



When Workers Want to Say No: A View into Critical Consciousness and Workplace Democracy in Data Work

CARL DISALVO, Georgia Institute of Technology, United States of America

ANNABEL ROTHSCHILD, Georgia Institute of Technology, United States of America

LARA L. SCHENCK, Georgia Institute of Technology, United States of America

BEN RYDAL SHAPIRO, Georgia State University, United States of America

BESTY DISALVO, Georgia Institute of Technology, United States of America

In this paper, we describe and reflect upon the development of critical consciousness and workplace democracy within an experimental workplace called DataWorks. Through DataWorks, we hire adults from communities historically minoritized in computing education and data careers, and train them in entry-level data skills developed through work on client projects. In this process, workers gain a range of skills. Some of these skills are technical, such as programming for data analysis; some are managerial, such as scoping and bidding projects; others are social, perhaps even political, such as the ability to say “No” to projects. In what follows, we describe a workshop series developed to build the workers’ critical literacy and consciousness about their data work, specifically regarding the use of data in machine learning systems. After that, we describe a data project the workers questioned and resisted because they determined the work to be harmful. In that process, they demonstrated and enacted a critical consciousness towards data and machine learning. Reflecting on this enactment of data-focused critical consciousness, we identify themes that characterize a democratic workplace, describe the work of designing for organizational action and institutional relations, and discuss how worker and researcher positionality affects this work. In doing so, we argue for enabling workers to resist and refuse harmful data work and challenge the standard power structures of academic research and data work.

CCS Concepts: • Social and professional topics;

Additional Key Words and Phrases: Data Work, Workplace Democracy, Refusal, Labor, Participatory Design

ACM Reference Format:

Carl DiSalvo, Annabel Rothschild, Lara L. Schenck, Ben Rydal Shapiro, and Besty DiSalvo. 2024. When Workers Want to Say No: A View into Critical Consciousness and Workplace Democracy in Data Work. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 156 (April 2024), 24 pages. <https://doi.org/10.1145/3637433>

1 INTRODUCTION

Data work is an often overlooked, undervalued, and exploitative aspect of contemporary computing and our economies. Numerous books, prominent journalism, and academic research are striving to address this situation, drawing our attention to data work as an important site of social action and topic of inquiry [6, 24, 48, 74, 78, 88, 97, 101, 102]. At the same time, higher education institutions are struggling to take a more responsive and responsible place in society. As they do so, we must be aware of the critiques of community-based work dating back decades, such as Ivan Illich’s polemic

Authors’ addresses: Carl DiSalvo, Georgia Institute of Technology, Atlanta, United States of America, cdisalvo@gatech.edu; Annabel Rothschild, Georgia Institute of Technology, Atlanta, United States of America, arothschild@gatech.edu; Lara L. Schenck, Georgia Institute of Technology, Atlanta, United States of America, lschenck@gatech.edu; Ben Rydal Shapiro, Georgia State University, Atlanta, United States of America, bshapiro@gsu.edu; Besty DiSalvo, Georgia Institute of Technology, Atlanta, United States of America, bdisalvo@cc.gatech.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

ACM 2573-0142/2024/4-ART156

<https://doi.org/10.1145/3637433>

"To Hell with Good Intentions" [54] and Sara Ahmed's recent scholarship detailing the hypocritical and manipulative aspects of some institutional diversity, equity, and inclusion efforts [3, 5]. It is within these contexts that we created DataWorks.

DataWorks is an experimental workplace set within a university. As a data services firm, DataWorks hires adults from communities historically minoritized in technology education and careers, and trains them in entry-level data skills developed through client data projects. We use the term minoritized, informed by the work of José Esteban Muñoz [73], and other computing researchers [77, 79] to emphasize that these communities may not be the minority in numbers but are systematically oppressed. One motivation for DataWorks is to develop approaches to community engagement and community-based research that centers those we collaborate with, respecting their lived experience, and prioritizing their perspectives. Instead of the performative and extractive practices familiar to much engagement and research, we are working with community partners, attempting to redistribute the wealth of the university beyond its standard beneficiaries. Another motivation for DataWorks is to co-construct a democratic workplace where the labor of data work is respected and workers have agency in their jobs, such that they can resist exploitative work practices [66]. Coupled with those motivations is the desire to broaden participation in computing beyond the usual patterns of computing as a field and career characterized by elite professionalism.

Most DataWorks projects involve data wrangling: cleaning, standardizing, correcting, integrating, and formatting data sets. These projects are usually more complicated than automated cleaning methods would address, and so workers use manual row-by-row cleaning in conjunction with basic scripting to automate repetitive portions of the wrangling. Some other projects are data labeling, and a few involve data collection through surveys and interviews. The workers are paid a salary with standard U.S.-style benefits, including health insurance. The salary and benefits come from contract work with businesses and non-profits, supplemented by financial support from the university. Research funding also contributes to DataWorks by supporting graduate students and faculty participating in the program. However, client project compensation does not support faculty and graduate students, and the research funding does not support the workers. Such boundaries are essential to maintain the work of DataWorks as more than merely a research project: it is a workplace providing services to clients and contributing to the development of the workers. For the workers, DataWorks is a way to gain skills to help them explore and pursue different careers. For many of the workers at DataWorks, this is their first office-based role, while for others DataWorks is a pathway to a career change. Some of the skills they develop are technical, such as programming for data analysis; some are managerial, such as scoping and bidding projects; others are social, perhaps even political, such as the ability to say "No" to certain work, such as work they determine to be harmful—which is a focus of this paper.

In attempting to co-construct a democratic workplace, DataWorks centers the workers, their lived experiences, and their professional development. In the broader context of the contemporary economy, everyday data tasks (similar to those comprising client projects at DataWorks) have slowly transformed from office-based labor to platform-based digital pieceworking, which are comparatively precarious for workers [8]. In particular, platform-based digital piecework rarely allows for upward mobility within an organization in the context of a worker's professional development [67]. We developed DataWorks, in part, to demonstrate that mid-skill data work can become the basis for a professional career. DataWorks also upends the current paradigm of data work, in which data workers and those who request their labor are separated by a series of barriers to communication and platform workers and requestors know almost nothing about one another [68]. In contrast many client-worker meetings at DataWorks happen either in person at the DataWorks office or using virtual conferencing systems. Critically, not only are the data workers active partners on client projects, but they can also interface with the client to address confusion

and ascertain the goals and implications of the project. The result of this is a collaboration between workers and clients, with the datasets meeting client expectations and also enabling workers to see and understand the larger context of their work. The DataWorks paradigm, thus, responds to what Sambasivan et al. observe in the dearth of respect accorded to data workers that both hampers their work and forces or coerces worker operation in unjust labor environments [88]. It is in this context that the workshop series and client project we discuss in this paper take place. We argue that the tensions that arise around workplace democracy in DataWorks is both the product of, and integral to, its organizational design. Such intentional and experimental structures and processes of DataWorks provide the opportunity to explore ways of restructuring data work, as well as restructuring the role of the university in the community, and the activities and identities of researchers as we endeavor to sustain DataWorks.

This paper is a description and reflection on the development of critical consciousness among the workers at DataWorks through a series of workshops, and a subsequent episode in which they surfaced concerns with a client project and eventually re-negotiated the project's design and execution. We begin by setting the research context for DataWorks across participatory design, engaged scholarship of labor and data in computing-supported cooperative work, and justice-centered approaches to computer science education, critical data literacy, and critical consciousness. We then describe a workshop series developed to build the workers' critical literacy and consciousness about the data work they perform, particularly regarding the creation and use of data in machine learning. After that, we describe a data annotation project that the workers questioned and resisted because they determined the work to be harmful. In that process they demonstrated a critical consciousness of data and machine learning, building upon their lived experiences and existing knowledge of racism. From this, we identify themes that characterize a democratic workplace, describe the work of designing for organizational action and institutional relations, and discuss how positionality affects this work. This work contributes to a broader goal of advancing more just forms of tech labor and scholarship, by providing ethnographic accounting and informed reflection that can inform research and practice. We hope this paper offers an empirical and inspiring account of workplace democracy, highlighting the importance of enabling workers to resist and refuse harmful data work. As we strive toward more just and fair data work and broaden participation in computing, it is crucial that we challenge the standard power structures of academic research and data work.

2 RELATED WORK: SETTING THE RESEARCH CONTEXT FOR DATAWORKS

As a research project, DataWorks is informed by and integrates participatory design, engaged scholarship of labor and data in computer-supported cooperative work, as well as critical perspectives on data and learning from the learning sciences. We see DataWorks as part of a history and practice of critical computing because core to the project is questioning the hegemony of computing as a field and profession, and attempting to refigure computing as we know it, including our research methods and community engagement. DataWorks, at least in part, can also be seen as a response to the provocation of Catherine D'Ignazio and Lauren Klein in *Data Feminism* to imagine and pursue more varied forms of data science [31]. In that sense, we strive to "do data work otherwise." In tandem with the technical work, DataWorks addresses concerns about respect for data workers, their skills, and their rights, as surfaced by Sambasivan et al. and Miceli et al.[68, 88]. Shared across the disciplines we draw from is a collaborative approach to knowledge production and direct action through research.

