

# The Data-Production *Dispositif*

MILAGROS MICELI\*, DAIR Institute, TU Berlin, and Weizenbaum Institute, Germany & USA  
 JULIAN POSADA\*<sup>†</sup>, Yale University, USA

Machine learning (ML) depends on data to train and verify models. Very often, organizations outsource processes related to data work (i.e., generating and annotating data and evaluating outputs) through business process outsourcing (BPO) companies and crowdsourcing platforms. This paper investigates outsourced ML data work in Latin America by studying three platforms in Venezuela and a BPO in Argentina. We lean on the Foucauldian notion of *dispositif* to define the *data-production dispositif* as an ensemble of discourses, actions, and objects strategically disposed to (re)produce power/knowledge relations in data and labor. Our *dispositif* analysis comprises the examination of 210 data work instruction documents, 55 interviews with data workers, managers, and requesters, and participant observation. Our findings show that discourses encoded in instructions reproduce and normalize the worldviews of requesters. Precarious working conditions and economic dependency alienate workers, making them obedient to instructions. Furthermore, discourses and social contexts materialize in artifacts, such as interfaces and performance metrics, limiting workers' agency and normalizing specific ways of interpreting data. We conclude by stressing the importance of counteracting the data-production *dispositif* by fighting alienation and precarization, and empowering data workers to become assets in the quest for high-quality data.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing**; • **Information systems** → **Crowdsourcing**; • **Applied computing** → *Annotation*; • **Computing methodologies** → **Artificial intelligence**.

Additional Key Words and Phrases: data production, data work, machine learning, data labeling, platform labor, crowdsourcing

## ACM Reference Format:

Milagros Miceli and Julian Posada. 2022. The Data-Production *Dispositif*. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 460 (November 2022), 37 pages. <https://doi.org/10.1145/3555561>

## 1 INTRODUCTION

Many machine learning (ML) models are built from training data previously collected, cleaned, and annotated by human workers. Companies and research institutions outsource several of these tasks through online labor platforms [64] and business process outsourcing (BPO) companies [53]. In these instances, outsourcing organization and their clients regard workers as independent contractors, considering them factors of production, and their labor a commodity or a product subject to market regulations [87]. They are paid as little as a few cents of a dollar per task, usually lack social protection traditionally tied with employment relations, and are subject to systems of control and surveillance [30, 37, 81]. Their assignments broadly comprise the interpretation and classification of data, and their work practices involve subjective social and technical choices that

\*Equal contribution.

<sup>†</sup>Also with University of Toronto, Canada.

Authors' addresses: Milagros Miceli, [m.miceli@tu-berlin.de](mailto:m.miceli@tu-berlin.de), DAIR Institute, TU Berlin, and Weizenbaum Institute, Germany & USA; Julian Posada, [julian.posada@yale.edu](mailto:julian.posada@yale.edu), Yale University, New Haven, CT, USA.



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

© 2022 Copyright held by the owner/author(s).  
 2573-0142/2022/11-ART460  
<https://doi.org/10.1145/3555561>

influence data production and have ethical and political implications. Workers interpreting and classifying data do not do so in a vacuum: their labor is embedded in large industrial structures and deeply intertwined with naturalized profit-oriented interests [44].

This paper presents an investigation of data production for ML as carried out by Latin American data workers mediated by three platforms operating in Venezuela and a business process outsourcing (BPO) company located in Argentina. To study *data work* for machine learning, which we define as the labor involved in the collection, curation, classification, labeling, and verification of data, we lean on Foucault’s notion of *dispositif* and apply the method of *dispositif analysis* [39]. A *dispositif* is an ensemble of objects, subjects, discourses, and practices as well as the relations that can be established between them [25]. Examples of *dispositifs* include prisons, police, and academia. These heterogeneous ensembles of discursive and non-discursive elements constitute what is perceived as reality and, as such, what is taken for granted.

The decision to lean on Foucault’s notion of *dispositif* is methodological, rather than theoretical. This notion and the method of *dispositif analysis* enables the study of data production as embedded in social interactions and hierarchies that condition how data is constructed and how specific discourses are reproduced. As we will describe in Section 3.1, this method also allowed us to integrate diverse qualitative data and focus on the relationships between them. We define the *data-production dispositif* as the network of discourses, work practices, hierarchies, subjects, and artifacts comprised in ML data work (see Figure 3) and the power/knowledge relationships that are established and naturalized among the three elements. The *data-production dispositif* determines the realities that ML datasets can reflect and the ones that remain erased from them. It has a crucial effect on the outputs that ML models will consider to be true. *Dispositif analysis* interrogates means of reality making, with a special focus set on the meanings that become dominant and those that are marginalized — “the said as much as the unsaid” [25]. Our *dispositif analysis* explores the sites where the production of ML data is outsourced. It comprises the investigation of (1) *linguistically-performed elements* (what is said/written), (2) *non-linguistically performed practices* (what is done), and (3) *materializations* (how linguistically and non-linguistically performed practices translate into objects) [39].

These elements and research questions relate specifically to the outsourcing of ML data-production tasks and can be structured as follows:

- **Linguistically performed elements:** *What discourses are present in task instructions provided to outsourced data workers?* (RQ1)  
→ We analyzed a corpus of 210 instruction texts for data-related tasks requested by ML practitioners and outsourced to data workers.
- **Non-linguistically performed practices:** *How do outsourced data workers, managers, and requesters interact with each other and instruction documents to produce data?* (RQ2)  
→ To explore how linguistically performed elements translate into practice, we conducted 41 interviews with data workers and inquired how they interpret the instructed tasks. In addition, we conducted interviews with six managers and eight ML practitioners (in their role as data-work requesters).
- **Materializations:** *What artifacts support the observance of instructions, and what kind of work they perform?* (RQ3)  
→ Through participant observation, we account for some of the material elements in which the *data-production dispositif* manifests, such as platforms and interfaces, tools to surveil workers, and documents that record the decisions made between service providers and service requesters.

To summarize, this paper, its contribution, and the extensive analysis it comprises can be described as follows: We start by exploring Foucault's notion of *dispositif* and defining key related concepts such as power, knowledge, and discourse. Then, we review previous investigations that have discussed ML data work and further define the scope of the data-production *dispositif*. After offering an overview of the *dispositif* analysis method, informants, and fieldwork sites, we present our findings. They are organized around the three elements that form the *dispositif*.

The findings show that, instead of seeking the “wisdom of crowds,” where a diverse and independent group cooperates to solve a problem, requesters use task instructions to impose predefined forms of interpreting, classifying, and sorting data that respond primarily to profit-oriented interest. Managers in BPOs and algorithms in labor platforms are in charge of overseeing the process. Poverty and dependence in the areas where data work is outsourced leaves workers with no other option but to obey and avoid questioning instructions. Documents, tools, and interfaces constitute some of the *dispositif*'s materializations. Given these findings, we outline some implications and propose three ways of counteracting the current composition of the data-production *dispositif* and its effects by fighting workers' precarization, alienation, and surveillance. Finally, we discuss the limitations of our investigation.

## 2 DEFINING KEY CONCEPTS

### 2.1 *Dispositif* and Other Foucauldian Concepts

Foucault proposes a relational conception of *power* [25] and argues that its exercise takes place in networks of relations rather than being placed in a specific social location [14]. He defines power as “a whole series of particular mechanisms, definable and defined, that seem capable of inducing behaviours or discourses” [24]. Therefore, power is not held by or exercised over individuals but works through the impersonal relations of force and strategy that connect subjects [14]. Power operates through practices that act upon subject's present or future actions. It is effective as long as it is normalized, that is, taken for granted and perceived as the inevitable way things are.

Present throughout Foucault's power analysis is the implicit relationship between power and knowledge, whereby one implies the other. He understands *knowledge* as entangled with discursive power and describes it as the power to define others and to produce truth through discourses. Knowledge and power are integrated with one another because “it is not possible for power to be exercised without knowledge, it is impossible for knowledge not to engender power” [25].

In a related manner, Foucault uses the term *discourse* to refer to a historically contingent system that produces knowledge and meaning [22]. Discourse is not only a way of organizing and presenting knowledge, but also, it can structure social practices and the relations that emerge through the collective understanding of reality encoded in discourse [21]. Discourses encode power in the sense that they can determine reality. Subjects are active participants in reality-making processes as co-producers of discourses, which puts explicit and implicit knowledge at their disposal [39]. Thus, discursive power is not exercised *on* subjects but flows *through* them. However, discourses still have a disciplinary effect on subjects as they assure the prevalence of certain knowledge of what can be said, done, and thought.

Power, knowledge, and discourse finally converge in *dispositif*, a notion that expands discourse to include non-discursive practices and artifacts. Foucault defines *dispositif* as “a thoroughly heterogeneous ensemble consisting of discourses, institutions, architectural forms, regulatory decisions, laws, administrative measures, scientific statements, philosophical, moral and philanthropic propositions – in short, the said as much as the unsaid . . . The [dispositif] itself is the system of relations that can be established between these elements” [25]. There is a multiplicity of *dispositifs* that influence each other and have strategic functions within power relationships. As Foucault puts it,

dispositifs respond to an “urgent need” that is bound to specific historical and geographical contexts. Link [48] draws attention to the etymological root shared by the French words “*disposition*” and “*dispositif*.” In colloquial French, “disposition” is used in the sense of being “at someone’s *disposal*.” This way, Link highlights the power element comprised in *dispositif* as the separation between those who are “at the disposal” (those who are instrumental) and those who have influence to determine the strategy used to meet a need [9]. Those who are “at the disposal” and those who “dispose” are part of the *dispositif*’s strategy.

In sum, the concept of *dispositif* comprises the knowledge that is built into linguistically performed practices (what is said, written, and thought), non-linguistically performed practices (what is done), and materializations (the objects) [39, 41]. A *dispositif* can therefore be defined as a constantly changing network of objects, subjects, discourses, and practices that shape each other, producing new knowledge and new power.

Previous HCI and CSCW research has engaged with Foucauldian theory. For instance, Harmon and Mazmanian [34] follow Foucault’s understanding of *discourse* to explore how US residents talk about smartphones and smartphone users. Kou et al. [46] use the Foucauldian concepts of *power*, *knowledge*, and *self* to explicate human-technology relationships. And Bardzell et al. [4] draw on Foucault’s Theory of Identity to investigate social practices within the virtual world Second Life. In terms of methodology, Kannabiran et al. [42] use Foucauldian Discourse Analysis (FDA) to study the rules and mechanisms involved in HCI discourses on sexuality, while Spiel [78] combines Actor-Network Theory with FDA into a “critical experience” framework to evaluate how children in the autistic spectrum interact with technologies. Despite the wide application of Foucauldian theory in HCI and CSCW, the *dispositif* notion and analysis has not been applied to the study of data production and data work. To that end, we believe that our contribution could be seen as methodological, in the sense that we intend to show a novel and comprehensive mode of analysis to approach data work.

## 2.2 Data Work for Machine Learning

The data-production *dispositif* analyzed in this paper comprises the “infrastructure” that enables the (re-)production and circulation of specific discourses in and through ML data work. As Foucault argues, the emergence of each *dispositif* responds to an “urgent need.” The data-production *dispositif* responds to the growing demand for data and labor in the AI industry.

We define *data work* as the human labor necessary for data production, in this case, for machine learning. Data work involves the collection, curation, classification, labeling, and verification of data. Users, developers, and outsourced workers carry out these tasks at any point in the development and deployment of AI systems. For example, medical professionals in the case of AI for healthcare [6, 55, 79], education professionals [50], or internet users when answering ReCAPTCHA tests [40]. This paper will employ the term “data work” to refer exclusively to the labor outsourced through crowdsourcing platforms and specialized business process outsourcing (BPO) companies, instead of the broader data work carried out by other professionals and users, while acknowledging the role of the former within the *dispositif* as requesters.

Platforms are one of the two significant ways of outsourcing data work. The rise of alternative forms of work different from traditional employment [43] and the expansion of the “gig economy,” or casual employment mediated through platforms [89], gave rise to “crowdsourcing,” “crowdwork,” or “digital piecework” platforms where geographically dispersed workers are allocated many fragmented tasks, which are carried out online from their homes. Platforms are hybrid organizations that combine traits of firms and multi-sided markets [11]. They serve as infrastructures that “facilitate and shape personalised interactions among end-users and complementors, organised through the systematic collection, algorithmic processing, monetisation, and circulation of data”

[63]. Platforms thrive in digital environments because they respond to deficiencies in markets and enterprises that fail to extract and appropriate data and allocate resources efficiently [11].

The second primary form of outsourced data work for ML is provided by business process outsourcing (BPO) companies. Conversely to crowdsourcing platforms, where hierarchies are primarily managed by algorithms, BPOs show rather traditional management structures. BPO is a form of outsourcing that involves contracting a third-party service provider to carry out specific parts of a company's operations, in the case of our investigation, data-related tasks. These service providers often specialize in one type of ML data service (e.g., semantic segmentation) or application domain (e.g., computer vision), contrary to platforms specializing in one or a few application domains but with more diverse data services. While prices per piece are significantly higher than those offered by platforms, many machine learning companies prefer to outsource their data-related projects with BPOs because of the perceived higher quality of data [54]. This is due to the companies' domain specialization and traditional managerial structures that allow more direct and personal communication.

