Products of Positionality: How Tech Workers Shape Identity Concepts in Computer Vision

Morgan Klaus Scheuerman
Morgan.Scheuerman@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Jed R. Brubaker
Jed.Brubaker@colorado.edu
University of Colorado Boulder
Boulder, Colorado, USA

Abstract

There has been a great deal of scholarly attention on issues of identity-related bias in machine learning. Much of this attention has focused on data and data workers, workers who do annotation tasks. Yet tech workers—like engineers, data scientists, and researchers—introduce their own "biases" when defining "identity" concepts. More specifically, they instill their own positionalities, the way they understand and are shaped by the world around them. Through interviews with industry tech workers who focus on computer vision, we show how workers embed their own positional perspectives into products and how positional gaps can lead to unforeseen and undesirable outcomes. We discuss how worker positionality is mutually shaped by the contexts in which they are embedded. We provide implications for researchers and practitioners to engage with the positionalities of tech workers, as well as those in contexts outside of development that influence tech workers.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → **Artificial intelligence**; • **Social and professional topics** → **User characteristics**; **Computer supported cooperative work**.

Keywords

Tech work, identity, positionality, computer vision, work studies, machine learning

ACM Reference Format:
Morgan Klaus Scheuerman and Jed R. Brubaker. 2024. Products of Positionality: How Tech Workers Shape Identity Concepts in Computer Vision. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3613904.3641890

1 Introduction

Artificial intelligence (AI)—or machine learning—has already become deeply ingrained in everyday life. Computer vision, a specific domain of machine learning focused on visual pattern recognition, has, for better or worse, permeated social media (e.g., [2, 22]), art (e.g., [41, 74]), hiring ([37, 60]), health (e.g., [62, 70]), advertising (e.g., [35]) and more. As a technology heavily reliant on making classifications about humans and human culture, computer vision has already been heavily critiqued for its propensity for identity-related bias (e.g., [4, 72]) and harmful representations, such as propagating stereotypes or reifying negative belief systems (e.g., [7, 65]).

Issues surrounding identity concepts in computer vision come down to a perspective on technology design common in human-computer interaction: computers are designed by people. Given computer vision is, then, designed by people, identity categories are not simply neutral and bias is not simply a mistake, but each is the result of the intentional decisions made by human actors. Increasingly, HCI scholars are exploring the way human actors influence the outcomes of computer vision artifacts (e.g., [20, 27, 45]). There has been a particular focus on how data workers, the (often contingent and invisible) workforce behind collecting and annotating data for computer vision (e.g., [18, 44, 53]). Such scholarship highlights opportunities to better understand not only how computer vision is shaped by people, but how people's individual and subjective perspectives influence their decisions. Implicit in this body of work is the acknowledgment that people's positionalities—the identities they occupy in the world and how those identities shape their perspectives—influence how they approach designing identity categories.

In this work, we explore how industry tech workers situated across a variety of roles, from engineering to research, are responsible for the design of enterprise-level computer vision systems. More specifically, we investigate how workers' positionalities—including the industrial contexts they are situated in, their own values, experiences, and perspectives, and their negotiations with their colleagues—influence the design of identity concepts in computer vision technologies. We address answers to the following research questions:

(1) How do the positionalities of tech workers impact the development of identity in computer vision products?
(2) How do tech workers negotiate their own positionalities with the positionalities of their colleagues and within the context of their organizations?
(3) What failures occur when tech workers fail to account for other positional experiences?

To answer these questions, we conducted semi-structured interviews with twenty-four industry practitioners who work on computer vision products. Participants worked at companies ranging from small startups to big tech; they worked in a variety of roles, such as data science, engineering, research, business, and project management. Further, they worked on various types of computer vision products, from video-based interviewing to facial

CHI '24, May 11–16, 2024, Honolulu, HI, USA

ACM ISBN 979-8-4007-0330-0/24/05.
https://doi.org/10.1145/3613904.3641890

demographics to gesture recognition. Interviews were designed to elicit descriptions of participants' company environments, their relationships with their fellow workers, and their personal experiences and values.

Findings showcase how the positionalities that workers inhabit influence the way that computer vision artifacts are designed. Workers seek to impact product design given their own positional perspectives about identity, while also being constrained by their fellow workers' differing perspectives and broader company-level contexts like regulation and company vision. Further, issues of representation and experience arise when workers fail to account for different positional perspectives than their own.

We discuss how worker positionalities are relational, rather than individualistic; they operate within larger contexts in which workers are embedded and their relationships with other actors within those contexts. We conclude with implications for attending to positionality in tech work, at a higher-level contextual level and at a lower-level actor level.

