

# Advancing NLP Data Equity: Practitioner Responsibility and Accountability in NLP Data Practices

Jay L. Cunningham<sup>1</sup>, Kevin Zhongyang Shao<sup>2</sup>, Rock Yuren Pang<sup>2</sup>, Nathanael Elias Mengist<sup>2</sup>

<sup>1</sup>DePaul University  
Chicago, IL USA

<sup>2</sup>University of Washington  
Seattle, WA USA

jgunni37@depaul.edu, kshao918@uw.edu, ypang2@uw.edu, mengin@uw.edu

## Abstract

While research has focused on surfacing and auditing algorithmic bias to ensure equitable AI development, less is known about how NLP practitioners, those directly involved in dataset development, annotation, and deployment, perceive and navigate issues of NLP data equity. This study is among the first to center practitioners’ perspectives, linking their experiences to AI governance frameworks and advancing participatory recommendations that bridge technical, policy, and community domains. Drawing on a 2024 questionnaire and focus group, we examine how U.S. based NLP data practitioners conceptualize fairness, contend with organizational and systemic constraints, and engage emerging governance efforts such as the U.S. AI Bill of Rights. Findings reveal persistent tensions between commercial objectives and equity commitments, alongside calls for more participatory and accountable data workflows. We critically engage debates on data diversity and “diversity-washing,” arguing that improving NLP equity requires structural governance reforms that support practitioner agency and community consent.

## Introduction

As language technologies permeate every facet of society, the need for responsible and inclusive dataset practices has never been more critical. Recent scholarship has documented persistent challenges in natural language processing (NLP) systems, including automated speech recognition (ASR), text classification, and language generation (Koencke et al. 2020; Sap et al. 2019a; Deas et al. 2023). Technologies such as smart virtual assistants and large language models often perform unfairly, reflecting biases rooted in dataset production practices, including limited diversity in training data and in the technologists who develop these systems (Bender and Friedman 2018; Goyal et al. 2022). In data science for machine learning (ML), dataset production involves multiple stages—curation, collection, annotation, and implementation (Muller et al. 2019); each shaped by human value judgments that can encode implicit biases toward or against certain language varieties.

These practices can have serious real-world consequences. Recent cases include low-income older Americans

being denied health benefits because of algorithmic decisions and a pregnant Black American woman wrongfully arrested due to flawed facial recognition technology. In response, fairness research communities within CHI, FAccT, and AIES have emphasized the need for direct participation by underrepresented socio-cultural groups in NLP data production (Denton et al. 2021). Community-driven initiatives such as DataWorks have demonstrated participatory approaches by annotating African American English (AAE) datasets to address representational harms (Rothschild et al. 2022). Policymakers have also advanced AI governance frameworks, including the U.S. AI Bill of Rights, to safeguard against socio-technical harms and promote socially responsible design (Lazar and Nelson 2023).

While prior work has examined fairness across AI/ML pipelines (Holstein et al. 2019; Miceli, Schuessler, and Yang 2020; Thakkar et al. 2022; Widder and Nafus 2023), NLP presents distinct equity challenges that are comparatively underexplored. The rapid deployment of large language models has transformed practitioners’ roles, intensifying pressures to address equity within complex organizational contexts. Yet there remains limited empirical understanding of how NLP practitioners themselves perceive and navigate data equity challenges, despite their direct influence on dataset integrity and, ultimately, system behaviors (Gebru et al. 2021-12; Suchman 2002).

We define NLP data equity as the practice of ensuring that NLP datasets are representative, contextually accurate, and co-produced in ways that acknowledge and mitigate structural inequities affecting marginalized groups. This requires both diversifying data sources and addressing power asymmetries in collection, annotation, and model deployment to support fair outcomes and epistemic justice (Barocas, Hardt, and Narayanan 2020; Hanna et al. 2020; Denton et al. 2021).

Our study includes practitioners working with both text-based and speech-based NLP systems, including projects such as ASR, virtual assistants, sentiment analysis, machine translation, and language generation<sup>1</sup>. We therefore treat data equity as a cross-cutting concern across modalities. Where relevant, we note that certain issues—such as linguistic bias, representational harm, and dialect recognition—may manifest differently in speech versus text sys-

<sup>1</sup>(See the Appendices (C) in extended paper version)

tems. Insights from this study may inform both domains.

This study investigates how NLP practitioners engage with data equity within their professional contexts. We focus on their awareness, decision-making processes, and perceived constraints, as well as their perspectives on governance and policy.

**We ask:**

1. How do NLP practitioners understand and navigate NLP data equity challenges within their workflows, and how are these shaped by organizational constraints and accountability structures?
2. What strategies, training, and resources do practitioners report using or lacking to identify and mitigate bias in NLP datasets?
3. How do practitioners perceive the role of governance, policy, and multi-stakeholder engagement in advancing equitable NLP data practices?

While past work has examined practitioner awareness of fairness in AI/ML broadly (e.g., Holstein et al., 2019; Deng et al., 2023), few studies center on NLP-specific data practices and practitioners’ agency in navigating fairness-related constraints in industry and research settings. Our study bridges policy and practice by examining how practitioners perceive and implement AI governance frameworks.

*\*An extended version of this paper, which includes supplementary materials, technical appendices, and expanded methodological details, is available on arXiv<sup>2</sup> under a Creative Commons Attribution (CC BY) license.*

## Related Work

### Critical Approaches in ML Data Practices

Critical approaches have been interrogating the foundational assumptions that shape the ML dataset production, and advocating for strategies to reduce bias and ensure the inclusion of socio-culturally diverse data in ML applications (Mothilal, Guha, and Ahmed 2024). For instance, *Data Feminism* calls for a power redistribution in data practices to achieve greater equity. Jo and Gebru (2020) highlights the limitations in many existing datasets, proposing more responsible data practices that prioritize consent, inclusion, transparency, power dynamics, and privacy. Additional recommendations included assembling diverse pools of annotators, incorporating iterative feedback processes, and fine-tuning models to account for ethnic dialects, alongside improving documentation protocols (e.g., datasheets for datasets) (Gebru et al. 2021-12; Ziman et al. 2018; Kapania, Taylor, and Wang 2023; Miceli et al. 2022; Pang et al. 2023; Raji et al. 2021). While data equity has been extensively studied in AI/ML, nuances on NLP-specific issues, such as linguistic representation, dialect variance, and cultural annotation complexity, are often overlooked (Deng et al. 2024). Moreover, the emergence of LLMs has changed how NLP practitioners interact with data curation workflows.

Among these strategies, participatory design (Muller and Kuhn 1993) has been used to improve ML systems by involving community stakeholders directly in the dataset production process, ensuring that the resulting data accurately represent diverse perspectives (Sloan et al. 2020; Cunningham 2023). Human-centered methods, such as community workshops, co-design sessions, and feedback mechanisms, are recommended as research tools to align with participatory modes of NLP dataset production. These methods promote a more inclusive approach to data creation by integrating input from the communities whose language and perspectives are represented in the data, thus contributing to more equitable and accurate models. Our work also intersects with HCI scholarship on data infrastructures and activism, which highlights how socio-technical participation is constrained and enabled by institutional, material, and political conditions (Irani 2016; Dourish 2017).

### Challenges of NLP Data Equity

As we defined, NLP data equity refers to a practice ensuring that NLP datasets themselves are representative and contextually accurate, as well as co-produced in ways that acknowledge and mitigate inequities affecting marginalized groups. However, achieving this ideal faces challenges, particularly given the rapid pace of AI development that often prioritizes speed over equitable data practices. For example, Paullada et al. (2021) delineates the limitations in historical and contemporary NLP data practices, focusing on the challenges of dataset design, development, and introspection, alongside the broader practices surrounding dataset usage. In particular, tensions between the NLP system performance and the data of linguistically diverse users expose the mis-alignments between the creation of the dataset and the real-world linguistic diversity, presenting complex socio-technical challenges that require critical reflection on current data practices (Mathur et al. 2022; Joshi et al. 2020).

