

Documenting Data Production Processes: A Participatory Approach for Data Work

MILAGROS MICELI, DAIR Institute, Technische Universität Berlin, and Weizenbaum Institute, Germany TIANLING YANG, Technische Universität Berlin & Weizenbaum Institute, Germany ADRIANA ALVARADO GARCIA^{*}, IBM Thomas J. Watson Research Center, United States JULIAN POSADA[†], Yale University, United States SONJA MEI WANG, Technische Universität Berlin, Germany MARC POHL, Independent Researcher, Germany ALEX HANNA, DAIR Institute, United States

The opacity of machine learning data is a significant threat to ethical data work and intelligible systems. Previous research has addressed this issue by proposing standardized checklists to document datasets. This paper expands that field of inquiry by proposing a shift of perspective: from documenting datasets towards documenting data production. We draw on participatory design and collaborate with data workers at two companies located in Bulgaria and Argentina, where the collection and annotation of data for machine learning are outsourced. Our investigation comprises 2.5 years of research, including 33 semi-structured interviews, five co-design workshops, the development of prototypes, and several feedback instances with participants. We identify key challenges and requirements related to the integration of documentation practices in real-world data production scenarios. Our findings comprise important design considerations and highlight the value of designing data documentation based on the needs of data workers. We argue that a view of documentation as a boundary object, i.e., an object that can be used differently across organizations and teams but holds enough immutable content to maintain integrity, can be useful when designing documentation to retrieve heterogeneous, often distributed, contexts of data production.

CCS Concepts: • Human-centered computing \rightarrow Empirical studies in collaborative and social computing; *Collaborative and social computing design and evaluation methods*.

Additional Key Words and Phrases: dataset documentation, data production, data work, data labeling, data annotation, machine learning, transparency

ACM Reference Format:

Milagros Miceli, Tianling Yang, Adriana Alvarado Garcia, Julian Posada, Sonja Mei Wang, Marc Pohl, and Alex Hanna. 2022. Documenting Data Production Processes: A Participatory Approach for Data Work. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 510 (November 2022), 34 pages. https://doi.org/10.1145/3555623

Authors' addresses: Milagros Miceli, m.miceli@tu-berlin.de, DAIR Institute, Technische Universität Berlin, and Weizenbaum Institute, Berlin, Germany; Tianling Yang, tianling.yang@tu-berlin.de, Technische Universität Berlin & Weizenbaum Institute, Berlin, Germany; Adriana Alvarado Garcia, adriana.ag@ibm.com, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, United States; Julian Posada, julian.posada@yale.edu, Yale University, New Haven, CT, United States; Sonja Mei Wang, s.wang.1@tu-berlin.de, Technische Universität Berlin, Berlin, Germany; Marc Pohl, marc@zweipohl.de, Independent Researcher, Berlin, Germany; Alex Hanna, alex@dair-institute.org, DAIR Institute, United States.



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

510

© 2022 Copyright held by the owner/author(s). 2573-0142/2022/11-ART510 https://doi.org/10.1145/3555623

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

^{*}Also with Georgia Institute of Technology, United States.

[†]Also with University of Toronto, Canada.

1 INTRODUCTION

In recent years, research has turned its attention toward the datasets used to train and validate machine learning (ML) models [28, 29, 78, 91, 94, 95, 98]. Work in this area [6, 10, 37, 38, 46] has proposed frameworks to document and disclose datasets' origins, purpose, and characteristics to increase transparency, help understand models' functioning, and anticipate ethical issues comprised in data. Frameworks such as Datasheets for Datasets [37], Dataset Nutrition Label [46], and Data Statements [10] have become increasingly influential. These forms of documentation are useful for providing a snapshot of a dataset's state at a given moment. Still, one persistent challenge is that datasets are living artifacts [8, 51, 94] that evolve and change over time. Moreover, their production involves actors with varying amounts of decision-making power. Capturing such variations requires, as we will argue, a shift of perspective: *from documenting datasets to documenting data production processes.*

In this paper, we understand *documentation* as a tool and the process of making explicit the contexts, actors, and practices comprised in ML data work, as well as the relationships among these elements. With the term *data work*, we refer to tasks that involve the generation, annotation, and verification of data [68]. We understand *data* as a process [72], not a fact, and lean on CSCW and HCI work that has focused on data work in the ML pipeline [49, 69, 90, 91], studying the socio-economic situation of data workers [34, 43, 87] and the power relations involved in the labor process [50, 64, 89]. Moreover, this paper expands our previous work [67, 69, 70] that studies a segment of data production that is outsourced through business process outsourcing companies (BPOs) that employ data workers from marginalized populations, mostly in developing countries. Unlike crowdsourcing platforms, where algorithms manage the labor process [86, 102, 116], BPOs are characterized by traditional and localized managerial structures and business hierarchies [67].

Building on that work, the focus of the present paper is exploring ways of *making distributed processes of data production more reflexive and participatory through documentation*. Promoting reflexivity through documentation means fostering collaborative documentation practices, including feedback loops, iterations, and the co-creation of tasks. Such considerations can help produce documentation that serves to interrogate taken-for-granted classifications and hierarchies in data work [70].

We report on 2.5 years of research comprising 33 semi-structured interviews, five co-design workshops, and several feedback instances where preliminary findings and prototypes were discussed with participants. We draw on participatory design to work closely with two BPO companies located in Argentina and Bulgaria. We focus on the perspectives of data workers at both research sites and identify the challenges and requirements of integrating documentation in real-world data production scenarios.

To that end, our research questions are:

- (1) How can documentation reflect the iterative and collaborative nature of data production processes?
- (2) How can documentation practices contribute to mitigating the information asymmetries present in data production and ML supply chains?
- (3) What information do data workers need to perform their tasks in the best conditions, and how can it be included in the documentation?

Our findings show how the participants prioritized documentation that is collaborative and circular, i.e., documentation that is able to transport information about tasks, teams, and payment to data workers and communicate workers' feedback back to the requesters. We discuss design considerations related to the integration of documentation practices as integral part of data production, questions of access and trust, and the differentiation between the creation and use of

documentation. We reflect on these observations and the ideas of our participants through the acknowledgment that documentation is a *boundary object* [100, 101], i.e., an artifact that can be used differently across organizations and teams but holds enough immutable content to maintain integrity. We further discuss a series of design and research implications that the acknowledgment of documentation as a form of boundary object entails.

This investigation extends previous work on dataset documentation in two aspects. First, instead of documenting datasets, we focus on the *documentation of data production processes*, particularly the intra- and inter-organizational coordination comprised in data production, taking into account the iterative character of such collaboration and its impacts on data practices and datasets. This type of documentation can help break with patterns of information gatekeeping and top-down communication in data production by creating documentation that enables and reflects feedback loops and iterations. Second, our investigation aims at producing *documentation for and with data workers*. In this sense, we explore documentation forms that enable data workers' participation in the discussion of data production processes. Through this approach, we aim to open up paths for designing documentation that captures the needs and practices of data workers in contexts where their voices and labor tend to remain silenced [42, 50, 64, 68, 69, 81, 89, 108].

2 RELATED WORK

2.1 The Documentation of ML Datasets and Data Pipelines

What counts as documentation is linked to the disciplinary background of each research community. For instance, while the field of open science envisions the documentation of data provenance and protocols toward reproducibility [39, 40], the Fairness, Accountability, and Transparency in the machine learning community is concerned with documenting ML models to prevent bias and improve transparency [71]. In this paper, we approach the documentation of data production by leaning on previous work that has advocated for transparency by proposing innovative frameworks for documenting ML datasets [6, 10, 37, 38, 46]. These frameworks vary in the forms of documentation, such as datasheets, dataset nutrition labels, and data statements, and have different prioritized goals and intended field of use in their design. They share similarities in requiring a standardized set of information for documenting the motivation [37], curation rationale [10], composition [37], provenance [6, 10, 37, 46], collection process and context [10, 37], actors involved [10, 38] and tools used [37, 38]. Some of these groundbreaking frameworks have been successfully integrated and implemented in corporations such as Google and IBM as pilot projects [1, 3], and several ML conferences that now require new dataset submissions **to be** accompanied by a datasheet, as evidenced by the recent NeurIPS Dataset and Benchmark track [2].

The foregoing frameworks focus on the documentation of *the datasets themselves*, offering a useful snapshot of the dataset's current state to dataset producers, users, and researchers. However, as previous work [68] has argued, two fundamental aspects of ML datasets represent a challenge to such standardized documentation frameworks: (1) The collaborative nature of data work and data production [69, 73, 75], including the often messy relationship between stakeholders [70], which challenges the strict differentiation between data subjects, data users, and data producers assumed by standardized documentation frameworks; and (2) the fact that datasets are not static entities but rather evolving projects with flexible boundaries [8, 94].

In view of these challenges, recent research in data documentation has raised attention toward a consideration of data production pipelines, including the multiplicity of workflows and actors involved in producing ML datasets [8, 30, 70, 82, 90]. Hutchinson et al. [47] put forward a set of

questions structured by interconnected stages of the data life cycle. Conceiving this cycle as nonlinear, the authors frame datasets as technical infrastructures and their production processes as "goaloriented engineering," which favors adopting methodologies with deliberation and intentionality to improve transparency and accountability over requiring only post-hoc justification. Relatedly, Balayn et al. [8] introduce a process-oriented lens to dataset documentation and expand the focus to the design of data pipelines and the reasoning behind them. This process lens can help identify and locate knowledge gaps in the pipeline and uncover places and procedures that are worth more rigorous documentation and careful reflection on their potential impact and harm. Other studies [33, 51, 94] have called for the professionalization of data work and care (i.e., the careful consideration of the dataset in terms of "the domain setting where it originates, and the potential questions modeling that data might answer or problems it might solve" [113]), the publication of data documentation, the implementation of institutional frameworks and procedures to promote documentation, and the consideration of data work as an own subfield of research [51, 91, 94].

These investigations shift the perspective to regard datasets as living artifacts that are formalized by actors with different positions and decision-making power. From this perspective, documentation could capture data pipelines and production contexts as well as various actors that shape datasets. At its ideal state, documentation should also be able to foster "a documented culture of reflexivity about values in their [data] work" [94]. As goals, interests, and needs vary across different contexts and organizations, documentation should be flexible to accommodate local specificities and constraints and, at the same time, keep certain common identities. In line with Pushkarna et al. [82], whose work in data documentation reflects the need for disclosure documents to be responsive to a larger body of stakeholders, including annotators, data curators, and users, we argue that recognizing documentation as a *boundary object* can be helpful to approach and reflect on such challenge.

2.2 Boundary Objects

The notion of boundary object, defined as "objects that both inhabit several communities of practice and satisfy the informational requirements of each of them" [15], is a useful lens to study how documentation operates across diverse contexts in the ML supply chain. Not only do boundary objects facilitate communication among different groups, but they also translate and transfer knowledge to be accessible to all potential stakeholders. This way, boundary objects can meet stakeholders' specific information needs and work requirements, but they do not require them to have a panoramic grasp and full comprehension of every detail. The variety of maps of a city is an example of a boundary object. The maps have the same geographical boundary, namely the city, but different groups use them differently. For example, consumption-minded citizens may search information about local parks and restaurants, physical geographers may pay attention to mountains and rivers, botanists to certain horticultural areas and botanical gardens, and administrators may keep an eye on political boundaries. Suchman [105] emphasizes the appropriability of objects into use contexts and working practices. In her analysis of a civil engineering project, the author shows how heterogeneous artifacts were brought together and integrated into everyday work practices [104]. When such "artful integration" [105] involves continuous and relatively stable relationships among different communities and shared objects jointly produced by them, "then boundary objects arise"[15].

Within CSCW and HCI, there is a longstanding use of the boundary object notion to investigate how artifacts can be used across varying contexts and to facilitate cooperation and collaboration, such as in the medical order [118], in the coordination among different actors in hospitals [13, 14], in aircraft technical support [60, 61], in disease management at home [4], and in issue tracking systems [11]. This notion is also used to study organizational memory [5] and common information space [9, 96]. Lee [58, 59] further coins the term *boundary-negotiating artifacts* to explore non-standardized,

fluid informal objects that not only travel across but also destabilize and push boundaries. Moreover, some participatory design approaches consider prototypes and models as boundary objects that facilitate the exchange of information and knowledge, and make it possible to involve different stakeholders in the design process by establishing a shared syntax or language [16, 22, 32, 44].