2 Related Work

In this paper, we focus on computer vision products developed and deployed in industrial contexts. Thus, in this section, we first describe prior work which has focused on how practitioners approach implementing machine learning in industry, including approaches to implementing ethical and fair machine learning. Yet, as we will demonstrate in our Findings, a crucial aspect for how identity characteristics becomes embedded into computer vision is through the positionalities of those developing it. Thus, the second area of related work we discuss is focused on defining positionality and discussing how positional values are embedded into machine learning approaches, broadly.

2.1 Machine Learning in Industrial Contexts

Computer vision is a subfield of machine learning. Machine learning, often referred to as AI, is a branch of computer science focused on using data to "teach" pattern modeling to computer systems. Most major tech companies have dedicated cloud-based computer vision platforms for purchase (e.g., Google, Microsoft, IBM). Certainly, every major tech company has dedicated resources to growing machine learning, prompting an AI arms race between companies [71] and even between countries [50].

Given the power industry has over the AI landscape [32], it is unsurprising that scholars have also focused on the implications of corporate machine learning. Corporate models are not only more powerful than research models, given the economic power of big tech companies, they are also deployed in real-world scenarios (e.g., [14, 29, 61, 66]). Through examinations of these models, many researchers have discovered troubling outcomes (e.g., poor gender classification performance for people of color [9] and transgender and non-binary individuals; biases against women even when gender is not explicitly encoded into model outputs [11]). Many other machine learning biases have also been discovered by users themselves, such as the notorious example of Google Photos labeling Black faces as "gorillas" [5]. Notably, these biases and negative experiences are acutely tied to issues of identity.

Building on the rich history of research on corporate settings in computing (e.g., [34, 39, 73]), many scholars have begun to examine the practices of machine learning practitioners in corporate settings to better understand how issues of bias occur and how practitioners attempt to mitigate them. For example, Holstein et al. identified the technical and organizational barriers preventing industry practitioners from effectively improving machine learning fairness; industry practitioners trying to implement fairness often had to engage in unrecognized additional labor which was not incentivized by their organizations or well received by colleagues [31]. Rakova et al. similarly identified constraints, like a lack of accountability and poor incentive structures which result in reactive, rather than proactive, approaches to enacting fairness initiatives; they use their insights into these constraints to offer aspirational future processes to better enable effective initiatives [49]. Working directly with industry practitioners is a fruitful method for uncovering their practices and identifying unknown challenges.

However, perhaps due to issues of access to industry settings, particularly large company settings, research on fairness practices in companies is still sparse. The study at hand buildings on the growing body of work focused on industrial-level machine learning practices. However, rather than focusing solely on the concept of preventing bias or improving fairness in developed products, we focus on how industry practitioners conceptualize identity characteristics for computer vision throughout the development life cycle.

2.2 Positional Values in Machine Learning

Positionality describes how the "position" a person occupies—including identities like gender, race, nationality, sexuality, and more—shapes the way they view and interact with the world around them [16]. The positions that each person occupies are often evolving and mutually reinforced through that individual's relationship with others and themselves [12, 15]. In computing research, positionality is largely attended to through reflexivity—a methodological process of self-reflection on how researcher and research mutually construct one another (e.g., [38]). Understanding how research subjects express their positionalities is often implicit. Increasingly, scholars like Cambo and Gergle are contributing approaches to actively considering reflexivity in occupational practices like data science [10].

In the realm of machine learning, many scholars have examined how human values have become embedded in and shape data categories. For example, Scheuerman et al. document the types of values driving the creation of datasets for computer vision [54]. Metcalf et al. found that many practitioners at the forefront of ethics in industry tolerated corporate values, like market fundamentalism and technological solutionism, for the sake of minimal ethical impact [43]. Numerous scholars have also critiqued the underlying values governing identity characteristics in machine learning. Much like Suchman argues in [63], language categories are explicitly designed to maintain current status quo social orders, not challenge them or provide space for social action. Scheuerman et al. explicate how the construction of race and gender categories for computer vision datasets reflects normative beliefs that identity categories are "insignificant, indisputable, and apolitical" [57]. Hanna et al. similarly argue that the treatment of race as categorical erases the reality that race is socially constructed and meaningful [28]. Such expressions of values are reflective of the larger positional standpoints occupied by the designers of machine learning artifacts, like datasets.