Implicit biases perpetuated in core NLP data practices such as transcription and annotation have been identified as key contributors to inadequate AI models, potentially exacerbating societal inequalities (Hutchinson et al. 2020). One illustrative example involves the Perspective API, which assigned disproportionately high toxicity scores to tweets featuring vernacular elements and linguistic structures associated with African American English (AAE) (Sap et al. 2019b). Such discrepancies highlight the inadequacies of models that fail to account for sociolinguistic diversity, often leading to suboptimal performance for users who communicate in minority language varieties (Kamikubo et al. 2022). This challenge is particularly pronounced for speakers of AAE, where existing NLP models have been shown to inadequately due to insufficient linguistic representation (Tatman 2017; Cunningham et al. 2024).

Scholarly discourse has also identified a misalignment between industry metrics for successful performance and the realities of socially situated, culturally nuanced human communication. For instance, Kapania, Taylor, and Wang (2023) identified primary obstacles to incorporating diversity into data practices: insufficient information about annotators, operational fragmentation, and the competing priorities in ML

<sup>2</sup>navigate to <https://jaylcunningham.com/work> for the arXiv web-link.

development. To address these issues, some researchers propose a participatory framework in which practitioners collaborate with annotators familiar with the linguistic varieties (Kapania, Taylor, and Wang 2023; Sloane et al. 2020).

Our study builds on this body of work by empirically exploring the perspectives of NLP practitioners regarding these challenges, while further identifying the potential for collaborative co-creation of datasets (Kapania, Taylor, and Wang 2023; Paullada et al. 2021; Deas et al. 2023). Additionally, it expands upon previous research by examining the integration of stakeholder participation across multiple stages of data practices, beyond the scope of annotation alone (Kapania, Taylor, and Wang 2023; McConvey and Guha 2024). To this regard, we integrate critiques of synthetic diversity and diversity-washing (Whitney and Norman 2024; Dalal, Hall, and Johnson 2024), highlighting how superficial diversity initiatives can obscure systemic inequities. Ultimately, this study builds on prior work (Holstein et al. 2019; Deng et al. 2023) by focusing specifically on how NLP practitioners interpret ML fairness frameworks, regulatory guidelines, and equity mandates in real-world contexts.

### **AI Policies on Addressing Algorithmic Disparities**

The introduction of policies and legislation has emerged as a response to the growing societal challenges posed by NLP technologies, as exemplified by initiatives such as the EU AI Act (Act 2024), the U.S. White House AI Bill of Rights (U.S. 2022). In 2023, the National AI Research Resource Task Force released a report providing recommendations for the establishment of a public resource to support AI research in the United States, emphasizing the need to consider societal impacts during the review of proposals that seek access to public resources. The U.S. AI Bill of Rights further emphasizes that AI-driven automated systems must actively and continuously protect the public from algorithmic discrimination. It stipulates that data used in system development or evaluation must be representative of the local communities where the AI will be deployed, and that this data should be scrutinized for bias within its historical and societal context to identify and mitigate potential harms.

Although these policies largely concentrate on downstream applications of AI systems, questions remain regarding the practical implementation of the proposed measures by technologists. In particular, NLP training datasets, which are central to the accuracy of downstream tasks, are shaped by data science practices that play a critical role in shaping model outcomes. This study explores how NLP practitioners perceive the applicability of the U.S. AI Bill of Rights' guidelines in promoting diversity and safeguarding against algorithmic harms within their own work. The findings contribute to a broader understanding of the real-world challenges associated with applying policy frameworks to technical data workflows in the development of AI systems.

### **AI Governance and Practitioner Responsibility**

AI governance has emerged as a critical area of scholarly and policy attention, encompassing the mechanisms, institutions, and practices through which the development, deployment, and oversight of AI systems are regulated (Florida

et al. 2018). While much attention has focused on technical standards, policy instruments, and high-level regulatory frameworks (Jobin, Ienca, and Vayena 2019), emerging scholarship highlights the need for multi-scalar governance strategies that engage actors across the AI development life-cycle, including practitioners, communities, industry stakeholders, and regulators (Drahos 2017).

In the context of NLP systems, however, practitioner responsibility remains under-theorized, with limited attention to how data workers, engineers, and analysts are positioned within governance debates. Recent critiques have also called for expanding governance beyond compliance frameworks, addressing how inclusion, diversity, and fairness discourses in AI can risk reinforcing extractive or colonial data practices if not critically interrogated (Whitney and Norman 2024; Dalal, Hall, and Johnson 2024). This critique is particularly salient for NLP practitioners navigating data curation and annotation tasks, where decisions about dataset composition, representational adequacy, and community engagement carry significant governance implications.

In parallel, recent scholarship has raised concerns about the limits of current governance approaches and diversity initiatives in AI development. While many efforts emphasize fairness principles or public-private partnerships, regulatory scholars argue for more expansive, multi-scalar governance frameworks that engage actors across policy, institutions, and the development pipeline (Drahos 2017; Jobin, Ienca, and Vayena 2019). These perspectives foreground the need to embed structural accountability and participatory oversight within AI systems development. Relatedly, emerging critiques of “diversity-washing” caution that efforts to increase dataset representativeness can obscure underlying power asymmetries or bypass community consent, particularly when synthetic or compliance-driven diversity metrics are used (Whitney and Norman 2024; Dalal, Hall, and Johnson 2024). Our work builds on these insights by focusing on how NLP practitioners themselves perceive their roles within these governance and equity tensions—highlighting both the constraints they face and the opportunities for more accountable, community-centered data practices.

### **Methods**

This study examines how some NLP practitioners advance data equity by reflecting on practices that produce representative, contextually accurate, and justice-oriented datasets to address technical misalignments and social inequities among diverse users. The analysis focuses on practitioner perspectives and experiences, rather than direct audits of NLP systems, to understand how organizational contexts, governance frameworks, and individual agency shape equity-oriented data practices.

We used a mixed-methods design, including a survey ( $N = 47$ ) and a focus group ( $N = 5$ ), to capture both breadth and depth in practitioner perspectives. Surveys have been effective in mapping practitioner attitudes in NLP fairness research (Sambasivan et al. 2021), while focus groups enable richer, context-specific insights (Miceli et al. 2022). Our single focus group used purposive sampling of “critical practitioners” to probe reasoning processes behind survey

trends; this approach prioritizes transferability over statistical generalization (Lincoln 1985).

Survey responses were analyzed descriptively in Microsoft PowerBI using frequency distributions, cross-tabulations, and percentages. Open-ended responses and focus group transcripts were coded thematically to identify patterns in bias mitigation experience, governance engagement, and organizational constraints.

All study procedures were reviewed and approved by the University of Washington Institutional Review Board (IRB Protocol).

## Questionnaire

The goal of the questionnaire was to understand on current NLP data practices in a lightweight and accessible manner. Specifically, the questionnaire was used to understand the strategies employed by NLP professionals to mitigate bias and ensure representation within datasets. Given the semi-structured, exploratory nature of the study, some questions were informed by topics of related works and developed through unvalidated measures (Hill et al. 2022) to address our research questions. The questionnaire was administered via the online platform (Typeform) in late 2023, and responses from 47 eligible participants were analyzed.