2.1 Participatory Design

The history of participatory design is rooted in workplace democracy. That history and commitment are the basis of this paper and one of the pillars of DataWorks. From its earliest projects, participatory design methods were meant to sustain workers' self-determination and skills [15, 36, 59, 86]. The values of workplace democracy and corresponding participatory design methods are grounded in the belief that workers should be able to contribute to shaping their tasks, processes, and environments, and that labor is a practice worthy of respect. These values clash with much contemporary data work, especially crowdwork and piecework. As such, we see DataWorks as a project to contribute to collaborative practices of shaping data work to be fair to workers. Randy Trigg and Karen Ishimaru's work with the Global Women's Fund [99] is a vital touchpoint in the history of participatory design and the quest for workplace democracy that influences DataWorks. Over more than a decade, Trigg, Ishimaru, and collaborators established an integrated infrastructure of participatory processes in the Global Women's Fund, exploring how participatory design could become part of the culture and practice of an organization. Their work sets an aspirational model for how participatory design might become integral to an organization. In addition to being aspirational, their work is crucial for understanding the challenges of integrating participatory design into a non-profit organization. The difficulty and importance they identify of moving between Designing Work Systems, Designing Organizational Frameworks for Action, and Designing the Industrial Relations Context [45] are key to our discussion of workplace democracy in DataWorks.

It is sometimes a challenge in participatory design projects, whether scholarly or applied, to identify the effects of participatory design. As Tone Bratteteig and Ina Wagner ask, "What is a Participatory Design result?" [17]. Over the past several decades, participatory design scholars have explored varied approaches to evaluation. In their 2016 paper "Evaluation in Participatory Design: A literature survey," Bossen, Dindler, and Iverson provide a thorough overview of almost 25 years of participatory design scholarship [16]. Through their analysis, they identify how, among other themes, workplace democracy has been assessed and call for more explicit and systemic evaluations, and the use of evaluation to democratize, empower, or enhance mutual learning. In their survey, they also raise the issue of criteria – by what and whose criteria do we evaluate the outcomes of a project? Spiel et al. have explored how to evaluate participatory design with autistic children, and through this work, raise crucial issues of how we can evaluate our research from the perspective of being an ally, realizing the need for evaluation while also remaining critical towards evaluation [95]. Hansen et al. proposed to use program theory as a lens for participatory design evaluation, focusing on causal relations in participatory processes and subsequent effects. While we did not use program theory in our evaluation, it does offer a compelling approach [49]. In our work, we hope to contribute to this literature on evaluating participatory design endeavors and, specifically, to provide empirically derived criteria for workplace democracy in data work.

Part of the challenge of evaluating workplace democracy is that definitions of workplace democracy are often vague: suggestive of worker choice but short of either operationalized or philosophical descriptions. There is, in fact, considerable scholarship on workplace democracy. Some scholarship is philosophical and concerned with whether there is a legitimate moral claim for workplace democracy [65]. Other scholarship comes from management and organizational studies and is concerned with how democracy is or is not manifested in various kinds of firms [41]. One contribution of this paper to participatory design is to name several characteristics of workplace democracy that other researchers and practitioners can use in their interpretations of the workplace. In line with the tradition of participatory design, we see workplace democracy as a practical issue. We take workplace democracy, generally, to be something good that is indicated by the capacity for workers to meaningfully influence their tasks and work environment. Identifying characteristics of that

capacity so they can be applied to designing, interpreting, and evaluating workplaces contributes to both the robustness of the concept and computing and design scholarship.

Another touchpoint from participatory design are projects with extended engagements between designers, organizations, and communities that question prevailing assumptions about work, democracy, and participation. At times, participatory design can be hegemonic by uncritically adopting norms. But other forms of engagement are possible. For example, Anna Seravalli's multi-year work with immigrant communities in Malmö, Sweden, demonstrates a participatory design practice that meets communities where they are and is responsive to their customs and values [51, 91]. Her and her colleagues' work calls attention to how fundamental participatory design concepts, including "participation", are culturally situated. Through her work with an immigrant community in Sweden, we witness how expectations and beliefs about who can or should participate and what participation entails are meaningfully different across cultures. We are similarly committed to questioning cultural assumptions about "participation," "data," and "work" through DataWorks—including our own beliefs as they are shaped by race, class, and gender.

2.2 Engaged Scholarship of Labor and Data in Computer-Supported Cooperative Work

In addition to participatory design, there are many modes of engaged scholarship in computer-supported cooperative work. Action research is one such example of scholars, practitioners, and communities working together to use computing to address pressing needs and concerns [50]. Intersectional perspectives examining the entwining of race, class, and gender are another prominent site of scholarship in computing research [25, 38, 39]. Over the past decade, there has been an increasing interest in labor as a site of engaged research, combining multiple methods and theories, including those of participatory design. Lilly Irani and Six Silberman's work on Turkopticon is a germinal example of using computing and design to address inequities in piecework [55]. Recently, significant attention has been put towards the gig economy and developing tools and processes to address wage theft as part of a broad social justice agenda in human-computer interaction design and research [35]. In their discussion of anti-oppressive design, practitioner-scholars Jill Dimond and Thomas Smyth highlight worker cooperatives as a model of alternate modes and sites of technology and knowledge production [94], and research on cooperatives, unions, and other forms of organizing labor is proliferating [62]. Drawing from participatory design and building from an ongoing engagement with issues of ethics and labor, Miceli et al., in collaboration with data workers, are exploring how documentation practices can be designed that foreground the practices and needs of workers [67]. Across these varied projects and perspectives is a shared agenda of pursuing research as a practice of allyship with workers towards more fair labor conditions.

Such engaged scholarship is another basis for DataWorks, and we continue to be inspired by those projects. We are attempting to address some of the same root issues by making an experimental workplace that centers the workers from the start. As such, organizational development is a significant part of our engaged scholarship. In this way, DataWorks is akin to projects such as Kotturi. et al.'s, help desk for local entrepreneurs [60]. In this project, the researchers constructed a physical "Help Desk" in a local community center to field technical questions for residents starting, or working to sustain, small businesses. The physical site transformed into a virtual place during the COVID-19 pandemic; nonetheless, such projects exemplify the work of building organizational platforms for engaged scholarship. We are similarly inspired and influenced by the work of Tawanna Dillahunt and her collaborators as they are exploring the design and use of technologies to support job-seekers, and especially their recent Community Tech Worker's project, through which they collaborate with community members to build fair and sustainable employment opportunities and resources [32–34]. The choice to use the university as a site of engaged scholarship brings into question the institution itself and its role: reflecting on, and sometimes refiguring, our institutions

are a vital part of engaged scholarship. From the start, it was important that DataWorks provides actual material support for residents employed as data workers. In this way, we see DataWorks as a way of redistributing the university's resources, and, subsequently, the resources of the state. To be clear, this is not a stated intent of those organizations that provide support for DataWorks, but it is an effect of that support.

2.3 Justice-Centered Approaches to Computer Science Education, Critical Data Literacy and Critical Consciousness

Fundamental to DataWorks is a justice-centered approach to computer science education and critical data literacy. A justice-centered approach to computer science education seeks to engage learners in the ethical consideration of computing, respond to learners' social and civic identities, and highlight the historical inequalities in computing [103]. While a handful of projects have begun to address computing education in K-12 setting with a justice centered approach [29, 80, 82, 89], there is less engagement with justice-centered workplace learning. We also draw from scholars who characterize critical data literacy as a key component of broadening participation in computing. Such scholars identify technical and social components of data literacy that are complementary and indivisible, noting how the technical skills required to analyze and work with datasets should be paired with an iterative process embedded in local context [100]. These scholars echo the work of Bhargava et al. [11] who assert that 1) data literacy must focus on fostering adaptive capacities and resilience instead of teaching platforms and technical languages, and 2) data literacy must empower people in meaningful and effective ways. Bhargava et al. also identify critical challenge to data literacy: Understanding the importance of context in data literacy.

Introduced by Paulo Freire in the 1960s, critical consciousness refers to the ability to interpret and resist how seemingly benign elements of the world reinforce systems of oppression [42, 56]. Friere developed this concept of critical consciousness as part of his broader program of emancipatory education in the context of work with adult laborers. There are three general components of critical consciousness: reflection, motivation, and action [30, 42]. Reflection speaks to learning to question social conditions regarding how those structure and enforce power relations that are often oppressive. Motivation speaks to developing a sense of agency to address oppressive conditions. Action speaks to working to change oppressive conditions, either individually or collectively [30, 42]. In the context of DataWorks, we use critical consciousness to refer to an awareness of how data perpetuates existing power structures and how data work itself can be a site of oppression, and also transformation. Although it is not the same as critical consciousness, critical reflection is another term and practice that aims to support an awareness of the social and political contexts of work and empower workers (at all levels) to take action in response to that awareness [19, 70]. Similar to critical consciousness, critical reflection is a process of learning. Much of the work on critical reflection in the workplace is grounded in the management and organizational literature and overlaps with scholarship on workplace learning. In contrast, much of the work on critical consciousness is grounded in education research, and specifically the praxis of critical pedagogy. Our focus is on critical consciousness as a concept and process, however we also draw from the scholarship on critical reflection, and we generously meld these literatures.

A common theme throughout the research on workplace learning, critical consciousness, and critical reflection is the immense difficulty of cultivating critical consciousness in the workplace because of the friction it often produces due to misalignments with the operational aspects of work [23, 40, 84]. This is not surprising; developing an awareness of oppressive systems in the workplace is likely to produce tension—indeed, that is a purpose of critical consciousness. For the worker, emerging critical consciousness can result in cognitive dissonance that drives a perspective transformation [83]. While this cognitive dissonance can produce a perspective shift within an

individual, without explicit organizational supports, there is "no inevitable flow between personal transformation and an organizational or societal level" [83]. Indeed, if the essence of critical consciousness is to facilitate emancipatory action, this is likely to be at odds with the capitalistic underpinnings of a typical workplace. Organizational supports that truly enable workers to effect this kind of change will necessarily disrupt the flow of work. Particularly for workers in the lower end of an organizational hierarchy, reconciling criticality and the nature of repetitive work or strict training processes may result in "unresolvable tension[s]" [18, 23] or punitive action [40].