As Davis writes, a "dataset is a worldview" [19]. The worldview that a dataset holds is the result of the worldviews of the humans working to produce it. The presence of identity characteristics in computer vision—like gender classifications, labels applied to images, and inferred racial demographics from textual data—indicates that, at some point, identity characteristics are developed for machine learning. Thus, beyond building on studies of industrial practice, we also contribute to burgeoning research on structuring and defining identity categories for machine learning. Specifically, we extend this area of inquiry by examining how industry practitioners embed their own perspectives about identity in the process of defining it.

3 Methods

In this section, we describe the methods we used to understand how the positionalities of tech workers influence identity in computer vision. We first discuss our approach to using semi-structured interviews. We then describe our participants, as well as the difficulties we had recruiting them. Next, we describe our approach to analyzing our interview data. Finally, we describe our own positionalities as researchers and how they influenced our study.

3.1 Interviews

To understand the role positionality plays in implementing identity in human-centric computer vision products, we conducted semi-structured interviews with full-time employees at technology companies. We chose to use interviews to gather rich descriptions of the perspectives, beliefs, and experiences of tech workers, as it allowed opportunities to ask clarifying questions, seek out specific examples, and tailor interview questions in real-time to contextual responses [58]. Participants worked on diverse projects, and many were in completely different companies. As such, the interview protocol was designed to be flexible towards participants' individual roles and products.

During interviews, we elicited descriptions of participants' roles and the products they worked on, how identity characteristics were embedded into those products, how decisions shaped those characteristics, and the constraints and difficulties they faced in implementing those characteristics. Our conceptualization of identity meant anything to do with human characteristics or culture. Though we did not prescribe any specific definition of identity during interviews, identity was largely characterized by participants as demographics (e.g., race, gender, age). Thus, much of what we discuss in our findings relates to demographic characteristics, presentations (e.g., gender presentations, skin tones), and human affect (e.g., emotion classification, sentiment analysis). Example interview questions included: Can you describe how the products you work on go about classifying people? Are there any characteristics you intentionally do not include and why? Are there any negative outcomes you are personally concerned about and why? Broader questions about roles and products were introduced first so that we could then delve more deeply into personal experiences and beliefs about identity classifications.

All interviews were conducted remotely using video conferencing software and recorded with participant consent. Interviews lasted, on average, about 73 minutes. This study was approved by the authors' institutional IRB.

3.2 Participants

We recruited 24 participants in tech industry contexts who worked on human-centric computer vision products (see Table 1). Participants worked on computer vision either as their primary responsibility or as major part of their job (e.g., some participants also worked on natural language processing tools). They also worked on characterizing aspects of human identity for various products, including: tools for explicit identity classifications (e.g., gender classification); tools that did not conduct classifications but required identity to be considered at the data level (e.g., gesture recognition models which required diverse examples of hands); and approaches to bias mitigation which required explicit classifications for evaluation (e.g., labeling race at the data level to test for biased model outputs in products).

Participants held a variety of roles at differing levels of seniority (from intern to C-level). We were interested in interviewing individuals in both "technical" roles (job roles responsible for implementation, measurement, and testing products) and "non-technical" roles (job roles focused on ideas, management, and research). While the boundaries between these two roles were fuzzy and overlapping, with many individuals engaging in both traditionally "technical" and "non-technical" responsibilities, the distinction aided us in targeted recruitment. We wanted to ensure that we did not have a balance of participants heavily skewed towards only research roles, for example. We similarly sought to interview participants at different companies, from startups to big tech companies, to get perspectives from differently-resourced employees.

To recruit participants, we largely employed an ad hoc sampling approach [64]. We chose ad hoc sampling because computer vision is a relatively narrow subfield of machine learning. Thus, identifying potential participants to recruit was difficult and response rates were low. We located computer vision companies largely through search tools on Google and LinkedIn, but identifying employees of those companies and whether they had direct engagement with computer vision products was opaque. In particular, those in more technical roles, like data scientists and software engineers, had less web presence than those in research or C-level roles. We largely contacted potential participants directly, via email or LinkedIn messaging.

Following interviews with some participants recruited via ad hoc sampling, we also recruited some participants via snowball sampling [64]. Participants put us in contact with others they felt would be a good fit for the study. The first author also maintains contact with many individuals working in industry, which aided in an ad hoc approach to direct recruitment and a snowball approach to sourcing new contacts.