**Questionnaire Recruitment and Respondents** The study targeted self-identified NLP data practitioners with professional experience in text and speech NLP systems, encompassing roles such as data engineers, data scientists, data annotators, data researchers, and crowdworkers. Recruitment strategies included outreach via online social networking platforms (e.g., X (Twitter), LinkedIn, Reddit, GroupMe), direct emails, and mailing lists from both internal and external organizations. Additionally, professional networks based on prior collaborations with academic and industry AI teams (e.g., Google, Microsoft, Apple, and Meta) were leveraged. The questionnaire started with an information statement detailing the objectives, benefits, and incentives of participation, followed by a disclosure of study procedures and an acknowledgment of informed consent. Eligibility screener questions ensured participants were at least 18 years old, resided and worked in the United States, and had professional experience in NLP dataset production.

**Questionnaire Design** The questionnaire comprised a total of 27 questions<sup>3</sup>, categorized into quantitative (i.e., 11 multiple-choice, 5 multi-select, 3 rating scale questions) and qualitative: (i.e., 8 open-ended questions). Subsequently, the questionnaire explored a variety of themes that were mapped to the study’s research questions. The questionnaire concluded with questions on demographic information and their interest in engaging in subsequent focus group studies.

**Participant Demographic Details** We collect and report context regarding demographics of study participants as a way to evaluate claims of data practitioner identity on the quality and performance of ML systems. Out of the 60 respondents who began the questionnaire, 13 were deemed ineligible or declined to complete the survey, resulting in 47 eligible responses for analysis. These respondents, with

an average age of 29 (ranging from 19 to 40), held various professional roles in NLP data practices, including researcher, scientist, and engineer positions. The majority of them (48.94%) have an intermediate level of expertise for NLP data practices. Beginners account for 10.64%, Advanced participants represent 25.53%, and Experts make up 14.89% of the surveyed group. Demographic data collected included gender identity, race and ethnicity, languages spoken, and educational credentials<sup>4</sup>. Specifically, the sample comprised 32 men, 13 women, one non-binary/non-conforming individual, and one participant who preferred not to disclose their gender identity. For those interested in participating in a drawing for a 50 USD Amazon gift card, email addresses were collected and aggregated into a participant interest pool. Participant demographic details are provided in (Table 1).

| Race/Ethnicity                    | # of respondents |
|-----------------------------------|------------------|
| Black or African American         | 20               |
| White or Caucasian                | 13               |
| Asian                             | 10               |
| American Indian or Alaskan Native | 1                |
| Hispanic or Latin Origin          | 1                |
| Other                             | 2                |

Table 1: Survey participant demographic distribution ( $N = 47$ ) by race/ethnicity, highlighting the sample diversity.

## Focus Group

This study included a single focus group of five critically engaged NLP practitioners. While this limits the breadth of perspectives represented, the focus group was designed as an exploratory complement to the broader survey, offering depth rather than statistical saturation. Our intention was not to reach thematic saturation across the field but to capture detailed narratives that could illuminate the organizational and governance contexts shaping equity-related decision-making. The insights are thus transferable to similar professional and institutional settings, providing a foundation for future, larger-scale qualitative research.

**Focus Group Moderation** In March 2024, two members of the research team facilitated a data analysis focus group involving five practitioners to extend the investigation and data analysis beyond the initial survey results. The focus group aimed to garner more nuanced qualitative data through insights gained from these practitioners supplementary to data from the initial questionnaire. By moderating an expert group discussion, we aimed to corroborate questionnaire insights and involve participants in data analysis through collective dialogue and consensus, further affirming informed strategies for enhancing fairness and broadening representation within NLP dataset production. The lead researcher directed participants to their provided guidebook ,

<sup>3</sup>(See the Appendices (A) in extended paper version)

<sup>4</sup>(See the Appendices (C) in extended paper version)

then transitioned the group into the questionnaire insights review activity. Participants were initially given 10-minutes to navigate through the PowerBI report on their own and were asked to review the report and jot notes of any interesting insights or analytics, in addition to any questions that they may have about the data to be discussed further.

**Recruitment and Participants** The inclusion criteria for the focus group mirrored those employed for the survey, ensuring consistency in participant selection. Out of 35 respondents who expressed interest in a 1.5-hour virtual expert focus group, five participants were recruited based on the subjective quality of their survey responses, including factors such as compelling content, evident interest in the topic, and nuanced contextual insights. The professional backgrounds of these participants were diverse, with two engaged in academic research and others specializing in data analysis, language engineering, and economic planning. Their areas of expertise spanned responsible AI, biotechnology, and economic planning. Each focus group participant received a \$75 USD gift card as an incentive for their successful participation. The focus group session was audio recorded using Jitsi, a free, open-source, and encrypted online conferencing software. Participants were instructed to disable their video cameras and avoid using profile images during the session to ensure anonymity. In contrast, the facilitating researchers' video feeds remained active to foster trust and confidence among participants. To further safeguard the study's integrity and participant privacy, each participant was assigned a code name (Table 2) based on their professional title (e.g., researcher 1, engineer 2, annotator 3) for use throughout the focus group discussion.

## Study Data Analysis

Our research team analyzed and visualized the questionnaire data using Microsoft Power BI, summarizing the data through quantitative and qualitative measures. The quantitative descriptive analysis approach facilitated a better understanding of the questionnaire-data distribution and variability through calculations of means, percentages, and number of responses. For instance, for questions involving multiple choice, multi-select options, and rating scales, we employed calculations alongside data visualizations to illustrate the distributions. We performed inferential analysis to make interpretations or inferences about the population based on the survey data (Marshall and Jonker 2011).

For open-ended questions, we employed a grounded-theory approach to code responses and thematically analyze them (Braun and Clarke 2006-01; Olson and Kellogg 2014). We conducted multiple rounds of open coding involving multiple members of the research team to connect the broader qualitative excerpts with overarching themes of the survey study areas, thus informing the findings. To effectively communicate sentiments from participants, we include direct quotes from in the Findings section, with the survey respondents mentioned broadly as “respondent,” and focus group participants by their assigned participant ID (e.g., R1) to differentiate them in the analysis.

## Findings

The findings presented here are not intended to be statistically generalizable across all NLP practitioners. Instead, they are analytically transferable to contexts where practitioners operate under comparable organizational conditions and governance constraints. This framing aligns with qualitative research principles that emphasize the depth and contextual richness of insights over broad representativeness.

Survey findings establish broad patterns in practitioner experiences and perspectives, while focus group insights add depth from critically engaged participants. The survey data provides a wide view of how NLP professionals encounter and interpret data equity challenges, while the focus group offers rich, contextual accounts of equity motivations, organizational constraints, and governance engagement. Together, these methods illuminate both the prevalence of certain views across the field and the underlying structural and institutional factors that shape them.

In this section, survey questions are referenced using the notation “SQ” followed by the question number (e.g., SQ-10c). Focus group excerpts are labeled “FG” with a sequential number (e.g., FG-3). These labels correspond to survey or focus group questions.<sup>5</sup>

### Practitioners' Applied Data Practices are Harms Neutral

The initial set of survey questions sought to understand the range of experiences of NLP practitioners in organizing, collecting, transcribing, and analyzing text or speech data for NLP datasets. The survey study starts with surfacing the various areas of experience among practitioners in [SQ-6a]. The analysis primarily indicated that the data work of most NLP practitioners was centered on data collection or annotation for training of the text / speech model (76.60%), research on data practices or methods (46.81%) and organizing and managing data annotation projects (42.55%). These activities were particularly relevant to NLP technologies such as automated speech recognition (ASR), virtual assistants, smart home devices, and sentiment analysis. The full spectrum of NLP technologies that NLP practitioners reported working on.<sup>6</sup> However, we find that practitioners had the least experience mitigating bias in data training sets (23.40%) and developing evaluation approaches to measure the effectiveness of training data (31.91%).