In order to think beyond the apparent paradox of criticality in the workplace, other scholars argue for thinking outside normative ideals that position critical consciousness as adherence to another singular vision of the world [23, 37, 92]. Building on political organizing literature [7, 9], Cotter re-conceptualizes critical reflection in the workplace as "reflexive spaces of appearance" that engender emancipatory praxis in the workplace [23]. Through dialogue, these spaces are characterized by participants' commitments to honesty, critique, and the generation of different ideas and starting points. In these spaces, the goal is not to ignite a critical flame then disperse to perform an action, the goal *is* an action: compromise, i.e., "the price of praxis" [7] according to some. While moving from reflection to motivation and action—or rethinking this paradigm entirely—is a persistent challenge across domains, little work on critical consciousness has examined tech labor generally, or data work specifically. We believe the framework of critical consciousness is useful for engaged scholarship of labor and data because it provides a process and vocabulary of learning that seems to match with the aspirations of research as a practice of allyship with workers towards more fair labor conditions.

3 METHODS

The structure of DataWorks creates a novel research environment. Participatory design and ethnographic methods are embedded within the workplace [13, 14, 93]. The workers collaborate in the research to varying degrees. What's important to emphasize is that **the workers are not individual research subjects** in any traditional sense. Instead, we are exploring the conditions and processes of data work within DataWorks. The faculty who participate in DataWorks are not just participating as researchers. They are also involved in supporting the regular operations of DataWorks in different ways: sometimes by offering training and sometimes through project management. Indeed, all of the faculty involved in DataWorks work directly and regularly with the workers and managers, across the full range of activities that make DataWorks possible, from coordinating hiring with university human resources, to providing technical instruction to the workers, to soliciting clients and writing project bids together with the workers, management, and staff, to sometimes assisting directly in the data work. This approach falls short of action research in that we are not (yet) collaboratively forming research questions with the workers. Still, we are inspired and informed by action research methods and commitments to being fully engaged in the research context. At times, workers directly participate in the research, including co-authoring papers. The workers are not included as authors on this paper because, unlike other efforts, we did not set out to study these episodes as research, and the workers did not participate in these episodes as researchers. Rather, we are reflecting and theorizing on these episodes after the fact.

The project described in this paper took place over approximately nine months. The lead author was not directly involved in the activities that will be described. Instead, they approached these events from an ethnographic perspective, drawing on methods of participant observation [96]. More specifically, the lead author attended weekly staff meetings which included all of the workers, the manager, and researchers (30 hours), weekly meetings with the researchers in which this project and others were discussed (30 hours), six hours of meetings related to the development and content of the workshops, and eight hours of meetings specific to this project (all times are approximate). In

addition to attending these meetings, the lead author met with the manager to discuss this project on four occasions during the project, each for an hour and again after the conclusion of the project for another hour. After the project, they also met with the researcher responsible for developing the workplace learning program to review in detail the program's structure and to document recollections and reflections from that researcher. Because DataWorks is, in part, a research project, faculty, managers, and workers are encouraged to keep notes on the work environment and their activities. As part of this particular project, the workers completed weekly reflections. All of these materials are shared to inform the ongoing development of DataWorks, and also contributed to this paper, providing documentation and reflections on the project from multiple perspectives. The materials analyzed and interpreted include the aforementioned weekly reflections, meeting agendas and notes, email communications, and project documentation by the faculty, managers, and workers. Altogether, this comprised a robust corpus for analysis and interpretation. Finally, because DataWorks is an ongoing project, the lead author could (and did) return to the workers, managers (both former and current), faculty, and graduate students throughout the analysis, interpretation, and writing process to ask additional questions.

To analyze these diverse sources of data, we followed a grounded theory approach [20, 47]. The materials described above were printed, read, and hand-coded by the lead author. Next, we conducted a thematic analysis to identify initial categories that were further discussed with the larger research team. Using constant comparison, the lead author elaborated these initial categories and iteratively developed, refined, and tested them in relation to the research literature. Emerging conceptual categories were also continually tested and refined in discussions with the larger research team to create more stable conceptual categories. Throughout this process the lead author also collaborated with the manager, workers, and researchers involved in the project as additional questions were raised and further information was needed.

4 A VIEW INTO CRITICAL CONSCIOUSNESS AND WORKPLACE DEMOCRACY IN DATAWORKS

4.1 What is DataWorks?

DataWorks is a novel workplace set within a university. The idea for DataWorks came from a small cadre of researchers who spanned the learning sciences, design, and computing and had extensive experience in community-based research and efforts toward broadening participation in computing. From prior experience working with non-profit and community-based organizations and city governments, they realized there was a need for basic data work. Many organizations and municipalities want to use data, and some have data, but they lack the time, expertise, or funding to make their data actionable. At the same time, from prior experience with community-based learning programs, it became clear that there was a need to provide vocational training to adults. So often, in efforts towards broadening participation in computing, the focus is on youth. That is essential work. However, adults also need technical skills and literacy. DataWorks sought to combine these: to create a program that provides free or low-cost data work for non-profits, community-based organizations, and government agencies and programs through a workforce of trained and paid adults to do data wrangling. Over time, we expanded to offer data services to academic researchers. This was in direct response to the preponderance of academic researchers that use crowdwork and piecework platforms, not always in ways that are fair to the workers. Our aim was, and is, to demonstrate to academic researchers that they can have data cleaned, formatted, and labeled through services such as DataWorks that provide fair compensation to the workers their research depends upon.

The workers for DataWorks are recruited through online employment postings and announcements of opportunities shared with local non-profits and community partners. Workers apply through the university employment office and are interviewed by the academic researchers and the DataWorks manager. To date, DataWorks has hired 13 workers: 6 men, 6 women, 1 non-binary person, aged 18 to 50. Once hired, the workers go through a training sequence using tools such as Excel and OpenRefine for basic data wrangling and also are introduced to Python for more advanced procedures. Much of the work is done in person, in a dedicated space at the university, and each worker has a desk and computer. The entire staff meets weekly to discuss projects. In those discussions, workers raise questions about current projects, request specific training needed for projects, and discuss upcoming projects. A sub-set of researchers, the manager, and 1-2 senior data workers meet weekly to discuss new clients and new projects, personnel or training issues, and general strategy for the growth of DataWorks.

From the start, we wanted DataWorks to be an experimental workplace that was worker-centered: that recognized and appreciated the value of the lived experience of the workers and sought to co-construct processes that enabled the workers to have a say in their work. From our perspective, having a say in their work requires a critical understanding of data and data work. In other words, from our perspective, critical consciousness and workplace democracy go hand-in-hand.

As part of our effort to create and sustain a worker-centered environment, faculty have crucial but limited involvement in the day-to-day operations of DataWorks. Instead, the day-to-day operations are directed by a manager. In the following scenario, we hired the manager separately from the workers. However, for the past two years, the manager has been promoted from within – workers can choose to move to management after spending a year in DataWorks. In collaboration with the manager, the faculty solicit work, interview and hire workers, and participate in weekly staff and project meetings, providing feedback, and when needed, direction and strategy. When called upon, the faculty also assist with client and personnel issues. But for the most part, the manager makes the day-to-day and week-to-week decisions that structure DataWorks, from scoping projects, to assigning tasks, to reviewing work, to billing.

In what follows, we describe an arc of learning and action that demonstrates the emergence of a critical consciousness of data and data work and the enactment of workplace democracy. This arc occurred over approximately nine months, happening in punctuated episodes. As with any ethnographic account, we have pruned the narrative. Across these moments of learning and action, the workers, research staff, and clients confronted the potential for harm in and through data directly and, subsequently, developed responses to those conditions. Through that process, we witness the workers' capacity to set the terms of work in a way that demonstrates the agency to meaningfully influence their tasks and work environment.

4.2 Developing Critical Data Literacy for Machine Learning

As DataWorks took on clients, we saw an opportunity in conducting work on datasets needed to develop and train artificial intelligence (AI), machine learning (ML), natural language processing (NLP), and computer vision (CV) models. The work of labeling, annotating, and standardizing data sets is tedious, yet a skill that can be learned by people new to technical work, which makes it apt for DataWorks. In addition, we believed there were situations when the knowledge and perspectives of the workers could be useful in processing the data: situations when they would bring a view that would add value to the data. We reached out to researchers at multiple institutions, describing DataWorks and promoting our services, and several projects quickly materialized. One such project involved clips of infants lying on their backs and labeling when they kicked their legs. This was used to investigate a system for early detection and intervention for infants at risk for motor impairment. Another involved labeling clips of videogame play to train an algorithm to identify pairings between

player strategies and game mechanics. After the first few such projects, we developed a workshop series to familiarize the workers with basic ML concepts and mechanisms, combining our expertise in participatory methods and the learning sciences. These workshop activities did not require prior experience in programming or computer science. They were intentionally designed as non-computational approaches to understanding ML that presented the larger concepts in the context of the everyday. They scaffolded an understanding of ML among novice data workers by connecting to their work experiences and current affairs, that is, their individual and collective lived experiences. The purpose of this workshop series was thus twofold: to contribute to the workers' contextual understanding of data and data work and to help the workers screen prospective client projects and assist them with refining their goals for the data task. In other words, we intended the workshop series to provide both critical data literacy skills and contribute to their ongoing technical and organizational skill development. These activities, and the outcome of a critical consciousness towards data and ML that resulted from these activities, affected subsequent DataWorks projects and processes in profound ways.