All Participants				
Alias	***Role***	***Company***	***Company Size***	***Location***
Jeremy	Software Engineer	Aqueous	Large	United States
Coleman	Principal Data Scientist	Aqueous	Large	United States
Kaleigh	Program Manager	Aqueous	Large	United States
Vasuda	Project Manager	Aqueous	Large	United States
Ethan	Senior Principal Research Manager	Aqueous	Large	United States
Callia	Principal Research Manager	Aqueous	Large	United States
Jacqueline	Lead UX Researcher	Maelstrom	Large	United States
Elliot	Research Scientist	Maelstrom	Large	United States
Madison	Lead Research Scientist	Maelstrom	Large	United States
Macy	UX Researcher	Exodia	Large	United States
Beiwen	Machine Learning Research Intern	Zeta	Large	United States
Nitesh	Data Engineer	Inoculus	Medium	United States
Irina	CEO	EnVision Data	Small	Bulgaria
Zephyr	Chief Impact Officer	EnVision Data	Small	Bulgaria
Thalia	Chief Operations Manager	EnVision Data	Small	Bulgaria
Samuel	Chief Commercial Officer	EnVision Data	Small	Bulgaria
Lynn	Head of Data Operations	MultiplAI	Small	United States
Kenny	Vice President of Business Development	MultiplAI	Small	United States
Nicholas	Chief IO Psychologist	Resoom	Small	United States
Lydia	Head of Data Science	Resoom	Small	United States
Kelly	Developer Advocate	Phrenx	Small	United States
Solange	AI Product Manager	SensEyes	Small	France
Siddhartha	Computer Vision Scientist	Sybil	Small	United States
Aishwarya	Computer Vision Research Intern	Verus	Small	United States

Table 1: The table lists the 24 participants in this study. The table is first organized by company size. Small company has 500 or fewer employees. Large company has 10,000 or more employees. Medium has between 501 and 9,999 employees. It is then organized by the number of participants per company. It is lastly organized alphabetically by company alias. Participant aliases were created using the same cultural origins as the participants' real names. Company aliases were randomly generated.

3.2.1 *Participation Concerns*

Recruitment for this study was difficult; response rates from participation requests were low. Throughout the process of both recruitment and conducting interviews, we realized several aspects that made recruitment for this study difficult. Two major participant concerns arose: (1) concerns tech workers had about accidentally violating their NDAs and (2) concerns about purposeful or accidental identity leaks. Though many of our recruitment emails went unanswered, some participants who had agreed to participate backed out before the interview due to legal concerns surrounding fresh controversy at their company. A participant who had initially declined to participate, but later participated after the first author had built a relationship with her, informed us that our initial recruitment email had caused a great deal of "backchanneling" (secret conversations that did not involve the first author) about whether participation was too risky without having the research team sign an NDA. Much like similar studies involving industry stakeholders (e.g., [31, 67], several participants expressed a distrust of researchers and speaking about AI due to fears that their personal and company identities would be leaked to the press. A distrust of journalists amidst a wave of articles covering AI ethics was salient among participants. Some participants also expressed concerns that the academic community was "reactionary" towards industry. Participants' concerns highlighted our positions as not only outsiders to tech workers but as perceived threats. This shaped how we approached research with tech workers. We became more upfront about our intentions during recruitment and actively reassured them that we had no intention of publishing information to harm them or their companies. It was also yet another factor that shaped this research; participants likely held back or altered the way they spoke out of concern. In the future, we plan to contribute a more in-depth paper focused on barriers to accessing traditional tech workers focused on AI.

3.3 Analysis

This study was conducted and analyzed from a constructivist epistemology [46]. Analysis was conducted through a series of coding and memoing exercises as themes coalesced [52]. The first author conducted open coding on each interview transcript to gain more intimate familiarity with the data and observe what themes emerged among individual participants. The first author wrote detailed memos about each participant's perspective on implementing identity. Since participants did not necessarily explicitly discuss their positionalities or express reflexivity, the first author employed his own lens as a researcher to understand how participant positions influenced their work. The first author then began clustering themes from participant-level memos into larger theoretical memos about how worker positionalities inform identity implementation. While the second author did not code or memo data, the first author and the second author met to discuss themes and examples, using each other's knowledge and expertise to interpret the data and solidify themes.

3.4 Researcher Positionality

Given this study is directly focused on how participants' positionalities influence the processes and outcomes of identity characteristics in computer vision, it is imperative that we also reflect on how our positionalities mutually shaped the outcomes of this work. The trust that we were able to build through our relationships and reputation with certain technology companies awarded us opportunities to interview individuals who would otherwise be inaccessible. Both authors are United States citizens who speak English and work in a highly-regarded department at a reputable institution, which helped instill trust in our research credentials. Both authors brought positions as English-speaking academics in the United States that aided in accessing participants and resources to conduct this work. The first author's background as a lower-class, queer, non-binary first-generation college graduate and the second author's background as queer also lent a specific perspective from which to analyze and consider the implications of identity-based positionalities when implementing identity concepts in computer vision. We do note that at times we find people's perspectives differ from our own. However, differing perspectives are important to consider when portraying how an individual's positional perspectives translate into the decisions that individual makes. We encourage readers to consider the role that our positionalities as researchers shaped the outcomes of this study.