This context indicates that while NLP practitioners are extensively engaged in the applied operational facets of data handling, they can be neutral to emergent harms, exposing deficiencies in practices related to recognizing and mitigating potential biases, along with evaluating the outcomes of training datasets for users and systems. This discrepancy directed our focus toward understanding the underlying reasons for the limited engagement in harms reducing practices.

**Barriers to Bias Mitigation** The focus group discussions provided deeper insight into why might NLP practitioners to be less experienced in mitigating bias [FG-6a], revealing

<sup>5</sup>(See the Appendices (B) in extended paper version)

<sup>6</sup>(See the Appendices (C) (Table 3) in extended paper version)

| Codename              | Gender Identity | Age | Job Title            | Industry/Field                 |
|-----------------------|-----------------|-----|----------------------|--------------------------------|
| Researcher 1 (R1)     | He/Him          | 25  | Academic Researcher  | PhD research                   |
| Researcher 2 (R2)     | She/Her         | 24  | Academic Researcher  | Responsible AI/Tech, Education |
| Data Analyst 1 (DA1)  | He/Him          | 29  | Biotech Data Analyst | Bio tech                       |
| Data Engineer 1 (DE1) | She/Her         | 29  | Language Engineer    | Natural language processing    |
| Data Analyst 2 (DA2)  | He/Him          | 33  | Economic Planner     | Economics Planning             |

Table 2: Profile of focus group participants.

both practical and philosophical challenges that may hinder efforts in this area. Participant DE1 expressed concerns over the potential repercussions of identifying biases, including financial costs associated with mitigation strategies and negative impacts on organizational reputation, which may discourage proactive measures in bias mitigation. In addition, the mention of “short contracts don’t meet basic needs” highlights another practical issue—temporary or short-term visions in NLP projects might not provide the stability or incentives needed for workers to engage deeply with complex issues like bias mitigation.

### Circumstances of Integrating Diverse Data and Inclusion Conscious Practices

Participants reflected on how the increasing reliance on LLMs reduces attention to equity-oriented data goals and raises the risk of flattening sociolinguistic nuance. Some voiced skepticism toward diversity as a compliance goal rather than a community-driven value. Our findings give insight into NLP practitioners’ efforts for achieving diversity and representation within data science practices. 9 out of 47 survey respondents elaborated on their efforts to strive for diverse representation and inclusion in their curated datasets [SQ-7a]. Notably, 28 respondents indicated that striving for diverse variations in their data collections to improve users inclusion and system performance were either not applicable to their roles or were not considerable factors.

The findings from this analysis indicate that a marginal percent of NLP practitioners are actively taking steps to promote diversity and inclusion in their datasets. Four primary themes emerged from the responses in relation to circumstances contributing to NLP data integration:

- **Seeking data from diverse sources:** Respondents emphasized their effort to acquire data from various sources representing different cultures, demographics, and communities, including online communities, minority-specific resources, and diverse media outlets.
- **Ethical and responsible data efforts with diverse teams:** Respondents mentioned ethical and responsible data collection methods to ensure representativeness and reduce bias. They also highlighted the role of diverse teams in identifying and mitigating biases within datasets.
- **Robust data quality and auditing practices** Ensuring data quality through regular reviews and updates was

identified as essential for maintaining dataset relevance and reflecting evolving cultural nuances of language. Techniques such as pre-processing, feedback loops, and verifying data accuracy across different demographics were mentioned as methods to mitigate bias and ensure accurate representation of target populations.

- **Cultivating a technical culture that prioritizes inclusion:** Respondents highlighted the importance of prioritizing inclusive awareness with the data-centric product team, particularly when project priorities vary based on stakeholder and business interests.

**Limitations in Sourcing Diverse Datasets** 21 survey respondents mentioned challenges in sourcing diverse datasets. To further understand this, we explored [FG-7a] the limitations and obstacles faced by NLP practitioners in the focus group discussion. Several barriers were revealed including: 1) Data volume and representation (e.g., “not all populations might ... provide the same ‘volume’ of data, to be considered across scale.”), 2) Considerations for diverse data variations (e.g., “data should be diverse and cover all races and all areas, not one region only.”), and 3) Cost and practical constraints (e.g., “a company will look at how they can produce the most data for the least cost first before considering ... diversity search.”).

**Consideration of Social Demographics and Cultural Nuances** With 17 respondents indicating that social demographics and cultural nuances were not applicable to their roles in [SQ-7b], the remaining 30 responses suggest that a growing number of NLP practitioners recognize its importance in ensuring diverse and representative datasets (e.g., “consider diverse social demographics and nuances by ensuring a broad representation...”). In addition, a few specific aspects of social demographics and cultural nuances were intentionally taken into considerations by participants. These include: 1) Race/Ethnicity; 2) Language: Accents, dialects, and English varieties; 3) Cultural background and sensitivities; 4) Regional variations based on country, state, and geographic areas; 5) Other aspects such as age, gender, income-level or socioeconomic status, accessibility challenges or disabilities (e.g., stuttering, stammering), and sexual-identity (LGBTQAI+). It was also noted that certain social identity aspects were more relevant to some NLP technologies than others. For example, one respondent noted “these features weren’t included when we applied sentimental analysis,” whereas another stated, “in my own work

within ASR research, in considering speakers of AAVE, we consider ethnic, cultural, and linguistic aspects.”

### **Defining and Measuring Success in NLP Dataset Annotation**

Practitioners described challenges in annotating culturally relevant text, curating datasets for diverse dialects, and aligning with internal metrics. These issues are unique to NLP systems where meaning is deeply context-dependent.

#### **Assessing the Implications of Data Annotation Practices**

Our survey showed the diverse criteria used by NLP practitioners to assess the performance of data annotators and their practices [SQ-8a]. A substantial majority of the respondents identified the quality of annotations (68.09%) and the precision of annotations (55.32%) as primary indicators of success, which underscores that producing reliable and precise annotations is essential for effective NLP model training. The respondents also indicated other applicable options: compliance with annotation guidelines (44.68%), feedback from other annotators (34.04%), and speed of annotation completion (31.91%). Another two respondents mentioned alternative factors such as participation in “internal annotator agreements” and comprehensive data collection efforts (e.g., “recruiting and paying participants, collecting their data/paperwork, and recording the tasks.”).

*Quality of Annotation and Accuracy of Annotations* were among the highest rated factors based on survey analysis, so our focus group discussions centered on how is quality and accuracy defined for data annotators and who develops metrics of success for data annotators [FG-8a]. One key finding is that the notions of “quality” and “accuracy” in annotation are inherently subjective. While factors like adherence to guidelines, consistency, and the overall usability of the annotated data are commonly cited, there is no universal definition that applies across all projects.

The focus group responses highlighted the tension between speed and quality as performance metrics. DE1 noted, “I spent two-ish years as an annotator and speed was considered a much more important factor as meeting deadlines was more important than having more quality data...” This underscores some projects prioritize meeting deadlines and high-volume data production, while others emphasize annotation precision. This tension varies with the project’s scale, data type, and goals (R1). Approaches to develop metrics for annotator success are diverse. Some approaches involve collaborative input from annotators, while others rely on a top-down approach by managers or clients (DE1). Based on participant responses, clear and precise guidelines are universally valued for steering annotation efforts and defining success metrics.