The intervention of humans in processes of data production has been addressed by a large body of CSCW and HCI research [19, 26, 53, 57, 58, 61, 62, 77, 80]. Some investigations have explored the role of worker subjectivity on datasets [7, 12, 29, 84] and have proposed ways to recognize and address worker bias [3, 28, 36, 84]. In contrast, other researchers have documented the "practices, politics, and values of workers in the data pipeline" [73, 74] and the sociotechnical organization of data work that privileges speed, scale, and scalability over worker wellbeing [37], low wages [18, 33], dependency [71], and the power asymmetries vis-à-vis requesters [38, 52, 53, 72].

As we argue, the "urgent need" addressed by the data-production *dispositif* is the exponential need for cheaper and more profitable data, which is also the exploitation of surveillance [90], natural resources [15], and other types of labor [13]. Previous research has highlighted the role of these elements to guarantee a façade where AI is seen as neutral, unbiased, and efficient due to the lack of human intervention — and error — while keeping workers and factors of production hidden from the public lens [10, 30, 37]. These elements show the wide extension of the "heterogeneous ensemble" that constitutes the discursive, non-discursive, and material elements of data production. Because the data-production *dispositif* is too vast to explore in one academic paper, we circumscribe its analysis around outsourced data work for ML as one of its crucial components.

### 3 METHODOLOGY

#### 3.1 *Dispositif* Analysis of Data Production

We lean on *dispositif* analysis to investigate the discourses implicit in annotation instructions, the non-discursive practices involved in the production of ML datasets, and how both materialize in artifacts. Often described as an extension of discourse analysis [9], *dispositif* analysis expands the field of inquiry beyond texts to include actions, relationships, and objects. *Dispositif* analysis rests on the notion of knowledge (and power) as the connecting force between discursive and non-discursive components. It accounts for hierarchies and power structures in societal fields and organizations that shape the construction of meaning in discourse [68]. Thus, our *dispositif* analysis crucially focuses on the relationship between discourse, practice, and objects in data production, and the power created through their interaction.

Foucault never outlined an explicit methodology of *dispositif* analysis. Several authors, most prominently Sigfried Jäger [39, 41], have explored ways of operationalizing the complex Foucauldian notion of *dispositif* into a method of inquiry. *Dispositif* analysis has thus been in constant evolution since the mid-1980s. Caborn [9] mentions four steps comprised in this methodology: (1) identifying the elements that constitute the *dispositif*, (2) determining which discourses they embody and their

entanglement with other discourses, (3) interrogating power by “considering who or what is *at the disposal* of whom”, and (4) analyzing non-discursive practices associated with the dispositif’s discourses. However, to our knowledge, a comprehensive guide or method of *how to* conduct a dispositif analysis has not yet been developed. As Jäger and Maier describe, dispositif analysis remains “a flexible approach and systematic incitement for researchers to develop their analytic strategies, depending on the research question and type of materials at hand” [39].

The study presented in this paper follows the experimental spirit of Jäger and Maier’s invitation to develop our analytical strategy, in this case, to study data production. Here, we combine methodological elements discussed by several authors in terms of the operationalization of power, knowledge, and discourse [8, 9, 47, 48, 60], apply them to our fieldwork on data production through platforms and at a BPO, and follow the examples provided by previous research that has successfully applied variations of dispositif analysis [9, 31, 51, 85, 86]. We followed the four steps outlined by Caborn mentioned above and based our analysis on the three-dimensional framework described by Jäger and Maier [39] as follows:

- The *analysis of linguistically performed elements*: which aims at reconstructing the knowledge built into what is said and written through discourse analysis. In terms of our investigation, this phase comprised an examination of the discourses encoded in the instruction documents received by data workers.
- The *analysis of non-linguistically performed practices*: which aims at reconstructing the knowledge that underlies linguistically performed practices and how they translate into action. In this phase, we investigated how workers make sense of the task instructions and their work in general, the interactions between workers, managers, and clients, and the labor conditions that structure these practices. We studied these elements through interviews conducted with data annotators who perform tasks guided by such instructions, machine learning practitioners who compose annotation instructions, and managers who oversee the process.
- The *materializations*: This phase of analysis consisted of identifying the knowledge that is built into physical and digital artifacts, i.e., discursive materialization, whose existence is coherent with the discourses they encode. Through this lens, we set the focus on the platforms and interfaces used to perform data work, documents (as artifacts and not as texts) that record decision-making processes, and tools used to surveil workers and quantify their performance. Our analysis of these materializations is based on participant observations and the above-mentioned interviews.

### 3.2 Researcher Positionality

Making researchers’ positionality explicit is key to situate the standpoint from which an investigation has been conducted. Positionality statements are relevant to all types of studies, specially qualitative and exploratory investigations such as this one. Moreover, given the flexible character of dispositif analyses, it seems appropriate to disclose some elements of the authors’ backgrounds that might have informed the analysis presented in this paper.

Both authors are multiracial researchers born in different countries of Latin America. Both are first-generation academics working in institutions located in the Global North, where they live under immigrant status. Both have a background in Sociology and Communication. Their first language is Spanish. One of the authors identifies as female and the other as male. Both are cisgender. Despite being born and raised within worker-class families and in the same regions as the data workers interviewed, the authors acknowledge that their class-related experiences differ from those of the interview partners and that their position as researchers living and working

in the Global North provides the authors with privilege that the study participants do not hold. Throughout data collection, analysis, and while considering the implications of this investigation, the authors have put much effort in remaining reflexive and acknowledging their position regarding the study participants and field of inquiry.

### 3.3 Data Collection and Analysis

This investigation comprises several weeks of participant observation, a total of 55 interviews, and the analysis of 210 instruction documents. These data were collected during several months of fieldwork from 2019 to 2021, online and in person, at two sites (see Table 1):

- virtually studying three crowdsourcing platforms operating in Venezuela and the experiences of platform workers, and,
- in a hybrid format, at a business process outsourcing company (BPO) located in Buenos Aires, Argentina, where data workers perform tasks related to the collection and labeling of data for machine learning.

At both fieldwork sites, we conducted participant observations and semi-structured interviews. To complement these data, we conducted a series of expert interviews with managers at other BPOs and with ML practitioners in their role of data-work requesters (see Section 3.3.2 and Table 3 for a detailed account of the interview participants).

#### 3.3.1 *Fieldwork.*

Fieldwork in Venezuela was carried out virtually between July 2020 and June 2021 due to restrictions related to the coronavirus pandemic. For the first phase of this research, we signed up and completed tasks for the platforms to understand the tasks available, working conditions, and interfaces. While some of these platforms presented similar tasks, they differed considerably in their general availability, the interfaces, labor process, and task applications. Initially, we contacted the platform workers using convenience sampling since this population is invisible, meaning that it is difficult to approach them without the support of the platforms, which we did not have for the study. We sought permission from the moderators of the most popular worker groups on Facebook and Discord to post a call for study participants. Thanks to this initial approach, we were able to use snowball sampling to contact further participants. We conducted in-depth interviews and asked workers about their experience working for the platforms. Additionally, we asked workers if they could share information about the instructions they received. Some workers also shared guides created by colleagues to understand and answer the tasks efficiently. We include those in our analysis as well. Moreover, we also searched the internet to find additional annotation instructions online. Our approach, notably the use of convenience and snowball sampling, present several limitations in terms of reproducibility and bias towards participants belonging to similar social circles. We have mitigated these issues by comparing our results with similar studies (see Section 2.2) and with the workers at the BPO company.

Fieldwork at Alamo, the Argentine business process outsourcing (BPO) company, was carried out in-person between May and June 2019 in Buenos Aires and continued online between August 2020 and February 2021. At the time of this investigation, this company is a medium-sized organization with branches in several Latin American countries. Alamo conducts, data collection, data annotation, content moderation and software testing projects. The company is an impact sourcing type of BPO, which refers to a branch of the outsourcing industry that purposely employs workers from poor and marginalized populations to offer them a chance in the labor market and to provide information-based services at lower prices. We contacted the company via e-mail to request field access. After several months of inquiry, a meeting with the company's management took place in which the researcher on site signed a non-disclosure agreement that specified several elements that

Table 1. Fieldwork sites: Studied data work BPO and platforms

Entity	Type	Primary Tasks	Applications
ALAMO	BPO	Data collection and annotation Content moderation	E-commerce
CLICKRATING	Platform	Data entry Algorithmic verification	Online search engine
TASKSOURCE	Platform	2D/3D image classification 2D/3D semantic segmentation	Self-driving vehicles Internet of things
WORKERHUB	Platform	2D image, text and video classification 2D semantic segmentation Text transcription	Content moderation E-commerce

we are not allowed to disclose in this or other papers. Most of these elements concern the identity of clients and specific details about their ML models. After this meeting, fieldwork was allowed to commence and we were able to observe several projects related to the collection and annotation of data for ML. Apart from shadowing workers, we were granted access to team meetings, meetings with clients, workers' briefings, and QA analysis related to three projects carried out by the company in 2019, involving the collection and labeling of image data. We complemented the observations with in-depth interviews with data workers, managers, and QA analysts.

### 3.3.2 Instruction Documents.

In total, we collected 210 annotation instruction documents from the platforms and the BPO. The analysis of the instructions was carried out by both authors. We used critical discourse analysis [39] to explore the instruction texts. The analysis comprised three stages: (1) the structural analysis of the corpus, (2) a detailed analysis of discourse fragments, and (3) a synoptic analysis. These steps (especially the synoptic analysis) included several iterations that allowed us to discover connections between different levels of analysis, collect evidence to support our interpretations, and develop arguments. Table 2 offers an overview of the codes used for the discourse analysis of the instruction documents, their evolution throughout the three phases of analysis, and explanatory memos that reflect our understanding of each code.

The goal of the *structural analysis* is to code the material to identify significant patterns and recurring themes and sub-themes comprised in the instructions. By the end of the structural analysis phase, we were able to identify elements of the text structure, regular tasks, and stylistic devices that appeared in the instruction documents. These elements helped us identify "typical" texts and representative discourse fragments for the following analysis step.

The *detailed analysis* comprised an examination of selected text fragments. We focused on identifying typical representations and their variations and interrogated the elements highlighted in the instruction documents and the contextual knowledge that is taken for granted and, thus, neglected in them. We also paid special attention to binary reductionisms, presupposition and attribution, examples, and visualizations. A critical aspect of the analysis focused on the taxonomies that structure the labels instructed by requesters.

Finally, the *synoptic analysis* included the overall interrogation of the observations that emerged from the structural and detailed analyses. This phase included an intensive exchange between both authors to reflect upon our shared understanding of the identified discourse strands. We contrasted the selected fragments and the identified elements with the interview and observation material. This approach helped us understand the role of work and managerial practices in legitimizing specific discourses.



Table 2. Evolution of codes throughout the three phases of discourse analysis as applied to the instruction documents

Structural Analysis	Detailed Analysis	Synoptic Analysis	Memo
Document characteristics	Format	Documents as constraints	Document elements that hinder the correct understanding or completion of tasks. Elements that constrain workers' interpretation of data.
	Document language		
	Language barriers		
	Document version	Document as artifacts	Evidence of documents evolving / being modified by requesters. Different versions of same task.
1st Person			
Descriptions	Project description	Worker alienation	Evidence of workers kept in the dark about the ML pipeline, made to feel foreign to the products of their labor. Workers forced to automatize their outputs.
	AI description		
	Use of tech jargon		
Task type	Data collection	Data generation	Data scraping, collection. Tasks that include taking or intervening pictures and/or generating texts.
	Classification		
	Involves text		
	Data labeling	Data annotation	Data segmentation and classification. Task related to labeling and keywording.
	Segmentation		
	Involves images		
	Keywording	Algorithmic verification	Assesment of algorithmic output by workers. Testing of ML systems, rating of search engine queries, moderation of content flagged by algorithm.
	Evaluation		
	Testing		
	Rating	Algorithmic impersonation	Workers are instructed to act as an AI and rendered invisible in the process.
Content moderation			
Taxonomies	Categories	Multiclass classification	Classifications that include more than two options. Rationale behind classes. Normalization.
	Classes	Binary classification	Binary classifications and simplification of complex phenomena.
	Attributes	Other/Unknown	Use of a third label, especially in relation to binary classifications. Otherness.
	Labels definition	Ambiguity/Exceptions	How ambiguity and marginal cases are dealt with.
Examples	Example description	Errors/Discrepancies	Cases where instructions contain errors. Discrepancies between instruction and interface.
	Counter-example	Explicit	What is explicitly described. Rationale behind taxonomies and examples made explicit.
	Clarifications		
	Interface description	Implicit	What remains unsaid. What is considered self-evident. Implicit rationale behind taxonomies.
Use of images			
Disturbing content	Content warnings	Exposure to disturbing content	Tasks that include dealing with that is sexual or violent in nature.
	Content inclusion		
Worker skills	Experience	Language	Discrepancies between instructions language and workers' native language. Problems and strategies. Use of Google Translate. Guides developed by workers in Spanish.
	Language skills		
	Technical skills	Quantification and surveillance	References to how the performance of workers is measured and surveilled. Consequences for low scores.
Warnings	Score	Worker obedience	How the unquestioning obedience is fostered in task instructions. How workers are prompted to think in terms of what the requester wants.
	Ban		
	Speed	Worker precarization	Precarious labor conditions made explicit in instruction documents. Threats of being banned from the task. References to piece rate pay. Arbitrary definitions of worker accuracy.
	No payment		
Accuracy			

### 3.3.3 Interviews.

To reconstruct the knowledge that underlies the practices that constitute the data-production dispositif, we turned to the experiences of those actors who interact with the instructions regularly. With this aim, we conducted a total of 55 interviews with dataworkers located in Venezuela and Argentina, BPOs managers and founders, and ML practitioners who regularly outsource data-related tasks. Table 3 shows a detailed overview of the interview partners, including their role within their organizations, location, type of interview, and language.