The workers began the workshops with an acute awareness of systemic racism from their lived experience. As Black citizens in the United States, they face racism every day. They are well aware of not only experiences of racism such as slurs, but also of how White people's perceptions of them affect their ability to get jobs and their treatment in the workplace. They are also well aware of how, for some of the workers, their neighborhoods were historically and consistently provided less services, from bus routes to grocery stores. We share this because we want to be clear that we are not claiming that the workers had no critical consciousness before these workshops. On the contrary, racism was ordinary in their lives [28, 58, 75]. Instead, our attention is on a critical consciousness about data, algorithms, and, more generally, computational systems and technologies. Building off the critical consciousness and lived experience workers already held in different domains, the goal of this workshop was to help workers analyze potential uses of data-intensive systems in conjunction with their role on those projects. This intentional effort to promote worker agency is one element of DataWorks that sets it apart from other data work sites. Critical perspectives on data were not as familiar to the workers before the workshop. This is unsurprising. Until recently, discussions of bias in data and algorithms were primarily academic, with scant coverage in popular media.

The workshop series took place once a week over approximately three and a half months. Each week, the workers and a researcher would meet for about 90 minutes. We built each workshop activity around a critical data literacy concept and a data set or algorithm through which the workers would explore that concept in practice. By anchoring the concepts in data and using the techniques of data analysis and operations that were part and parcel of their client work, we wanted to materially and experientially demonstrate the connections between data, practices of data work, and issues of justice. The workshop was separated roughly into two halves; the first half focused more on understanding ML, while the second examined basic data skills in conjunction with issues of equity and justice.

Through this, we expected these workshop activities would contribute to a critical consciousness about data, within which the workers would come to recognize the potentials, consequences, and limits of data work. For example, the first activity of the workshop series was a discussion of evidence-based assertions, that is, how data can be used to advance a narrative or support a position. The key to this was to develop an awareness that data is not "objective," but can be used, and even distorted, to present a specific narrative. The case study for this example was a map of the COVID-19 cases in the Southeastern United States (where DataWorks is located). When the map and accompanying data set was introduced, we asked the data workers to spot what was "wrong" or "off" about the graphic. After several minutes and some nudging, they spotted that the scale

of the graphics made no sense — the same color scale (used for two identical maps, to showcase progression over time) was matched to different case values. The workers then explored a dataset about COVID cases in neighboring states. The researcher then posed a series of questions. The questions were formed around understanding what arguments, or narratives, might be developed from the data. These included questions such as, *Could the data be used to make the argument that COVID-19 affects women more than men?* and *Are people with existing comorbidities at a higher risk of death?* and *Could the data be used to make the argument that people in Southern states are unhealthier than in other parts of the country?* These questions were designed to demonstrate the complexity of situations represented by data, while also drawing on workers' lived experiences as Southerners in the COVID-19 pandemic. For example, the question about Southern states, which was depicted as having a higher case load, overlooks the context of those states, which are generally a more rural and underfunded part of the United States that, in some places, lack in both basic preventative and emergency healthcare infrastructure. Here, the objective was to demonstrate how data can be used to further an agenda, and how a critical examination of what the data shows can help uncover co-contaminate social and technical phenomena. The researcher also asked reflective questions that prompted considering the limits of this data set, such as *What claims would you feel comfortable making based on this dataset?* *What information about the dataset would you want to have, but don't?* and *What reservations would you have about working with this dataset?* After answering these questions individually, they were discussed as a group. Through that, they continued to collaboratively articulate a collective understanding of this particular data set and develop critical perspectives on data more generally.

More technical terms and processes were introduced as the weeks progressed. At the same time, we also introduced more social and political connections between data, ML, and everyday life, by drawing from news stories and integrating videos from organizations such as Data for Black Lives that directed our attention to multifaceted aspects of data: from being responsible stewards of data, to the weaponization of data and algorithms, to the pro-social potential of data and data work. The workers were less familiar with these situations, particularly regarding data. For example, while they had experienced racism in job hiring, they had not previously considered how data and algorithms might perpetuate and amplify racist hiring practices. Once realized, the implications were unsurprising to them, as yet another experience of racism.

Throughout, we designed the weekly activities to support worker agency in developing critical data literacy. To this end, we believed it was not enough to make workers aware of the themes of critical data literacy and the potential harms of data and algorithms. It was also necessary to assist the workers in developing skills that they could use to address harms caused by data or algorithms and take direct action through their work. For example, an activity on ML applications included developing a list of questions to ask whenever starting work with a new dataset about its contents and potential applications. Workers practiced this list on hypothetical scenarios like a medical company that develops a system to diagnose different kinds of bug bites without having patients interact with a doctor in person; instead, patients submit photos of their bug bites for automated classification via an app. Potential concerns for this scenario discussed in the workshop include issues of medical ethics (removing in-person care), privacy (storage and processing of user photos), and fairness (whether the system has been trained to recognize potential variance of bites on different skin tones), among others. We reiterated the real-world implications of these projects; for example, the following week after the bug bite recognition system, we began with a discussion of a news story that had emerged in the interim, in which Google was launching a dermatology AI app in the European Union.

One outcome of this workshop series is that workers expanded their vocabulary for talking about data, data work, and issues of justice. That is, they developed a critical consciousness and

capacities to express that critical consciousness. For example, one workshop focused on the task-work systems often used in data labeling projects, like projects they were doing. At the start of the workshop, the workers recalled earlier experiences with such systems, saying they had not always been compensated and that the work seemed unfair. After the workshop activities and discussions, the workers could speak about these systems and issues with critical insight. As one worker expressed in a written reflection referring to a task work platform, *“This system is set up to exploit workers because there is a low level of transparency and virtually no workers’ rights. There should be regulations in place to protect [crowd] workers.”* Such a reflection reveals an astute understanding of both the systems and practices of data work and potential courses of action to address problematic conditions. Significantly for relating critical consciousness to workplace democracy, the worker casts this problem not merely as a technical issue, but also as a labor rights issue. Seeing the connection between data, their work, and their experience of oppression was overwhelming for one of the workers. After viewing a Data for Black Lives video, she left the room in tears. A while later, she returned and stated that she was overwhelmed, and that it was *“us against them”*, and saw herself as a data worker in a position to act against the systemic racism and oppression against herself and her community.

4.3 Expressing Critical Consciousness of ML In A Data Annotation Project

A few months after completing the critical data literacy and ML workshops, a computing researcher from another institution approached us about a labeling project involving African American Vernacular English (AAVE), which is also referred to as Black Vernacular English. The client used the term AAVE, so we are using that term in describing the project. After discussion, we thought this project might be a significant opportunity for the workers, who at that time were all Black Americans, to engage in a project that built upon their positionality and demonstrated the value of their expertise. From a design justice stance of “nothing about us without us,” such a project was an opportunity for the workers to contribute to shaping technology that was about them and would affect them [22]. Following our usual process, we agreed to an exploratory foray into the data set and accompanying work led by the workers. Very quickly, the project revealed complexities and potentials of supporting critical consciousness in the workplace as the workers questioned, resisted, and in due course, re-negotiated the project’s design and execution in meaningful ways. What had been hypothetical scenarios in the prior workshops materialized as real-world events as the workers, researchers, and clients engaged the issues of bias and harm in data sets and data work.

The project was to examine short statements and annotate the statements as examples of AAVE or not. Some statements were taken from social media channels, such as Twitter or online forums, such as Quora. Other statements had been algorithmically generated based on a model of AAVE. In other words, some of these statements were not expressions of AAVE authored by people but algorithmically generated AAVE statements. The purpose of the task was to assess the capacity of the algorithm to generate AAVE, knowing already that the algorithm was flawed. From that assessment, the client believed they could modify the algorithm. The client sought out DataWorks specifically because they wanted to approach this project responsibly and ethically. Rather than distributing the work through an online crowdwork platform, they wanted to use workers who were culturally situated for the annotation task to produce data, which would, ostensibly, enable them to create a more accurate and appropriate computational model of AAVE.

Numerous problems quickly arose. To begin with, the workers were uncomfortable with the notion of a singular African-American vernacular. They pointed out that their experience of a Southern vernacular was distinct from what might be a vernacular in Chicago, Los Angeles, or Boston. If they were experts, it was only regarding a Southern Black vernacular. Then there were issues with the data itself. Both the algorithmically generated phrases and those taken from

social media and online forums were often offensive, leading to discomfort among the workers when annotating the phrases. Most of the offensive material was sexual in nature, some of it was homophobic. Some of it was also racist, in particular expressing anti-Asian and anti-Asian-American sentiments. The workers immediately identified these phrases, and the condition of having to work with them, as inappropriate and unacceptable. They requested that the manager reach out to the client, make the client aware of the issue, and ask for a remedy. That recognition and request, as simple as it might seem, is an important expression of workplace democracy, as the workers had the capacity to identify work conditions as inappropriate and the agency to request those conditions be addressed. This is in contrast to the dynamic at many data work sites, such as content moderation firms, where workers are neither encouraged nor supported in voicing concerns about their work [85].

In addition to the content of the data being offensive, the workers began to find the very concept of the project problematic. In daily check-ins and conversations, the workers began to question the premise behind the project and the project's actual outcomes. They raised concerns such as "*Who really wants to be able to identify Black speakers?*" and "*Once someone has been identified as a Black speaker, then what?*" These questions are similar in kind and subject to the questions raised in the prior workshop series. As these questions emerged, the workers seemed to be applying the critical perspectives from those previous conjectural activities to this client work. As the questions and conversations continued, several of the workers arrived at a specific concern: "*It seems like this could be misused in a way that will hurt my community.*" And then, they requested a meeting with the client. In the meantime, we collectively agreed to put the work on hold until they had expressed their concerns and the client had adequately addressed those concerns to the workers.