Additionally, survey insights highlighted the significance of time which is typically allocated to handle conflicts of opinion among annotators and their processes [SQ-8b]. The survey indicated that 30 of 47 respondents allocated a moderate to substantial amount of time to resolve potential conflicts of opinion. Active conflict resolution practices can help identify potential biases and create a more inclusive dataset.

### **Data Science Considerations Toward Improving Data Inclusion and Quality**

In our survey study, participants indicated their applicable factors that should be considered to improve representation and inclusion in NLP datasets [SQ-9a]: annotation guidelines (63.83%), identity inclusion datasets (61.70%), data collection sources (57.45%), model evaluation metrics (48.94%), and privacy review (42.5%).

In addition, respondents suggested a few additional consideration factors, such as the need to expose models to diverse cultures and backgrounds, the incorporation of non-algorithmic metrics to assess bias and harm, questioning dataset ownership and access rights, and advocating for alternative annotation methods to enhance accuracy.

Focus group discussions delved deeper into the reasons and challenges behind the relatively lower emphasis on model evaluation metrics [FG-9a]. Participant R1 noted, “A sentiment analysis algorithm might see a word like “hate” to be negative, but in the context of a statement like ‘I hate how this show ended on a cliffhanger’ it is not negative”. Furthermore, the lack of industry-wide standards in defining data quality and the success of NLP systems exacerbate this challenge (DA1, DA2). DE1 also emphasized the difficulty of designing metrics for voice-activated systems, which must account for accent and dialect diversity.

The survey also explored the perceived responsibility of data science teams in developing systems that accurately represent users’ experiences [SQ-9b]. Respondents rated the importance on a 5-point Likert scale (0-5). Less than 5% of NLP practitioners rated it slightly important, while a substantial majority (76.6%) rated it either important or very important.

### **Opportunities for Public Participation in the NLP Dataset Production**

Our findings offer insight into the role in which the public can play when involved within various stages of the NLP data production pipeline. By examining survey responses and focus group discussions, we specifically focus on the implications for broadened public participation in enhancing the integrity of NLP datasets and system outcomes.

**Community Involvement in Data Production** Out of 47 respondents, 28 respondents (60%) considered involving diverse community members with varying demographic backgrounds (gender, age, race/ethnicity, geo-location, etc.) in the NLP dataset production pipeline [SQ-10a], while 19 respondents did not take it into account. Out of those 28 respondents, 18 detailed specific stages and methods to involve diverse communities. Key areas identified include: 1) Early stage involvement (e.g., planning, project/study design), 2) Full-scale inclusion through the entire data production pipeline, 3) Specific stages for community input (e.g., collection, model training and evaluation, deployment and feedback, and 4) innovative equity-centered practices (e.g., community advisory board). Several respondents also highlighted challenges or the absence of diverse group involvement, often citing practical constraints related to their roles

or existing project structures (e.g., “Not so far in my work, since data is usually already collected.”).

**Benefits of Diverse Community Involvement** Focus group participants highlighted benefits of involving diverse external communities in NLP work [FG-10a]. DEI noted that more diverse internal and external perspectives are often better equipped to address unique challenges, saying, “To have the most well rounded product you need a well rounded team,” while R1 emphasized co-annotation with end-users and other stakeholders, where community members collaborate with NLP practitioners to label, correct, and validate data, thereby enhancing data quality, increasing representation, and building users trust & transparency. Participants noted the learning opportunities for NLP practitioners through community participatory innovation, including gaining deeper insights into different language use across communities; heightening awareness of ethical considerations (e.g., bias mitigation, data privacy); and learning effective ways to incorporate diverse feedback to ensure project success and relevance. By directly engaging diverse communities, as indicated in [SQ-7b] in NLP dataset production, practitioners and technology development organizations could ensure that the data collected is as representative and unbiased as possible. In fact, 96% of survey respondents thought that the demographic background (gender, age, race, nationality, etc.) of data practitioners can affect the outcome of NLP datasets [SQ-10b]. This consensus underscores the pivotal role that diverse perspectives play in shaping data quality and model performance.

**Challenges in Engaging Community Data Agents** The survey evaluated the feasibility of integrating external communities into their NLP data annotation workflows on a 5-point Likert scale (0-5) [SQ-10c]. 16.02% of NLP practitioners rated less feasible or not feasible, which was potentially caused by practical challenges or perceived structural, regulatory, or technical barriers. However, a significant majority (83.98%) indicates some degree of feasibility (e.g., fairly feasible, feasible, very feasible) to involve external communities in the annotation process.

Respondents identified several potential challenges faced by external community members if involved in NLP data annotation [SQ-10d]: Access to data sources (44.68%); Length of time to productively annotate (44.68%); Understanding metrics of success (42.55%); Understanding the application or technology in which data is used for (40.43%); Annotation quality (23.40%). 7 respondents (14.89%) indicated other challenges, including recruitment difficulties, logistical considerations, and integration with existing teams. Additional insights from open responses further underscored the complexity of these challenges and the practicalities of engaging external annotators or collaborators, such as model over-tuning, cultural and linguistic barriers, privacy and security concerns, and undervaluing external community experts knowledge. The focus group elaborated on some underlying perceptions and structural issues in NLP data annotation that can impact the involvement and valuation of external data agents’ skills and knowledge. Participant D1 stated “annotation and data work is considered and sold as

low skill work gig/contract work that anyone could do.”

## Perspectives on AI Policy and Regulation

To gain perspective on the Blueprint for an AI Bill of Rights (U.S 2022), which is aimed at making automated systems work for the American people, we gauged practitioners’ familiarity with the Algorithmic Discrimination Protections (ADP) [SQ-11a]. We found that 59.57% of surveyed NLP practitioners reported low familiarity (e.g., slightly familiar, not familiar at all) and only a combined 14.89% being very or extremely familiar, indicating a substantial knowledge gap regarding these protections and a need for broader educational initiatives within the industry.

Despite this gap, a substantial (93.62%) of respondents rated the necessity of an AI Bill of Rights as at least moderately important, with a majority (62.71%) leaning towards important or very important for preventing algorithmic discrimination.

**Importance of the AI Bill of Rights and Regulation** Focus group participants underscored the importance of an AI Bill of Rights from both a technologist’s perspective and the broader public’s viewpoint. DEI points out the practical and immediate consequences of AI advancements (e.g., creations of deepfakes, privacy invasions). They emphasize the need for regulation to keep people safe and mitigate harms already evident as AI becomes more pervasive and capable. DA2 advocates for the Bill as crucial for protecting consumer privacy and ensuring that data usage respects individual rights “regardless of their class of life.” Furthermore, DA2’s call for equitable protections under AI regulations reinforces the role of AI policy and governance in promoting fairness and preventing discrimination.

**Necessary Algorithmic Discrimination Protections (ADP)** 27 survey respondents provided additional insights on the AI Bill of Rights Algorithmic Discrimination Protections [SQ-11e]. They highlighted their practical applicability, importance, and areas for improvement in NLP data-centric systems: 1) *Significance and Support for ADP*: Respondents’ sentiments recognize ADP as fundamental in maintaining fairness and ethics in NLP practices (e.g., “ADP are very important...ensure that NLP systems are used in a way that is fair and equitable for all users”). Some responses emphasize the evolving nature of ADP and suggest that while current measures are a good start, ongoing refinement is necessary. 2) *Practical Applicability and Implementation Concerns*: Multiple respondents mention that implementing ADP policies is both crucial and challenging, advocating for robust, actionable strategies such as “continuous auditing, transparency in model decisions, and proactive fairness measures in data collection” to preemptively mitigate issues. 3) *Skepticism and Concerns of Exploitative Practices*: Respondents highlighted a) trust issues on the effectiveness of government-led initiatives and b) inconsistent international protection standards.