Socio-material perspectives point out that the features of physical objects impact the ways they behave and operate as boundary objects [45, 48, 66]. Documents, as a specific form of boundary objects, can become part of communicative practices in three ways: they can be an object of evaluation, a medium to facilitate communication, and/or the arena where communicative practices are shaped [48, 66, 119]. Contexts of data production, where actors often pursue different goals and interests, pose their own challenges in terms of fostering shared values and effective communication and, finally, promoting common documentation practices. Given these challenges, documentation should adapt to the needs of different stakeholders while remaining stable enough not to lose its function across different settings. The acknowledgment of data documentation as a boundary object [82] can provide a useful lens to approach and address some of the tensions arising from coordinating and collaborating across actors and organizations in data production. Specifically, the boundary object notion can add to existing work on data documentation in two ways. First, the acknowledgement of data documentation as a boundary object helps to make visible and specify the variety of actors involved in data production contexts with their own priorities, workflows, and needs. Secondly, this acknowledgment can add a practice-oriented perspective to the formal standardization of documentation frameworks and make work more visible. This includes work not only in the accommodations of standardized documentation frameworks to local needs and specificities, but also in the adaptation and improvement of existing work practices and workflows. Looking at documentation through the lens of boundary objects can be helpful in fostering documentation practices and design documentation frameworks that capture the practical, economic, and organizational decisions involved in data production and encoded in datasets.

3 METHOD

This investigation comprises 2.5 years of research in collaboration with two BPO companies located in Buenos Aires, Argentina (Alamo), and Sofia, Bulgaria (Action Data), specializing in the collection and annotation of data for ML. Our engagement with both organizations included fieldwork at their offices, interviewing phases, moments of feedback and discussion, the collaborative development of prototypes, and a series of co-design workshops. In this section, we provide details of our engagements with the participants and the methods used. Figure 1 shows how the research process evolved over time.



Fig. 1. Timeline of the iterative phases of data collection and analysis comprised in this research.

3.1 Participatory Design

At the core of Participatory Design (PD) is the active involvement of the communities that will use technologies and systems in their design. Moreover, PD recognizes the importance of the context in which technologies are used and the situated practices within that context [84]. In CSCW and HCI, participatory design methodologies have been deployed to explore various fields of inquiry, ranging from public services [7, 25, 26, 36, 54, 93, 97], to healthcare [56, 79, 88], to organizational settings [62, 106, 117]. Additionally, PD has been used in the field of ML Fairness, Accountability, and Transparency to guide algorithmic accountability interventions [52], learn about the concerns of communities affected by algorithmic decision-making systems [21, 83], and elucidate causal pathways to be input into complex modeling [65]. Regardless of the discipline, what these previous efforts have in common is their commitment to building reciprocal partnerships to inform the design of systems and technologies based on the values of justice and participation.

In this investigation, we use PD to explore the development, implementation, and use of documentation in real-world data production scenarios that involve many stakeholders that are geographically distributed. Similar to Wong-Villacres et al. [115], we understand participatory design as a long-term iterative process in which exploring participants' insights is as essential as tangible design outcomes. We use PD and co-design as synonyms [27].

Throughout the research and design process, we focused on three core elements of PD: *having a say, mutual learning*, and *co-realization* [17, 74]. We were intentional about listening to what the different actors involved in data production (the data workers, their managers, and the clients) had to say about dataset documentation in order to derive design considerations that would adapt to their needs. As the investigation evolved, we especially focused on the experiences of data workers because we realized they were the ones that usually *did not have a say* in how workflows and processes are designed in ML supply chains. We discuss this shift of perspective in Section 4.

While working on the research design, we put effort in remaining sensitive to the "politics of participation" [27] by trying to anticipate power asymmetries that could emerge among participants and between researchers and participants. In this sense, the interviewers and workshop facilitators were at all times intentional about *making room* for data workers' voices, for instance, by insisting on conducting in-depth interviews with them or preventing managers from hogging the mic during the workshops. Moreover, Bratteteig and Wagner [18] argue that "language is power" and by speaking one's own language, discourse can be expanded and *model monopoly*, i.e., the adoption of external mental models, can be avoided [17]. In this sense, it was important to us to let the participants maintain their preferred language — both verbally and visually. To that end, we provided interpretation and facilitation in the participants languages and worked largely with visual symbols [18]. In section 5.2, we reflect upon our research and highlight a few implications for participatory research in terms of power, negotiation, and participation in corporate settings [99].

3.2 The Participating Organizations

Two companies dedicated to the collection and labeling of ML data participated in this investigation:

- Alamo (pseudonym), located in Buenos Aires, Argentina.
- Action Data (pseudonym), located in Sofia, Bulgaria.

Both BPOs are impact sourcing companies. *Impact sourcing* refers to a branch of the outsourcing industry that proposes employing workers from poor and marginalized populations with the twofold aim of offering them a chance in the labor market and providing information-based services at lower prices. In both organizations, data workers perform tasks related to the collection and annotation of data. These tasks are requested by external clients located in various regions of the world, who

seek to outsource the production of datasets used to train and validate ML models. The relationship between requesters and data workers is mediated by managers, reviewers, and quality assurance (QA) analysts that form a hierarchical structure where data workers occupy the lowest layer [69].

3.3 Data Collection

3.3.1 Semi-Structured Interviews.

A total of 33 interviews were included in the present investigation. 19 were conducted in person in 2019. Due to the on-going COVID-19 pandemic and the resulting travel restrictions, 14 interviews were conducted remotely between August 2020 and September 2021.

Table 1. Overview of Participants and Research Methods. Some participants were interviewed more than once.

Method	Organization	Description	Participants	Medium	Language
S	Alamo	BPO in Argentina	6 data workers	In-Person	Spanish
Ε	Alamo	BPO in Argentina	2 managers	In-Person	Spanish
52	Action Data	BPO in Bulgaria	9 data workers	In-Person	English with
LEI					Arabic interpre-
z					tation
<u>A</u>	Action Data	BPO in Bulgaria	2 managers	In-Person	English
IRE	Alamo's	ML Practitioners request-	1 requesters	Virtual	Spanish
Ľ	clients	ing data work from Alamo			
nc	Action Data's	ML Practitioners request-	5 requesters	Virtual	English
IR	clients	ing data work from Action			
S-I		Data			
EM	Other Clients	ML Practitioners request-	4 requesters	Virtual	English
S		ing data work from other			
		BPOs			
	Action Data	BPO in Bulgaria	2 managers		
PS	Action Data	BPO in Bulgaria	2 data workers		English with
ЮН	Ranua	Action Data's Syrian part-	1 manager	Virtual	Arabic
KS		ner organization			interpretation
OR	Ranua	Action Data's Syrian part-	13 data workers		
À		ner organization			
S	Action Data's	ML Engineers at Scandina-	3 requesters		
ESIG	clients	vian University			
IQ-	Alamo	BPO in Argentina	13 data workers	Virtual	Spanish
8	Alamo	BPO in Argentina	4 managers	virtual	opanish

All interviews were semi-structured. Those conducted in 2019 were in-depth interviews that explored the lived experiences of data workers and aimed at understanding the contexts in which data work is carried out [69]. The average duration of those interviews was 65 minutes. For the analysis comprised in this paper, we selected only the in-depth interviews that revealed insights about documentation practices. In 2020 and 2021, we carried out a second round of interviews focused on the topic of documentation. Each one of those interviews had an average duration of 50 minutes. We interviewed actors who were actively involved in the documentation of projects at both companies. In addition, we interviewed machine learning practitioners with experience outsourcing data work to platforms and BPOs. Some of these practitioners were direct clients of Alamo and Action Data. Table 1 provides an overview of the interviews and participants included in this work.

The interviews allowed us to explore documentation practices carried out in different organizations and identify perceived challenges, experiences, and needs. Some interview partners were interviewed more than once at different stages of the design process. For each interview, the informants received a gift card for €30 (or the equivalent in their local currency) for an online marketplace. The audio of the interviews was recorded and, later on, transcribed. The name of the interview partners and any other elements that would help identify them were changed. The names that appear in the excerpts included in this paper are pseudonyms that the interview partners chose for themselves.

3.3.2 Co-Design Workshops.

Given the multiplicity of actors involved in processes of data production [69, 70], we saw value in inviting data workers, their managers, and their clients to a series of co-design workshops to discuss documentation practices. Table 1 offers an overview of the organizations and individuals participating in each workshop series.

The workshop series with each company took place on different days with a similar structure but was tailored to each setting and based on the ongoing projects and clients of the respective BPO. Due to the COVID-19 pandemic, all sessions took place online via Zoom and the activities were conducted in parallel on the visual tool Miro. As compensation, all workshop participants received €15 per hour of participation. The sessions were video-recorded and, later, transcribed. In addition, note-takers registered the activities and interactions and one of the authors produced a real-time graphic recording on Miro. The graphic recording was kept visible to participants as it evolved throughout the workshop sessions.



Fig. 2. The ground rules for participation were visible at all times throughout the co-design sessions with Action Data's workers. A Spanish version of the same rules was introduced at the workshops with Alamo.

The workshops with the Argentine company Alamo were conducted in Spanish, the native language of both participants and the co-authors who served as facilitators. The sessions took place on October 26 and 28, 2021. Each session comprised two hours. 17 participants attended the workshops, including four managers and 13 data workers. We suggested Alamo's management

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

Documenting Data Production Processes: A Participatory Approach for Data Work

include one of the company's clients in the workshops but this request was repeatedly declined. Our proposal to include a third workshop session for data workers only (without the presence of management) was also declined due to scheduling issues according to the company's management.

The 3-day workshop series that we organized with Action Data took place on November 3, 5, and 8, 2021. On day 1, we hosted a 90-minute session and invited Action Data's management (two participants) located in Bulgaria and 13 of the data workers that Action Data manages through its partner organization, Ranua, located in Syria. For day 2, we prepared a 3-hour workshop and had three members of one of Action Data's client organizations — a Nordic university — join in. Finally, day 3 was conceived as a 2-hour round table only for the Syrian data workers to discuss their impressions without the presence of managers and clients. This was motivated by the acknowledgment that the presence of managers can influence the level of comfort workers feel when expressing their opinions. In all sessions, the co-authors who served as main workshop facilitators spoke in English with simultaneous interpretation. For the breakout rooms, we offered facilitation in English and Arabic, depending on the needs of each group.

The power differentials present among the participant groups motivated many of our decisions in terms of how to design the respective sessions with Alamo and Action Data. The participation rules introduced at the beginning of the workshops reflect our stance (see Fig. 2). The acknowledgment of such power differentials also demanded much attention and flexibility on our side to quickly adapt and reformulate activities in view of difficulties identified while conducting the workshops.

3.4 Data Analysis

The collected data were analyzed through reflexive thematic analysis (RTA) as developed by Braun and Clarke [19] and Braun et al. [20]. RTA is a type of thematic analysis that places importance on researcher subjectivity and sees the researcher as playing an active role in the knowledge production process, a value which is emphasized by the word "reflexive." We describe the position from which we have conducted this research in the next subsection, 3.5.

To analyze the data collected in the present investigation, two of the co-authors explored the interview and audio transcriptions first. They started with a first coding round after which two further authors revised and added codes. A meeting followed where the four coders exchanged impressions and discussed disagreements around the codes. Additionally, they revisited video recordings of the workshops as well as the activities captured on Miro boards.

We constructed candidate themes after revisiting coding categories. Braun et al. [20] argue that themes "do not emerge fully-formed" and that there are first candidate themes that are developed from earlier phases, before settling on the final themes. It took us several iterations until arriving at a (more-or-less) final set of themes. Table 2 provides an overview.

The transcriptions of the data collected at Alamo were in Spanish. The data collected at Action Data was in English, with the exception of three interviews and all workshop sessions that included English and Arabic-speaking participants. In those cases, the English interpretations as performed by the interpreters were transcribed. Coding the transcriptions in both Spanish and English was possible because all authors are proficient in English and three of them are native Spanish speakers. Some of the excerpts included in the Findings sections were originally in Spanish and were translated into English only after the analysis was completed.