We reached out to the client and the first meeting went poorly. The graduate student who was leading the work met with the workers. He was not a person of color and demonstrated little experience talking about race. He listened to the workers and told them about the purpose of the study: to demonstrate that current algorithms for AAVE were inadequate and that their work would help substantiate that claim. However, that explanation, and its presentation, did not build confidence in the workers or adequately address their concerns. Two Black American Ph.D. students also attended the meeting—we thought to help alleviate power imbalances. This proved unnecessary, as the workers directly, and without hesitation, expressed their concerns with the data work. The workers made multiple suggestions, including that a Black researcher would have better insights and sensibilities to conduct this work, and that rather than reviewing an algorithm already known to be flawed, it might be more productive to consider how to more directly generate AAVE better. The graduate student was flustered and returned to the lab to share these concerns with the lead investigator. In the meantime, the workers requested that the project remain on hold until their concerns were addressed.

Several weeks later, the graduate student met again with DataWorks, with three changes to the project, each in an attempt to address the concerns raised by the workers. The first was to use a filter to remove profanity from the data set before sending it to DataWorks, to address the problem that the algorithmically generated statements were at times offensive. The second was to change the response form to include more nuance besides correct or incorrect. Workers could instead express if a translation was correct, incorrect, or concerning. If a translation was marked as concerning, the research team would follow-up to understand why it was labeled as such. Finally, an open entry field was added to the form used to process each phrase, to enable the workers to provide ideas about why the translation was wrong, and express other concerns. With this new responsibility, what had previously been a task defined and dictated by researchers had, in effect, become a co-creation task between the data workers and the researchers. It is important to note that during this process, the workers continued to be paid their regular hourly wage and were

encouraged to discuss their concerns while "on the clock." The workers were compensated not only for performing the annotation work, they were employed and accorded due respect as consultants on the project, even during the stoppage of the annotation work.

With these changes in place, the workers resumed the task and completed the first batch of labeling. However, several months later, the student researcher sent over the second batch of AAVE translations for labeling. New workers were on staff who were concerned with what they heard from the veteran data workers about the previous batch of data. In a staff meeting, the data workers expressed their concerns after taking a pass at the new data set. The data still contained profanity and inappropriate content. Moreover, the new workers had no context for how this work could be helpful to their community—they only could see how it would be harmful to the Black community. One of the DataWorks research staff, a faculty member, reached out to the lead investigator on the project and advisor to the graduate student involved. They talked about the data workers' concerns. They brainstormed how to implement the suggestions previously made by the workers, including changing the lead student researcher to a Black Ph.D. student and focusing on translations rather than annotations based on the workers' previous feedback. Subsequently, the lead investigator met with the workers. She answered the workers' questions about the study's goals, assuring them that it sought to create technology inclusive of dialects and vernaculars beyond Standard American English. She related how other dialects were also part of the study, including language variants associated with her positionality as a person of color and an immigrant. She understood their concerns, and while she was not Black, she could empathize with how language is used to advance racist perspectives. And she sought out compromises.

The first was to review the algorithmically generated phrases for offensive keywords with a stricter protocol and delete offensive entries from the data set—in effect, to do what the workers had asked for already but which the client had failed to actually follow through on. The second was to ask the data workers for translations in addition to annotations. Notably, the data workers rejected this idea; they felt it would produce biased data as they knew only a Southern dialect of AAVE, and they *"could not represent all Black people."* The third was to change the primary graduate student investigator to a Black student. One of the new data workers was insulted by this suggestion. They told us it signified their grievances were about the race of the researcher rather than the data and the purpose of the data. However, the workers who had previously requested this and researchers explained the prior reasoning for this, and subsequently followed through with this change. Finally, the lead investigator suggested a data use agreement that the workers would write. This data use agreement would state the intended uses of the data and what the data could not be used for. Once the data workers wrote the agreement, the lead investigator would sign the use agreement and commit to it. Only after the offensive data was removed and the data use agreement was signed did the workers continue and complete the project.

Of course, data use agreements, unless bound to policy, regulation, or law, are nebulous and nearly impossible to enforce [63]. This data use agreement depends upon a social contract and goodwill. While that is problematic, we should not dismiss the importance of the worker agency in this situation and the critical consciousness that developed. It is important that they identified consequences of this data work. It is crucial that they were able to act upon that, to choose to stop the project, to require accountability from others, and then to determine under what conditions they were willing to act. This is essential to fair and just approaches to data work.

5 DISCUSSION

The workshop series and subsequent AAVE project demonstrate an emergent critical consciousness and provide insight into a workplace democracy. In the following sections we reflect upon this, connecting back and building upon the research literature and themes that motivate this work. First,

we name two characteristics of workplace democracy that other researchers and practitioners can use in their interpretations of the workplace: refusal and agonism. Then we discuss the difficulty and importance of moving between designing work systems, designing organizational frameworks for action, and designing the industrial relations context. Finally, we reflect upon the institutional labor of workplace democracy, how the positionality of the workers, clients, and researchers affected this project, and the changes we have made to DataWorks based on this experience.

5.1 Refusal and Agonism in the Workplace

As the workers developed more understanding of the implications of data, they approached data work with increasing skepticism. In the AAVE project, upon realizing it was dangerous to their community they chose to halt the work. This caused a disruption to the flow of work within DataWorks and prompted challenging engagements with the client. **Such trouble is appropriate and necessary.** Critical consciousness should be disruptive to the status quo. Such trouble is undoubtedly what was intended in the formulation of critical consciousness by Friere [42]. Critical consciousness finds its apogee in action, and action against an oppressive status quo, intentional or not, will be disruptive. In the workplace, for that disruption to be meaningful, it must be coupled with the capacity of workers to stop, or at least pause, work. That is to say, **workplace democracy includes the possibility of refusal**, and that possibility increases as we develop our critical consciousness. The possibility of refusal, thus, is a characteristic of workplace democracy that other researchers and practitioners can use in their interpretations of the workplace.

Of late, refusal has emerged as a vibrant theme in feminist data and data science approaches [43, 44, 104], with the Feminist Manifest-No as a prime example [21]. Refusal is also appearing in participatory design research and practice [87]. Data for Black Lives advocates for abolishing big data [69], which we might consider another form of refusal. In the context of DataWorks and similar programs, the disruption to the status quo brought about by critical consciousness and refusal will affect the program, the work, relationships, and the research. Accepting that disruption, and working within it, is necessary to sustain a commitment to a fair work environment that centers the workers.

What the workers at DataWorks engaged in is a tentative refusal, akin to a work stoppage. They did not categorically refuse to work on AAVE projects. Instead, the actions of the workers are those of the killjoy [4] in the context of research that Parvin and Pollack [76] call out as necessary: to not give in to the aspirations of researchers or technology advocates, or accept so-called “unintended consequences” but to call attention to problematic situations and push back. In the face of perceived harm, the workers did not simply abide by the researchers’ good intentions nor the possibility that these data sets and algorithms might be put to good use. Rather, the workers at DataWorks refused to continue working on the project until a series of conditions they authored were met. This turn of events is vital to understanding and fostering critical consciousness because it demonstrates the move from reflection to motivation and action: a move from awareness (which was achieved in the prior workshops) to developing a sense of agency to address oppressive conditions, and then working to change oppressive conditions.

From the outset, the goal of the workshops was to assist the workers in developing critical perspectives on data and algorithms in relation to their daily work at DataWorks. While the development of such perspectives might not be entirely surprising, the expression of this critical consciousness through halting a project and requesting the client make changes to a project to lessen harm was unexpected. What is important and distinctive is the move beyond awareness to action, from pointing out the significant problems with fundamental assumptions underlying the project to engaging in the act of refusal.

In demonstrating an arc and expression of critical consciousness, the AAVE project also brought to the surface the tensions that critical consciousness creates in the workplace. While some research has examined how critical consciousness and critical reflection may complicate the workplace [23, 84], less scholarship addresses how such complications might be productive for democracy. If we take democracy seriously (which we do), we must recognize and appreciate that a lack of tension does not characterize democracy. Indeed, while such tensions might be disruptive to the normal operations of the workplace, and therefore cast unfavorably by some, such tensions are essential to democracy. Following from perspectives of agonistic pluralism, democracy is characterized by the capacity for contestation [52, 71, 72]. Over the past two decades, theories of agonism have significantly influenced Participatory Design; troubling research and practice by challenging assumptions about democracy and participation [61]. If the goal of democracy is not to arrive at consensus, then instead of designing tools or environments to hasten consensus, research suggests there is value in constructing events where contestation and dissensus can foment and be expressed. This is a different mode of participation than was previously standard to participatory design. Much of the work on agonism in participatory design focuses on community and government settings [12, 46, 90]. However, it could just as well be explored within the workplace, especially given the longstanding commitment of participatory design to work. Another characteristic of workplace democracy that researchers and practitioners can use in their interpretations of the workplace is agonism: a democratic workplace is one that allows for agonism. The question then becomes what tools and processes we might co-create to foster work environments where workers have the agency to challenge and reshape those environments?

5.2 Blending Context and Action

In discussing the issues of long-term and embedded participatory design in the Global Women's Fund, Trigg and Ishimaru [99] draw upon a set of distinctions first put forward by Gärtner and Wagner [45]. Those are the distinctions between Designing Work, Designing Organizational Frameworks for Action, and Designing the Industrial Relations Context. What is often meant by Designing Work is cooperatively designing the activities of work through cooperatively designing the tools used to perform the work. In the project described in this paper, supporting workplace democracy occurred through an oscillation between the other two categories: Designing Organizational Frameworks for Action and Designing the Industrial Relations Context. This, in turn, connects us with the extension of participatory design into institutions and organizations as sites of design.