Additionally, there were concerns that regulations might potentially drive exploitative practices or hinder open-source innovation: “it’s just going to force the people who don’t agree with these rights to go underground, which leads

to black market involvement.” These insights indicate that while there is broad support for the ADP, there is a clear need for more concrete guidelines and effective implementation methods in the broader global AI landscape.

These findings underscore that while NLP practitioners are often expected to uphold fairness and equity standards, they operate within environments that frequently lack the organizational, technical, and governance support necessary to do so. The challenges they articulated—ranging from dataset limitations to institutional constraints—highlight a gap between ethical aspirations and practical realities. These tensions invite deeper examination of how responsibility is distributed across the AI development pipeline and raise important questions about the structural conditions required to meaningfully advance equitable NLP systems. We explore these implications further in the following discussion.

## Discussion

Our study examines how NLP practitioners advance NLP data equity by reflecting on practices that produce representative, contextually accurate, and justice-oriented datasets to address technical mis-alignments and social inequities among diverse users. This study offers three key contributions. First, it is among the few to center the voice of NLP practitioners – capturing the lived realities of those embedded in the data production pipeline and their perceptions of equity, bias, and governance. Second, it connects these practitioner experiences directly to a multi-scalar governance framing, linking day-to-day decisions to institutional, national, and international policy contexts. Third, it advances participatory recommendations that bridge academic, industry, and civic domains, offering actionable pathways for enhancing practitioner agency and accountability.

### Expanding Metrics and Accountable Approaches to Fostering Diversity in Data Collection

Developing robust strategies and metrics for measuring dataset diversity is essential. Further research emphasis could be placed on techniques for mitigating bias introduced during data collection. For instance, Cunningham et al. (2024) has proposed creating human-centered evaluation methods for better assessing AI’s sociocultural consequences in language technologies on underrepresented language communities. How are NLP professionals currently measuring the diversity of their datasets? Are there limitations to these metrics? Can new metrics be developed to provide a more comprehensive picture of dataset diversity? Simulations and user experience case studies can be used to evaluate the effectiveness of different bias mitigation strategies in various NLP data collection scenarios.

Subsequently, approaches to addressing the lack of diverse data sources (RQ1) can be extended to practices that enable transparency and equity in developing more representative NLP systems. Considering how might user-controlled mechanisms to collecting data improve the user experiences and performance of NLP systems. For instance, community-based governance and participatory data works can be integrated to broaden perspectives integral for humanis-

tic data, especially from people of disproportionately affected communities. Community-based data governance and sovereignty empowers communities to actively participate in the data collection process and ensure that their experiences and perspectives are accurately represented (McNealy 2020; Ada Lovelace Institute 2021; Hovmand and Hovmand 2014). Furthermore this enables specific communities-of-interest to reclaim authority and control over their data, ensuring self-determination in how their information is gathered, owned, stored, and utilized. This concept has proven to be especially crucial for First Nations or Indigenous peoples (Lilley et al. 2024; West, Hudson, and Kukutai 2020). Technical and policy frameworks for data co-operatives or data trusts (McNealy 2020; Reed and Ng 2019) offer a way for stewards of community data to collectively own and govern their data, ensuring that they receive fair equity for their contributions and that their data are used in ways that benefit them (Pentland, Lipton, and Hardjono 2021; McNealy 2020).

### Promoting Community-collaborative Computing

While participation is recognized as essential for fairness and accountability in NLP, practical mechanisms for community participation remain limited. One promising model is the use of Community Advisory Boards (CABs), which have supported equitable algorithm development with underrepresented groups (Farmer et al. 2022; Mitchell et al.). In our study, one respondent proposed a CAB to incorporate Indigenous conceptions of health in clinical NLP. This also echoes recent participatory efforts to mitigate racial disparities in speech technologies (Cunningham 2023).

Yet, in industry settings, such collaboration often depends on project leaders or even organizations who may deprioritize public engagement due to perceived effort, costs, or misconceptions about skills necessary for data work. Investing in infrastructure to upskill community-based experts is often not considered. Although companies like Google, Meta, and Microsoft have launched expert-based external community-facing advisory boards (Smith-Loud et al. 2023), these mechanisms rarely involve experts whose language communities have been disproportionately impacted by NLP tools among other ML technologies.

In contrast, academic institutions face fewer barriers to sustained community collaboration. Initiatives like the Black Community Research Council at the University of Washington demonstrate how institutional research can meaningfully partner with marginalized communities to shape technology equity and design (Cunningham 2024).

### Equipping Practitioners for Equity-Centered NLP Work

Despite interest in fairness, many NLP practitioners lack experience in core equity practices. Only 23.4% of survey respondents reported any prior engagement with dataset bias mitigation, revealing substantial gaps in both training and institutional support. This gap reflects both workplace dynamics and the absence of critical education in mainstream technical training.

To address this, academic and professional training programs must more explicitly integrate social and ethical dimensions of computing. For example, “CompSci 240: Race, Gender, Class, & Computing” at Duke University trains students to interrogate how identity, power, and systems of inequality shape the development and deployment of AI (University 2021). Industry initiatives, such as Google’s Product Inclusion Summit, represent a promising model for skills-building in inclusive design.

However, training alone is not sufficient. Many participants expressed skepticism toward checklist-style diversity efforts that emphasize representational inclusion without substantive engagement. These concerns align with critiques of “diversity-washing” (Whitney and Norman 2024; Dalal, Hall, and Johnson 2024), which caution that superficial inclusion efforts can obscure structural inequities or bypass community consent. Practitioners emphasized that equity must extend beyond demographic representation to include epistemic justice, participatory design, and reciprocal engagement with affected communities. Taken together, these insights underscore the importance of coupling practitioner capacity-building with structural reforms. Equity in NLP is not only a technical or moral challenge, but an institutional one—requiring both governance support and practitioner empowerment.

### **Resituating Practitioner Agency in Multi-Scalar Governance**

Our findings reveal that while NLP practitioners play a crucial role in shaping equitable AI outcomes, their efforts are often constrained by organizational priorities, short-term incentives, and a lack of structural support or accountability mechanisms. These tensions highlight the need to move beyond framing NLP data equity as an individual responsibility and instead embed it within broader AI governance ecosystems.

Drawing on Drahos’ Regulatory Theory (Drahos 2017), we argue for multi-scalar approaches to AI governance—those that span policy regimes, institutional norms, and technical infrastructures. Survey responses, including from items [SQ-11a] and [SQ-11d], showed limited familiarity with algorithmic discrimination protections but a strong consensus that public policymakers, rather than private companies, should lead in ensuring fairness and safety (Lander and Nelson 2021; Araya and Nieto-Gómez 2020). Calls for international alignment further resonate with UNESCO’s recommendation for globally harmonized AI ethics.

Public-private collaboration has a role to play, but only if accountability and public interest are centered (Crawford 2021). Civic organizations (e.g., ACLU) can help bridge this gap by facilitating community engagement beyond corporate priorities (Costanza-Chock 2018; brown 2017). Governance frameworks must explicitly support practitioner agency—not simply by prescribing compliance, but by providing the institutional scaffolding needed to make equity work viable and meaningful within teams and workflows.