It is worth mentioning that some of the initial interview data had been previously coded for an earlier study conducted by some of the authors [69] using constructivist grounded theory [23], a method that was consistent with the exploratory spirit of the fieldwork and the first interviews conducted in 2019. For the present paper, we revisited some of those interviews and re-coded them along with the newly collected data using RTA.

Domain Summaries	Codes	Themes		
	Standardization			
	Foster Workers Growth	- - - - - - DOCUMENTATION AS COMMUNICATION		
FUNCTION OF	Quality Control & Metrics			
DOCUMENTATION	Work Enabler			
DOCUMENTATION	Inter-org. Accountability			
	Knowledge Preservation			
	Communication Medium	= DOCUMENTATION AS COMMUNICATION		
	Auditor	– ENADLER ACROSS DOUNDARIES – –		
	Managers			
ROLES AND	Service Owners			
COLLABORATION	Labelers			
	BPO's Teams			
	Clients			
	Wiki			
DOCUMENTATION	Kick-off Document			
DOCUMENTATION	Scope of Work (SoW)			
FORMATS	Dataset Documentation			
	Postmortem Report	HOW DOCUMENTATION CAN SUPPORT		
	Impact and Responsibility	WORKERS' AGENCY		
WORK ETHIC AND	Trust	-		
LADOD CONDITIONS	Security			
LABOR CONDITIONS	Labor Conditions			
	Confidentiality			
	Integration in Workflows			
LIMITATIONS OF	Use of Documentation	CHALLENCES DESIGNING FOR		
CURRENT DOC	Lack of Information	WODVEDS IN CODDOD ATE		
DRACTICES	Access	- WORKERS IN CORFORATE		
FRACTICES	Feedback	- ENVIRONMENTS		
	Updates			
	Ownership			
	Synthesis /Simplicity			
	Interactivity			
POSSIBLE DESIGN	Information Type	DESIGN IDEAS EMERGED FROM THE WORKSHOPS		
SOLUTIONS	Integration			
	Centralizing			
	Automation			
	Visualization			
	Lessons from Fieldwork	WHAT DADTICIDANTS & ODCANIZEDS		
LESSONS	Lessons from our Feedback	- WHAT FARTICITAN IS & ORGANIZERS		
LESSONS	Lessons from the Interviews	— PROCESS		
	Lessons from the Workshops			

Table 2. Data Analysis: Summary of codes and themes

3.5 Positionality Statement

Our ethical stance throughout this investigation was centered on respecting workers' expertise [27, 85, 92] and creating spaces where they could describe work practices and conditions, challenges, and needs. A central motivation — and certainly a challenge — of this research was to empower data workers to speak up and participate in shaping their own workflows and tools [103].

Our position and privilege vis-à-vis our study participants have been in constant interrogation and discussion throughout data collection, analysis, and while considering the implications of this investigation. Opting for an analysis method that would help us to acknowledge and make our positionalities as researchers explicit was essential to us. Following the idea of conducting a *reflexive* thematic analysis, it seems appropriate to disclose some elements of the authors' backgrounds that might have informed the findings we will present in the next section.

All authors except one are academics working in institutions located in the Global North. Most of us carry passports that are foreign to the places where we live and work. Four of us are sociologists, one a designer, and two are computer scientists. We all work in the field of Human-Computer Interaction and Critical Data Studies. Three of us are Latin American, two were born in China, one of us is German, and another is a first-generation US-American. All of us are first-generation academics. Our position as researchers living and working in the Global North provides us with privilege that our study participants do not hold. Despite being born and raised within working-class families, some of us in the same country as some of the participants and having too experienced migration, we acknowledge that our experiences differ from that of the participants in the sense, for instance, that we have never experienced war or had to flee our home countries to become refugees. This is especially relevant vis-à-vis the Syrian data workers that participated in this investigation. The acknowledgment of such privilege has guided several of the decisions that we have made, for instance, in terms of mitigating power differentials in data collection. In this sense, we must mention that, despite our best efforts, some power asymmetries could not be mitigated and we could only acknowledge and reflect upon them during data analysis and in this paper (see Sections 3.3.2 and 5.2).

4 FINDINGS

Our findings are presented in a way that reflects the diverse priorities, opinions, positions, and power of our participants. For each research site (Alamo in Sect. 4.1 and Action Data in Sect. 4.2), we first provide details of our engagement with the respective partner organization and the preliminary knowledge obtained through the interviewing phases. We then move on to describe salient discussions that took place at the workshops. In Section 4.3, we present a summary of findings including commonalities and differences across sites.

Long-term engagements with participants are key to participatory research [24, 110] and, as Sloane et al. [99] argue, it is important "to make the tensions that characterize the goal of long-term participation in ML visible, acknowledging that partnerships and justice do not scale in frictionless ways, but require constant maintenance and articulation with existing social formations in new contexts." We believe that there is a contribution to be made by placing our findings in conversation with the vicissitudes of conducting participatory research, especially with the inter-mediation of partner organizations. This is why we show how our observations evolved over 2.5 years and across sites and methods. We are convinced that our findings would not be the same had we stopped after the interviews. In the same way, without a description of the interviewing process and how our relationship with the participant organizations evolved, the ideas emerging from the co-design workshops would lose some of their nuances.

4.1 Case 1: Alamo

The Buenos Aires-located company that we will call *Alamo* is, at the time of this investigation, a medium-sized organization with around 400 data workers distributed in three offices located in Argentina and two other Latin American countries. Its main client is a large regional e-commerce corporation. 90% of Alamo's projects are requested by this client and include image segmentation and labeling, data collection, and content moderation. The Buenos Aires office employs around 200 data workers, mainly young people living in very poor neighborhoods or slums in and around the city. Looking for workers in poor areas of Buenos Aires is part of Alamo's mission as an impact sourcing company. Alamo's data workers are offered a steady part- or full-time salary and benefits,

which contrasts with the widespread contractor-based model observed in other BPOs and data work platforms [80]. Even so, Alamo's workers receive the minimum legal wage in Argentina (the equivalent to US\$1.50/hour in 2019) and their salaries situate them below the poverty line.

4.1.1 Investigating Alamo's Documentation Practices.

During fieldwork in 2019, one of the observations that stood out was the company's stated commitment to fostering its workers' professional growth [69]. In contrast with this mission, Alamo's emerging documentation practices seemed, at the beginning of this investigation, almost completely oriented toward surveilling workers and measuring their performance. Noah, one of the data workers that we interviewed in 2019 described that documentation process as follows:

For instance, for this project, we have three metrics for individual performance and now we're about to add one more. Mostly, we quantify and document errors or rollbacks reported by the requester.

Nancy, who was one of Alamo's project managers in 2019, explained what was done with the documented performance metrics:

Within the company's structure, QA [the quality assurance department] then takes those metrics and puts together something like a "scoreboard" that shows everyone's performance and that is visible to the whole company.

After our first fieldwork period and once we had analyzed the collected data, we provided Alamo with a report summarizing our observations. We found that there was a contradiction between the company's mission of fostering worker growth and its emphasis on quantifying individual performances, and suggested reorienting the purpose of documentation to include information that could help workers understand the background and purpose of their work. We have discussed those observations and their implications at length in previous work [67, 69]. After presenting our observations, we engaged in a conversation with Alamo's co-founder Priscilla who, as shown by the following interview excerpt, agreed with our suggestions and expressed her desire to establish documentation practices based on the needs of data workers:

In my opinion, this is what we need to do. Listen more to their needs and ask them what they need to perform better and do their work better. And we must generate and provide the information they need. We need to make time for that.

In September 2020, we resumed contact with Alamo's management and invited the company to participate in our investigation of data production documentation, starting with a series of followup interviews. By then, the company was already in the process of restructuring its documentation practices. Nati, one of Alamo's former QA analysts, had just been promoted to "continuous improvement analyst" and put in charge of bringing together the company's quality standards and its goals in terms of promoting workers' professional growth. Nati's main task was re-designing Alamo's documentation practices. She described this process as follows:

During the past year, we made an effort to ask everyone "what would you like to see reflected in documentation? What do you need?" After that, we started dropping old practices that didn't make sense anymore. We threw them away, and started anew. We learned a lot about empathy too, you know? This work is very much about empathy [...] because we assign workers in each team according to their competencies and their skills, and they need to know what projects are about to actually make the most out of those skills.

What came out of the process led by Nati was a series of documentation practices strongly oriented toward *knowledge preservation*, i.e., the documentation of workflows, tools, specific details about the conduction of projects, best practices, and lessons learned [70]. In Nati's words,

Documenting Data Production Processes: A Participatory Approach for Data Work

documentation was envisioned as an educational resource for Alamo's workers and its aim was to mitigate information loss due to fast worker turn around. A special team led by Nati was put in charge of developing and implementing this vision. With the documentation guidelines and templates designed by them, each team leader would document the projects carried out by their data workers. Moreover, following one of our suggestions, the company saw value in documenting contextual information about the ML products that are trained or validated using data produced at Alamo. Management then started to seek clients' involvement in documentation practices. The first step was sharing the project documentation with the project's requester and asking for feedback. These collaborative practices progressed to the point that some clients became co-producers and users of the documentation produced by Alamo, as Pablo, leader of one of Alamo's data labeling teams described:

I mean, it depends on the client. Luckily this requester is very committed and is always looking at our documentation, providing feedback, checking if the information is up to date. Sometimes they just jump in and add information, they collaborate on the best practices manual. It's also good for them to see if their information is updated too, not just for us but for the requester's own team, because they share these documents within the requester's company.

While this arrangement was beneficial for the clients, Alamo was providing the infrastructure and most of the workforce that made documentation possible. Alamo's workers started to feel the burden, as team lead Pablo described:

We definitely noticed the change: from not documenting at all to documenting everything. It also depends on the project. For instance, our main client [a large regional e-commerce corporation] brought many more requirements for us to implement in our documentation in terms of what they wanted to see reflected. There are new documents and new guidelines all the time.

We followed up with continuous improvement analyst Nati in December 2020 as she continued to review and revise documentation practices within Alamo. By then, Alamo's documentation revolved around a format they called *the Wiki* that was conceived as a collaborative document to enable the conduction of projects (see Fig. 3). Each Wiki comprised information, collected in many cases in direct collaboration with the requester, about processes, tools, and best practices concerning each specific project. Nati's main concern back then was reducing the time and effort that were flowing into documentation and, at the same time, finding efficient ways of keeping documentation updated.

After these interviews with Nati, Pablo, and some other data workers, we got back to Alamo's management with feedback. Among other things, we highlighted the need to reduce the amount of work put on Alamo's team in terms of producing and maintaining documentation and distributing responsibilities more equitably. The idea of producing exhaustive metadata for each project was commendable but the price of such an effort was being paid by the data workers. Moreover, we found that the original idea of the Wikis, namely, that of preserving the knowledge generated in each project and offering more information about the ML supply chain to foster workers' professional growth, had shifted toward a constant consideration of what clients might want to see reflected in documentation, as Sole, leader of another data labeling team at Alamo, described:

We need to know what we want to document first. Who we are documenting for. That's a challenge at the moment. Because right now I ask myself: who am I doing this for? What for? [...] Right now, we're focusing on the clients: what is the client asking for? Because I may want many things, but what is the client interested in? What is the client expecting from us as service providers?



Fig. 3. The wiki-based documentation prototype created by Alamo to brief and train workers.

In view of these challenges, we proposed the idea of conducting a series of co-design workshops with Alamo's data workers and managers with the purpose of moving the focus *away from "what clients want" and re-imagining documentation based on the needs of workers.* In June 2021, we started conversations with Alamo. Negotiation points included participants' compensation and attendance, i.e. how many and which workers would Alamo allow to participate in the workshops. We agreed to use the Wiki as the prototype we would discuss and re-imagine at the workshops. After several months of negotiation, the workshops finally took place in October 2021.

4.1.2 Co-Designing Documentation with Alamo's Workers.

The workshops with Alamo workers and managers started with a discussion of the Wiki template. The discussion focused on a recent project from Alamo's main client, an e-commerce corporation that operates primarily in Latin America. The client instructed Alamo to produce a training dataset with images of false identity cards from three Spanish-speaking Latin American countries. This dataset was later used by the client to train an algorithm to screen IDs. The instructions for workers involved the collection of images of authentic identification cards, their modification to render them invalid, and the labeling of the resulting photos.