Due to the restrictions related to the COVID-19 pandemic, the interviews with platform workers were conducted online through video calls, while those with BPO workers were conducted in person before the pandemic. The interviews with the data workers were conducted in Spanish, which is the native language of the interviewers and the participants. We conducted the expert interviews in English. While most platform workers had tried several platforms, they usually focused on one, except for two workers who worked simultaneously for Tasksource and Workerhub. All interview partners were asked to choose a code name or were pseudonymized post-hoc to preserve their identity and that of related informants.

The goal of the interviews was to reveal practices and perceptions and obtain additional information about the organizational relations and structures that inform how data-related tasks come about, how instructions are communicated, and how workers execute them. For instance, hierarchical structures can have an essential effect on meaning-making practices as enacted through the annotation instructions without being referred to explicitly or implicitly in the instruction documents. The in-depth interviews with data workers include accounts of specific work situations involving the interpretation of data. Moreover, they cover task descriptions, widespread routines, practices, working conditions, lived experiences, and general views on their work and the local labor market.

It is essential to mention that the differentiation between in-depth and expert interviews refers to the interview method chosen for each situation and informant and was not based on informants' occupational status or position. Our priority was engaging in in-depth conversations with data workers to discuss and learn from their experiences in and beyond data work. Conversely, we used the expert-interview method to conduct focused exchanges with actors that possessed a broad overview of the machine learning pipeline. The expert interviews covered the topics of data work and the relationship between BPO/platform and requesters.

Dispositif analysis allowed us enough flexibility to obtain valuable insights from the interviews by combining inductive and deductive coding. Some of the topics that we identified through discourse analysis in the instruction documents helped us build categories to code the interviews. This form of deductive coding was oriented towards finding additional evidence for phenomena identified in the instruction texts and understanding the contexts in which instructions are formulated and carried out. In addition, there was room for inductive category formation so that several codes could emerge directly from the interviews during coding. This approach helped us identify valuable observations that otherwise have been lost. Through this form of analysis, we aimed at identifying patterns. Those patterns were later confronted with the elements identified in the instruction texts and complemented with participant observations. The development of coding schemes for the analysis and the coding process itself was carried out in iterations involving cross-coding between both authors. The interview transcripts were analyzed in their original language (Spanish or English). The excerpts included in Section 4 were translated by us when writing this paper and only after the analysis phase. Our emergent understanding evolved throughout numerous discussions and several iterations until reaching the set of findings that we present in Section 4.

Table 3. Overview of interview partners and interview characteristics

	ORGANIZATION	INTERVIEW METHOD	MEDIUM	LANG.	INFORMANTS
<b>WORKERS</b>	Alamo (BPO in Argentina)	In-depth interview	In-person	Spanish	10 data workers
	Tasksource (Platform in Venezuela)	In-depth interview	Zoom	Spanish	8 data workers
	Workerhub (Platform in Venezuela)	In-depth interview	Zoom	Spanish	19 data workers
	Clickrating (Platform in Venezuela)	In-depth interview	Zoom	Spanish	6 data workers
<b>MANAGERS</b>	Data processing company (Bulgaria)	Expert interview	Zoom	English	1 company founder
	Data processing company (Iraq)	Expert interview	Zoom	English	1 general manager 1 program manager
	Data processing company (Kenya)	Expert interview	Zoom	English	1 country manager
	Data processing company (India)	Expert interview	Zoom	English	1 director of ML services 1 project manager
<b>REQUESTERS</b>	Computer vision company (Germany)	Expert interview	In-person	English	1 data protection officer 1 co-founder 1 product manager 1 CV engineer
	Machine learning company (USA)	Expert interview	Zoom	English	1 product engineer
	Machine learning company (Spain)	Expert interview	Zoom	Spanish	1 lead engineer
	Computer vision company (Bulgaria)	Expert interview	Zoom	English	1 co-founder 1 CV engineer

### 3.3.4 Observations.

Through fieldwork at the BOP and the platforms, we were able to observe interactions among data workers and between them and clients using in-person observation in the Argentinian case and digital ethnography in the Venezuelan one [35]. Furthermore, we observed workers' interactions with crowdsourcing platforms and the software interfaces used to complete annotation tasks. Special attention was paid to the interaction of workers with task instructions. It is important to mention that the instruction documents underwent a twofold form of analysis: On the one hand we analyzed instruction documents as texts through discourse analysis as described in Section 3.3.2. On the other hand, we used the observations conducted to analyze these documents as artifacts or materializations of the data-production dispositif. For the latter form of analysis, the focus was set on documents' function, provenance, and the interactions they allow or constrain.

The level of involvement regarding observations varied from shadowing to active participant observations. In some cases, we had the opportunity to observe and try the interfaces and perform data annotation tasks for several hours. All observations were recorded as jottings taken in real-time. Those jottings comprised descriptions of briefings, meetings, tasks, documents, communication channels, and interfaces, as well as their advantages and limitations. In parallel, reflections on the researchers' impressions and perceptions, including explicitly subjective interpretations, were noted. Simple sketches and, when permitted, photos helped to complete the observations registered. The information gathered in keywords or bullet points was later transformed into complete texts and integrated into more consolidated field notes. In the analysis phase, we combined these field notes with the interview transcripts and coded them following the steps described in the previous subsection.

## 4 FINDINGS

The presentation of our findings comprises a descriptive subsection (4.1) and three analytical parts (4.2, 4.3, and 4.4).

In 4.1, we present several examples of the different tasks carried out by data workers at each one of our four fieldwork sites. We include details of how task instructions are formulated and

Table 4. Types of tasks in outsourced data work based on Tubaro et al. [81]

Task Type	Description	Examples (based on our fieldwork)
Data Generation	The collection of data from the worker's environment	"You can earn \$2.5 by completing the task 'Do you wear glasses?' Upload a picture of a document with your prescription values now."
Data Annotation	The classification of data according to a predefined set of labels	"Based on the text in each task, select one of these three options: Sexually Explicit, Suggestive, Non-Sexual."
Algorithmic Verification	The evaluation of algorithmic outputs	"You'll be shown two lists of up to eight search suggestions each. Your task is to indicate which list suggestion is better."
AI Impersonation	The impersonation of an artificial agent	As the assistant, the "user will initiate the conversation... you need to use the facts to answer the user's question."

communicated to workers and how workers follow or interrogate instructions. Through these descriptions, we seek to locate our analysis in specific settings with specific ways of doing things.

Next, we move into dissecting the described tasks, instructions, and practices while outlining specific characteristics of the data-production dispositif. In 4.2, 4.3, and 4.4, we center the findings of the dispositif analysis around our three research questions. Following RQ1, we describe *discursive practices* such as those involved in the taxonomies used to collect and classify data and the warnings and threats included in them. Following RQ2, we describe *non-discursive practices and social contexts* such as the obedience to instructions, the dependence of Latin American workers to precarious work, and moments of interrogation and solidarity among workers. Finally, and following RQ3, we describe some of the *dispositif's materializations*, such as documents, work interfaces, and tools to measure workers' performance and surveil them.

#### 4.1 Different Tasks, Different Instructions

Before moving towards answering our three research questions, we will describe in this section the different tasks available in data work and explored in this study. We use the framework proposed by Tubaro et al. [81] to differentiate the tasks (see Table 4). While this framework was initially conceived to analyze digital platform labor, we think it can also be applied to BPOs in the broader field of data work due to the similar types of tasks available.

Tubaro et al. define three moments in outsourced AI production: "artificial intelligence preparation," "artificial intelligence verification," and "artificial intelligence impersonation." The authors divide AI preparation into the collection of data and its annotation. AI verification involves the evaluation of algorithmic outputs. Finally, AI impersonation, often seen in the corporate and AI-as-a-service sector [59], refers to the non-disclosed "human-in-the-loop" principle that makes workers hardly distinguishable from algorithms" [81].

##### 4.1.1 Data Generation.

Platform workers are directed to collect data from websites or to produce media content (e.g., text, images, audio, and video) from their devices. For example, a task on Clickrating instructed workers to find information online from companies in the United States, including their address and telephone number. Workerhub required workers to take photos of themselves in certain poses or pictures of family members (including children) and enter attributes of the subjects in these images, including their age and gender. While tasks involving data collection from the web were paid a few cents per assignment, those that involved capturing photos, video, and audio were compensated with a few dollars per file generated. Interviewed workers found the latter type of task attractive because they were the best remunerated. In this case, financial need overcame any privacy concern.

At the BPO, one of the data generation projects produced an image dataset to train a computer vision algorithm capable of identifying fake ID documents. For this purpose, workers were instructed to use their IDs and those of their family members. They took several pictures of the documents and used some of those pictures to create different variations of the ID document (changing the name, the address, the headshot). This way, they produced imagery of authentic as well as fake IDs. The requester of this task was a large e-commerce corporation and Alamo's most important client. The task is just one of many data-related projects Alamo has conducted for this corporation in the last four years. Because of their ongoing service relationship, the client invested considerable time and money in training Alamo's workers for each project. One particular characteristic of this relationship is that instruction documents are not created by the client unilaterally but co-drafted with Alamo's project managers and team leaders. Managers and leaders then serve as support to answer data workers' questions, should they arise after reading the instructions or while completing the tasks. Because of the sensitivity of the data involved in the ID project, a special interface with several security measures was created to work on the images and store them. Furthermore, all workers had to sign a consent form, allowing the use of their ID document. The request to use their ID and those of family members caused some unease and raised several questions among workers, as one of Alamo's managers reported. Then, it was the managers' job to "convince them that their IDs were not going to be used for anything bad, no crime or something. And to do that without revealing too much of the client's idea because we had signed an NDA."

#### 4.1.2 *Data Annotation.*

The most common task in data work is data annotation. Projects of this kind were available in all of the studied platforms and the BPO.

Tasksource and Workerhub provided many image segmentation and classification tasks as platforms specializing in computer vision projects. Some of these tasks included the classification and labeling of images of people according to gender, race, or age categories. Since segmentation tasks required around one hour depending on the size of the image and the number of labels, these tasks are paid more than classifying entire pictures according to a set of categories. Most of the segmentation tasks were designed for training self-driving cars and devices that are part of the internet of things. At the same time, image classification has multiple uses, including content moderation (notably for hate speech and sexual content), healthcare, facial recognition, retail, and marketing. For instance, the same BPO workers that collected pictures of ID documents were subsequently asked to classify and label them as "authentic" or "fake." In addition, workers had to segment the "fake" ID documents and mark the part of the image they had modified. Finally, they had to annotate the type of modification the image had undergone (e.g., "address has been modified" or "headshot is fake").

Outside of computer vision, workers of Workerhub were asked to identify hate speech and sexual content in text, notably for social media. For example, in an assignment titled "Identify Racism," workers were asked to read social media posts and identify whether or not the content included racism or if this judgment was not possible. In another task titled "\$exxybabe69," workers had to judge if usernames included examples of child exploitation, general sexual content, or none of the previous categories. Video and audio annotation were also present in this platform. For example, in the task "How dirty is this," workers had to identify if there was sexual or "racy" content in different media types, including audio files and videos. In section 4.2, we will revisit some of these tasks to describe how workers navigated the different taxonomies they encountered to perform data annotation tasks.

#### 4.1.3 *Algorithmic Verification.*

Algorithmic verification involves the assessment of algorithmic outputs by workers. This type of task was observed primarily on the annotation platform Tasker and accessed through Clickrating. Tasker is the internal platform of a major technology company that develops a search engine. They use Clickrating to recruit their workers who, depending on the project, have to sign special contracts and non-disclosure agreements. Algorithmic verification tasks, for example, include assessing how the search engine has responded to a user query, the objects that accompany a search result (e.g., images, maps, addresses), or whether the search result contains adult content or not. In many cases, these assessments include comparing search results with a competitor search engine and assessing which one is more accurate and substantial.

In another example of algorithmic verification, one of the tasks conducted by the BPO Alamo for its largest client (the same e-commerce corporation behind the “ID project”) consisted in verifying the outputs of a model used by the client to moderate user-generated content in their marketplaces. In this case, the task consisted of reviewing content flagged as inappropriate by an algorithm and confirming or correcting the output. For this purpose, the client had provided handbooks that contained each marketplace’s terms and conditions and examples of the specific forms a violation could take. For workers, this task often involved being exposed to disturbing images and violent language, which several interview partners described as “tough.”

#### 4.1.4 *AI Impersonation.*

Impersonation is the rarest type of task, and it was only observed once in the platform Clickrating. Tubaro et al. [81] describe it as a task that occurs “whenever an algorithm cannot autonomously bring an activity to completion, it hands control over to a human operator.” The task that we encountered, developed by a major social media company, asked workers to dialogue with users and respond to their queries according to a set of predefined “facts, history, and characteristics.” If the worker couldn’t answer the user’s query, they were asked to say, “Sorry, I don’t know about that, but can I tell you about...” and then they had to “insert fact related that may be of interest to user.” The platform instructed workers to complete dialogues in the least amount of time, and they have to be logged into the platform in specific 3-hour sessions.