4 Findings

We present Findings on how workers approached identity concepts in computer vision and how those approaches were influenced by the positions they occupied. First, we introduce how identity is being defined in industrial contexts, including the challenges practitioners face in defining it. Of course, worker positionality did not occur in isolation, at solely the individual level, and was influenced by the contexts in which they worked and their interactions with fellow employees. Thus, we then present how the company context shapes how individual workers are able to approach identity. Next, we show how workers consider their own values and identity affinities in their approaches, and how they must often negotiate differing perspectives with their colleagues. Finally, we describe how gaps in positionality arise due to workers having their own limited viewpoints during the development process, resulting in undesirable outcomes that can become embedded into products. Throughout our Findings, we prioritize sharing participants positional perspectives to highlight how their positionality is connected to their interpretations of identity concepts.

4.1 How Identity Is Defined in Industrial Contexts

As Siddhartha said: *"Identity is very important, right?"* Before building a computer vision product, workers define what identity should look like and how it should be scoped, in terms of its categories and data representations. For example, a gender classification model often uses the categories "male" and "female" for gender, and data representations include face images annotated with those categories. We saw identity show up in computer vision products in three ways: in explicit identity classifications (e.g., a gender classification model); in implicit identity classifications (e.g., for fairness audits); and in non-human objects which were still imbued with socio-cultural meaning (e.g., immaterial concepts like "racy," which might have embedded cultural beliefs about gender, sexuality, and purity, for content moderation).

The process of defining identity varied depending on the company participants worked for and the products they were working on. As Elliot described: *"When you're trying to incorporate information into a model during training ... then you end up needing to do much more rigid things."* The *"rigid things"* Elliot is referring to are the categories workers define for computer vision products—such as gender categories for demographic classifications. Elliot expressed that in order for supervised machine learning products to work, the categories must be made into something rigid. Something like gender must be turned into discrete categories to be classified.

In practice, workers often struggled with how to best scope identity, given the vast possibilities for categorizing identity characteristics. For example, Lynn explained how *"intimidating"* it was to determine how to measure racial categories for bias testing:

> *The kind of resounding sentiment from the team was like, we have to constrain this problem because it's an impossibility criterion if we just allow ourselves to think about every single phenotype and every single appearance of human facial features … So leaving those out entirely, was just based on skin color, I guess, right? And so, you think about, like, where is there maybe a standardized thing that I can steal from and then like, well, there's the U.S. Census, which is highly problematic.*

Workers who felt identity *should be* incorporated into the systems they worked on often struggled to decide on *how* to best represent identity for core product needs, like testing, and often made decisions based on necessity even if that meant using resources that are historically flawed (e.g., the U.S. Census). Moreover, Lynn's experience demonstrates that representing identity was a given. It was not questioned whether identity attributes need to be included, because they are seen as necessary.

This necessity was similarly apparent in Lynn's approach to gender in a wedding classification model. Lynn explained that they did not include *"non-binary"* or *"androgynous"* as categories, because they felt only *"masculine"* and *"feminine"* expressions could be classified. Lynn's explanation also revealed her underlying perspective that gender is always inherent on the face in a specific binary way. Her perspective clashed with those of other participants, who viewed binary gender as not *"biological"* but social. Madison explained how her team created guidelines for annotating gender that reflected gender classification as a visual interpretation made by others. *"I've been calling [self-annotated gender] 'first-party' gender, and … 'third-party' gender, when it's about what's being perceived by someone else or by a system."*

Finally, certain representations were much more difficult to get data for. For example, Jeremy worked on gesture recognition and described how difficult it is to get data on hands which are missing digits. Certain types of identities, like disability, were often seen as untenable, because making them into *"rigid things"* was much more difficult due to the vast diversity of disabilities. Jeremy explained that accounting for the spectrum of human diversity is difficult to imagine: *"When you design the dataset, you can't anticipate every type of failure, you try to vary it up as much as possible."*

Scoping identity for a computer vision product was a complex and tangled process—influenced by everything from "common sense" to technical constraints and product needs. Identity is defined as a result of the interplay between worker positionality and product requirements. As we detail in the next section, positionality is often in tension with the overriding constraints of product needs. This tension always occurs within the broader context of development, which includes differing personal values amongst workers, economic constraints, client demands, and regulatory frameworks.