The managers participating in the session mentioned that the main benefit of the Wikis is for data workers to take an interest in existing projects, explore them through the respective Wikis, and use that information for their education and growth. In the words of a project manager:

We've been working with growth metrics for each person working at Alamo this year. For example, if someone wants to learn programming skills, they can look at the Wiki, see what tools are available, the projects, the languages used, and know what skills to develop.

However, while managers want to encourage workers to have an overview of the undergoing projects at the company, they are also mindful of their clients' demands in terms of confidentiality and security. For this reason, either much of the information in the documentation lacks details and appears general and superficial, or the access to such information is restrained. Several of the data workers argued that the restrictions and the lack of access to the Wikis as well as difficulties keeping the information updated represented a hurdle for their work:

One of the limitations is access. That we all know the information is there and can access it. How can we make that happen? Not everyone has access. How can we

guarantee that the information is updated? It takes a lot of time to write down the documentation.

In addition, the data workers asked the managers for trust in accessing complete versions of the documentation, but the managers argued that the information is sensitive and clients often request not to make it available to all. This tension, coupled with the difficulty of keeping documents updated as projects evolve, makes it hard for workers to incorporate feedback in the documentation due to the lack of information.

Moreover, the workers acknowledged that they were less inclined to use the information on Wikis for reference. For one thing, some workers were not aware that the Wikis were available for them to see other projects and had to be reminded of their presence. For another, the documentation was perceived as "on the machine" and "not close at hand" and, thus, it was not the workers' initial responses to check if relevant information could be found on the Wikis. By the end of day 1, it became evident that there is a discrepancy between how management thinks workers should use the Wikis and the latter's perceived usefulness vis-à-vis the time and effort required to maintain these documents.



Fig. 4. One of the workshop activities aiming to understand roles, workflows, and collaboration within Alamo.

We started Day 2 with a focused discussion of the tensions identified during the previous session: *how to keep the Wikis up to date, how to guarantee access and promote use, how to provide details while keeping confidentiality, and how to integrate documentation practices in existing workflows.* When discussing how to update the documentation efficiently, one of the data workers pointed out that gathering required information often took weeks due to the organizational structure and dispersed nature of teams. In this discussion, possible forms of documentation were explored. A former data analyst, now team leader, argued that automating the production of metadata was probably not the best solution because much information exchange and knowledge transfer took place in daily conversations. He argued that *the potential of documentation lies in integrating it into communication processes:*

I think that's where the greatest possibility of improving communication lies. To leave at least a record. It does not necessarily have to be so precise to set up a whole process, but a spreadsheet where I document things, an e-mail, a chat. I don't know. There are many ways of communication. I don't know if the best thing is automation, because we are quite beyond automation. Yes, it is a matter of communication.

Finally, the last workshop activity asked workers to imagine possible approaches to the problems discussed earlier in the session. We asked them to think of a documentation process that would guarantee simplicity, accessibility, and integration in workflows. These requirements were derived from the previously conducted interviews. In their own way, all teams focused on the use of documentation rather than the act of documenting. They all agreed that documentation should be centralized rather than dispersed in different locations and be easy to edit and interpret by any team member. In this sense, three themes were salient. First, the documents should have better navigation tools for readers. Some proposed additions include a history of editions (such as in other Wikis like Wikipedia) and a glossary to enable translation between languages and illustrate technology jargon. Second, documentation should be more legible. In this sense, participants suggested including visual material, especially in task instructions, such as graphics or video tutorials, and reducing text. Some participants also proposed more interactive documentation. For example, teams can benefit from more precise structures instead of documents without format. Third, participants suggested three ways to *increase participation in the documentation*: making everybody in the organization in charge of documenting, automating the passage between text input and the graphic output, and adding elements of gamification to encourage and measure participation.

4.2 Case 2: Action Data

The Bulgarian company *Action Data* specializes in image data collection, segmentation, and labeling. Its clients are computer vision companies and academic institutions, mostly located in North America and Western Europe. Action Data offers its data workers contractor-based work and the possibility to complete their assignments remotely, with flexible hours. Workers are paid per piece (image/video or annotation), and payment varies according to the project and its difficulty. At the time of our first visit in July 2019, Action Data operated with one manager and two coordinators in salaried positions handling operations and a pool of around 60 freelance data workers. As part of its impact-sourcing mission, Action Data recruited its data workers among refugees from the Middle East who had been granted asylum in Bulgaria. That approach changed in 2021 and now Action Data outsources its projects to partner organizations located in Syria and Iraq. In fact, in November 2021, 90% of Action Data's projects were conducted by data workers in Syria using a re-intermediation model [41] in which a third-party organization that we will call *Ranua* recruits the workers while Action Data manages the projects with clients. This model of using a third party to recruit workers has been documented in the crowdsourcing and gig economy sectors, in which some platforms recruit workers to do data work for other platforms [109].

4.2.1 Investigating Action Data's Documentation Practices.

In 2019, as we began this investigation, Action Data's projects and workers fluctuated frequently and communication seemed to adapt to the needs of each project. Consequently, Action Data did not have standardized documentation guidelines. This was critical considering the fact that most of the workers worked remotely after the initial training at the company's office.

After that first fieldwork period and once we had analyzed the collected data, we provided Action Data with a report summarizing our observations — most of them related to labor conditions, highlighting their effects on the quality of the datasets produced by the company. In addition, we provided recommendations regarding the need to educate data workers about the ML pipeline, the

urgency of considering labor conditions an essential point in the "ethical AI" discourse at the core of the firm's marketing strategy, and the importance of documenting workflows and practices [53]. The company's management found that our suggestions were useful and actionable. By September 2020 they proceeded to implement some of them, starting with an education program for data workers. The program included a machine learning course that covered technical and ethical aspects of ML as well as an overview of the ML supply chain.

It was around that time that we invited Action Data to participate in a series of follow-up interviews centered around documentation practices. An important development was the company's implementation of an internal database of all the projects they had conducted since the beginning of their operations. However, some of the information contained in it was considered confidential or sensitive, which resulted in the documents only being shared with Action Data's management and the respective clients. The data workers did not have access to the project database. We provided feedback on this issue and highlighted the importance of sharing information with the data workers.

The recently hired operations officer at the time, Tina, was then put in charge of creating a series of documents targeted at data workers, that contained guidelines and information about the different projects carried out by Action Data. Apart from our recommendation, two factors contributed to the company's decision to develop these documents. First, Action Data had managed to establish itself in the market of data services for machine learning — especially data labeling. This means that, by 2020, the company was conducting long-term projects for a more or less steady set of clients, which resulted in the need to keep better records and continually train workers. Second, because of its growth, the company had started outsourcing some requests to partner organizations located in the Middle East. This development revealed the need for clear — written — guidelines for a geographically distributed workforce that, in most cases, did not speak English.

NAME OF COMPANY/CLIENT		NAME OF PROJECT	
NAME OF CONTACT PERSON	POSITION	EMAIL	PHONE NUMBER
PROJECT TYPE	PLATFORM	TYPE OF DATA	VOLUME
TRIAL START DATE DATE	TRIAL END DATE	PROJECT START DATE	PROJECT END
PROJECT BACKGROUND AND models will be trained on the	PURPOSE: Why are we data?	annotating this data? How	will it be used? What
PROJECT BACKGROUND AND models will be trained on the CLASSES: What is the class lis	PURPOSE: Why are we data?	annotating this data? How y	classes exactly?
PROJECT BACKGROUND AND models will be trained on the classes: What is the class lis	PURPOSE: Why are we data? t, tags, and attributes?	annotating this data? How t	viii it be used? What

Fig. 5. One of the documentation templates prototyped by Action Data: The Scope of Work form (SoW).

Since the inclusion of third-party organizations in the conduction of projects and because of the disadvantaged position of service providers vis-à-vis their clients, Action Data's perceived usefulness of documentation started to shift toward the idea of *inter-organizational accountability* [69]. The company started to develop a documentation template that was able to reflect its business relationship with other organizations, i.e. clients and partner organizations in the Middle East to which projects were being increasingly outsourced. This form of documentation would make explicit each organization's responsibilities in terms of work, payment, and deadlines and would reflect iterations and updates. Moreover, keeping clear records of the task instructions provided by clients and documenting instructed changes could provide proof that tasks were carried out as instructed, especially in cases where clients are not satisfied with the quality of the service provided or decide to demand more, as Eva explained:

We also keep the client accountable so that they don't come up with a new requirement or something that we haven't mentioned before. So, documents are also for accountability of us toward the client as well so that the client can have a document where they can keep track of what the arrangement is and so on beyond our contract.

toward the beginning of 2021, Action Data began developing three documentation templates. The prototypes were designed as forms where information would be filled in by Action Data's management. When sending the prototypes to us to request feedback, operations manager Tina explained via email:

On the first page is the *SoW* (Scope of Work) that we regularly use in order to scope clients' projects before we start working. On the second one, there is a new idea Eva is piloting for *Dataset Documentation*, so some information we will fill out during and after projects so that the client can use it for transparency purposes if they wish. And then the third page is a *Post Mortem* template that we can use to address issues and challenges after the project is completed.

These three prototypes were examined and interrogated at the co-design workshops conducted with Action Data workers that we describe in the next subsection.

4.2.2 Co-Designing Documentation with Action Data's Workers.

The co-design workshops took place in three different sessions facilitated over Zoom.

Day 1 worked as an introduction and focused on getting to know Action Data's general structure and workflows, and discussing how stakeholders use the three documentation prototypes ("Scope of Work," "Dataset Documentation," "Postmortem Report"). To this end, participants mapped the actions and relationships of each of the stakeholders in the data labeling process, from clients to Action Data's intermediation, to their partner organization Ranua (pseudonym) in Syria, and the Syrian data workers (see Fig. 6).

Day 2 introduced the perspectives of the clients — three researchers at a Nordic university. They were working on automating a safety system at the harbor front that would release a signal, for instance, if people fall in the water. For this client, the workers in Syria annotated thermal videos from CCTV footage that would serve as training data. The task was coordinated by Action Data and outsourced to Ranua's workers in Syria.

The first group activity of the day was a roleplay (see Fig. 7). We asked each group to reflect on the knowledge required to conduct data-production projects in their assigned roles by filling information cards. Afterward, each team presented their cards in a discussion round during which each group gave their feedback on how accurately it had been roleplayed by the others. The clients in the role of data workers generated many cards comprising questions about labeling type and classes, edge cases, desired output, and annotation precision (see Fig. 7). However, the feedback of

510:18

Documenting Data Production Processes: A Participatory Approach for Data Work



Fig. 6. Stakeholder map showing the three groups participating in the second workshop session, the nature of their relationship, and the workflows that characterize it. The map was developed with the workshop participants. The real names of the organizations and individual actors, originally included in the stakeholder map, have been changed to be included in this paper.

the actual data workers uncovered two further issues that the clients generating the questions had not contemplated. First, workers wanted to know how much they would earn. And second, it was important to them to include information about the ethical implications of the data they had to work with in terms of potential harms to data subjects as well as concerns about the mental health of data workers dealing with violent, offensive, or triggering data.

In the same discussion, workers considered the extent of annotation precision required an important information type because it defines the difficulty of the task and can have negative effects on the received wages. This requirement is usually negotiated between requesters and managers without the participation of data workers. The dispersion and intermediation — and power — of actors like Action Data and Ranua also render pricing negotiation inaccessible to workers. These negotiations occur primarily between the client and Action Data with lesser participation of managers from Ranua.

Following these discussions, participants acknowledged that one of the main problems of the documentation prototypes developed by Action Data is that they only concerned managers and clients while ignoring data workers. Moreover, instruction documents for workers are translated into Arabic by Ranua's management due to workers' limited knowledge of the English language. In this sense, the English-based documentation prototypes proposed by Action Data remain inaccessible to most Syrian workers.

Finally, the last activity of day 2 prompted participants to imagine new forms of documentation by thinking of its format, structure, information, and the stakeholders involved in its development and access. Similar to the Alamo team, Action Data's workers prioritized a view of *documentation as a communication medium*. Mostly, they acknowledged that documentation should allow the co-creation of task instructions with different stakeholders, particularly workers. In the words of one of the Syrian workers:

510:19

Communication and feedback should be the highest priorities. The biggest loser when communication is lacking is the annotators because they spend more effort and time adapting their work each time.