## 4.2 Linguistically-Performed Elements

*RQ1: What discourses are present in task instructions provided to outsourced data workers?*

To explore our first research question, we analyze the task instructions *as text*. Our analysis focuses on the categories and classes used for collecting, sorting, and labeling data as contained in the task instructions. We describe three recurrent elements: (1) the normalization of conventions from the Global North and oriented towards profit maximization, (2) the use of binary classifications and the inclusion of residual categories such as “other” or “ambiguous,” and (3) the discursive elements that aim at constraining data workers’ agency in the performance of data-related tasks.

### 4.2.1 *Normalized Classifications.*

Taxonomies are the main component of task instructions for data work. They consist of classification systems comprising categories and classes used to collect, sort, and label data. Definitions, examples, and counterexamples usually accompany taxonomies. The number of classes depends on the task and varies according to each platform or company. For instance, assignments on the platform Workerhub usually present a smaller set of labels, ranging from two to a dozen maximum. At the same time, jobs in Tasksource usually feature dozens of classes classified in several categories (e.g., for the semantic segmentation of a road, the category “cars” included labels like “police car” or “ambulance”). The taxonomies that we observed in instructions carry self-evident meanings to the clients but are not necessarily relevant to the annotators or communities affected by the ML

system. For instance, the label “foreign language” refers to languages other than English, and the category “mail truck” only comprises examples of USPS vehicles.

One of the projects conducted at the BPO Alamo consisted in analyzing video footage of a road. A particularity of this project is that the requester did not predefine the labels, but workers were asked to come up with mutually exclusive classes to label vehicles in Spanish. One of the main difficulties of the project, however, lay in the nuances of the Spanish language: Probably oriented towards targeting a broader market, the requester wanted the labels to be formulated in *español neutro* (“neutral Spanish”), i.e., without the idioms that characterize how Argentines and most Alamo workers speak. This contrast led to many instances of discussion among workers and managers about which vehicles’ designations the client would consider “neutral.”

The mismatch between the classifications that requesters and outsourced data workers consider self-evident becomes critical in cases of social classification. For instance, in the task shown in Example 1, workers were asked to label individuals’ faces for facial recognition according to predefined racial groups that included “White, African American, Latinx or Hispanic, Asian, Indian, Ambiguous,” where the last category should be selected “ONLY if you cannot identify the RACE of the person in the image.”

#### Example 1

In this task you will be determining the **race** of the persons in the images.

You should select only **one** of the following categories:

- White
- African American
- Latinx or Hispanic
- Asian
- Indian
- Ambiguous

Beyond the already problematic situation of being asked to “guess someone’s race,” for many interviewed Latin American workers, this type of classification did not make sense, as it was conceived by a US company with a US-centric conception of racial classification, i.e., with little regard for the cultural and racial complexities in Latin America. Similarly, Gonzalo, a data worker for Workerhub, declared having trouble with a task that asked him to label “hateful speech” in social media posts:

They give you many examples of what they consider “hateful” and not. But, once you’re doing the task, you don’t encounter basic examples, and it’s up to you as a worker to interpret the context and decide what counts as “hate.” Clearly, [the requesters] have their parameters and, if they don’t consider something hateful, they will mark your work as wrong. ... For example, a sentence like “kick the latinos out” would not be regarded as “hateful,” and you will not interpret it the same way as a Latino.

These examples of social classification and conceptualization are not just about cultural differences between requesters and data workers, but they reflect the prevalence of worldviews dictated by requesters and considered self-evident to them. Furthermore, the taxonomies respond to the commercial application of the product that will be trained on the data that outsourced workers produce. For example, the instructions to categorize individuals for facial recognition technologies in Example 1 are based on a US-centric definition of “protected group.” Such definitions suggest requesters’ efforts to prioritize the mitigation of legal risks, neglecting the safety of users from

social minorities and groups facing discrimination in other contexts, such as ethnic groups with defined caste systems or social categories not protected by US law, such as economic class.

In another example that shows the inscription of requesters' profit orientation, workers of a major social media app working through Clickrating, were instructed to evaluate if a post was "building awareness." Here, "awareness" referred to posts with commercial content and discarded anything personal or political in nature (see Example 2).

#### Example 2

**Building awareness** means that the purpose of the [post] is to give information about a brand, product, or other [sic] object. For example, "@Restaurant is awesome for karaoke and the food is delicious!" is building awareness. A story is not building awareness if it's primarily about the author's life. A story that mentions the name of a business or service without providing much additional context, or a story that refers to a product in passing while its author is sharing one of their experiences, is not building awareness. For example, if the author says, "I'm eating in @restaurant", the story is not building awareness for the restaurant.

Requesters' profit orientation is implicit in task instructions and gets inscribed in the ways data workers approach their tasks. For instance, when asked what the procedure would be if they were unsure about annotation instructions, most of the workers at the BPO answered that they would abide by the requesters' opinion because "their interpretation is usually the one that makes more sense as they know exactly what kind of system they are developing and how they plan to commercialize it."

#### 4.2.2 Binary and Residual Categories.

While complex taxonomies are common in the BPO, the most common classification tasks in platforms are binary. For example, for tasks described as "Does this link title use sensationalist phrasing or tactics?", "Who is the author of this story?," and "You will be looking at a bounding box and choosing whether it is around a primary face," the possible labels were "yes or no," "human or non-human," and "yes or no" respectively. As mentioned above, many tasks, especially in Workerhub, are based on categories protected by the United States legislation when defining what counts as hate speech, racism, and other forms of discrimination (See Example 3). However, the classification of whether a text contains hate speech is often reduced to a binary decision without considering the context.

#### Example 3

In this task, you will be identifying messages that contain **hate speech**.

Based on the text, you must select:

- **Hate Speech:** if the username contains hateful content
- **None:** if there is NO hateful or abusive language in the given [sic]

#### **Definition:**

Select **Hate Speech** if the text contains **any** of the following:

- Discrimination, disparagement, negativity, or violence against a person or group based on a **protected attribute**
- References to hate groups, or groups that attack people based on a **protected attribute**

This form of binary classification often ignores ambiguity or uncertainty, such as when workers are confronted with contexts that are ambiguous or separate from their cultural setting. Moreover, given the impossibility of platform workers to send feedback to requesters, many omissions



(involuntary or not) remain unquestioned. For instance, a task on Workerhub asked workers to classify images as “racist” or not; we observed an image representing several copies of the “crying Wojak” meme wearing kippahs inside a heating oven while the meme “Pepe the Frog” is watching outside. The text above reads: “Changed the wooden doors today frens [sic], this is actually working as intended now!” While this image was, for us, a clear example of antisemitism, a form of racism condemned worldwide, including by the United Nations [83], requesters instructed workers not to consider this type of hateful content as “racist.”

In several of the cases where binary categories were prescribed, we encountered a third label, usually called “other” or “ambiguous,” not to designate an additional class that would break the binary classification but to merge errors or instances where the worker cannot apply one of the two labels (see Example 4). Workers are also encouraged to ignore ambiguity altogether. Some Clickrating tasks acknowledge the limits of this binary classification and urge workers to ignore other possible attributes when categorizing data. For example, in an assignment where workers tagged queries for a search engine, the instructions referred to “Queries with Multiple Meanings” for “queries [that] have more than one meaning. For example, the query [apple], in the United States might refer to the computer brand, the fruit, or the music company.” In this case, the task instructed workers to “keep in mind only the dominant or common interpretations unless there is a compelling reason to consider a minor interpretation.” The “dominant” or “common” interpretation of a term in the US may be different to the one in Latin America. Still, we repeatedly encountered similar instructions to deal with instances of ambiguity in many tasks.

#### Example 4

##### Overview

Select:

- **Male:** if the boxed face is a male
- **Female:** if the boxed face is a female
- **Other:** if there is no face in the box

Very often, instruction documents were dated, and requesters provided several updates. For instance, in the task “How dirty is this image/video” on the platform Workerhub, workers were initially instructed to label adult content in images. Later on, the requesters provided videos without updating the documents, confusing some workers because they could not apply the original instructions to the video footage. As a result, the requesters had to update the instructions to include details about video annotation. In another example, a document that asked workers to classify elements in a road revised its definition of “Pedestrians sitting on the ground” to include also those “laying [sic] on benches and laying [sic] or sitting on the ground” (see Example 5). This example also shows that the “other” category that we described above does not represent “everything else” included in the classification but that there are elements deliberately or accidentally left out. In this latter example, the requesters failed to see that people in the streets are not necessarily always “walking” or “sitting,” but that a segment of the population lies or sleeps on them.

## Example 5

**Pedestrians sitting on the ground**

Use the “pedestrian” label if a pedestrian is sitting on the ground, bench, ledge, then use the “pedestrian” label.

**UPDATE!!**

Use the “**PEDESTRIAN LYING DOWN**” label for pedestrians laying [sic] on benches and laying [sic] or sitting on the ground.

## 4.2.3 Warnings.

As we will argue in the following sections, the influence and preferences of powerful actors in data work are stabilized through narrow task instructions, specially tailored work interfaces, managers and quality assurance analysts in BPOs, and algorithms in crowdsourcing platforms. Some of these processes are part of the tacit knowledge workers have about their position (i.e., *it goes without saying* that workers must carry out tasks according to the preferences of requesters). However, task instructions often make explicit reference to the power differentials between workers and requesters as they include threatening warnings such as “low quality responses will be banned and not paid” or “accurate responses are required. Otherwise you will be banned” (see Examples 6–8).

Platform workers risk being banned and even expelled from the platform if they fail to follow task instructions. At the BPO Alamo, the communication of instructions is mediated by project managers and team leaders. For this reason, such warnings are not explicit in written documents but are present in reviewing instances and evaluating workers’ performance. Here, any concerns expressed at the workers’ end are filtered through hierarchical managerial structures and hardly ever reach requesters.

## Example 6

This is a **high paying job**, a special job, but to gain access to it and to keep access to it after passing the qualification test, we require patience and **VERY careful [sic] thought out and accurate responses**.

*Otherwise, you will, unfortunately be banned from the job :(*

## Example 7

The picture above shows SHIFTING DATA which means that the LiDAR points for stationary objects move or slide around throughout the scene.

**Any [Project] Tasks with shifting data are not usable by the customer & have to be cancelled. You will NOT get paid for working on task [sic] with shifting data!!!!!!!!!!**

Every time you get a [Project] Task (before you start working) always turn on dense APC and look around the entire scene to check for shifting data. You will be able to tell that the shift is big enough to be cancelled if it makes any object 0.3m plus larger than it’s [sic] normal size (or if it makes a flat wall 0.3m plus thick) and effects [sic] multiple cuboids.

## Example 8

We value your individual opinion and review each result, so please provide us with your best work possible. We understand that this can be a tiring task, so if you are in any way unable to perform your best work, please stop and come back once you are refreshed. You may also see multiple queries with the same kind of visual treatment.

Please keep your judgments consistent UNLESS you feel that there is some difference in the two that would result in a change of overall score.

**Judges providing low quality responses will be banned and not paid.**

These messages encode a precise definition of “accurate responses”: Accuracy is classified according to what the client believes to be an accurate truth value, while divergence from that value is considered inaccurate. Here, too, the classifications that make sense to requesters have prevalence. This is why workers at the Argentine BPO are permanently encouraged by management to think in terms of “what the client might want and what would bring more value to them.”

Given the social and economic contexts in which the outsourcing of data work occurs, warnings and threats of being banned or fired reinforce the hierarchical structure of the annotation process and compel workers to follow the norms as instructed or risk losing their livelihood. In the next section, we will present evidence of how the social contexts of workers and the fear of losing their job shapes how assignments are carried out.

### 4.3 Non-Linguistic Practices and Social Context

*RQ2: How do outsourced data workers, managers, and requesters interact with each other and instructions to produce data?*

This section explores our second research question. Here, we focus on analyzing interviews to describe the contexts in which data-related tasks are carried out and the interactions they enable. We describe (1) the social contexts of Latin American workers that lead to their dependence on data work regardless of the labor conditions, (2) the elements that contribute to the unquestioning obedience to instructions, and (3) moments of subversion of rules as well as workers’ organization and solidarity.

#### 4.3.1 Poverty and Dependence.

Being an impact sourcing company, Alamo employs around 400 young workers who live in slums in and around Buenos Aires. As stated on its website, Alamo specifically recruits workers from poor areas as part of its mission. As Natalia, one of the BPO’s project managers, describes, this is a population that does not receive many opportunities in the Argentine labor market:

They are very young, and a bit, you know... Alamo works with people another company wouldn’t hire, so people who live in areas... slums with difficulties, with a very low socioeconomic level. That’s something the company pays attention to when it comes to recruiting, and if during the interview we detect that the candidate could have an opportunity somewhere else, we prefer not to hire that person and hire someone else.

One particularity of Alamo is that it provides workers with a regular part- or full-time salary. This form of employment contrasts with the widespread piece-wage model in platforms. The salary Alamo’s workers received in 2019 was the equivalent of 1.70 US dollars per hour which was the minimum legal wage in Argentina. Despite the low wages and exhausting tasks such as semantic segmentation or content moderation, all interviewees were satisfied as the company offers better conditions than previous experiences. According to a report published in November 2020 by the national Ministry of Production [76], the unemployment rate in Argentina is 10%, and 35% of the

employed labor force is not registered. Argentina has a long tradition of undeclared labor. This way, employers avoid paying taxes while workers remain without protection or benefits.