### **Study Limitations**

While this study employed a mixed-methods approach to provide both breadth and depth of insight, several limitations should be noted. First, the survey and focus groups were conducted in late 2023 and early 2024, ensuring relevance to shifts in NLP data practices following the proliferation of large language models (LLMs). Second, our participant sample over-represents Black practitioners, reflecting targeted outreach within AI fairness and equity networks. While this offers critical insights from often underrepresented voices, the findings are qualitative and exploratory rather than statistically generalizable. We acknowledge the potential for response bias, as participants may have been predisposed to engage with NLP data equity issues. Third, while our study focused on practitioners’ individual perspectives, we recognize that organizational structures and workplace policies play a significant role in shaping NLP data equity practices. A deeper analysis of these dynamics was beyond the study’s scope expanded in the discussion with reference to prior research. Finally, the study is primarily U.S. centric and findings should not be assumed to reflect global NLP practitioner contexts.

### **Conclusion**

This study examined how NLP practitioners navigate NLP data equity in the design and development of language technologies. Through a mixed-methods study of U.S.-based practitioners, we surfaced how individuals responsible for dataset development and annotation engage with ethical frameworks, fairness goals, and systemic constraints. Our findings reveal persistent tensions between commercial imperatives and equity commitments, as well as practitioner skepticism toward superficial inclusion efforts.

We argue that equitable NLP systems cannot be achieved solely through technical interventions or individual initiative. Instead, they require structural realignments—governance frameworks that support practitioner agency, participatory infrastructure for community engagement, and accountability mechanisms embedded across the AI development lifecycle. Drawing on critiques of diversity-washing and multi-scalar governance theory, we advocate for a shift from representational metrics to relational approaches grounded in consent, context, and community-defined equity. By foregrounding practitioner perspectives, this work contributes to ongoing efforts to rethink responsibility and accountability in AI development—not as abstract principles, but as situated practices shaped by institutional conditions and socio-technical dynamics.

### **Acknowledgements**

This work was supported by the National Science Foundation Graduate Research Fellowship, the Google Award for Inclusive Research, and the Social Science Research Council Just Tech Fellowship. The first author gratefully acknowledges the commitment of these organizations to advancing research at the intersection of technology equity and social justice.

## References

2022. Blueprint for an AI Bill of Rights (OSTP).
- Act, E. A. I. 2024. The eu artificial intelligence act. *European Union*.
- Ada Lovelace Institute. 2021. Participatory Data Stewardship.
- Araya, D.; and Nieto-Gómez, R. 2020. Renewing multilateral governance in the age of AI. *MODERN CONFLICT AND ARTIFICIAL INTELLIGENCE*, 6.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2020. Fairness and machine learning. *Recommender systems handbook*, 1: 453–459.
- Bender, E. M.; and Friedman, B. 2018. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics*, 6: 587–604.
- Braun, V.; and Clarke, V. 2006-01. Using thematic analysis in psychology. 3(2): 77–101.
- brown, a. m. 2017. *Emergent Strategy: Shaping Change, Changing Worlds*. AK Press. ISBN 978-1-84935-261-1.
- Costanza-Chock, S. 2018. Design Justice: Towards an Intersectional Feminist Framework for Design Theory and Practice. (ID 3189696).
- Crawford, K. 2021. The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence.
- Cunningham, J.; Blodgett, S. L.; Madaio, M.; Iii, H. D.; Harrington, C.; and Wallach, H. 2024. Understanding the impacts of language technologies’ performance disparities on African American language speakers. In *Findings of the Association for Computational Linguistics ACL 2024*, 12826–12833.
- Cunningham, J. L. 2023. Collaboratively Mitigating Racial Disparities in Automated Speech Recognition and Language Technologies with African American English Speakers: Community-Collaborative and Equity-Centered Approaches Toward Designing Inclusive Natural Language Systems. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI EA ’23, 1–5. Association for Computing Machinery. ISBN 978-1-4503-9422-2.
- Cunningham, J. L. 2024. *Methods of Designing Justice-oriented Interactive AI Systems*. Ph.D. thesis, University of Washington.
- Dalal, S.; Hall, S. M.; and Johnson, N. 2024. Provocation: Who benefits from” inclusion” in Generative AI? *arXiv preprint arXiv:2411.09102*.
- Deas, N.; Grieser, J.; Kleiner, S.; Patton, D.; Turcan, E.; and McKeown, K. 2023. Evaluation of African American Language Bias in Natural Language Generation. 2305.14291 [cs].
- Deng, W. H.; Yildirim, N.; Chang, M.; Eslami, M.; Holstein, K.; and Madaio, M. 2023. Investigating Practices and Opportunities for Cross-functional Collaboration around AI Fairness in Industry Practice. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’23*, 705–716. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.
- Deng, Y.; Qiu, X.; Tan, X.; Pan, J.; Jue, C.; Fang, Z.; Xu, Y.; Chu, W.; and Qi, Y. 2024. Promoting Equality in Large Language Models: Identifying and Mitigating the Implicit Bias based on Bayesian Theory. *arXiv preprint arXiv:2408.10608*.
- Denton, R.; Díaz, M.; Kivlichan, I.; Prabhakaran, V.; and Rosen, R. 2021. Whose Ground Truth? Accounting for Individual and Collective Identities Underlying Dataset Annotation. 2112.04554 [cs].
- Dourish, P. 2017. *The Stuff of Bits: An Essay on the Materialities of Information*. MIT Press.
- Drahos, P. 2017. *Regulatory theory: Foundations and applications*. ANU Press.
- Farmer, N.; Osei Baah, F.; Williams, F.; Ortiz-Chapparo, E.; Mitchell, V. M.; Jackson, L.; Collins, B.; Graham, L.; Wallen, G. R.; Powell-Wiley, T. M.; and Johnson, A. 2022. Use of a community advisory board to build equitable algorithms for participation in clinical trials: a protocol paper for HoPeNET. 29(1): e100453.
- Floridi, L.; Cows, J.; Beltrametti, M.; Chatila, R.; Chazerand, P.; Dignum, V.; Luetge, C.; Madelin, R.; Pagallo, U.; Rossi, F.; et al. 2018. AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds and machines*, 28: 689–707.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021-12. Datasheets for datasets. 64(12): 86–92.
- Goyal, N.; Kivlichan, I.; Rosen, R.; and Vasserman, L. 2022. Is Your Toxicity My Toxicity? Exploring the Impact of Rater Identity on Toxicity Annotation. 2205.00501 [cs].
- Hanna, A.; Denton, E.; Smart, A.; and Smith-Loud, J. 2020. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT\* ’20*, 501–512. Association for Computing Machinery. ISBN 978-1-4503-6936-7.
- Hill, J.; Ogle, K.; Gottlieb, M.; Santen, S. A.; and Artino Jr, A. R. 2022. Educator’s blueprint: A how-to guide for collecting validity evidence in survey-based research. 6(6): e10835. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/aet2.10835>.
- Holstein, K.; Wortman Vaughan, J.; Daumé III, H.; Dudik, M.; and Wallach, H. 2019. Improving fairness in machine learning systems: What do industry practitioners need? In *Proceedings of the 2019 CHI conference on human factors in computing systems*, 1–16.
- Hovmand, P. S.; and Hovmand, P. S. 2014. *Group model building and community-based system dynamics process*. Springer.
- Hutchinson, B.; Prabhakaran, V.; Denton, E.; Webster, K.; Zhong, Y.; and Denuyl, S. 2020. Social Biases in NLP Models as Barriers for Persons with Disabilities. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5491–5501. Online: Association for Computational Linguistics.