In addition, by increasing communication between different stakeholders, the documentation can potentially *integrate feedback* to make data-production processes clearer by reducing redundant work. Furthermore, it was discussed that a post-mortem report aiming at reflecting on the finished projects allows producing knowledge for future tasks. Such a document could include missing information, met goals, fairness in pricing and wages, the quality of the work environment, difficulties, and learned lessons.



Fig. 7. Roleplay activity in breakout groups. Here the requester group had to generate questions (the yellow cards) assuming the role of the data labelers. Afterward, the data labelers provided feedback and added their own questions (in the blue cards.)

Finally, on day 3, we hoped to enable a space for data workers to express their opinions and wishes without the presence of managers. Although all workshop sessions were designed to be about documentation practices, there were some more pressing needs that were the focus of the data workers: to earn a living wage and be able to provide for their families. This way, we opened up the floor to discuss these issues and talk about possible approaches. After the workshops, we helped facilitate the dialogue [12] between the Syrian data workers and Action Data management. One of the main issues described by the workers was that *they received the same payment regardless of their effort or time spent on the tasks.* One of the data workers explained the general feeling among her colleagues:

Some feel that the profit or the earnings they're making after the end of the job are not enough compared to the effort they have put into it.

Another issue with payment was the lack of room for negotiation by workers about their wages. For example, one practice that workers perceived as unfair was that *management divided the client's* payment equally among all team members regardless of individual efforts. Furthermore, the workers

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

are not consulted in decisions about team composition and cannot influence how many workers will share the payment:

When we divide the amount of money that we have received for that certain project, when it's divided by the number of people working on that project, so it turns out to be something completely unfair.

Tensions were observed during the discussion on the tools used to perform data work. The limitations of tools used by workers require them to share accounts and keep a shared spreadsheet with the number of annotations per task that each worker had done. This makes individual performances visible, accessible, and modifiable to all. One of the workshop participants explained that this was a practical choice due to the diverse working conditions and location of workers and the complexity of supervising and managing their work. He further emphasized mutual trust between the management and workers and among workers. While acknowledging the existence of honest behavior and mutual trust within the team, the other participants did not regard trust as a satisfying justification for the limitations of tools and privacy bridging, and pointed out that the problem was beyond trust issues:

I get the question of trust, but I believe that some of the issues that have been raised don't really have much to do with trust, but maybe with technical problems or the adequacy of free tools that, you know, make work easier or more difficult.

For one thing, current tools and arrangements posed difficulties in keeping track of and securing the tasks they individually performed. For another, they did not allow adequate respect for privacy, which could potentially undermine team solidarity, as some participants explained. Therefore, other workers advocated for better tools and arrangements to facilitate work and enhance privacy and security. In this sense, as a potential solution related to documentation efforts, the workers suggested that documentation should reflect how much workers are being paid and the purpose of their work. However, the individual earnings of each worker should remain private.

4.3 Salient Considerations and Summary of Findings

Before the workshops, we observed two different orientations in the ways in which the BPOs implemented our feedback to create the documentation prototypes. While Alamo's documentation prototype (*the Wiki*) was oriented toward the *preservation of knowledge* that could be useful to train (future) workers, Action Data had oriented the design of their documentation templates (especially *the Scope of Work*) toward keeping *inter-organizational accountability*, i.e., clearly stating clients' expectations and workers' responsibilities [70].

After the workshops, however, we learned that *the needs of data workers at both organizations were not significantly different in terms of what they wanted to see reflected in the documentation.* The apparent diverging orientations adopted by Action Data and Alamo were based on the corporate needs of each organization because, up until that point, the design process had been piloted by the companies' founders and managers. Engaging in hands-on design sessions directly with data workers gave us the possibility to explore their challenges and needs beyond corporate priorities and gave rise to a different set of considerations regarding documentation.

To summarize our findings, we present here five design considerations derived from our longterm engagement at both research sites. Taking steps toward approaching actionable documentation guidelines, we add concrete questions to illustrate how these considerations could flow into a documentation framework based on data workers' needs. We present these documentation items as questions because that is how the data workers expressed them at the co-design workshops as we prompted them to re-imagine documentation (see Fig. 8). For the most part, these observations can be generalized to both BPOs. We added specific notes where differences between the research sites were observed.

• The documentation of data production should be collaborative. Many actors take part in data production processes and their documentation. The roles can vary according to the context. There are, however, two patterns that can be derived from our findings. First, most of the information about tasks comes from the requesters and is documented by managers. The managers act as gatekeepers, only granting access to information according to the client's confidentiality requests. Second, data workers receive only minimum information that allows them to complete tasks according to the client's instructions. To break with these patterns, workers must be acknowledged as collaborators and be granted access to information.

 \rightarrow *Who is the requester?* The data workers would like to see information about the client, its products, and its relationship history with the data processing company included in the documentation. This information could help workers get an idea of "what to expect" from each project and produce strategies to collaborate more effectively.

 \rightarrow *Who else is working on this project*? Information about team composition is helpful for the coordination and division of labor. Moreover, at Action Data where the project budget is divided among the team members, it is essential for data workers to know how many people will be working on the project because their earnings depend on that.

• The documentation of data production should enable communication. It should transport information that is useful and intelligible to diverse actors and organizations. Language could present a barrier for geographically distributed teams, as reported by Action Data workers. Moreover, communication should not just flow top-down. In this sense, data workers at both sites emphasized the need to enable communication channels for feedback from workers to reach clients.

 \rightarrow What kind of data is this? Where does it come from? Data workers and managers at both BPOs argued that requesters should communicate key facts about the data, i.e., data provenance, dataset's previous use, and proof of compliance with local regulations, especially in the case of personal data. Moreover, Action Data participants suggested including a warning to inform workers if projects involve violent or offensive material.

 \rightarrow What have we learned from this project? Participants at both sites argued for the inclusion of a post-mortem report after project completion to reflect upon edge cases, lessons learned, and possible mismatches between time, effort, and payment. Alamo workers considered this to be a useful way to preserve knowledge for their professional growth and for future projects. Instead, Action Data workers saw such reports as a useful communication channel to provide feedback to requesters, especially on whether pricing and wages were fair vis-a-vis the labor involved in the task.

• The documentation of data production should not be seen as a one-time action. Data production is often messy and iterative. ML datasets evolve and change over time. Therefore, data production requires living, evolving documentation.

 \rightarrow How has this project evolved over time? Participants at both sites favored the inclusion of an update history log, including the date, time, and name of the person who documented each update. Given the constant negotiation and modification of task instructions and the adaptation of data work projects, a history log could keep important decisions recorded and task instruction and requirements updated.

 \rightarrow What is the data going to be used for? Data workers at both sites expressed the need to receive more information about the ML pipeline and dataset maintenance and stewardship. Specifically, the participants wanted documentation to include a description of the ML product that will be trained on the basis of the data produced by them. Regarding the involvement of sensitive or personal data, Action Data workers suggested asking clients for a written commitment that the data will not be used for harmful purposes. In comparison, workers in Alamo pointed to the need for clear instructions on how to process and store sensitive data to enhance privacy.



Fig. 8. Template for one of the co-design exercises carried out at the workshop series with Action Data and Alamo. The activity prompted data workers to deconstruct and re-imagine documentation practices.

• The documentation of data production should be integrated in existing workflows and routines. This could help foster a view of documentation as a constitutive part of data production and not as extra work. Action Data workers expressed that documentation could help them understand projects and their requirements better, which would ease their everyday tasks. Alamo workers mentioned that the preservation of information might come in handy and help them save time in future projects or to train future workers.

 \rightarrow What should be done with the data? The participants asked to receive a detailed description of tasks, including categories and classes for data annotation and collection, edge cases, and the required quality of annotations. One important point mentioned by Action Data workers was to include information about annotation precision and acceptance and rejection standards which could negatively impact their wages.

 \rightarrow What tools are used? This includes clear guidelines in terms of communication tools, the software used to work on the data, and therepository where the data is to be stored. Alamo participants mentioned that the use of certain tools (e.g. Photoshop) has sometimes increased the interest of workers to join certain projects because they wanted to learn that specific software tool.

• The documentation of data production should be adaptable to stakeholder needs.

We base the design considerations presented in this paper on the needs and wants of data workers and acknowledge that other stakeholders such as requesters, data subjects, or policymakers might have different requirements. Moreover, even if the suggestions and needs of the data workers at both BPOs were similar, some subtle differences could be observed. For instance, the workers at Action Data attached value to the evaluation of project workload and payment details, which would help them assess and decide the extent of commitment to the project, based on the perceived fairness of the wages. Instead, workers in Alamo placed more emphasis on documenting the evolution of projects because the company runs periodical evaluations based on individual and project-based performance metrics.

 \rightarrow How difficult is the task? How labor-intensive is this data? Action Data workers suggested including a "difficulty rating" for workers to quickly estimate the workload and time required for each project. In a similar way, Alamo workers mentioned that an overview of the tools used in each project would help workers understand the difficulty and the knowledge required for specific tasks.

 \rightarrow How much does this project pay each data worker? Clear information about wages, payment structure, and rules was the most highlighted need at the workshops with Action Data workers because they are paid according to the project and the task. Instead, Alamo workers, who receive a fixed salary independently from the project, prioritized information about client expectations and quality standards that would help them complete tasks "according to what each client has in mind," which could have a positive impact on workers' performance reviews.

5 DISCUSSION

Our findings show that data workers are expected to use the documentation to perform better and learn new skills. However, they receive only the bare minimum information to allow them to complete tasks according to clients' instructions. Workers' feedback is rarely considered and only occasionally integrated into task instructions. In addition, access to documents is not always granted (as reported by Alamo's workers), and language translations can constitute a barrier (as discussed by Action Data's workers). The production of ML datasets involves many iterations and the collaboration of actors working in different organizations that are often geographically distributed. Moreover, datasets are frequently updated, relabeled, and reused even after being produced. Therefore, the challenge consists of creating documentation that is able to capture the evolving character of datasets and the intricacies of data work. To address this challenge, we advocate for a shift of perspective in the development and implementation of data documentation frameworks: from documenting datasets toward documenting data production.

Data production documentation should travel across multiple stakeholders and facilitate their communication. In this sense, documentation should register iterations in production processes and enable feedback loops. *Documenting data production processes means making collaboration instances and discussions explicit.* By including detailed information about tasks and wages and enabling the feedback of workers, documentation could make precarious labor conditions in outsourced data work explicit and, therefore, contestable. Previous data documentation initiatives in ML were rooted in the value of transparency and motivated by the need to inform consumers and the general public about dataset characteristics and composition. Our research shifts the focus away from the notion of transparency and toward promoting reflexivity [31, 70, 94] in terms of how tasks are laid out and how datasets are produced.

Documenting Data Production Processes: A Participatory Approach for Data Work

The acknowledgment that this form of documentation *works as a boundary object* [15, 101] can be helpful in approaching the challenges arising from coordinating and collaborating across organizations [82]. In this sense, we recognize two contributions that the study of data documentation as a form of boundary object can bring to the line of research on developing documentation frameworks for datasets. First, while previous work on data documentation often assumes a single and consistent notion of "dataset producer," the boundary object notion helps to highlight the multiplicity of involved actors and needs. This is especially important in data production, as actors that are often globally distributed may have different – and even divergent – goals, priorities, and work requirements. Second, this acknowledgment helps draw more attention to the dynamics between formal standardization and local accommodations, making the local tailoring of documentation practices visible. As described by Star [100], we believe that this local adaptation is also a form of invisible work. Useful as the existing data documentation frameworks are, they remain relatively stable and standardized. In this regard, making visible how a documentation framework accommodates local specificities and requirements, which can be observed in how it operates as a boundary object, has values for the design of documentation frameworks.

Three aspects identified and explored by Star [100] and Star and Griesemer [101] could be the starting point for valuable considerations regarding the functioning of documentation as a boundary object. The first aspect is interpretive flexibility, understood as the adaptability of boundary objects to be interpreted, used, and mobilized differently depending on the communities that use them. In our findings, this is manifested in the workers' view to conceive data documentation primarily as a medium to facilitate the co-creation of task instructions, while managers see documentation as an important site for preserving knowledge and improving performance. Moreover, companies use documentation to keep each other accountable in the event of discrepancies [70] and researchers to promote transparency and reproducibility [37, 38]. The second aspect is the material and organizational structures of boundary objects, which arise from the information needs and work requirements of different communities. Such material and organizational structures could include differential access to information, the ownership of information among workers, and the incorporation of feedback instances. In this sense, the questions included in the previous section reflect the needs and requirements of data workers regarding documentation. The third aspect is the tension between the loosely structured common use and more strongly structured local use of boundary objects. In view of the ideas expressed by our participants, such tension could be eased through the incorporation of text-based approaches to document processes, visual elements such as video and images to present the documentation, the use of interactive features to retrieve them, and a translation aid to facilitate the process.