4.2 How Context Influences Worker Positionality

To better understand how workers embed tacit knowledge informed by their positionality into computer vision products, it is necessary to understand the organizational context in which workers are situated. Participants described three characteristics of the organizational context they were in: economic constraints, regulation and policy, and company values.

Economic constraints heavily influenced how workers were able to approach identity in computer vision development. Workers were enabled or constrained to approach identity in specific ways depending on the company they were working in. In particular, workers discussed how the economic power of their company shaped what the focus of their work was. Many felt that lower economic resources at their companies hindered their ability to conduct more in depth or expansive identity work. The workers discussing economic factors were those in small companies where money was a constraint, whereas those in larger companies did not discuss economic costs or disadvantages. For example, Kenny, the Vice President at a small computer vision startup, attributed the ability to engage more deeply with data collection, annotation, and research to economic power. Kenny explained that because MultiplAI is so small, they take a "business first" approach to building products: *"The reality is that we build products based out of market need, right?… And in the early stages of the company's infancy, there was a significant amount of inbound demand for gender identification."* A client's vision for what identity should be in a product—like gender classification for marketing purposes—drove how identity was then designed.

On the other hand, Kaleigh, who worked for a large company with vast economic resources, did highlight that economic resources do not rid them of all challenges:

> *It is absolutely spot on that, yes, big companies have a lot more resources. But we still need and welcome help with more practices about how to do this well, how to measure things well, how to mitigate well, how to do participatory design well, because I think a lot of that work could directly translate into product changes.*

Kaleigh described that, even with economic resources, her company and her team did not necessarily know the best way to approach identity. Such approaches were still social enterprises, with humans making informed decisions guided by best practices. Beyond needing more guidance, Kaleigh also described how, even in a large company, teams had to request budget allocations. In order to approach identity in computer vision from a responsible standpoint, Kaleigh expressed the need for more money and more team members. Yet requesting an increased budget also comes with the burden of proving the money was necessary. She had to spend the past year documenting the progress her team was making on fairness and responsible AI initiatives. She also had to showcase what made these initiatives important. Even though her company had vast amounts of economic power in comparison to MultiplAI, it did not mean focusing on identity was a company priority.

Regulation and policy also influenced the context in which workers approached identity categories. Regulation occurred in both the strict legal sense and through localized company policy. For example, Lydia, the Head of Data Science at a small company focused on providing AI interview tools, explained that the company is bound to fair employment laws in the United States. Given that the company's approach to identity is built on the concept of fair hiring practices, the categories used are derived from a legal perspective.

On the other hand, many companies also have their own internal policies about identity approaches, separate from legal requirements. For example, Madison described the ethical AI policies her company is expected to follow. She explained that people *"use [the policy for] decision making around what should and shouldn't be released."* Formal company policies can empower workers to push back on products that do not align with ethical principles, because they have been formalized in ways that implicit company values are not.

However, some workers might fall back on regulatory guidelines for demographic categories rather than expanding them. Lydia's company only attends to EEOC frameworks for gender through the lens of male and female, because the law does not currently protect non-binary people in hiring. Therefore, workers like Lydia did not attend to expanding gender categories beyond the binary. While regulation and policy might maintain certain legal and ethical standards for identity that could go otherwise overlooked, they also risk becoming a constraint that narrowed worker thinking to more rigid categories.

Finally, more implicit **company values** influenced worker perspectives on identity. Workers expressed instances where their values seemed to align or differ with their views of their company's values. Some participants seemed to express an alignment and adoption of company values. For example, Nicholas told the underlying story of why the CEO founded his company. He said the CEO had been rejected by a large banking company because he did not go to a prestigious school that the company recruited from, which led him to create his own business focused on giving people a more meritocratic chance at landing a job. Nicholas communicated an alignment and identification with the goals of the company, which provides computer vision software for video interviews. He expressed that incorporating demographics into AI hiring software for bias mitigation was an improvement on in-person interviews. Specifically, he believed in his company's approach to first measure demographics and then obfuscate them for automated interview analysis, including sentiment and affective assessments.

Siddharta also expressed joining his company because he felt that the vision of the company was "ethical:" *"You know, it is one of the reasons I joined Sybil, and not the other places I had offers from."* Siddhartha contrasted this with other companies, which he felt used computer vision for "scary" purposes, like highly targeted advertising. This contrast was intriguing because Sybil provides affective classification for targeted advertising, a use case that Siddharta expressed was a violation of privacy and something he personally disliked. However, because he was working on research on computer vision for autonomous vehicles, he was distanced from his company's main product.