- Irani, L. 2016. Hackathons and the making of entrepreneurial citizenship. *Science, Technology, & Human Values*, 41(5): 705–740.
- Jo, E. S.; and Gebru, T. 2020. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT\* '20, 306–316. Association for Computing Machinery. ISBN 978-1-4503-6936-7.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- Joshi, P.; Santy, S.; Budhiraja, A.; Bali, K.; and Choudhury, M. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. *arXiv preprint arXiv:2004.09095*.
- Kamikubo, R.; Wang, L.; Marte, C.; Mahmood, A.; and Kacorri, H. 2022. Data Representativeness in Accessibility Datasets: A Meta-Analysis. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '22, 1–15. Association for Computing Machinery. ISBN 978-1-4503-9258-7.
- Kapania, S.; Taylor, A. S.; and Wang, D. 2023. A hunt for the Snark: Annotator Diversity in Data Practices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, 1–15. Association for Computing Machinery. ISBN 978-1-4503-9421-5.
- Koenecke, A.; Nam, A.; Lake, E.; Nudell, J.; Quartey, M.; Mengesha, Z.; Troups, C.; Rickford, J. R.; Jurafsky, D.; and Goel, S. 2020. Racial disparities in automated speech recognition. 117(14): 7684–7689.
- Lander, E.; and Nelson, A. 2021. Americans need a Bill of rights for an AI-powered world. *Wired*.
- Lazar, S.; and Nelson, A. 2023. AI safety on whose terms?
- Lilley, S.; Oliver, G.; Cranefield, J.; and Lewellen, M. 2024. Māori data sovereignty: contributions to data cultures in the government sector in New Zealand. 0(0): 1–16. Publisher: Routledge. eprint: <https://doi.org/10.1080/1369118X.2024.2302987>.
- Lincoln, Y. S. 1985. *Naturalistic inquiry*, volume 75. sage.
- Marshall, G.; and Jonker, L. 2011. An introduction to inferential statistics: A review and practical guide. *Radiography*, 17(1): e1–e6.
- Mathur, R.; Rentsch, C. T.; Venkataraman, K.; Fatumo, S.; Jobe, M.; Angkurawaranon, C.; Ong, S. E.; Wong, A. Y.; and Siddiqui, M. K. 2022. How do we collect good-quality data on race and ethnicity and address the trust gap? *The Lancet*, 400(10368): 2028–2030.
- McConvey, K.; and Guha, S. 2024. "This is not a data problem": Algorithms and Power in Public Higher Education in Canada. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–14.
- McNealy, J. 2020. An Ecological Approach to Data Governance. *Bulletin of the American Physical Society*, 65.
- Miceli, M.; Schuessler, M.; and Yang, T. 2020. Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision. 4: 115:1–115:25.
- Miceli, M.; Yang, T.; Alvarado Garcia, A.; Posada, J.; Wang, S. M.; Pohl, M.; and Hanna, A. 2022. Documenting Data Production Processes: A Participatory Approach for Data Work. 6: 510:1–510:34.
- Mitchell, J.; Perry, T.; Rorai, V.; Ilardo, J.; Lichtenberg, P. A.; and Jackson, J. S. 2022. Building and Sustaining a Community Advisory Board of African American Older Adults as the Foundation for Volunteer Research Recruitment and Retention in Health Sciences. 30: 755–764.
- Mothilal, R. K.; Guha, S.; and Ahmed, S. I. 2024. Towards a Non-Ideal Methodological Framework for Responsible ML. *arXiv preprint arXiv:2401.11131*.
- Muller, M.; Lange, I.; Wang, D.; Piorkowski, D.; Tsay, J.; Liao, Q. V.; Dugan, C.; and Erickson, T. 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*, 1–15.
- Muller, M. J.; and Kuhn, S. 1993. Participatory design. *Communications of the ACM*, 36(6): 24–28.
- Olson, J. S.; and Kellogg, W. A. 2014. *Ways of Knowing in HCI*, volume 2. Springer.
- Pang, R. Y.; Cenatempo, J.; Graham, F.; Kuehn, B.; Whisenant, M.; Botchway, P.; Stone Perez, K.; and Koenecke, A. 2023. Auditing Cross-Cultural Consistency of Human-Annotated Labels for Recommendation Systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 1531–1552. Association for Computing Machinery. ISBN 9798400701924.
- Paullada, A.; Raji, I. D.; Bender, E. M.; Denton, E.; and Hanna, A. 2021. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11).
- Pentland, A.; Lipton, A.; and Hardjono, T. 2021. *Building the new economy: Data as capital*. MIT Press.
- Raji, I. D.; Bender, E. M.; Paullada, A.; Denton, E.; and Hanna, A. 2021. AI and the Everything in the Whole Wide World Benchmark. 2111.15366 [cs].
- Reed, C.; and Ng, I. 2019. Data trusts as an AI governance mechanism. Available at SSRN 3334527.
- Rothschild, A.; Meng, A.; DiSalvo, C.; Johnson, B.; Shapiro, B. R.; and DiSalvo, B. 2022. Interrogating data work as a community of practice. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–28.
- Sambasivan, N.; Arnesen, E.; Hutchinson, B.; Doshi, T.; and Prabhakaran, V. 2021. Re-imagining Algorithmic Fairness in India and Beyond. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, 315–328. Association for Computing Machinery. ISBN 978-1-4503-8309-7.
- Sap, M.; Card, D.; Gabriel, S.; Choi, Y.; and Smith, N. A. 2019a. The Risk of Racial Bias in Hate Speech Detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1668–1678. Association for Computational Linguistics.
- Sap, M.; Card, D.; Gabriel, S.; Choi, Y.; and Smith, N. A. 2019b. The risk of racial bias in hate speech detection. In

*Proceedings of the 57th annual meeting of the association for computational linguistics*, 1668–1678.

Sloan, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2020. Participation is not a design fix for machine learning (pp. 1–7). In *Proceedings of the International Conference on Machine Learning, Vienna, Austria*.

Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2020. Participation is not a Design Fix for Machine Learning. 7.

Smith-Loud, J.; Smart, A.; Neal, D.; Ebinama, A.; Corbett, E.; Nicholas, P.; Rashid, Q.; Peckham, A.; Murphy-Gray, S.; Morris, N.; et al. 2023. The Equitable AI Research Roundtable (EARR): Towards Community-Based Decision Making in Responsible AI Development. *arXiv preprint arXiv:2303.08177*.

Suchman, L. 2002. Located accountabilities in technology production. 14: 16.

Tatman, R. 2017. Gender and Dialect Bias in YouTube’s Automatic Captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 53–59. Association for Computational Linguistics.

Thakkar, D.; Ismail, A.; Kumar, P.; Hanna, A.; Sambasivan, N.; and Kumar, N. 2022. When is machine learning data good?: Valuing in public health datafication. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–16.

University, D. 2021. CompSci 240: Race, Gender, Class, & Computing. Course offered at Duke University.

West, K.; Hudson, M.; and Kukutai, T. 2020. Data Ethics and Data Governance from A Māori World View. In George, L.; Tauri, J.; and Te Ata o Tu MacDonald, L., eds., *Indigenous Research Ethics: Claiming Research Sovereignty Beyond Deficit and the Colonial Legacy*, volume 6 of *Advances in Research Ethics and Integrity*, 67–81. Emerald Publishing Limited. ISBN 978-1-78769-390-6 978-1-78769-389-0.

Whitney, C. D.; and Norman, J. 2024. Real Risks of Fake Data: Synthetic Data, Diversity-Washing and Consent Circumvention. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’24*, 1733–1744. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704505.

Widder, D. G.; and Nafus, D. 2023. Dislocated accountabilities in the “AI supply chain”: Modularity and developers’ notions of responsibility. *Big Data & Society*, 10(1): 20539517231177620.

Ziman, K.; Heusser, A. C.; Fitzpatrick, P. C.; Field, C. E.; and Manning, J. R. 2018. Is automatic speech-to-text transcription ready for use in psychological experiments? *Behavior research methods*, 50: 2597–2605.