While this paper centers on the needs and requirements of data workers, future research could approach the perspectives of other stakeholders, the relationships and communication among them, and how they shape the development of documentation in intra- and inter-organizational contexts through the lens of boundary objects.

5.1 Design Implications: Documentation for and with Data Workers

Our findings show that designing documentation based on data workers' needs and ideas requires careful consideration of the unique conditions present in machine learning supply chains and in specific data production settings. For instance, the fact that Action Data is in the European Union while Ranua and the data workers are in the Middle East requires higher coordination efforts than data projects at the Argentina-based Alamo. Similarly, while all workers in Alamo are native Spanish speakers, in the case of Action Data, English is used as a vehicular second language. Given that most data workers do not speak English, the language choice renders coordination – and

documentation – efforts challenging. Such variations entail several implications for the design of data production documentation that we list in the following.

First, access to information must be guaranteed for documentation to support collaboration. However, different groups might have diverging needs regarding privacy and security [55]. In this sense, trust in collaboration [76] must be supported by the documentation framework to foster a view of data workers as essential collaborators who need access to information to produce better data [68]. Some of the approaches mentioned by our participants included promoting the ownership of information among workers and differentiating personal and group information when restricting access.

Second, requesters could benefit from the early feedback of data workers regarding task design and instructions. This feedback should also address possible mismatches between payment and effort. For such an exchange to be possible, documentation frameworks must enable feedback loops and transparent information about tasks, intended uses of the datasets as well as wages and payment conditions.

Third, to facilitate the integration of documentation in data production workflows, the information input should be simple. In this sense, the workers favored standardized text-based options such as checklists or multiple-choice questionnaires. Conversely, visual forms were favored in the documentation presentation. Some of the participants' ideas included integrating images and short videos to illustrate instruction. Moreover, conciseness, good structure, and the integration of a search function and interactive elements were considered crucial to the use of documentation.

Fourth, the mismatch between a documentation input that should be text-based and a documentation output that should be visually appealing must be considered. The participants discussed the integration of some form of translation aid, in this case, to turn a concise checklist into a compelling and visual document. Translation was also considered in linguistic terms, for instance, in the case of Action Data, where most of the information is documented in one language (English), but the documentation users need it in another (Arabic).

Fifth, the design of a boundary object such as documentation is not a neutral process [45, 48]. In fact, boundary objects are shaped within power relations that inform how they are understood and enable specific interpretations out of a general interpretative flexibility [45, 107]. The acknowledgment of such power differentials is at the core of our decision to focus on the needs of data workers to re-imagine data documentation and is a critical issue to consider when designing documentation frameworks.

Finally, boundary objects, especially documents, can comprise different, and even divergent, perspectives of involved participants but do not necessarily require reaching a consensus [100, 101]. Attempts to reach consensus may reinforce hegemonic aspirations and favor the worldviews of certain groups over others [45, 48]. In this regard, any form of documentation based on workers' needs should aim to preserve moments of different and even conflicting viewpoints rather than striving for consensus, as is commonly sought in crowdsourcing tasks.

An initial low-fidelity prototype showing one of the many ways in which our findings and these design considerations could be integrated into a documentation framework is available at dataworksheet.com.

5.2 Research Implications: Challenges of Participatory Research

Leaning on the considerations argued by Le Dantec and Fox [57], in what follows, we reflect upon our work's limitations and outline the negotiations and interactions through which we constructed the relationship with our participants. By doing so, we bring to light our position as researchers within the power relations present at the research sites. We believe that the challenges we discuss here are not exclusive to our investigation and can be understood as more general implications for participatory research in CSCW.

First, conducting participatory research with organizations does not mean having full access to observe everything within them. The challenges of negotiating access to corporate sites are not new within CSCW [35, 77, 111, 114]. In many cases — including this research — the participating organizations influence what researchers are allowed to see and to whom they are allowed to talk. This affects the study sample since, often, individual participants are selected by the organizations' management. In this sense, our observations were enabled but also constrained by the organizations that hosted us.

Second, co-designing with partner organizations does not always mean getting to design with the intended *community*. Often, the gatekeeping tendency of organizations constitutes a barrier for researchers to establish rapport with participants and meaningfully engage in participatory processes with them. In our case, even though we had engaged in in-depth interviews with data workers and their views were prioritized in our feedback to the BPOs, the development of the first documentation prototypes was piloted by the companies' founders and managers, and, therefore, based on corporate priorities instead of workers' needs. After the co-design workshops, we realized that enabling spaces for workers to lead the design process would reveal ideas and requirements that had not surfaced before due to the intermediation of the management.

Third, including culturally and geographically diverse communities in participatory research and design is as necessary as it is challenging [63, 112]. For instance, in our investigation, working across language barriers came with its own set of challenges that, at times, counteracted the dynamic of the workshops. Such cases demand large amounts of researcher adaptability, flexibility, and creativity.

Finally, being critical while maintaining a relationship with partner organizations requires researchers to balance their commitments [99]. We acknowledge that a continuous challenge of our relationship with the BPOs has been finding the right balance between being explicit and critical with our feedback without losing access to the research sites [114]. This also surfaces an issue that many participatory projects face, namely, that of balancing contributions to the research community and to the community that is participating in the design process [57]. Keeping such balance required us to compromise and adapt while remaining reflexive and true to our commitment to the data workers.

6 CONCLUSION

In the context of industrial machine learning, discretionary decisions around the production of training data remain widely undocumented. Previous research has addressed this issue by proposing standardized checklists or datasheets to document datasets. With our work, we seek to expand that field of inquiry toward a documentation framework that is able to retrieve heterogeneous (and often distributed) contexts of data production.

For 2.5 years, we conducted a qualitative investigation with two BPO companies, Alamo and Action Data, dedicated to producing ML datasets. We engaged in a participatory process with data workers at both sites to explore the challenges and possibilities of documentation practices. The process involved phases of interviewing, data analysis, presentation and discussion of preliminary findings, the collaborative development of prototypes, feedback rounds, and a series of co-design workshops.

Previous practices at Alamo and Action Data reflect a constant of linear and top-down communication through documentation. However, as our findings show, designing a documentation framework based on the needs and ideas of data workers involves a different set of considerations. The data workers prioritized documentation that is able to transport information about tasks, teams, and payment. They also expected documentation to be able to reflect workers' feedback and communicate it back to the requesters. In this sense, task instructions could be co-created through iterations and communication enabled by documentation. Our findings include design considerations related to communication patterns and the incorporation of feedback, questions of access and trust, and the differentiation between the creation and use of documentation.

Together with our participants, we have re-imagined data documentation as a process and an artifact that travels among actors and organizations across cultural, social, and professional boundaries and is able to ease the collaboration of geographically distributed stakeholders while preserving the voices of data workers. Based on this, we have identified and approached several tensions and challenges that emerge from such a form of boundary object. Breaking with top-down communication and gatekeeping patterns by enabling feedback loops and promoting reflexivity in the documentation process have been, in this sense, helpful design considerations.

ACKNOWLEDGMENTS

Funded by the German Federal Ministry of Education and Research (BMBF) – Nr 16DII113, the International Development Research Centre of Canada, and the Schwartz Reisman Institute for Technology and Society. Our deepest gratitude goes to the workers and organizations that participated in this research. Thanks to Christopher Le Dantec, Marisol Wong-Villacres, and the anonymous reviewers for their comments and suggestions.

REFERENCES

- [1] [n.d.]. AI FactSheets 360. https://aifs360.mybluemix.net/
- [2] [n.d.]. Call For Datasets Benchmarks. https://neurips.cc/Conferences/2021/CallForDatasetsBenchmarks
- [3] [n.d.]. Google Cloud Model Cards. https://modelcards.withgoogle.com/about
- [4] Rikke Aarhus and Stinne Aaløkke Ballegaard. 2010. Negotiating boundaries: managing disease at home. In Proceedings of the 28th international conference on Human factors in computing systems - CHI '10. ACM Press, Atlanta, Georgia, USA, 1223. https://doi.org/10.1145/1753326.1753509
- [5] Mark S. Ackerman and Christine Halverson. 1998. Considering an organization's memory. In Proceedings of the 1998 ACM conference on Computer supported cooperative work - CSCW '98. ACM Press, Seattle, Washington, United States, 39–48. https://doi.org/10.1145/289444.289461
- [6] Shazia Afzal, C Rajmohan, Manish Kesarwani, Sameep Mehta, and Hima Patel. 2021. Data Readiness Report. In 2021 IEEE International Conference on Smart Data Services (SMDS). IEEE, Chicago, IL, USA, 42–51. https://doi.org/10.1109/ SMDS53860.2021.00016
- [7] Mariam Asad, Christopher A. Le Dantec, Becky Nielsen, and Kate Diedrick. 2017. Creating a Sociotechnical API: Designing City-Scale Community Engagement. 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, 2295–2306. https://doi.org/10.1145/3025453.3025963
- [8] Agathe Balayn, Bogdan Kulynych, and Seda Gürses. 2021. Exploring Data Pipelines through the Process Lens: a Reference Model for Computer Vision. (2021), 8.
- [9] Liam Bannon and Susanne Bødker. 1997. Constructing Common Information Spaces. In Proceedings of the Fifth European Conference on Computer Supported Cooperative Work. Springer Netherlands, Dordrecht, 81–96. https: //doi.org/10.1007/978-94-015-7372-6_6
- [10] Emily M. Bender and Batya Friedman. 2018. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics* 6 (2018), 587–604. https://doi.org/10.1162/tacl_a_00041
- [11] Dane Bertram, Amy Voida, Saul Greenberg, and Robert Walker. 2010. Communication, collaboration, and bugs: the social nature of issue tracking in small, collocated teams. In *Proceedings of the 2010 ACM conference on Computer* supported cooperative work - CSCW '10. ACM Press, Savannah, Georgia, USA, 291. https://doi.org/10.1145/1718918. 1718972
- [12] Susanne Bødker and Morten Kyng. 2018. Participatory Design That Matters—Facing the Big Issues. ACM Trans. Comput.-Hum. Interact. 25, 1, Article 4 (feb 2018), 31 pages. https://doi.org/10.1145/3152421
- [13] Claus Bossen, Lotte Groth Jensen, and Flemming Witt. 2012. Medical secretaries' care of records: the cooperative work of a non-clinical group. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work* -

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

CSCW '12. ACM Press, Seattle, Washington, USA, 921. https://doi.org/10.1145/2145204.2145341