In opposition, some participants were explicitly against what they felt their own companies' values and priorities were. Lynn, who initially joined MultiplAI because she felt it was a *"social good"* company lamented how she felt the company had changed. Lynn's values were in opposition to the company's market-based approach that Kenny described, building any product they legally can. Lynn actually left the company due to these differing values. Such instances represent how the company context might drive workers with certain values—in Lynn's case, a focus on social good over profit—away and potentially lead to a company full of workers that operate from a status quo perspective.

Irina, the CEO of an ethical data company, was in a unique position where her own *personal* values could drive *company* values. She developed a method for screening potential clients. She would assign each potential client an impact score, ranging from 1 (the project contributes to social good) to 4 (the project causes active harm, like military projects or content moderation projects). Her policy was to never accept projects with a score of 4, and to carefully consider how projects with a score of 3 (the project indirectly benefits society) might implicitly contribute to social goods even if was not the purpose of the client's project. She described how, in cases where clients came to her and she felt the projects were too harmful, she went about rejecting them:

> *There was this really problematic [proposal] about intelligent weapons … You know, you don't want to offend the people and tell them that they're horrible people and they shouldn't be building this AI. Usually what I say is that we're a social enterprise, and given that the majority of our workforce comes from conflict-affected countries, we're not able to perform such type of labeling. We have some Palestinians as well and this is also a personal preference of mine that I've instilled in the company to reject any project from an Israeli company, so in that case we tell them we work with … Palestinian refugees so we prefer not to work with Israel.*

Evident in Irina's example is how the social and cultural context *outside* of companies was also core to worker perspectives. Ethan emphasized that the overall context of research and development occurs in relation to the broader societal context companies sits within:

> *"[Corporate work] doesn't happen in a vacuum. And just like any, you know, academic research or anywhere else, the kinds of questions you're asking, there's a reason you're asking those questions. It's driven by societal concerns, by company concerns, by what you can get funding for, by … all of these kinds of things. They all come together."*

4.3 How Worker Positionality Influences Product

Now that the contextual factors in which workers are situated have been established, we present findings on how individual workers apply their own personal perspectives to their work and how they negotiate those perspectives with other workers and clients. In this section, we first present Findings focused on how workers described their own personal interest—or lack thereof—in approaches to implementing identity characteristics. We then describe instances where workers discussed disagreeing with or clashing with the positions of their colleagues. Finally, we conclude this section by describing approaches participants took to negotiating their own worldviews with the worldviews of those they worked with.

4.3.1 (Im)personal Stakes in Identity Work

Worker positionalities became particularly apparent when workers discussed their own personal stakes in defining identity characteristics for computer vision. Many participants attributed an interest in working on identity issues for computer vision to their own personal values and affinities with certain identities. Kaleigh, whose primary role is overseeing fairness initiatives across multiple computer vision products at Aqueous, described how the majority of resources come from people who are personally passionate about fairness. As a concrete example, Kaleigh was working on guiding product teams to update their approach to gender categories and concepts across computer vision. In doing that work, she focused on identifying and contracting researchers with gender-specific expertise *"who are passionate"* to bring product teams on board with the changes. Kaleigh worked to connect product teams to appropriate research resources in instances where a product needs to move from the status quo (*"what already exists"* in product) to what product should be (*"what we need for product"*).

Beyond value-driven interest, workers who held specific affinities with group identities played a major role in their approaches to their work. Vasudha (as a person of color), Kenzie (as biracial and a child of immigrants), Elliot (as non-binary), Callia (as a mother to a blind child), Lynn (as a wealthy white woman) and Madison (as a woman) all recognized how their own positionalities played a role in their work. Madison described how workers at Maelstrom prioritize certain projects:

> *Deciding what projects to focus on … is at least partially informed by people's identities and who they are, not only as a [worker] but also who they are as a person in life generally ... One example, in particular, is one of my coworkers on my team is gay … So, they felt a personal interest in this, as well as a professional interest. So that's one way that … that identity is manifesting. There's the corporate or tech kind of [way], like, 'What are the terms? How do we deal with them?' But that's also tied to the individuals doing [the work] and what they want to prioritize.*

Beyond driving this colleague to work on a project focused on LGBTQ identity, being LGBTQ also informed their approach to their work. This worker collected data at Pride from people directly, to both avoid online trolls and to work directly with the LGBTQ community. Madison felt this approach *"embraced identity by the horns"* rather than letting it be implicit or "neutral," as simply demographics or data points.