Behind the numbers are people like Nati, who did different types of precarious work before working for Alamo. She started at the BPO as a data annotator and quickly became a reviewer until being offered a position as an analyst in the company's QA department. Like other Alamo workers, she acknowledges the difficulty of securing a desk job somewhere else. Moreover, many of our research participants mentioned being proud of the work they do at Alamo because a desk job has "a different status." For several of them, working at Alamo means finally having a steady income and breaking with generations of informal gigs, for example, in the cleaning or construction sectors. As Nati explains, what Alamo offers is better than the alternatives:

That was the situation at home; we were going through a rough time. My mother was out of work because her former boss had found someone else to clean, and I had lost my job too. So I needed a job and when I found this one I was surprised to work at a friendly place for a change! Now I have a desk, a future, and I feel appreciated. This is new to me.

In the case of the platforms, they have thrived in the Venezuelan economy, which is characterized by the highest levels of inflation in the world, with an average of 3,000% in 2020 [1]. All participants that we interviewed from Venezuela stated that the "situación país" [country's situation] was the main reason they resorted to online work. Workers reported difficulty finding employment in the local labor market, especially for income that is not dependent on the national currency, the bolivar, which devalues quickly. For example, Rodrigo, a Clickrater worker, quit his job as an information technology consultant because online platforms were the only way he could earn US dollars. He explains the monetary situation of his country as follows:

There are two types of currency exchange rates: the official rate dictated by the government and the one used on the black market, which everyone uses. Everyone knows this black-market exchange rate. It's an Instagram profile that posts the average exchange rates of several independent currency exchange websites. They make this average and post the fluctuation several times per day, which is the exchange rate that we use today.

Platforms' low entry barriers make outsourced data work an attractive — and sometimes the only — source of employment during social, economic, and political crises. Data workers earn 15 to 60 US dollars per week, the average being around 20 dollars, which is substantially higher than the minimum wage in Venezuela reported by workers to be around 1 US dollar per month in March 2021.

Dependency on the platforms is exacerbated by the high unemployment levels and reduced government support during the COVID-19 pandemic. In this situation, workers have limited access to subsidies and pay for services such as healthcare from their income [65]. For example, Olivia, one of the Tasksource workers, was diagnosed with diabetes and has to self-fund the costs of insulin. For her, losing access to the platform is a life-threatening situation:

Imagine, with this pandemic, what can I do? My medical situation does not allow me to go outside and risk getting the [coronavirus disease]. If I get it, I'll die. For this reason, I cannot take the risk and expose myself to something worse; I can't risk this job either because this is my only source of income.

This dependency affects the labor process as well. Workers usually do not choose which tasks to perform, even if they disagree ethically with their assignments. When asked what criteria Carolina, a Clickrater worker, uses to choose a task over another, she answered:

My priority is to get the tasks that pay the best. But I don't even have that choice. The platform restricts which jobs are available here in Venezuela, so I have to make the most of it to earn the minimum and get paid as soon as I get one task.

By "minimum," Carolina refers to the minimum income workers can transfer out of the platforms, which is another form of creating dependency. Platforms establish a minimum of 5 to 12 US dollars before they make payments and, if a worker cannot achieve this threshold, they have to wait for a week before withdrawing their salary. This payment process is a form of institutionalized wage theft implemented by the platforms. Workers lose the money they have worked for if they get banned before reaching the threshold for payment.

#### 4.3.2 *Obedience to Instructions.*

In its outsourcing capacities, Alamo focuses on data-related services ordered mainly by machine learning companies. Even if they display some similarities, each of those projects is different from the previous ones, and workers need to be briefed regularly. Depending on the difficulty and extension of the task, briefings can be more or less sophisticated and involve more or fewer actors, meetings, and processes. Sometimes, the instructions for new projects are sent by the requester via email in a PDF document. One of the area managers receives that information and transmits it to a project manager, who would then put together a team and work closely with their leader. Depending on the degree of difficulty, one or more meetings with the team will be held to explain the project, answer questions, and supervise the first steps. When handling large projects from multinational organizations, Alamo invests a considerable amount of resources in the briefings. No matter how big or small the requester, briefings at Alamo consist of getting the workers acquainted with the expectations of the requesters and are a way of making sure workers are on the same page and thinking similarly:

The information from the client usually reaches the team leader or the project manager first, and, at that moment, what we do is to have a meeting for criteria alignment... that is generally what we do. The team meets to touch base and see that we all think in the same way. (Quality assurance analyst with Alamo)

These briefings give workers a framework for new projects and are instrumentalized by the company as the first instance of control, aiming at reducing room for subjectivity. Further control instances, aiming to ensure that data work is done uniformly and according to requesters' expectations, take place in numerous iterations where reviewers and team leaders review and revise data and go back to the instruction documents or contact the requester to clarify inconsistencies.

In companies like Alamo, *data quality* means producing data precisely according to the requester's expectations. According to Eva, a BPO manager in Bulgaria, this view on data quality is commonplace in data services companies. In the following excerpt, she summarizes the importance and main function of instructions and further instances of control, i.e., making sure that the workers interpret the data homogeneously:

Normally, issues in data labeling do not come so much from being lazy or not doing your work that well. They come from a lack of understanding of the specific requirements for the task or maybe different interpretations because a lot of the things, two people can interpret differently, so it's very important to share consistency and, like, having everyone understand the images or the data in the same way.

The interviews we conducted with requesters show that the priority behind the formulation of task instructions is producing data that fits the requester's machine learning product and the business plan envisioned for that product. What does not match the requester's instructions is

considered low-quality data or “noise.” Dean is a machine learning engineer working for a computer vision company in Germany. He reported on this widespread view as follows:

Dean: Noise is what doesn’t fit your guidelines.

Interviewer: And where do those guidelines come from?

Dean: We say, “actually we want to do this, we want to do that,” and then, of course, since the client is the king, we translate that business requirement into something like... into a requirement in terms of labels, what kind of data we need.

As described in Section 4.2.3, compliance with requesters’ views is made explicit in instruction documents in the form of warnings for workers. Those documents are usually the only source of information and training platform workers have to complete their tasks. However, the case of the platform Tasksource is slightly different, as it employs Latin American coaches to brief and explain to workers how to interpret instructions and annotate tasks. This approach is similar to the one used by the BPO Alamo and described above. However, at Tasksource, briefings take place in week-long unpaid digital courses called “boot camps” and later evaluation periods called “in-house.”

The use of the military and correctional term “boot camp” could be interpreted as reflecting this training’s purpose: conditioning workers to obey tasks without question. Ironically, even though the platform employed workers to help train artificial agents, they were supposed to behave like “robots,” according to a Tasksource worker named Cecilia:

When you start, they tell you: “To be successful in this job, you have to think like a machine and not like a human.” After that, they explain to you why it has to be like that. For example, you are teaching a [self-driving] car how it has to behave. When you segment an image, there is a police car, and you label it like a regular car, the [self-driving] car will think it’s a regular car and, if it crashes against it, something terrible can happen. The mistake was not of the car that crashed into the police vehicle, but it’s yours as a tasker, as a worker, who taught the car to behave like that.

Platform workers serve a similar role as BPO employees in reinforcing the primacy of instructions and requester intent to complete tasks effectively, producing data that fits model and revenue plan while shifting the responsibility for failures on workers. In this context, obedience to instructions is critical for data workers to keep their job and make a living. The fear of being fired, banned from the job, or not being paid for the task reinforces the disposition of workers to being compliant, even when instructions look arbitrary. This is what Rodolfo, Tasksource worker, reported:

That is why I don’t like that platform very much. Because they give us the instructions and we have to follow. And there are many cases where, if you don’t complete the task really to perfection, according to what they want or what they think is right, they just expel you. Just like that, even if you followed the instructions thoroughly.

#### 4.3.3 *Worker Solidarity and Organization.*

Not everything is imposition and obedience in data work. There are also several expressions of workers organizing to improve working conditions and help each other deal with tasks and make the most out of them.

For instance, Alamo’s employment model that includes data workers as part of the company’s permanent staff instead of having them as contractors results from workers organizing to demand receiving a fixed salary and benefits. As reported by one of Alamo’s workers, Elisabeth, in 2019, further workers’ demands were being negotiated with the company:

We asked for a couple of things like the possibility of home office and a better healthcare plan. We are organizing many things. It’s being negotiated.

In 2020, probably also motivated by the COVID-19 pandemic, Alamo’s data workers were finally allowed to work remotely. It is worth mentioning that before 2020, every other company’s department and management were allowed to work at least some days of the week remotely while the data workers could not.

In the case of the platforms, data workers organize in virtual groups and fora. The existence of virtual and local groups of workers that provide solidarity and support has been reported in other examples of platform labor [16, 67, 88]. In the case of Venezuelan data workers, we observed similar situations. Because we used convenience and snowball sampling and worker groups on social media as a starting point, all the interviewed participants were directly associated with them. Participants use these independent and worker-led spaces to exchange information about which tasks pay more and are less challenging to complete and warn each other about non-reliable requesters. One of the aspects that workers paid significant attention to was the presence of bugs in the tasks. When asked about their existence, Yolima, a worker with the platform Tasksource, said to us:

Errors occur all the time. But, since we are in groups on Facebook and Whatsapp, we alert each other and say, “Hey, don’t do this task because it has a bug. It will flag you as mistaken even if you have done everything ok.”

Some smaller groups, with high entry barriers to ensure privacy and trustworthiness among members, recommend specific tasks over others. For example, when describing tasks with sexual or violent content, Estefanía, one of Clickrating’s workers, stated:

I don’t like those tasks with pornographic content. I do them only when my friends from the groups say, “look, this is a good task, here’s the link.” I don’t have to look for good tasks, and that’s great. I just have to log into my account and do the annotation without worrying about which tasks to do.

Some users of these smaller groups also craft guides to explain the instructions to their peers. Most interviewed workers stated that their knowledge of the English language was limited. Since Tasksource and Clickrating only presented instructions in that language, and Workerhub provided automated translations with errors, these guides in Spanish are a fundamental tool for workers. They are written by workers for their peers and contain Spanish translations of taxonomies, definitions, and examples. They also provide further explanations about the contexts in which workers can apply the taxonomies, avoid being banned by the algorithm, and maintain high accuracy scores. For example, in the introduction of a guide for a task to annotate hate speech in text for Workerhub, a user wrote:

Example 9

**IMPORTANT INFORMATION**

What I’m sharing in this guide is based on my experience with the task. I’ll try to explain as best as I can the tips that I consider are the most important to avoid being banned and the essential information to understand the task.

**BE CAREFUL**

The task “No Hatred” is not available on all accounts. You must have been paid **AT LEAST ONCE**. **IT’S IMPORTANT THAT YOU CONSIDER THIS GUIDE FOR WHAT IT IS: A “GUIDE”** made for you to understand the task better. You must earn real experience by doing the task with perseverance and dedication.

Work practices in the data-production *dispositif* are not informed exclusively by the relationships between requesters, intermediaries (platform or BPO), and individual workers. They are also dependent on the networks formed by the latter group. This can be observed in BPOs where data

workers share the same office space and constantly consult and advise each other on conducting projects more quickly and easily. Among platform workers, online groups help to choose what tasks to carry out, and such decisions are influenced by recommendations and guides from peers who evaluate instructions from requesters.

#### 4.4 Dispositif's Materializations

*RQ3: What artifacts support the observance of instructions, and what kind of work they perform?*

In this section, we focus on the third research question. Based on the observations conducted at the crowdsourcing platforms and the BPO company, we present three of the many possible materializations of the data-production dispositif: (1) the function of diverse types of documents that embody the dispositif's discourses, (2) the platforms and interfaces that guide and constrain data work, and (3) the tools used by managers and platforms to surveil workers and quantify their performance.

##### 4.4.1 Documents as Artifacts.

In Section 4.2, we focused on the *content* of instruction documents to describe the discursive elements comprised in them. Here, instead, we look into a variety of documents – instructions included – to analyze them as *artifacts*, focusing on their form, function, and type of work they perform.

One common document related to data work at BPOs is that containing metadata and project documentation. Alamo, for instance, records the details of each project in several documents that vary in form and purpose according to the task and the requester. Often, that documentation aims to preserve the evolution of task instructions, registering changes requested by clients. Keeping this type of documentation functions as a form of “insurance” for Alamo and can help resolve discrepancies if requesters are not satisfied with the service provided. In those cases, project documentation serves as proof that data was produced as instructed. Documents containing project details can also serve the purpose of preserving situated and contingent knowledge that would otherwise get lost and could help improve future work practices [53]. Sometimes, these documents become artifacts that cross Alamo's boundaries and reach the requesters. For them, the documents might have a factual function (in terms of the information they want or need) or a symbolic one (to reassure clients that Alamo is *at their disposal*). Alamo's QA analyst Nati describes this as follows:

We send a monthly report to the clients, including what was done and problems we encountered; we set objectives for the following month and send an overview of the metrics. Some clients don't even look at the report but insist on receiving it every month. Others value it and use the information to report to leadership or investors within their organization.