- [14] Claus Bossen, Lotte Groth Jensen, and Flemming Witt Udsen. 2014. Boundary-Object Trimming: On the Invisibility of Medical Secretaries' Care of Records in Healthcare Infrastructures. *Computer Supported Cooperative Work (CSCW)* 23, 1 (Feb. 2014), 75–110. https://doi.org/10.1007/s10606-013-9195-5
- [15] Geoffrey C. Bowker and Susan Leigh Star. 1999. Sorting things out: classification and its consequences. MIT Press, Cambridge, Mass. https://mitpress.mit.edu/books/sorting-things-out
- [16] Eva Brandt, Thomas Binder, and Elizabeth Sanders. 2012. Tools and techniques: Ways to engage telling, making and enacting. In *Routledge Handbook of Participatory Design*, Jesper Simonsen and Toni Robertson (Eds.). Routledge, 145–181.
- [17] Tone Bratteteig, Keld Bødker, Yvonne Dittrich, Preben H. Mogensen, and Jesper Simonsen. 2012. Methods. Organising principles and general guidelines for Participatory Design projects. In *Routledge Handbook of Participatory Design*, Jesper Simonsen and Toni Robertson (Eds.). Routledge, 117–144.
- [18] Tone Bratteteig and Ina Wagner. 2016. Unpacking the Notion of Participation in Participatory Design. Computer Supported Cooperative Work 25, 6 (dec 2016), 425–475. https://doi.org/10.1007/s10606-016-9259-4
- [19] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3 (01 2006), 77–101. https://doi.org/10.1191/1478088706qp0630a
- [20] Virginia Braun, Victoria Clarke, Nikki Hayfield, and Gareth Terry. 2019. Thematic Analysis. In Handbook of Research Methods in Health Social Sciences, Pranee Liamputtong (Ed.). Springer Singapore, Singapore, 843–860. https://doi.org/10.1007/978-981-10-5251-4_103
- [21] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-Making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19*). Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300271
- [22] Paul Carlile. 2002. A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development. Organization Science 13 (08 2002), 442–455. https://doi.org/10.1287/orsc.13.4.442.2953
- [23] Kathy Charmaz. 2006. Constructing Grounded Theory: A Practical Guide through Qualitative Analysis. Sage Publications, London ; Thousand Oaks, Calif.
- [24] Michael Chibnik. 2020. Practical and Ethical Complications of Participatory Research. Annals of Anthropological Practice 44, 2 (Nov. 2020), 208–212. https://doi.org/10.1111/napa.12153
- [25] Eric Corbett and Christopher Le Dantec. 2019. Towards a Design Framework for Trust in Digital Civics. In Proceedings of the 2019 on Designing Interactive Systems Conference (San Diego, CA, USA) (DIS '19). Association for Computing Machinery, New York, NY, USA, 1145–1156. https://doi.org/10.1145/3322276.3322296
- [26] Terry Costantino, Steven LeMay, Linnea Vizard, Heather Moore, Dara Renton, Sandra Gornall, and Ian Strang. 2014. Participatory Design of Public Library E-Services. In Proceedings of the 13th Participatory Design Conference: Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium Papers, and Keynote Abstracts - Volume 2 (Windhoek, Namibia) (PDC '14). Association for Computing Machinery, New York, NY, USA, 133–136. https: //doi.org/10.1145/2662155.2662232
- [27] Sasha Costanza-Chock. 2020. Design Justice: Community-Led Practices to Build the Worlds We Need. The MIT Press, Cambridge, MA. https://design-justice.pubpub.org/
- [28] Kate Crawford and Trevor Paglen. 2019. Excavating AI: The Politics of Images in Machine Learning Training Sets. https://www.excavating.ai tex.ids: zotero-3263.
- [29] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society* 8, 2 (July 2021), 205395172110359. https://doi.org/10.1177/20539517211035955
- [30] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. 2020. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets. arXiv:2007.07399 [cs] (July 2020). http://arxiv.org/abs/2007.07399 arXiv: 2007.07399.
- [31] Catherine D'Ignazio and Lauren F. Klein. 2020. *Data feminism*. The MIT Press, Cambridge, Massachusetts. https://mitpress.mit.edu/books/data-feminism
- [32] Pelle Ehn. 2008. Participation in Design Things. In Proceedings of the Tenth Anniversary Conference on Participatory Design 2008 (Bloomington, Indiana) (PDC '08). Indiana University, USA, 92–101.
- [33] Jordan Famularo, Betty Hensellek, and Philip Walsh. 2021. Data Stewardship: A Letter to Computer Vision from Cultural Heritage Studies. (2021), 11.
- [34] Shaoyang Fan, Ujwal Gadiraju, Alessandro Checco, and Gianluca Demartini. 2020. CrowdCO-OP: Sharing Risks and Rewards in Crowdsourcing. Proceedings of the ACM on Human-Computer Interaction 4, CSCW2 (Oct. 2020), 1–24. https://doi.org/10.1145/3415203

- [35] Casey Fiesler, Jed R. Brubaker, Andrea Forte, Shion Guha, Nora McDonald, and Michael Muller. 2019. Qualitative Methods for CSCW: Challenges and Opportunities. In *Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing* (Austin, TX, USA) (*CSCW '19*). Association for Computing Machinery, New York, NY, USA, 455–460. https://doi.org/10.1145/3311957.3359428
- [36] Asbjørn Ammitzbøll Flügge. 2021. Perspectives from Practice: Algorithmic Decision-Making in Public Employment Services. In Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing. Association for Computing Machinery, New York, NY, USA, 253–255. https://doi.org/10.1145/3462204. 3481787
- [37] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. Commun. ACM 64, 12 (nov 2021), 86–92. https://doi.org/10.1145/3458723
- [38] R. Stuart Geiger, Kevin Yu, Yanlai Yang, Mindy Dai, Jie Qiu, Rebekah Tang, and Jenny Huang. 2020. Garbage in, Garbage out? Do Machine Learning Application Papers in Social Computing Report Where Human-Labeled Training Data Comes From?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (*FAT* '20*). Association for Computing Machinery, New York, NY, USA, 325–336. https://doi.org/10.1145/ 3351095.3372862
- [39] Yolanda Gil, Cédric H. David, Ibrahim Demir, Bakinam T. Essawy, Robinson W. Fulweiler, Jonathan L. Goodall, Leif Karlstrom, Huikyo Lee, Heath J. Mills, Ji-Hyun Oh, Suzanne A. Pierce, Allen Pope, Mimi W. Tzeng, Sandra R. Villamizar, and Xuan Yu. 2016. Toward the Geoscience Paper of the Future: Best practices for documenting and sharing research from data to software to provenance. *Earth and Space Science* 3, 10 (2016), 388–415. https: //doi.org/10.1002/2015EA000136 arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2015EA000136
- [40] Alyssa Goodman, Alberto Pepe, Alexander W. Blocker, Christine L. Borgman, Kyle Cranmer, Merce Crosas, Rosanne Di Stefano, Yolanda Gil, Paul Groth, Margaret Hedstrom, David W. Hogg, Vinay Kashyap, Ashish Mahabal, Aneta Siemiginowska, and Aleksandra Slavkovic. 2014. Ten Simple Rules for the Care and Feeding of Scientific Data. *PLOS Computational Biology* 10, 4 (04 2014), 1–5. https://doi.org/10.1371/journal.pcbi.1003542
- [41] Mark Graham, Isis Hjorth, and Vili Lehdonvirta. 2017. Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research* 23, 2 (may 2017), 135–162. https://doi.org/10.1177/1024258916687250
- [42] Mary L. Gray and Siddharth Suri. 2019. Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. Houghton Mifflin Harcourt, Boston.
- [43] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. 2018. A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. ACM, Montreal QC Canada, 1–14. https://doi.org/10.1145/3173574.3174023
- [44] Lisa Haskel. 2017. Participatory design and free and open source software in the not for profit sector: the Hublink Project. Ph.D. Dissertation. Bournemouth University.
- [45] Beverley Hawkins, Annie Pye, and Fernando Correia. 2017. Boundary objects, power, and learning: The matter of developing sustainable practice in organizations. *Management Learning* 48, 3 (July 2017), 292–310. https: //doi.org/10.1177/1350507616677199
- [46] Sarah Holland, Ahmed Hosny, Sarah Newman, Joshua Joseph, and Kasia Chmielinski. 2018. The Dataset Nutrition Label: A Framework To Drive Higher Data Quality Standards. arXiv:1805.03677 (2018). http://arxiv.org/abs/1805.03677
- [47] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. ACM, Virtual Event Canada, 560–575. https://doi.org/10.1145/3442188.3445918
- [48] Isto Huvila. 2011. The politics of boundary objects: Hegemonic interventions and the making of a document. *Journal of the American Society for Information Science and Technology* 62, 12 (2011), 2528–2539. https://doi.org/10.1002/asi.21639 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.21639
- [49] Lilly Irani. 2015. The cultural work of microwork. New Media & Society 17, 5 (2015), 720–739. https://doi.org/10. 1177/1461444813511926
- [50] 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 (CHI '13). Association for Computing Machinery, Paris, France, 611–620. https://doi.org/10.1145/2470654.2470742
- [51] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: strategies for collecting sociocultural data in machine learning. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. ACM, Barcelona Spain, 306–316. https://doi.org/10.1145/3351095.3372829
- [52] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and P. M. Krafft. 2020. Toward Situated Interventions for Algorithmic Equity: Lessons from the Field. In *Proceedings* of the 2020 Conference on Fairness, Accountability, and Transparency (Barcelona, Spain) (FAT* '20). Association for

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

Computing Machinery, New York, NY, USA, 45-55. https://doi.org/10.1145/3351095.3372874

- [53] Gunay Kazimzade and Milagros Miceli. 2020. Biased Priorities, Biased Outcomes: Three Recommendations for Ethics-oriented Data Annotation Practices. In *Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics,* and Society. (AIES '20). Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/ 3375627.3375809 tex.ids: kazimzade2020a, kazimzade2020b.
- [54] Sandjar Kozubaev, Fernando Rochaix, Carl DiSalvo, and Christopher A. Le Dantec. 2019. Spaces and Traces: Implications of Smart Technology in Public Housing. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300669
- [55] Jacob Leon Kröger, Milagros Miceli, and Florian Müller. 2021. How Data Can Be Used Against People: A Classification of Personal Data Misuses. https://papers.ssrn.com/abstract=3887097
- [56] Diana S. Kusunoki and Aleksandra Sarcevic. 2015. Designing for Temporal Awareness: The Role of Temporality in Time-Critical Medical Teamwork. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 1465–1476. https://doi.org/10.1145/2675133.2675279
- [57] Christopher A. Le Dantec and Sarah Fox. 2015. Strangers at the Gate: Gaining Access, Building Rapport, and Co-Constructing Community-Based Research. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (Vancouver, BC, Canada) (CSCW '15). Association for Computing Machinery, New York, NY, USA, 1348–1358. https://doi.org/10.1145/2675133.2675147
- [58] Charlotte P. Lee. 2005. Between Chaos and Routine: Boundary Negotiating Artifacts in Collaboration. In EC-SCW 2005, Hans Gellersen, Kjeld Schmidt, Michel Beaudouin-Lafon, and Wendy Mackay (Eds.). Springer-Verlag, Berlin/Heidelberg, 387–406. https://doi.org/10.1007/1-4020-4023-7_20
- [59] Charlotte P. Lee. 2007. Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work. Computer Supported Cooperative Work (CSCW) 16, 3 (June 2007), 307–339. https: //doi.org/10.1007/s10606-007-9044-5
- [60] Wayne G. Lutters and Mark S. Ackerman. 2002. Achieving safety: a field study of boundary objects in aircraft technical support. In Proceedings of the 2002 ACM conference on Computer supported cooperative work - CSCW '02. ACM Press, New Orleans, Louisiana, USA, 266. https://doi.org/10.1145/587078.587116
- [61] Wayne G. Lutters and Mark S. Ackerman. 2007. Beyond Boundary Objects: Collaborative Reuse in Aircraft Technical Support. Computer Supported Cooperative Work (CSCW) 16, 3 (June 2007), 341–372. https://doi.org/10.1007/s10606-006-9036-x
- [62] Michael A. Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376445
- [63] Henry Mainsah and Andrew Morrison. 2014. Participatory Design through a Cultural Lens: Insights from Postcolonial Theory. In Proceedings of the 13th Participatory Design Conference: Short Papers, Industry Cases, Workshop Descriptions, Doctoral Consortium Papers, and Keynote Abstracts - Volume 2 (Windhoek, Namibia) (PDC '14). Association for Computing Machinery, New York, NY, USA, 83–86. https://doi.org/10.1145/2662155.2662195
- [64] David Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a turker. In Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing. ACM, Baltimore Maryland USA, 224–235. https://doi.org/10.1145/2531602.2531663
- [65] Donald Martin Jr, Vinodkumar Prabhakaran, Jill Kuhlberg, Andrew Smart, and William S Isaac. 2020. Participatory problem formulation for fairer machine learning through community based system dynamics. arXiv preprint arXiv:2005.07572 (2020).
- [66] Laurie McLeod and Bill Doolin. 2010. Documents As Mediating Artifacts in Contemporary IS Development. In Proceedings of the 43rd Hawaii International Conference on System Sciences. IEEE, Honolulu, HI, USA, 1–10. https: //doi.org/10.1109/HICSS.2010.155
- [67] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. arXiv. https://doi.org/10.48550/arXiv.2205. 11963 arXiv:2205.11963 [cs] type: article.
- [68] 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, Article 34 (Jan. 2022), 14 pages. https://doi.org/10.1145/ 3492853
- [69] 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 (Oct. 2020), 1–25. https://doi.org/10.1145/3415186
- [70] Milagros Miceli, Tianling Yang, Laurens Naudts, Martin Schuessler, Diana Serbanescu, and Alex Hanna. 2021. Documenting Computer Vision Datasets: An Invitation to Reflexive Data Practices. In Proceedings of the 2021

