House. 2022. Blueprint for an AI Bill of Rights. https://www.whitehouse.gov/ostp/ai-bill-of-rights/. Accessed: 2024-7-3.

[68] David R Thomas. 2003. A general inductive approach for qualitative data analysis. (2003).

[69] Kristin Undheim, Truls Erikson, and Bram Timmermans. 2023. True uncertainty and ethical AI: regulatory sandboxes as a policy tool for moral imagination. *AI Ethics* 3, 3 (Aug. 2023), 997–1002.

[70] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed consent: Studying GDPR consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*. ACM, New York, NY, USA, 973–990.

[71] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA.

[72] Alissa Wilkinson. 2023. The looming threat of AI to Hollywood, and why it should matter to you. undefined. Accessed: 2023-7-26.

[73] Stephen Tze-Inn Wu, Daniel Demetriou, and Rudwan Ali Husain. 2023. Honor ethics: The challenge of globalizing value alignment in AI. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA.

[74] Angie Zhang, Alexander Boltz, Jonathan Lynn, Chun-Wei Wang, and Min Kyung Lee. 2023. Stakeholder-centered AI design: Co-designing worker tools with gig workers through data probes. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23, Article 859)*. ACM, New York, NY, USA, 1–19.

[75] Haiyi Zhu, Bowen Yu, Aaron Halfaker, and Loren Terveen. 2018. Value-Sensitive Algorithm Design: Method, case study, and lessons. *Proc. ACM Hum. Comput. Interact.* 2, CSCW (Nov. 2018), 1–23.

A Participant Demographics

Demographic Variables	N	Percentage	Demographic Variables	N	Percentage
Age			Employment Status		
18-20	0	0%	Employed full-time	5	41%
21-29	5	41%	Employed part-time	1	8%
30-39	4	33%	Self-employed	2	17%
40-49	3	25%	Student	3	25%
Gender			Unemployed (currently looking for work)	1	8%
Female	6	50%	Income		
Male	6	50%	$0	1	8%
Sexuality			1–9,999	2	17%
Heterosexual	10	83%	10,000–24,999	1	8%
Prefer not to say	2	17%	25,000–49,999	3	25%
Relationship Status			50,000–74,999	1	8%
Single (never married)	5	41%	75,000–99,999	3	25%
Married or Domestic Partnership	4	33%	$150,000 and greater	1	8%
In a relationship	3	25%			
Ethnicity					
African	2	17%			
Black	2	17%			
Hispanic	1	8%			
South Asian	2	17%			
Taiwanese	1	8%			
Indigenous	1	8%			
White	3	25%			

Table 1: Aggregate participant demographics.

Pseudonym	Role Description
Maddie	Non-profit Industry AI Researcher
Sean	Participatory AI Researcher in Agricultural Contexts
Tyler	Participatory AI Researcher Studying Online Communities
Scott	AI Developer Focused on Tools for Software Engineers
Falen	Expert Data Annotator for Commercial Mental Health AI
Sabrina	Social Computing Researcher and Volunteer Moderator
Fiona	Former Commercial Content Moderator and Current Content Moderators Activist
Katie	Former mTurk Data Worker and Current Data Rights Activist
Kayla	AI Ethics Researcher and Activist
Charles	AI Developer for Company
Anthony	AI Developer for Company
Spencer	AI Researcher Bulding Public Sector Tools

Table 2: Interview participants' pseudonym and their role description. All names in this section are pseudonyms from the top 100 US baby names and, therefore, do not indicate the culture or ethnic background of the participants. Company and institution names have been redacted unless the participant explicitly requested otherwise.