The documents produced by the BPO are tailored to be valuable for requesters. Conversely, the documents formulated by requesters often remain unintelligible for data workers, even if they are the primary addressees, as in the case of instruction documents. In many cases, language is the main issue hindering the intelligibility of documents: Most of the workers we interviewed have limited knowledge of the English language and reported using translation services, notably Google Translate, to understand the instructions provided by requesters. As mentioned in the previous section, one of the main reasons platform workers resort to guides written by peers is that they are written in Spanish. But beyond language differences, elements of the taxonomies used in documents can also be confusing, as explained by Tasksource's worker Yolima and described in Section 4.2.1:



For the [categories], they are made in the United States, I think. I don't know what they would call a laundry sink<sup>1</sup>, a shower, or parts of the bathroom. Most of the time, my mistakes were with parts of the bathroom, especially around the shower, the tap, and those things. That was confusing because that was a shower for me, but it was something else for [the platform].

The confusion produced by the different languages is not merely a matter of cultural bias. Looking for cheap labor, platforms and requesters target the Venezuelan market but ignore the language barriers and formulate instructions in English. Moreover, further documents that workers encounter in their work, such as privacy policies, contracts, and non-disclosure agreements, are also prepared in English and remain, partially or totally, unintelligible for them. Workers usually sign these documents without understanding the full scope of their contractual relationship with their employers. Along with instructions, these documents embody the data-production *dispositif*. They are a materialization of normalized discourses and practices that shape data workers to be dependent and, therefore, obedient, while their subjectivities as Spanish-speaking Latin American workers are ignored and erased.

#### 4.4.2 *Work Interfaces.*

In BPOs like Alamo, choices regarding which platform will host the data and will be used as a tool are made by clients. In many cases, the requester has developed an annotation software specifically tailored to the needs of their business and the dataset to be produced. In other projects, the company uses a commercial platform designed by a third party. In this case, the client would suggest the tool that best fits their needs among several choices available on the market.

The choice of a specific tool comes with limitations that, in one way or another, constrain data workers' agency to interpret and sort data. The most notorious one is that the taxonomies contained in instruction documents are also embedded in the software interfaces that workers use to collect, organize, segment, and label data. Workers usually interact with a drop-down menu containing all the classes or attributes they are allowed to apply to data. Most interfaces do not allow workers to add further options to the list of pre-defined labels that they receive from requesters. This is most prominent in software interfaces specially designed by requesters and tailored to specific projects. In those cases, the software interfaces that mediate between workers and data are designed to ensure that tasks are completed according to particular parameters pre-defined by requesters and made explicit in the instruction documents.

In the case of regular data work tools for commercial use, the impossibility of changing the predefined categories or adding more classes is perceived as a limitation that makes data work harder at BPOs and requires communication throughout hierarchical structures until the requester is reached. Jeff, one of the managers leading a BPO in Iraq, reports on this issue:

There was a limitation on the annotation tool that they were using. They were relying on an open-source platform that doesn't have that feature that lets you add or create predefined attributes, which makes the work many times easier.

Some of these generic tools give the project owner — generally the requester or a BPO's project manager — the faculty to allow workers to add further options to the classification system (see Figure 1). However, this does not seem to be a widespread practice at Alamo. Among the many projects that we had the opportunity to observe during fieldwork, only once were data workers allowed to co-create the taxonomy around which data was organized and annotated.

---

<sup>1</sup>We use this term to translate "lavadero," a commonplace in Latin American homes for the washing of clothes equipped with a washboard basin and sink.




Fig. 1. Commercial data annotation tool. Only the project owner can grant rights to data workers.

In the crowdsourcing platforms that we studied, only Clickrating presented external interfaces, meaning that workers had to log into internal annotation platforms of clients, notably in the case of tasks requested by major technology companies. For Tasksource and Workerhub, workers interacted with data annotation interfaces developed by these companies. In both cases, the screen displayed a top bar with an accuracy score or the percentage of tasks submitted by the worker that the platform judged accurate. For Workerhub, the top bar also showed the number of annotations completed for the assignment, the earnings, the time spent per task, and a button to display the instructions (see Figure 2). On both platforms, the labels were available in the right sidebar alongside tools to zoom in and out and configure the visibility of the data. In all three platforms, workers could not change the predefined labels or suggest changes.

The interfaces present in the BPO and the platforms feature gamification elements (scores and timing) to speed up the labor process and keep workers focused on the tasks at hand. The over-reliance on speed privileges action over reflection and increases the alienation between workers and the production process. Even tasks that ask for workers' judgment, such as those present in Clickrating, are timed and reward fast thinking. That said, unlike the other platforms, they offer room for comments to evaluate algorithmic outputs that can be substantial, creating more engagement for workers beyond narrow annotation tasks.

#### 4.4.3 Tools to Assess Worker Performance.

To differentiate itself in the very competitive market of outsourced data services, the BPO Alamo makes a selling point out of its performance metrics and quality assurance mechanisms. The company puts much effort into developing more and better ways of measuring performance and quality, and transforming those into numbers and charts the client may perceive as valuable. As a response to market demands, quality controls intensify, which results in pressure and surveillance for workers. Moreover, the need for quantifiable data to translate "quality" into a percentage exacerbates the standardization of work processes, which, once more, results in less room for workers' subjectivity.

Alamo has highly standardized processes that include a team leader and several reviewers per team and a quality assurance (QA) department using several metrics to ensure that projects are conducted in accordance with the requesters' expectations. In addition, team leaders and the QA department use metrics to quantify workers' labor. As a token of transparency, sometimes workers'



Fig. 2. Age-based classification of images on Workerhub's interface.

scores are shared with clients. Noah, one of the BPO's team leaders, describe the function of metrics within Alamo and concerning its clients:

We have metrics for everything. They can be individual, for personal output, or they can be general in the project. We have some to measure correct and incorrect output, there we see where we fail, where we can give more support to the team so that those errors are corrected, how we can solve those problems. In QA, what they do is metrics. Metrics, and ensure that the quality provided to the client is high.

In platforms, workers are also constantly evaluated with accuracy and speed metrics. Instead of being managed by company employees like in the case of Alamo, platform workers are assessed and controlled by algorithms. All platform workers we interviewed reported being often banned from tasks because the algorithms negatively evaluated their performance. Of course, this represents a serious obstacle to maintaining a stable income, especially when it is permanent. This is what Juan, a worker of Workerhub, reported:

Juan: The platform pays every Tuesday. Once they ban you, you lose all your credits, in the sense that, without an account under your name and email, you can't open a new account and access the money you've earned.

Interviewer: Did they tell you why?

Juan: No. I could have asked in the [platform managed] Discord channel, but if you ask anything, you get banned. They are the ones who command... they are the ones who decide. I was banned without cause because my accuracy was high. I never knew why they expelled me.

Interviewer: How did you realize you were banned?

Juan: One day, I couldn't access my account. ... I created another account with the same email, worked for a week, and they banned me again. They didn't pay me. Some

of my colleagues from the same neighborhood and cousins who work for the platform told me: “Don’t create an account with the same email because they won’t let you. They will let you open it, but then they won’t pay you.”

The algorithms that assess worker performance in the three platforms that we have studied follow exactly the same three-step process: First, workers have to work with the same data again after some time. For example, if the task involves categorizing photographs of flowers according to their colors, if a worker marks the same image differently, the algorithm will consider it “spam.” The second mechanism is to verify workers’ answers with previously labeled data. If there is a mismatch, the algorithm will assume that the worker is not performing their activities “accurately.” Finally, from interviews with workers and previous observations in Amazon Mechanical Turk [56], a platform that is not well established in Latin America outside of Brazil and, therefore, not the focus of this study, the third method used by algorithms is to compare workers’ answers with those of peers and assume that the most common answer is the correct one. Many of the workers’ groups that we encounter provide guides, so workers do not diverge from the responses of the majority and, thus, keep high levels of accuracy from the perspective of the algorithms.

## 5 DISCUSSION

In concurrence with previous work [27, 53, 54, 75], we have observed that workers collecting, interpreting, sorting, and labeling data do not do so guided solely by their judgment: their work and subjectivities are embedded in large industrial structures and subject to control. Artificial intelligence politics are inextricably connected to the power relations behind data collection and transformation and the working conditions that allow preconceived hegemonic forms of knowledge to be encoded in machine learning algorithms via training datasets. Labor conditions and economic power in the production of ML datasets manifest in decisions related to what is considered data and how each data point is interpreted.

While task instructions help data workers complete their tasks, they also constitute a fundamental tool to assure the imposition of requesters’ worldviews on datasets. Sometimes, the meanings and classifications comprised in data work instructions appear self-evident to workers, and a shared status quo is reproduced on the dataset. Often, however, the logic encoded in the instructions does not resonate with them. This could be due to cultural differences between requesters and data workers, lack of contextual information about the dataset’s application area, perceived errors that cannot be reported, or simply because the tasks appear ethically questionable to workers. In such cases, another form of normalized discourse persists: that of a hierarchical order where service providers are conditioned to follow orders because “the client is always right” and workers should “be like a machine.”

According to Foucault, discourse organizes knowledge that structures the constitution of social relations through the collective understanding of the discursive logic and the acceptance of the discourse as a social fact. A normalized discourse is, therefore, what goes without saying. This way, the prevalence of requesters’ views and preferences does not need to be explicitly announced to workers. Instead, such implicit knowledge influences how the outsourced data workers that we observed and interviewed perform their tasks: carefully following instructions, even when they do not make sense to them or when they do not agree with the contents and taxonomies in the documents. The context of poverty and lack of opportunities in the regions where data production is outsourced is also fundamental as it makes workers dependent on requesters and, thus, obedient to instructions.

Finally, artifacts such as narrow work interfaces with embedded predefined labels, platforms that do not allow workers’ feedback, and metrics to assess workers’ “accuracy” (understood as

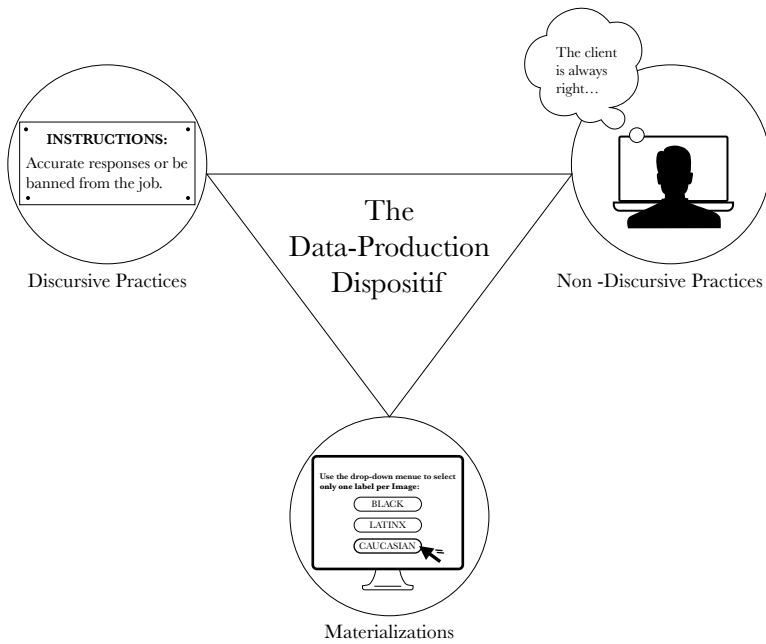


Fig. 3. The three components of the data-production dispositif based on the framework and figure proposed by Jäger and Maier [39]

compliance to requesters' views) constitute discursive materializations that, at the same time, ensure the perpetuation and normalization of specific discourses.

**All these elements combined – the predefined truth values encoded in instructions, the work practices and social positions of workers, and materializations such as interfaces – constitute the data-production dispositif.** Without any of these elements, the dispositif would not be able to function as such. As Foucault puts it, dispositifs respond to an “urgent need” [25] that is historically and geographically contingent. **The data-production dispositif responds to the voracious demand for more, cheaper, and increasingly differentiated data to feed the growing AI industry** [5, 13]. Its goal is to produce *subjects* that are compliant to that need.

The Foucauldian notion of *subject* has a twofold meaning, with subjects, on the one hand, being producers of discourse and, on the other hand, being created by and subjected to dispositifs. All subjects are entangled in dispositifs and have, therefore, tacit knowledge of *how to do things* within specific contexts. This tacit knowledge includes “knowing one’s place” and what is expected from each subject depending on their position. Thus, data workers know that subjects in their social and professional position are implicitly expected to comply with client’s requests. This way, dispositifs normalize and homogenize the subjectivities of those they dominate, producing power/knowledge relationships that shape the subjects within the dispositif according to certain beliefs, actions, and behaviors that correspond to the dispositif’s purpose [21, 23].

Following Foucault’s perspective, we argue that **the goal of the data-production dispositif is creating a specific type of worker, namely, outsourced data workers who are kept apart**

**from the rest of the machine learning production chain and, therefore, alienated. Data workers who are recruited in impoverished areas of the world, often under the premise of “bringing jobs to marginalized populations,” but are not offered opportunities to rise socially or professionally in terms of salary and education. Data workers who are surveilled, pushed to obey requesters and not question tasks, and who are constantly reminded of the dangers of non-compliance.** Data production cannot be a dignifying type of work if it does not provide workers with a sustainable future.

The implications of this data-production dispositif designed to constrain workers’ subjectivities and perpetuate their *alienation*, *precarization*, and *control*, will be unpacked in the following subsection.

## 5.1 Implications

As the generous corpus of research literature dedicated to mitigating bias in crowdsourcing suggests, controlling workers’ subjectivities is considered essential to avoid individual prejudices being incorporated in datasets and, subsequently, in machine learning models. However, as we have shown with our findings, unilateral views are already present at the requesters’ end in the form of instructions that perpetuate particular worldviews and forms of discrimination that includes racism, sexism, classism, and xenophobia.