ACM Conference on Fairness, Accountability, and Transparency. ACM, Virtual Event Canada, 161–172. https://doi.org/10.1145/3442188.3445880

- [71] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference* on Fairness, Accountability, and Transparency (FAT* '19). Association for Computing Machinery, 220–229. https: //doi.org/10.1145/3287560.3287596
- [72] Michael Muller and Angelika Strohmayer. 2022. Forgetting Practices in the Data Sciences. (2022), 30.
- [73] Michael Muller, Christine T Wolf, Josh Andres, Zahra Ashktorab, Narendra Nath Joshi, Michael Desmond, Aabhas Sharma, Kristina Brimijoin, Qian Pan, Evelyn Duesterwald, and Casey Dugan. 2021. Designing Ground Truth and the Social Life of Labels. (2021), 17.
- [74] Michael J. Muller. 2009. Participatory Design: the third space in HCI. In *Human-computer interaction*. CRC press, 181–202.
- [75] Samir Passi. 2018. Collaboration as Participation: The Many Faces in a Corporate Data Science Project. In The Changing Contours of "Participation" in Data-driven Algorithmic Ecosystems: Challenges, Tactics, and an Agenda' workshop in the 2018 ACM CSCW. https://www.samirpassi.com/pubs/working-papers/SamirPassi-CollaborationAsParticipation.pdf
- [76] Samir Passi and Steven J. Jackson. 2018. Trust in Data Science: Collaboration, Translation, and Accountability in Corporate Data Science Projects. Proc. ACM Hum.-Comput. Interact. 2, CSCW, Article 136 (nov 2018), 28 pages. https://doi.org/10.1145/3274405
- [77] Samir Passi and Phoebe Sengers. 2020. Making data science systems work. Big Data & Society 7, 2 (July 2020), 205395172093960. https://doi.org/10.1177/2053951720939605
- [78] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. 2020. Data and its (dis)contents: A survey of dataset development and use in machine learning research. arXiv:2012.05345 [cs] (Dec. 2020). http://arxiv.org/abs/2012.05345 arXiv: 2012.05345.
- [79] Laura R. Pina, Sang-Wha Sien, Teresa Ward, Jason C. Yip, Sean A. Munson, James Fogarty, and Julie A. Kientz. 2017. From Personal Informatics to Family Informatics: Understanding Family Practices around Health Monitoring. In Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (Portland, Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 2300–2315. https://doi.org/ 10.1145/2998181.2998362
- [80] Julian Posada. 2020. The Future of Work Is Here: Toward a Comprehensive Approach to Artificial Intelligence and Labour. Ethics in Context (2020).
- [81] Julian Posada. 2022. Embedded Reproduction in Platform Data Work. Information, Communication & Society (2022).
- [82] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI. arXiv:2204.01075 [cs] (April 2022). http://arxiv.org/abs/2204.01075 arXiv: 2204.01075.
- [83] Samantha Robertson and Niloufar Salehi. 2020. What If I Don't Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design. https://doi.org/10.48550/ARXIV.2007.06718
- [84] Toni Robertson and Jesper Simonsen. 2012. Participatory Design: An introduction. In *Routledge International Handbook* of *Participatory Design*. Routledge, 1–18.
- [85] T. Robertson and I. Wagner. 2012. Ethics: engagement, representation and politics-in-action. In *Routledge Handbook of Participatory Design*, Jesper Simonsen and Toni Robertson (Eds.). Routledge, 64–85.
- [86] Alex Rosenblat and Luke Stark. 2016. Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. International Journal Of Communication 10, 27 (2016), 3758–3784. https://doi.org/10.2139/ssrn.2686227
- [87] Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the crowdworkers?: shifting demographics in mechanical turk. In CHI '10 Extended Abstracts on Human Factors in Computing Systems. ACM, Atlanta Georgia USA, 2863–2872. https://doi.org/10.1145/1753846.1753873
- [88] M. J. Rothmann, D. B. Danbjørg, C. M. Jensen, and J. Clemensen. 2016. Participatory Design in Health Care: Participation, Power and Knowledge. In *Proceedings of the 14th Participatory Design Conference: Short Papers, Interactive Exhibitions, Workshops - Volume 2* (Aarhus, Denmark) (*PDC '16*). Association for Computing Machinery, New York, NY, USA, 127–128. https://doi.org/10.1145/2948076.2948106
- [89] Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbe, Kristy Milland, and Clickhappier. 2015. We Are Dynamo: Overcoming Stalling and Friction in Collective Action for Crowd Workers. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. ACM, Seoul Republic of Korea, 1621–1630. https://doi.org/10.1145/2702123.2702508
- [90] Nithya Sambasivan. 2022. All Equation, No Human: The Myopia of AI Models. Interactions 29, 2 (mar 2022), 78–80. https://doi.org/10.1145/3516515
- [91] 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

Proc. ACM Hum.-Comput. Interact., Vol. 6, No. CSCW2, Article 510. Publication date: November 2022.

2021 CHI Conference on Human Factors in Computing Systems. ACM, Yokohama Japan, 1–15. https://doi.org/10.1145/ 3411764.3445518

- [92] Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. CoDesign 4, 1 (2008), 5–18. https://doi.org/10.1080/15710880701875068 arXiv:https://doi.org/10.1080/15710880701875068
- [93] Devansh Saxena and Shion Guha. 2020. Conducting Participatory Design to Improve Algorithms in Public Services: Lessons and Challenges. In Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing. Association for Computing Machinery, New York, NY, USA, 383–388. https://doi.org/10.1145/ 3406865.3418331
- [94] Morgan Klaus Scheuerman, Emily Denton, and Alex Hanna. 2021. Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development. arXiv:2108.04308 [cs] (Sept. 2021). https://doi.org/10.1145/3476058
- [95] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R Brubaker. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. Proc. ACM Hum.-Comput. Interact. 4, CSCW1 (2020). https://doi.org/10.1145/3392866 Article 058.
- [96] Kjeld Schmidt. 2008. Taking CSCW Seriously: Supporting Articulation Work (1992). In Cooperative Work and Coordinative Practices. Springer London, 45–71. https://doi.org/10.1007/978-1-84800-068-1_3
- [97] Kristen M. Scott, Sonja Mei Wang, Milagros Miceli, Pieter Delobelle, Karolina Sztandar-Sztanderska, and Bettina Berendt. 2022. Algorithmic Tools in Public Employment Services: Towards a Jobseeker-Centric Perspective. In 2022 ACM Conference on Fairness, Accountability, and Transparency (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 2138–2148. https://doi.org/10.1145/3531146.3534631
- [98] Shilad Sen, Margaret E. Giesel, Rebecca Gold, Benjamin Hillmann, Matt Lesicko, Samuel Naden, Jesse Russell, Zixiao (Ken) Wang, and Brent Hecht. 2015. Turkers, Scholars, "Arafat" and "Peace": Cultural Communities and Algorithmic Gold Standards. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15). Association for Computing Machinery, New York, NY, USA, 826–838. https: //doi.org/10.1145/2675133.2675285 tex.ids= sen2015a.
- [99] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. Participation is not a Design Fix for Machine Learning. arXiv:2007.02423 [cs] (Aug. 2020). http://arxiv.org/abs/2007.02423 arXiv: 2007.02423.
- [100] Susan Leigh Star. 2010. This is Not a Boundary Object: Reflections on the Origin of a Concept. Science, Technology, & Human Values 35, 5 (Sept. 2010), 601–617. https://doi.org/10.1177/0162243910377624
- [101] Susan Leigh Star and James R. Griesemer. 1989. Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science* 19, 3 (Aug. 1989), 387–420. https://doi.org/10.1177/030631289019003001
- [102] David Stark. 2021. Algorithmic Management in the Platform Economy. 3, 2020 (2021), 47-72.
- [103] Marc Steen. 2013. Co-Design as a Process of Joint Inquiry and Imagination. Design Issues 29, 2 (04 2013), 16–28. https://doi.org/10.1162/DESI_a_00207 arXiv:https://direct.mit.edu/desi/article-pdf/29/2/16/1715163/desi_a_00207.pdf
- [104] Lucy Suchman. 2000. Embodied Practices of Engineering Work. Mind, Culture, and Activity 7, 1-2 (Jan. 2000), 4–18. https://doi.org/10.1080/10749039.2000.9677645
- [105] Lucy Suchman. 2002. Located accountabilities in technology production. Scandinavian journal of information systems 14, 2 (2002), 91–105. http://aisel.aisnet.org/sjis/vol14/iss2/7
- [106] Christine T. Wolf and Jeanette L. Blomberg. 2020. Ambitions and Ambivalences in Participatory Design: Lessons from a Smart Workplace Project. In Proceedings of the 16th Participatory Design Conference 2020 - Participation(s) Otherwise - Volume 1 (Manizales, Colombia) (PDC '20). Association for Computing Machinery, New York, NY, USA, 193–202. https://doi.org/10.1145/3385010.3385029
- [107] Pascale Trompette and Dominique Vinck. 2009. Revisiting the notion of Boundary Object. Revue d'anthropologie des connaissances 3, 1, 1 (2009), 3. https://doi.org/10.3917/rac.006.0003
- [108] Paola Tubaro and Antonio A. Casilli. 2019. Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics* (2019). https://doi.org/10.1007/s40812-019-00121-1 ISBN: 4081201900 Publisher: Springer International Publishing.
- [109] Paola Tubaro, Antonio A. Casilli, and Marion Coville. 2020. The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society* 7, 1 (2020). https://doi.org/10.1177/ 2053951720919776
- [110] Gabriela Vargas-Cetina. 2020. Do Locals Need Our Help? On Participatory Research in Anthropology. Annals of Anthropological Practice 44, 2 (Nov. 2020), 202–207. https://doi.org/10.1111/napa.12152
- [111] James R. Wallace, Saba Oji, and Craig Anslow. 2017. Technologies, Methods, and Values: Changes in Empirical Research at CSCW 1990 - 2015. Proc. ACM Hum.-Comput. Interact. 1, CSCW, Article 106 (dec 2017), 18 pages. https://doi.org/10.1145/3134741
- [112] Heike Winschiers. 2006. The challenges of participatory design in a intercultural context: designing for usability in Namibia. In PDC. 73–76.

- [113] Christine T. Wolf. 2019. Conceptualizing Care in the Everyday Work Practices of Machine Learning Developers. In Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion. Association for Computing Machinery, New York, NY, USA, 331–335. https://doi.org/10.1145/3301019.3323879
- [114] Christine T. Wolf, Julia Bullard, Stacy Wood, Amelia Acker, Drew Paine, and Charlotte P. Lee. 2019. Mapping the "How" of Collaborative Action: Research Methods for Studying Contemporary Sociotechnical Processes. In Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing (Austin, TX, USA) (CSCW '19). Association for Computing Machinery, New York, NY, USA, 528–532. https://doi.org/10.1145/3311957. 3359441
- [115] Marisol Wong-Villacres, Carl DiSalvo, Neha Kumar, and Betsy DiSalvo. 2020. Culture in Action: Unpacking Capacities to Inform Assets-Based Design. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376329
- [116] Alex J. Wood, Mark Graham, Vili Lehdonvirta, and Isis Hjorth. 2018. Good Gig, Bad Big: Autonomy and Algorithmic Control in the Global Gig Economy. Work, Employment and Society 00, 0 (2018), 1–20. https://doi.org/10.1177/ 0950017018785616
- [117] Allison Woodruff, Sarah E. Fox, Steven Rousso-Schindler, and Jeffrey Warshaw. 2018. A Qualitative Exploration of Perceptions of Algorithmic Fairness. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3173574.3174230
- [118] Xiaomu Zhou, Mark Ackerman, and Kai Zheng. 2011. CPOE workarounds, boundary objects, and assemblages. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, Vancouver BC Canada, 3353-3362. https://doi.org/10.1145/1978942.1979439
- [119] Carsten Østerlund. 2008. The Materiality of Communicative Practices. Scandinavian Journal of Information Systems 20, 1 (Jan. 2008). https://aisel.aisnet.org/sjis/vol20/iss1/4

Received January 2022; revised April 2022; accepted August 2022

510:34