Designing Organizational Frameworks for Action "refers to the explicit and tacit norms on which organizational actors base their broader decisions" [45]. In DataWorks, these norms included an emergent critical consciousness of data developed through the workshop. In other words, the organizational framework for action was—loosely—critical consciousness, and the design was in the workshop's structure, content, and enactment. At the start of the AAVE project, critical consciousness comprised awareness of the potential harms of data and algorithms and a sense of agency. But that awareness and agency had not been put into action because the workers at DataWorks had not yet encountered any substantially problematic jobs. It was only in the encounter with the AAVE project's offensive, biased, and harmful content that the move to action occurred, as the workers reacted to, and in particular against, that content. This reaction against is characteristic of agonism as described in the preceding section. This reaction against also surfaced a missing element of the workshops: a process for responding. Without such a process, the workers and management did what was most expedient and practical: to stop the work. The encounters and activities that followed, the meetings with the researchers, the authoring of requirements, and the request for a data use agreement blends into Designing the Industrial Relations Context. Through that process, the organizational framework for action that was loosely critical consciousness

becomes much more fitted: the general qualities of critical consciousness are made specific to the context of data work at DataWorks.

Designing the Industrial Relations Context refers to “where legal and political frameworks are negotiated that govern the relations among industrial sectors” [45]. The process of working through the problematic conditions of the AAVE project and the eventual resolution of those conditions is an example of this negotiation work. In the AAVE project, the (industrial) relations are between the workers, the research staff and directors, and the clients, which in this case were other researchers. Those relations are power relations, making this an explicitly political endeavor [10]. What was negotiated in the AAVE project were the standards of working conditions. This included specific actions of data preparation taken by the client to remove offensive terms and phrases, an in-depth explanation of the work and the intended value of the work from the client to the workers, and the client signing a data use agreement written by the workers to set boundaries for the use of the data. These requests and structures are essential because they assert respect and authority to the workers. Such respect and authority are usually absent from data work. Establishing these conditions as standards advances workplace democracy by acknowledging the workers’ value and expertise. The process, more generally, demonstrates a shift in power relations between workers and clients. Of course, this shift is bounded and tenuous, but it is nonetheless a moment of achievement through contention.

At the same time, we should acknowledge that describing this process as *design* is a stretch. We intended the workshops to develop critical perspectives, but we did not plan for the AAVE project to unfold as it did. Indeed, as mentioned, we did not have a plan for responding to such problematic conditions within the context of a DataWorks project. As we consider this move into Designing the Industrial Relations Context through DataWorks we must acknowledge how we—as researchers—are responsible for and imbricated in this context and the challenges this raises due to contrasting positionalities.

5.3 Institutions and Positionalities

As an experimental workplace, DataWorks is a space to explore how else we might structure data work, and what sorts of institutions and organizations we might create to support more varied approaches to data science and computing. We—as researchers—are part of that exploration; we are part of those institutions and organizations. So, reflecting on how our work cultures and practices align and contrast with DataWorks and the workers is appropriate and important. Such reflexivity follows work in participatory design that examines the roles of design in navigating and influencing institutions [53, 98] as well as work in computing research calling for attending to the positionality of researchers [58, 75, 81].

One challenge we faced throughout the AAVE project was the instability that critical consciousness introduces into the workplace. The starting, stopping, restarting, stopping again, and then restarting significantly disrupted the workflow in DataWorks. We had to reschedule client work and find other work to fill the days and weeks during the stoppages. In addition to this being an administrative difficulty, it also threatened to become a financial difficulty. Although there is a reserve operating budget we can draw upon so that we do not have to be on client work constantly, we do need to consistently complete contract work. As the work was delayed, we grew concerned, and the situation created a tension between the interests of the workers and the researchers. While that seemingly did not inhibit the expressions and actions of the workers, they were aware of the disruption, difficulty, and pressure this put upon DataWorks. Such problems are documented in the literature on critical reflection and critical consciousness in the workplace [23, 84]. Taking the cue to reflect critically upon our interpretation of the situation, we—the academics—were operating from a perspective that prioritized predictability. But if we are committed to democracy

in the workplace, such managerial perspectives must be balanced with those that acknowledge the contestation inherent to democracy and we must continue to hold space for that contestation.

The process of stopping the work, bringing in the academic client, and the negotiations laid bare power dynamics within academic research specifically, and more generally, between those who request and those who perform data services. Outside of established formal institutions such as the Institutional Review Board, researchers are rarely directly questioned about the ethics of their work in conducting their research. Questions may be raised upon publication or other forms of dissemination, but rarely as the work is being done. Thus, the workers' resistance to the data work challenged not just the content and activity of the work, but also the traditions and routines of academic research and the power hierarchies within those traditions and routines. To the credit of the academic client, they met these challenges with openness: one reason they chose to employ DataWorks for the labeling was to work with data workers who would bring their positionality to the interpretation of the data. This returns us to an initial motivation for DataWorks: to engage the university as a distinctive institutional site and in the process challenge, and attempt to shift, if ever so slightly, university cultures. In the context of higher education institutions, part of the work of workplace democracy is to push against the privilege of the university and the de facto authority given to academics, and instead work towards respecting the diversity of labor, skill, and expertise within and beyond the university.

While the data workers had the agency and motivation to act, they did not, *at the time*, have full authority or access in the university setting. For instance, before these projects, the workers interacted with clients but less regularly; most interactions with the clients were handled by a manager who was separate from the workers. So, in the events described in this paper, the research faculty moved the emergent critical consciousness of the workers into the organizational structure and culture of the institution and served as initial facilitators with the client. In contemporary participatory design literature, this work is referred to as “institutioning.” [53, 98] Huybrechts, Benesch, and Geib describe institutioning as “a practice of interweaving between—as well as producing—various insides and outsides... consolidating and challenging existing institutional frames as well as forming new ones” [53]. In the case of the AAVE project, this institutioning involved bringing the workers and clients together, holding the space for contention, and assisting in establishing a process that foregrounded the workers in that project and going forward. This is not to valorize the work of the research staff or management—we made many mistakes. Rather, it calls attention to the necessary transformation of the role of the researcher in contexts such as DataWorks. Such institutioning requires more than being an ally: it requires becoming an accomplice [1].

Throughout the workshops and the AAVE project, the Whiteness of most of the research staff was a persistent tension. We attempted to ground the critical data literacy workshops in the lived experience of the data workers, who at the time were all Black. Similarly, we approached this work as what we perceived to be an opportunity for design justice: an opportunity to approach data work and the development of sophisticated data sets and algorithms from the perspective of “nothing about us, without us” [22]. What we did not consider was how such experiences would be stressful, perhaps even traumatizing to the workers. **This was a failure of ours.** A few months after the AAVE project, the DataWorks manager left to take another job, and we replaced them by promoting one of the DataWorks workers, thereby putting a worker and person of color in a direct leadership role. As we iterate on the workshop materials, we are collaborating with the workers to select the sample data sets and examples that we draw upon, and future instantiations of the workshop will be co-taught with a worker. Acknowledging these problems and their consequences is crucial both to be more just in our research and also to further workplace democracy: we cannot solve

the problems of Whiteness through projects such as DataWorks. Still, we can, and we must, hold ourselves accountable for how our work furthers these problems [2, 27, 57].

5.4 Forever Changing and Moving Forward

As an experimental workplace, DataWorks is forever changing. We—the research faculty and the workers—are constantly learning from each other and our collective experiences with DataWorks, adjusting and adapting how we work and what work we do. The experiences of the AAVE project have had significant effects on DataWorks, impacting our structure and processes.

The AAVE project provided an opportunity to reflect on the manager's role and consider how we might continue to transition more authority and power to the workers. We decided from then onwards to promote a manager from among the workers. The first such manager was, in fact, one of the workers who had most strongly led in resisting the AAVE work. That manager left approximately six months later, after securing another position (a success for DataWorks). A new manager was then appointed from the current workers. This process of promoting data workers as managers puts the workers in direct control of their daily and weekly tasks, as well as putting a worker in the first line of decision making regarding new clients and new projects.

We have also formalized the workshop curriculum, making it a component of the training process for all new data workers. As new data workers are brought in DataWorks and are learning basic skills in data wrangling, that learning is coupled with the workshop activities introducing workers to critical perspectives on data and machine learning, and uses the AAVE project as an example. That workshop curriculum is freely available online at <https://dataworkforce.gatech.edu/data-tools/critical-data-literacy/>. As part of formalizing the curriculum, we have also changed the examples and presentation of examples to address the trauma they might elicit, and we have provided the space within the workshop to better support the affective dimensions of critical consciousness through discussion and reflection.

We have been less successful in creating a standardized review process or binding and enforceable contracts. When a client approaches DataWorks, we request an explanation of their intended use of the data. This is presented by the manager (now a former data worker) to the other data workers. The workers, together with the manager, can raise concerns about a project and choose not to pursue a project. For instance, since the project described in this paper, DataWorks has been presented with several other AAVE projects, which have been declined. Because each project is so different, checklist-style reviews are inadequate. Instead, we approach each project individually and assess it through discussion. Binding and enforceable contracts are also elusive. In large part, this is because there are scant data policies that would give substance to such a contract. Nonetheless, we do continue with data agreements, if for no other reason, because they are statements on behalf of the data workers at DataWorks of what they value as being the responsible use of data they collect and process.

6 CONCLUSION

Critical consciousness and workplace democracy can be nebulous concepts. This paper presented a case of the development of critical consciousness and workplace democracy within an experimental workplace. In doing so, we contributed ethnographic descriptions of these concepts in action and provided interpretations that connect to themes in participatory design and critical studies of data work. Specifically, we identified the role of agonism in the democratic workplace and described acts of refusal as part of workplace democracy, expanding the repertoire of agonistic data practices [26]. We also discussed the blending of designing organizational frameworks for action and designing the industrial relations context in data work. Finally, we addressed how multiple and varied positionalities affect this work. In future work, we plan to delve further into

these positionalities and the institutional labor of establishing and sustaining workplaces, such as DataWorks, that strive to be just and fair within compromised contexts [64].