Given its characteristics, the data-production dispositif is detrimental to data workers and the communities affected by machine learning systems trained on data produced under such conditions. To close this paper, we would like to make a call to dismantle the dispositif. However, before going into the implications of our call, it is crucial to consider that we never cease to act within dispositifs and, by dismantling the data-production dispositif, we would inevitably give rise to another one. Therefore, we discuss here ways of dismantling the data-production dispositif *as we know it today*, that is, by changing the material conditions in data work and making its normalized discourses explicit.

### 5.1.1 Fighting Alienation by Making the ML Pipeline Visible to Workers.

Substantial efforts in research and industry have been directed towards investigating and mitigating worker bias in crowdsourcing. Many of these initiatives portray data workers as bias-carrying hazards whose subjectivities need to be constrained to prevent them “contaminating” data. This widespread discourse within the data-production dispositif gives place to narrow instructions and work interfaces and the impossibility of questioning tasks. Workers are required to “think like a machine” to be successful in the job. Moreover, data workers are often kept in the dark about requesters’ plans and the machine learning models that they help train. Such conditions lead to workers’ alienation as they are kept apart from the rest of the ML production chain.

Researchers have often referred to data workers such as data labelers and content moderators, practicing *ghost work* [30] that remains “invisible” [70]. However, as Raval [69] accurately argues, it is worth asking *invisible for whom* and, most importantly, “what does this seeing/knowing—hence generating empathetic affect among Global North users—provide in terms of meaningful paths to action for Global South subjects (workers and others)?” Breaking with the alienation of data workers means much more than rendering them visible. It rather requires making the rest of the machine learning supply chain visible *to them*. It means providing information and education on technical and language matter that could help workers understand how their valuable labor fuels a multi-billion dollar industry. This also concerns questions of labor organization and unionizing: For instance, the recently-created Alphabet Workers Union has taken steps in this direction by including contractors — many of them outsourced data workers. To help counter their alienation,

researchers and industry practitioners need to regard data workers as tech workers as much as we do when we think of engineers.

*Why would requesters want to educate data workers and disclose technical or commercial information to them?*

As mentioned above, the design of the tasks that we encountered failed to acknowledge and rely on the unique ethical and societal understanding of workers to improve the annotations and, with them, models. We found that the BPO model generates a stronger employment relationship with workers compared to platforms, notably Workerhub and Tasksource, which translates into higher engagement with the tasks at hand. Furthermore, BPO workers interviewed by us said they wished they knew more about the requesters' organizations and products because this would help them understand their work and perform better. In this sense, expanding instructions to include contextual information about the task, its field of application, and examples that show its relevance for systems and users could improve data workers' motivation and satisfaction, and help them understand the value of their labor within ML supply chains.

### 5.1.2 *Fighting Precarization by Considering Data Workers Assets.*

One of the most pressing ethical and humanitarian concerns surrounding outsourced data work is the workers' quality of life. The data-production *dispositif* is designed to access a large and cheap labor pool and profit from workers' precarious working conditions. It is not a coincidence that, in Latin America, the platforms we encountered were established primarily in Venezuela, a country mired in a deep socio-political crisis exacerbated by the COVID-19 pandemic, and that the BPO company in Argentina recruited its workers from low-income neighborhoods.

While the arrival of these platforms and BPOs has allowed many workers to circumvent the limits of their local labor markets, the system of economic dependency and exploitation that they reproduce hinders efforts for sustainable development that include access to decent work, and economic growth [82]. Labor is an often overlooked aspect in discussions of ethical and sustainable artificial intelligence [66]. We argue that we cannot truly create fair and equitable machine learning systems if they depend on exploitative labor conditions in data work.

*Why would requesters want to improve labor conditions in outsourced facilities?*

All ML practitioners interviewed for this study had experience outsourcing data-related tasks with both crowdsourcing platforms and BPOs. They all agreed that platforms are cheaper than BPOs, but the latter offer higher quality. As argued by our interview partners, BPO teams remain more or less unchanged throughout the production project, which results in better quality. Moreover, direct communication with project managers allows for iterations and the incorporation of feedback. Several ML practitioners also report preferring not to outsource data-related tasks, especially in cases where a unique "feel for the data" [61], that can only be achieved with time and experience, was required. The evidence pointing to a negative correlation between cheap labor and the quality of data [49] described by the ML practitioners that we interviewed could be a strong argument for requesters to take measures and fight precarious work in outsourced facilities. Improving labor conditions might result in less expensive (and, perhaps, more effective) approach than investing in "debiasing" datasets after production.

### 5.1.3 *Fighting Workers' Surveillance and Control by Encouraging Interrogation.*

Our findings show that the widespread use of "protected categories" for human classification is bound to the cultural contexts and local jurisdictions that define what counts as a protected group. Moreover, even tasks that do not involve classifying humans, such as identifying objects in a road, can potentially have fatal consequences for individuals or groups, as in the case of the Tasksource requester who did not include labels for humans sleeping or lying on the streets. Making the rationale behind task instructions explicit can be difficult if categories are implicitly considered

commonplace for requesters, as they might not even notice the normativity behind instructed taxonomies. Moreover, data workers that are subject to surveillance and control and who risk being banned from tasks are less likely to question instructions. De-centering the development of taxonomies from an “a priori” (i.e., classifying exclusively based on personal experience) and data-based (i.e., classifying solely based on quantitative data) classification to one that derives from the context and experiences of those who may be affected by it could be a fruitful approach to this issue [17]. Data workers often perceive errors in task instructions or interfaces that remain unnoticed by the requesters. Even if this feedback could be valuable for requesters, the data-production dispositif is designed to silence workers’ voices.

We argue that the approach observed in Clickrating, where feedback from workers was encouraged, could be constructive here. However, expanding and implementing such an approach would require a general shift of perspective: from considering workers’ subjectivities a danger to data towards considering workers as assets in the quest for producing high-quality datasets. Fostering workers’ agency instead of surveillance and opening up channels for feedback could allow workers to become co-producers of datasets instead of mere reproduction tools.

*Why would requesters want to be questioned in their logic?*

While taxonomies respond to the commercial necessities of requesters, they also need to be built with equity and inclusivity in mind. This is not only an ethical issue, but it can quickly become a commercial one. Public scrutiny can have fatal consequences for a machine learning product that is perceived to be discriminatory or harmful [2, 32, 45]. Furthermore, instruction documents are living documents. We have observed how requesters update them by withdrawing the tasks, reinstating the instructions, and seeking data work again, a time-consuming and costly process. Thus, requesters could benefit from considering instructions as the product of exchanges with the different stakeholders contributing to data production and deployment. Data workers could play a key role in interrogating and improving tasks and, therefore, datasets and ML systems.

## 5.2 Limitations and Future Research

Our findings are bound to the platforms, companies, individuals, and geographical contexts covered by our study and our positionality as researchers, which has undoubtedly oriented but probably also limited our observations and interpretations. Because of the qualitative nature of our investigation, we have striven for inter-subject comprehensibility [20] instead of objectivity, which means making sure that our interpretations are plausible for both authors and the contexts observed. Furthermore, the use of multiple data sources allowed us to procure supporting evidence for observed phenomena. In addition, the use of expert interviews allowed us to confirm and discuss several of our initial interpretations.

This paper only covers some aspects of the data-production dispositif. This is because no dispositif works in isolation but is always entangled with other discourse, action, and materialization networks. To explicate the totality of the data-production dispositif would mean to analyze its relationship with, among many others, the scientific dispositif, the economic dispositif, and more specifically, the academic and the tech-industry dispositifs. Critical aspects of these relationships have been reported in these pages, but covering them all in one paper would be unfeasible. The fact that our analysis is bound to remain “incomplete” could be seen as a limitation. However, we consider it an opportunity for future research to expand our findings and interrogate ways of working with data that today seem commonplace. We think that a profound exploration into the tech-industry dispositif and its relationship with the data-production dispositif could be especially fruitful.



## 6 CONCLUSION

To explore how data for machine learning is produced through labor outsourced to Venezuela and Argentina, we have turned to the Foucauldian notion of *dispositif* and applied an adapted version of the *dispositif* analysis method outlined, among others, by Sigfried Jäger [39, 41]. Our investigation comprised the analysis of task instructions, interviews with data workers, managers, and requesters, as well as observations at crowdsourcing platforms and a business process outsourcing company. What we have called the data-production *dispositif* comprises discourses, work practices, and materializations that are (re)produced in and through ML data work. Our findings have shown that requesters use task instructions to impose predefined forms of interpreting data. The context of poverty and dependence in Latin America leaves workers with no other option but to obey. In view of these findings, we propose three ways of counteracting the data-production *dispositif* and its effects: making worldviews encoded in task instructions explicit, thinking of workers as assets, and empowering them to produce better data.

While the potentially harmful effects of algorithmic biases continue to be widely discussed, it is also essential to address how power imbalances and imposed classification principles in data creation contribute to the (re)production of inequalities by machine learning. The empowerment of workers and the decommodification of their labor away from market dependency, as well as the detailed documentation of outsourced processes of data creation, remain essential steps to allow spaces of reflection, deliberation, and audit that could potentially contribute to addressing some of the social questions surrounding machine learning technologies.

## ACKNOWLEDGMENTS

Funded by the German Federal Ministry of Education and Research (BMBF) – Nr 16DIII113, the International Development Research Centre of Canada, and the Schwartz Reisman Institute for Technology and Society. We thank Tianling Yang, Marc Pohl, Alex Taylor, Alessandro Delfanti, Paula Núñez de Villavicencio, Alex Hanna, Paola Tubaro, Antonio Casilli, and the anonymous reviewers. Special thanks to the data workers who shared their experiences with us. This work would not have been possible without them.

## REFERENCES

- [1] Agence France-Presse. 2021. Venezuela reports 2020 inflation of 3,000 percent. *ABS CBN News* (2021).
- [2] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. *ProPublica* (2016).
- [3] Ron Artstein and Massimo Poesio. 2005. Bias decreases in proportion to the number of annotators. In *Proceedings of FG-MoL 2005 : the 10th Conference on Formal Grammar and the 9th Meeting on Mathematics of Language, Edinburgh, 5–7 August, 2005*. 139–148.
- [4] Jeffrey Bardzell, Shaowen Bardzell, Guo Zhang, and Tyler Pace. 2014. The lonely raccoon at the ball: designing for intimacy, sociability, and selfhood. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Toronto Ontario Canada, 3943–3952. <https://doi.org/10.1145/2556288.2557127>
- [5] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *Conference on Fairness, Accountability, and Transparency (FAcT '21)* (mar 2021).
- [6] Claus Bossen, Kathleen H Pine, Federico Cabitza, Gunnar Ellingsen, and Enrico Maria Piras. 2019. Data work in healthcare: An Introduction. *Health Informatics Journal* 25, 3 (Sept. 2019), 465–474. <https://doi.org/10.1177/1460458219864730>
- [7] C. E. Brodley and M. A. Friedl. 1999. Identifying Mislabeled Training Data. *Journal of Artificial Intelligence Research* 11 (Aug. 1999), 131–167. <https://doi.org/10.1613/jair.606>
- [8] Andrea D. Bührmann and Werner Schneider. 2007. Mehr als nur diskursive Praxis? – Konzeptionelle Grundlagen und methodische Aspekte der Dispositivanalyse. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research* Vol 8, No 2: From Michel Foucault's Theory of Discourse to Empirical Discourse Research (May 2007). <https://doi.org/10.17169/FQS-8.2.237>