Of course, settings like DataWorks are different from most data work environments. DataWorks is an intentional exploration of how data work might be structured and performed otherwise, in ways that center the workers and their lived experiences. It is also set within a university, which is a particular institutional structure and culture. Nonetheless, the setting of DataWorks, and the lessons learned from such experimental workplaces are valuable to help us, collaboratively and collectively, understand what is possible, desirable, and undesirable. As researchers across the fields of human-centered computing, critical data studies, and design collaborate with workers in different ways to investigate and co-create more varied and equitable work environments, we hope that the descriptions, interpretations, and themes in this paper contribute to those efforts.

ACKNOWLEDGMENTS

Thank you to all of the workers at DataWorks, past and present. Support for this research was provided by the Public Interest Technology University Network Challenge Fund, a fiscally sponsored project of New Venture Fund, and the National Science Foundation under Grant No. 1951818 and the Constellations Center for Equity and Computing at Georgia Institute of Technology.

REFERENCES

- [1] Indigenous Action. 2014. Accomplices not allies: Abolishing the ally industrial complex. *Indigenous Action Media* 2 (2014), 2.
- [2] Sara Ahmed. 2007. A phenomenology of whiteness. *Feminist theory* 8, 2 (2007), 149–168.
- [3] Sara Ahmed. 2012. *On Being Included: Racism and Diversity in Institutional Life*. Duke University Press.
- [4] Sara Ahmed. 2016. CONCLUSION 2. A Killjoy Manifesto. In *Living a Feminist Life*. Duke University Press, 251–268.
- [5] Sara Ahmed. 2021. *Complaint!* Duke University Press.
- [6] Ifeoma Ajunwa, Sareeta Amrute, Lilly Irani, Winifred R Poster, and Meg Stalcup. 2021. Tech firms need Black AI scholars and labour rights. *Nature* 590, 7846 (2021), 389–390.
- [7] Saul Alinsky. 1989. *Rules for radicals: A pragmatic primer for realistic radicals*. Vintage.
- [8] Ali Alkhatib, Michael S. Bernstein, and Margaret Levi. 2017. Examining Crowd Work and Gig Work Through The Historical Lens of Piecework. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 4599–4616. <https://doi.org/10.1145/3025453.3025974>
- [9] Hannah Arendt. 2013. *The human condition*. University of Chicago press.
- [10] Eevi E Beck. 2002. P for political: Participation is not enough. *Scandinavian Journal of Information Systems* 14, 1 (2002), 1.
- [11] Rahul Bhargava, Erica Deahl, Emmanuel Letouzé, Amanda Noonan, David Sangokoya, and Natalie Shoup. 2015. Beyond data literacy: Reinventing community engagement and empowerment in the age of data. (2015).
- [12] Erling Björgvinsson, Pelle Ehn, and Per-Anders Hillgren. 2012. Agonistic participatory design: working with marginalised social movements. *CoDesign* 8, 2-3 (2012), 127–144.
- [13] Jeanette Blomberg and Helena Karasti. 2012. Ethnography: Positioning ethnography within participatory design. In *Routledge international handbook of participatory design*. Routledge, 86–116.
- [14] Jeanette Blomberg and Helena Karasti. 2013. Reflections on 25 years of ethnography in CSCW. *Computer supported cooperative work (CSCW)* 22, 4 (2013), 373–423.
- [15] Susanne Bødker and Morten Kyng. 2018. Participatory design that matters—Facing the big issues. *ACM Transactions on Computer-Human Interaction (TOCHI)* 25, 1 (2018), 1–31.
- [16] Claus Bossen, Christian Dindler, and Ole Sejer Iversen. 2016. Evaluation in participatory design: a literature survey. In *Proceedings of the 14th Participatory Design Conference: Full papers-Volume 1*. 151–160.
- [17] Tone Bratteteig and Ina Wagner. 2016. What is a participatory design result?. In *Proceedings of the 14th Participatory Design Conference: Full papers-Volume 1*. 141–150.
- [18] Stephen Brookfield. 2002. Overcoming alienation as the practice of adult education: The contribution of Erich Fromm to a critical theory of adult learning and education. *Adult Education Quarterly* 52, 2 (2002), 96–111.
- [19] Stephen Brookfield. 2018. Developing critically reflective practitioners: A rationale for training educators of adults. In *Training educators of adults*. Routledge, 317–338.

- [20] Kathy Charmaz. 2006. *Constructing grounded theory: A practical guide through qualitative analysis*. sage.
- [21] Garcia P. Cowan T.L. Rault J. Sutherland T. Chan A. Rode J. Hoffmann A.L. Salehi N. Nakamura L. Cifor, M. 2019. Feminist Data Manifest-No.
- [22] Sasha Costanza-Chock. 2020. *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
- [23] Richard J Cotter. 2014. Reflexive spaces of appearance: Rethinking critical reflection in the workplace. *Human Resource Development International* 17, 4 (2014), 459–474.
- [24] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- [25] Kimberlé Crenshaw. 1989. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *u. Chi. Legal f.* (1989), 139.
- [26] Roderic Crooks and Morgan Currie. 2021. Numbers will not save us: Agonistic data practices. *The Information Society* 37, 4 (2021), 201–213.
- [27] Girish Daswani. 2020. On the Whiteness of Academia. *Everyday Orientalism* (2020). <https://everydayorientalism.wordpress.com/2021/02/17/on-the-whiteness-of-academia/>
- [28] Richard Delgado and Jean Stefancic. 2023. *Critical race theory: An introduction*. Vol. 87. NyU press.
- [29] Kayla Desportes, Kathleen McDermott, Yoav Bergner, and William Payne. 2022. “Go [ing] hard... as a woman of color”: A case study examining identity work within a performative dance and computing learning environment. *ACM Transactions on Computing Education (TOCE)* 22, 4 (2022), 1–29.
- [30] Matthew A Diemer, Luke J Rapa, Adam M Voight, and Ellen H McWhirter. 2016. Critical consciousness: A developmental approach to addressing marginalization and oppression. *Child Development Perspectives* 10, 4 (2016), 216–221.
- [31] Catherine D’ignazio and Lauren F Klein. 2020. *Data feminism*. MIT press.
- [32] Tawanna R Dillahunt, Nishan Bose, Suleman Diwan, and Asha Chen-Phang. 2016. Designing for disadvantaged job seekers: Insights from early investigations. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 905–910.
- [33] Tawanna R Dillahunt, Matthew Garvin, Marcy Held, and Julie Hui. 2021. Implications for Supporting Marginalized Job Seekers: Lessons from Employment Centers. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–24.
- [34] Tawanna R Dillahunt, Alex Jiahong Lu, Aarti Israni, Ruchita Lodha, Savana Brewer, Tiera S Robinson, Angela Brown Wilson, and Earnest Wheeler. 2022. The Village: Infrastructuring Community-based Mentoring to Support Adults Experiencing Poverty. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [35] Lynn Dombrowski, Adriana Alvarado Garcia, and Jessica Despard. 2017. Low-wage precarious workers’ sociotechnical practices working towards addressing wage theft. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4585–4598.
- [36] Pelle Ehn. 2017. Learning in participatory design as I found it (1970–2015). In *Participatory Design for Learning*. Routledge, 7–21.
- [37] Elizabeth Ellsworth. 1989. Why doesn’t this feel empowering? Working through the repressive myths of critical pedagogy. *Harvard educational review* 59, 3 (1989), 297–325.
- [38] Sheena Erete, Aarti Israni, and Tawanna Dillahunt. 2018. An intersectional approach to designing in the margins. *Interactions* 25, 3 (2018), 66–69.
- [39] Sheena Erete, Yolanda A Rankin, and Jakita O Thomas. 2022. A method to the madness: Applying an intersectional analysis of structural oppression and power in HCI and design. *ACM Transactions on Computer-Human Interaction* (2022).
- [40] Tara Fenwick. 2005. Conceptions of critical HRD: Dilemmas for theory and practice. *Human Resource Development International* 8, 2 (2005), 225–238.
- [41] Janice R Foley and Michael Polanyi. 2006. Workplace democracy: Why bother? *Economic and Industrial Democracy* 27, 1 (2006), 173–191.
- [42] Paulo Freire. 1970. *Pedagogy of the oppressed*. Herder Herder.
- [43] Patricia Garcia, Tonia Sutherland, Marika Cifor, Anita Say Chan, Lauren Klein, Catherine D’Ignazio, and Niloufar Salehi. 2020. No: Critical refusal as feminist data practice. In *conference companion publication of the 2020 on computer supported cooperative work and social computing*. 199–202.
- [44] Patricia Garcia, Tonia Sutherland, Niloufar Salehi, Marika Cifor, and Anubha Singh. 2022. No! Re-imagining Data Practices Through the Lens of Critical Refusal. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–20.
- [45] Johannes Gartner and Ina Wagner. 1996. Mapping actors and agendas: Political frameworks of systems design and participation. *Human–Computer Interaction* 11, 3 (1996), 187–214.