- [9] Joannah Caborn. 2016. On the Methodology of Dispositive Analysis. *Critical Approaches to Discourse Analysis Across Disciplines* 1, 1 (2016), 115–123. <https://doi.org/10.5209/CLAC.53494>
- [10] Antonio A. Casilli. 2017. Digital labor studies go global: Toward a digital decolonial turn. *International Journal of Communication* 11 (2017), 3934–3954.
- [11] Antonio A. Casilli and Julian Posada. 2019. The Platformisation of Labor and Society. In *Society and the Internet* (vol. 2 ed.), Mark Graham and William H. Dutton (Eds.). Oxford University Press, Oxford.
- [12] Justin Cheng and Dan Cosley. 2013. How annotation styles influence content and preferences. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media - HT '13*. Association for Computing Machinery, Paris, France, 214–218. <https://doi.org/10.1145/2481492.2481519> tex.ids: cheng2013a.
- [13] Kate Crawford. 2021. *Atlas of AI. Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, New Haven, CT.
- [14] Ciaran Cronin. 1996. Bourdieu and Foucault on power and modernity. *Philosophy & Social Criticism* 22, 6 (Nov. 1996), 55–85. <https://doi.org/10.1177/019145379602200603>
- [15] Peter Dauvergne. 2020. *AI in the Wild. Sustainability in the Age of Artificial Intelligence*. MIT Press, Cambridge, MA.
- [16] Alessandro Delfanti and Sarah Sharma (Eds.). 2019. Log Out! The Platform Economy and Worker Resistance. *Notes from Below* 8 (2019).
- [17] Catherine D'Ignazio and Lauren F. Klein. 2020. *Data feminism*. The MIT Press, Cambridge, Massachusetts. <https://mitpress.mit.edu/books/data-feminism>
- [18] 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>
- [19] Melanie Feinberg. 2017. A Design Perspective on Data. In *CHI '17: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, Denver, Colorado, USA, 2952–2963. <https://doi.org/10.1145/3025453.3025837>
- [20] Uwe Flick. 2007. *Qualitative Sozialforschung: Eine Einführung* (10, erweiterte neuauflage ed.). Rowohlt Taschenbuch, Reinbek bei Hamburg.
- [21] Michel Foucault. 1971. Orders of discourse. *Social Science Information* 10, 2 (April 1971), 7–30. <https://doi.org/10.1177/053901847101000201> Publisher: SAGE Publications Ltd.
- [22] Michel Foucault. 1982. *The Archaeology of Knowledge: And the Discourse on Language*. Vintage, New York.
- [23] Michel Foucault. 1982. The Subject and Power. *Critical Inquiry* 8, 4 (1982), 777–795. <https://www.jstor.org/stable/1343197>
- [24] Michel Foucault. 1996. What Is Critique? In *What is Enlightenment?: Eighteenth-Century Answers and Twentieth-Century Questions*, James Schmidt (Ed.). University of California Press.
- [25] Michel Foucault and Colin Gordon. 1980. *Power/knowledge: selected interviews and other writings, 1972-1977* (1st american ed ed.). Pantheon Books, New York.
- [26] Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. 2017. Clarity is a Worthwhile Quality: On the Role of Task Clarity in Microtask Crowdsourcing. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media (HT '17)*. Association for Computing Machinery, New York, NY, USA, 5–14. <https://doi.org/10.1145/3078714.3078715>
- [27] 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>
- [28] Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. Are We Modeling the Task or the Annotator? An Investigation of Annotator Bias in Natural Language Understanding Datasets. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 1161–1166. <https://doi.org/10.18653/v1/D19-1107>
- [29] Bhavya Ghai, Q. Vera Liao, Yunfeng Zhang, and Klaus Mueller. 2020. Measuring Social Biases of Crowd Workers using Counterfactual Queries. Honolulu, HI, USA. <http://fair-ai.owlstown.com/publications/1424>
- [30] Mary L. Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt, Boston.
- [31] Julian Hamann, Jens Maesse, Ronny Scholz, and Johannes Angermuller. 2019. The Academic Dispositif: Towards a Context-Centred Discourse Analysis. In *Quantifying Approaches to Discourse for Social Scientists*, Ronny Scholz (Ed.). Springer International Publishing, Cham, 51–87. [https://doi.org/10.1007/978-3-319-97370-8\\_3](https://doi.org/10.1007/978-3-319-97370-8_3)
- [32] Karen Hao. 2020. In 2020, let's stop AI ethics-washing and actually do something. *MIT Technology Review* (2020). <https://www.technologyreview.com/2019/12/27/57/ai-ethics-washing-time-to-act/>

- [33] 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>
- [34] Ellie Harmon and Melissa Mazmanian. 2013. Stories of the Smartphone in everyday discourse: conflict, tension & instability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Paris France, 1051–1060. <https://doi.org/10.1145/2470654.2466134>
- [35] Heather A. Horst and Daniel Miller. 2012. *Digital Anthropology*. Berg. 196–213 pages.
- [36] Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. Understanding and Mitigating Worker Biases in the Crowdsourced Collection of Subjective Judgments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300637> tex.ids: hube2019a event-place: Glasgow, Scotland Uk.
- [37] Lilly Irani. 2015. The cultural work of microwork. *New Media & Society* 17, 5 (2015), 720–739. <https://doi.org/10.1177/1461444813511926>
- [38] 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>
- [39] Siegfried Jäger and Florentine Maier. 2016. Analysing discourses and dispositives: A Foucauldian approach to theory and methodology. *Methods of critical discourse studies* (2016), 109–136.
- [40] Brian Justie. 2021. Little history of CAPTCHA. *Internet Histories* 5, 1 (2021), 30–47. <https://doi.org/10.1080/24701475.2020.1831197>
- [41] Siegfried Jäger. 2007. *Deutungskämpfe: Theorie und Praxis Kritischer Diskursanalyse*. Springer-Verlag.
- [42] Gopinath Kannabiran, Jeffrey Bardzell, and Shaowen Bardzell. 2011. How HCI talks about sexuality: discursive strategies, blind spots, and opportunities for future research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Vancouver BC Canada, 695–704. <https://doi.org/10.1145/1978942.1979043>
- [43] Lawrence F. Katz and Alan B. Krueger. 2016. The Role of Unemployment in the Rise in Alternative Work Arrangements. , 10 pages. <https://doi.org/10.1257/aerp20171092>
- [44] 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. (AIIES '20)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3375627.3375809>
- [45] Granate Kim. 2019. Microsoft Funds Facial Recognition Technology Secretly Tested on Palestinians. *Truthout* (2019).
- [46] Yubo Kou, Xinning Gui, Yunan Chen, and Bonnie Nardi. 2019. Turn to the Self in Human-Computer Interaction: Care of the Self in Negotiating the Human-Technology Relationship. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–15. <https://doi.org/10.1145/3290605.3300711>
- [47] Valérie Larroche. 2019. *The Dispositif: A Concept for Information and Communication Sciences* (1 ed.). Wiley. <https://doi.org/10.1002/9781119508724>
- [48] Jürgen Link. 2014. Dispositiv. In *Foucault-Hanbuch*, Clemens Kammler, Rolf Parr, Ulrich Johannes Schneider, and Elke Reinhardt-Becker (Eds.). J.B. Metzler, Stuttgart, 237–242. [https://doi.org/10.1007/978-3-476-01378-1\\_27](https://doi.org/10.1007/978-3-476-01378-1_27)
- [49] Leib Litman, Jonathan Robinson, and Cheskie Rosenzweig. 2015. The relationship between motivation, monetary compensation, and data quality among US- and India-based workers on Mechanical Turk. *Behavior Research Methods* 47, 2 (June 2015), 519–528. <https://doi.org/10.3758/s13428-014-0483-x>
- [50] Alex Jiahong Lu, Tawanna R. Dillahunt, Gabriela Marcu, and Mark S. Ackerman. 2021. Data Work in Education: Enacting and Negotiating Care and Control in Teachers' Use of Data-Driven Classroom Surveillance Technology. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (oct 2021), 1–26. <https://doi.org/10.1145/3479596>
- [51] Katharina Manderscheid. 2014. The Movement Problem, the Car and Future Mobility Regimes: Automobility as Dispositif and Mode of Regulation. *Mobilities* 9, 4 (Oct. 2014), 604–626. <https://doi.org/10.1080/17450101.2014.961257>
- [52] 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>
- [53] 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>
- [54] 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>

- [55] Naja Holten Møller, Claus Bossen, Kathleen H. Pine, Trine Rask Nielsen, and Gina Neff. 2020. Who Does the Work of Data? *Interactions* 27, 3 (April 2020), 52–55. <https://doi.org/10.1145/3386389>
- [56] Bruno Moreschi, Gabriel Pereira, and Fabio G. Cozman. 2020. The Brazilian Workers in Amazon Mechanical Turk: Dreams and realities of ghost workers. *Revista Contracampo* 39, 1 (apr 2020). <https://doi.org/10.22409/contracampo.v39i1.38252>
- [57] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q. Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, Glasgow, Scotland Uk, 1–15. <https://doi.org/10.1145/3290605.3300356>
- [58] 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.
- [59] Gemma Newlands. 2021. Lifting the curtain: Strategic visibility of human labour in AI-as-a-Service. *Big Data & Society* 8, 1 (Jan. 2021), 205395172110160. <https://doi.org/10.1177/20539517211016026>
- [60] Magdalena Nowicka-Franczak. 2021. Post-Foucauldian Discourse and Dispositif Analysis in the Post-Socialist Field of Research: Methodological Remarks. *Qualitative Sociology Review* 17, 1 (Feb. 2021), 72–95. <https://doi.org/10.18778/1733-8077.17.1.6>
- [61] Samir Passi and Steven Jackson. 2017. Data Vision: Learning to See Through Algorithmic Abstraction. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, Portland, Oregon, USA, 2436–2447. <https://doi.org/10.1145/2998181.2998331>
- [62] 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 (Nov. 2018), 1–28. <https://doi.org/10.1145/3274405>
- [63] Thomas Poell, David B. Nieborg, and José van Dijck. 2019. Platformisation. *Internet Policy Review* 8, 4 (2019). <https://doi.org/10.14763/2019.4.1425>
- [64] Julian Posada. 2020. The Future of Work Is Here: Toward a Comprehensive Approach to Artificial Intelligence and Labour. *Ethics in Context* (2020).
- [65] Julian Posada. 2022. Embedded Reproduction in Platform Data Work. *Information, Communication & Society* (2022).
- [66] Julian Posada, Nicholas Weller, and Wendy H. Wong. 2021. We Haven't Gone Paperless Yet: Why the Printing Press Can Help Us Understand Data and AI. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AI/ES '21)* (2021).
- [67] Rida Qadri. 2020. Algorithmized but not Atomized? How Digital Platforms Engender New Forms of Worker Solidarity in Jakarta. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. ACM, New York, NY, USA, 144–144. <https://doi.org/10.1145/3375627.3375816>
- [68] Sverre Raffnsøe, Marius Gudmand-Høyer, and Morten S. Thaning. 2016. Foucault's dispositive: The perspicacity of dispositive analytics in organizational research. *Organization* 23, 2 (March 2016), 272–298. <https://doi.org/10.1177/1350508414549885> Publisher: SAGE Publications Ltd.
- [69] Noopur Raval. 2021. Interrupting invisibility in a global world. *Interactions* 28, 4 (July 2021), 27–31. <https://doi.org/10.1145/3469257>
- [70] Sarah T. Roberts. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press, New Haven, CT. 280 pages. <https://doi.org/10.1177/1461444819878844>
- [71] 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>
- [72] Niloufar Salehi, Lilly C. Irani, Michael S. Bernstein, Ali Alkhatib, Eva Ogbé, 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>
- [73] 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>
- [74] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/3411764.3445518>
- [75] Morgan Klaus Scheuerman, Emily Denton, and Alex Hanna. 2021. Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development. *arXiv:2108.04308 [cs]* (Aug. 2021). <https://doi.org/10.1145/3476058> arXiv: 2108.04308.

- [76] Daniel Scheingart, Martin Trombetta, and Gisella Pascuariello. 2020. Primas salariales sectoriales en Argentina. *Ministerio de Desarrollo Productivo de la Nación*. Centro de Estudios para la Producción XXI (Nov. 2020), 39.
- [77] Cathrine Seidelin, Yvonne Dittrich, and Erik Grönvall. 2018. Data Work in a Knowledge-Broker Organisation: How Cross-Organisational Data Maintenance Shapes Human Data Interactions. In *Proceedings of the 32nd International BCS Human Computer Interaction Conference (Belfast, United Kingdom) (HCI '18)*. BCS Learning Development Ltd., Swindon, GBR, Article 14, 12 pages. <https://doi.org/10.14236/ewic/HCI2018.14> (2018).
- [78] Katta Spiel. 2017. Critical Experience: Evaluating (with) Autistic Children and Technologies. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 326–329. <https://doi.org/10.1145/3027063.3027118>
- [79] Divy Thakkar, Azra Ismail, Pratyush Kumar, Alex Hanna, Nithya Sambasivan, and Neha Kumar. 2022. When is Machine Learning Data Good?: Valuing in Public Health Datafication. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)* (2022).
- [80] Divy Thakkar, Neha Kumar, and Nithya Sambasivan. 2020. Towards an AI-powered Future that Works for Vocational Workers. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376674>
- [81] 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>
- [82] United Nations. 2015. Sustainable Development Goals. <https://www.un.org/sustainabledevelopment/>
- [83] United Nations General Assembly. 1999. *Measures to combat contemporary forms of racism, racial discrimination, xenophobia and related intolerance*. Technical Report. United Nations.
- [84] Fabian L. Wauthier and Michael I. Jordan. 2011. Bayesian Bias Mitigation for Crowdsourcing. In *Proceedings of the 24th International Conference on Neural Information Processing Systems (NIPS'11)*. Curran Associates Inc., Granada, Spain, 1800–1808. <http://papers.nips.cc/paper/4311-bayesian-bias-mitigation-for-crowdsourcing.pdf>
- [85] Glen Whelan. 2019. Born Political: A Dispositive Analysis of Google and Copyright. *Business & Society* 58, 1 (Jan. 2019), 42–73. <https://doi.org/10.1177/0007650317717701>
- [86] Ricky Wichum. 2013. Security as Dispositif: Michel Foucault in the Field of Security. *Foucault Studies* (Jan. 2013), 164–171. <https://doi.org/10.22439/fs.v0i15.3996>
- [87] Alex J. Wood, Mark Graham, Vili Lehdonvirta, and Isis Hjorth. 2019. Networked but Commodified: The (Dis)Embeddedness of Digital Labour in the Gig Economy. *Sociology* (2019). <https://doi.org/10.1177/0038038519828906>
- [88] Alex J. Wood, Vili Lehdonvirta, and Mark Graham. 2018. Workers of the Internet unite? Online freelancer organisation among remote gig economy workers in six Asian and African countries. *New Technology, Work and Employment* 33, 2 (2018), 95–112. <https://doi.org/10.1111/ntwe.12112>
- [89] Jamie Woodcock and Mark Graham. 2020. *The Gig Economy: A Critical Introduction*. Polity Press, London. 160 pages.
- [90] Shoshana Zuboff. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, New York, NY. 691 pages.

Received January 2022; revised April 2022; accepted May 2022