- [46] Amanda Anne Geppert and Laura Ellen Forlano. 2022. Design for equivalence: Agonism for collective emancipation in participatory design. In *Proceedings of the Participatory Design Conference 2022-Volume 1*. 158–168.
- [47] Barney G Glaser and Anselm L Strauss. 2017. *Discovery of grounded theory: Strategies for qualitative research*. Routledge.
- [48] Mary L Gray and Siddharth Suri. 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- [49] Nicolai Brodersen Hansen, Christian Dindler, Kim Halskov, Ole Sejer Iversen, Claus Bossen, Ditte Amund Basballe, and Ben Schouten. 2019. How participatory design works: mechanisms and effects. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction*. 30–41.
- [50] Gillian R Hayes. 2011. The relationship of action research to human-computer interaction. *ACM Transactions on Computer-Human Interaction (TOCHI)* 18, 3 (2011), 1–20.
- [51] Per-Anders Hillgren, Anna Seravalli, and Mette Agger Eriksen. 2016. Counter-hegemonic practices; dynamic interplay between agonism, commoning and strategic design. *Strategic Design Research Journal*, 9 (2): 89–99 May-August 2016 (2016).
- [52] Bonnie Honig. 2016. Political theory and the displacement of politics. In *Political Theory and the Displacement of Politics*. Cornell University Press.
- [53] Liesbeth Huybrechts, Henric Benesch, and Jon Geib. 2017. Institutioning: Participatory design, co-design and the public realm. *CoDesign* 13, 3 (2017), 148–159.
- [54] Ivan Illich. 1968. To hell with good intentions. *Combining service and learning: A resource book for community and public service* 1 (1968), 314–320.
- [55] Lilly C Irani and M Six Silberman. 2013. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 611–620.
- [56] Alexis Jemal. 2017. Critical consciousness: A critique and critical analysis of the literature. *The Urban Review* 49, 4 (2017), 602–626.
- [57] Azeezat Johnson. 2018. An academic witness: White supremacy within and beyond academia. *The fire now: Anti-racist scholarship in times of explicit racial violence* (2018), 15–25.
- [58] Britney Johnson, Ben Rydal Shapiro, Betsy DiSalvo, Annabel Rothschild, and Carl DiSalvo. 2021. Exploring Approaches to Data Literacy Through a Critical Race Theory Perspective. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [59] Finn Kensing and Joan Greenbaum. 2012. Heritage: Having a say. In *Routledge international handbook of participatory design*. Routledge, 41–56.
- [60] Yasmine Kotturi, Herman T Johnson, Michael Skirpan, Sarah E Fox, Jeffrey P Bigham, and Amy Pavel. 2022. Tech Help Desk: Support for Local Entrepreneurs Addressing the Long Tail of Computing Challenges. In *CHI Conference on Human Factors in Computing Systems*. 1–15.
- [61] Helena Kraft. 2020. A critical exploration of agonistic participatory design. *The Design Journal* 23, 1 (2020), 31–48.
- [62] Airi Lampinen, Moira McGregor, Rob Comber, and Barry Brown. 2018. Member-Owned Alternatives: Exploring Participatory Forms of Organising with Cooperatives. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–19.
- [63] Karen Levy, Kyla E Chasalow, and Sarah Riley. 2021. Algorithms and decision-making in the public sector. *Annual Review of Law and Social Science* 17 (2021), 309–334.
- [64] Max Liboiron. 2017. Compromised agency: The case of BabyLegs. *Engaging Science, Technology, and Society* 3 (2017), 499–527.
- [65] Robert Mayer. 2001. Robert Dahl and the right to workplace democracy. *The Review of Politics* 63, 2 (2001), 221–247.
- [66] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 460 (nov 2022), 37 pages. <https://doi.org/10.1145/3555561>
- [67] Milagros Miceli, Julian Posada, and Tianling Yang. 2022. Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 34 (jan 2022), 14 pages. <https://doi.org/10.1145/3492853>
- [68] Milagros Miceli, Martin Schuessler, and Tianling Yang. 2020. Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–25.
- [69] Yeshimabeit Milner. 2019. Abolish big data. *University of California, Irvine* 8 (2019).
- [70] Jennifer A Moon. 2013. *Reflection in learning and professional development: Theory and practice*. Routledge.
- [71] Chantal Mouffe. 1999. Deliberative democracy or agonistic pluralism? *Social research* (1999), 745–758.
- [72] Chantal Mouffe. 2013. *Agonistics: Thinking the world politically*. Verso Books.
- [73] José Esteban Muñoz. 2019. Cruising utopia. In *Cruising Utopia, 10th Anniversary Edition*. New York University Press.

- [74] Labor Tech Reserch Network. 2023. Labor Tech Research Network.
- [75] Ihudiya Finda Ogbonnaya-Ogburu, Angela DR 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*. 1–16.
- [76] Nassim Parvin and Anne Pollock. 2020. Unintended by design: On the political uses of “Unintended Consequences”. *Engaging science, technology, and society* 6 (2020), 320–327.
- [77] Lucy Pei, Marisol Wong-Villacres, Sheena Erete, Daniela Rosner, Alex Taylor, and Mikael Wiberg. 2022. Who we are and what we have: Designing with minoritized communities. *interactions* 29, 5 (2022), 5–5.
- [78] B Perrigo. 2023. OpenAI Used Kenyan Workers on Less Than 2 Per Hour to Make ChatGPT Less Toxic. *Jacobin* (2023). <https://time.com/6247678/openai-chatgpt-kenya-workers/>
- [79] 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*. 1–11.
- [80] Yolanda A Rankin and Kallayah K Henderson. 2021. Resisting Racism in Tech Design: Centering the Experiences of Black Youth. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–32.
- [81] Yolanda A Rankin, Jakita O Thomas, and Nicole M Joseph. 2020. Intersectionality in HCI: Lost in translation. *Interactions* 27, 5 (2020), 68–71.
- [82] Brianna Richardson and Juan E Gilbert. 2021. A Framework for Fairness: A Systematic Review of Existing Fair AI Solutions. *arXiv preprint arXiv:2112.05700* (2021).
- [83] Clare Rigg and Kiran Trehan. 2004. Reflections on working with critical action learning. *Action Learning: Research and Practice* 1, 2 (2004), 149–165.
- [84] Clare Rigg and Kiran Trehan. 2008. Critical reflection in the workplace: is it just too difficult? *Journal of European Industrial Training* (2008).
- [85] Sarah T Roberts. 2019. *Behind the screen*. Yale University Press.
- [86] Toni Robertson and Jesper Simonsen. 2012. Participatory Design: an introduction. In *Routledge international handbook of participatory design*. Routledge, 1–17.
- [87] Sarah Robinson, Nicola J. Bidwell, Laura Maye, Nadia Pantidi, and Conor Linehan. 2020. Participation through substituting and refusing. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 143–147.
- [88] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [89] Rafi Santo, Leigh Ann DeLyser, June Ahn, Anthony Pellicone, Julia Aguiar, and Stephanie Wortel-London. 2019. Equity in the who, how and what of computer science education: K12 school district conceptualizations of equity in ‘cs for all’ initiatives. In *2019 research on equity and sustained participation in engineering, computing, and technology (RESPECT)*. IEEE, 1–8.
- [90] Nitin Sawhney and Anh-Ton Tran. 2020. Ecologies of contestation in participatory design. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 1*. 172–181.
- [91] Anna Seravalli. 2014. *Making Commons: attempts at composing prospects in the opening of production*. Ph.D. Dissertation. Malmö University.
- [92] Amanda Sinclair. 2007. Teaching leadership critically to MBAs: Experiences from heaven and hell. *Management Learning* 38, 4 (2007), 458–472.
- [93] Rachel Charlotte Smith and Ton Otto. 2020. Cultures of the future: Emergence and intervention in design anthropology. In *Design anthropological futures*. Routledge, 19–36.
- [94] Thomas Smyth and Jill Dimond. 2014. Anti-oppressive design. *Interactions* 21, 6 (2014), 68–71.
- [95] Katta Spiel, Christopher Frauenberger, and Geraldine Fitzpatrick. 2017. Experiences of autistic children with technologies. *International Journal of Child-Computer Interaction* 11 (2017), 50–61.
- [96] James P Spradley. 2016. *Participant observation*. Waveland Press.
- [97] JS Stan. 2018. Tech Workers Need to Keep Organizing. *Jacobin* (2018). https://jacobin.com/2018/11/tech-worker-organizing-google-union-walkout?fbclid=IwAR25-l3YuHgEAOS3Ewt34a8711amwkALTQtJqLXhe9VgfNerj2Tl1TF_KW8/
- [98] Maurizio Teli, Marcus Foth, Mariacristina Sciannamblo, Irina Anastasiu, and Peter Lyle. 2020. Tales of institutioning and commoning: participatory design processes with a strategic and tactical perspective. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 1*. 159–171.
- [99] Randy Trigg and Karen Ishimaru. 2012. Integrating participatory design into everyday work at the Global Fund for Women. *Routledge International Handbook of Participatory Design* (2012), 233–254.
- [100] Alan Freihof Tygel and Rosana Kirsch. 2016. Contributions of Paulo Freire for a critical data literacy: A popular education approach. *The Journal of Community Informatics* 12, 3 (2016).

- [101] Meredith Whittaker. 2023. Origin Stories: Plantations, Computers, and Industrial Control. *Logic(s)* 19 (2023). <https://logicmag.io/supa-dupa-skies/origin-stories-plantations-computers-and-industrial-control/>
- [102] Milagros Miceli Williams, Adrienne and Timnit Gebru. 2021. The Exploited Labor Behind Artificial Intelligence. *Neoma* (2021).
- [103] Aman Yadav and Marie K Heath. 2022. Breaking the Code: Confronting Racism in Computer Science through Community, Criticality, and Citizenship. *TechTrends* (2022), 1–9.
- [104] Jonathan Zong and J Nathan Matias. 2023. Data Refusal From Below: A Framework for Understanding, Evaluating, and Envisioning Refusal as Design. *ACM Journal on Responsible Computing* (2023).

Received January 2023; revised July 2023; accepted November 2023