Much like Madison's colleague, Vasudha brought specific expertise to the table due to her own positionality as a person of color. She explained that historically, product teams at Aqueous have approached evaluating racial biases in facial recognition models by using skin tone. Yet, given her own experience, those categories were ineffective because they were far too static:

> *Because of a brown skin tone and a lot more melanin content, if you leave me under the sun for 20 days, I have like five shades darker skin tone ... So, when you think about comparing those, skin tone didn't really make any sense ... We realized very quickly that skin tone wasn't something that we could test on, especially taking into consideration the aging aspect [that an individual's skin differs across time].*

Vasudha knew from her own personal experience that if they were to use skin tone as a metric, the system may perform poorly over periods of time as her skin tone fluctuated. This personal knowledge also informed user studies her team went on to do, which showcased a need for information beyond skin tone, including facial structures. Of course, this did not make the decision for which categories to use necessarily easier or more concretely correct. Vasudha continued: *"We pivoted to ancestry background ... [but] there are issues that come up with 18 demographics. Why not 24? Why not 36?"* Vasudha's personal identity made asking questions about the viability of skin tone more obvious to her, but she did not perceive herself as able to determine what the best course of action was.

Of course, the positionalities people brought to their work were not static and changed over time as workers were exposed to new ideas and adapted to their company context. As Vasudha explained, when she was in graduate school, focusing entirely on research, she was not considering a business context. She explained that simply relabeling and retraining models in an industrial context is often economically untenable, especially due to labor costs. Therefore, her own positional perspective was now grounded in *"thinking about dollars"* and *"ship timelines"*.

The above examples showcase how workers make subjective judgments about what projects they care about and how to approach defining identity in those projects. Such subjective judgments reflect their own positional perspectives, informed by their experiences and beliefs. They relied on their own familiarity with identity concepts—either because they personally valued concepts, like equality or ethics, because they identified with the identity attribute in question, or because they became exposed to values they later came to internalize.

4.3.2 Colleagues with Clashing Positional Perspectives

Participants also seemed to lament values or perspectives that differed from their own. There seemed to be an overwhelming perspective that those in heavily technical roles—–like engineering and data science—had an interest in tasks rather than social implications. Kaleigh described how machine learning research teams often explore novel problems out of personal interest, which then later become embedded into products. This approach was the case for one of the computer vision products under her purview as a program manager, a mobile application for real-time classification for accessibility purposes. The researchers at the time created the product to classify gender and age because they had *felt* it would be useful, though they did not assess utility in any empirical way. Coleman, who worked on a prototype of this accessibility product for a research project, described choosing human characteristics that seem "useful": *"We tried to make sure there are certain things, like person names ... Because otherwise a human might not find it useful."* Much like Kaleigh described, Coleman took a utilitarian approach to his work on identity in the product. Coleman described how the *"part"* of himself that is trained as a machine learning researcher desires to build new models and focus on improving methods, perhaps at the expense of fairness:

> *So, I mean, part of me is a researcher who wants to get basically whatever data I can get my hands on, toss it into the data grinder, and build models ... and ignore the fact that maybe I've just produced something which is very, very biased towards certain things, which have optimized my scores.*

Machine learning researchers may simply not view identity bias as relevant to their position and expect others in the company to handle it. For example, Nitesh described not knowing anything about how identity tags are selected or filtered in the image indexing system he worked on. He stated that it was not his responsibility, but rather the responsibility of the data science team to consider how identity is represented and to mitigate biases. Much like Coleman's motivation to pursue state-of-the-art modeling, Nitesh's major motivation for pursuing a career in computer vision was solving technical problems, not social ones. Nitesh was so out of touch with how identity was implemented in the model he was engineering that he was entirely unaware that slurs were associated with targeted subgroups on his company's public-facing website.

Elliot critiqued the technical approach to scoping identity. In particular, Elliot was concerned with how categories like race and gender are represented—implicitly, because of a personal stake in gender as a non-binary person: *"Honestly, it's a bunch of ... you know, predominantly white, cis, male engineers who have not thought too much about identity."* Elliot explicitly questioned the positions that many engineers inhabit—white, cis, male—and assumed they lend to a less thoughtful approach to identity.

While workers like Elliot, Siddharta, Madison, and Lynn viewed gender as difficult to define from visuals, Kenny viewed it simply as *"a logical human decision"* on behalf of the data workers annotating gender in images. Callia expressed that gender representations in computer vision were not as *"big a deal"* as some others were making it out to be—her focus was on accessibility for blind people, and she thus advocated a binary gender. These clashing perspectives also highlight that personal affinities and experiences with certain positions motivate workers to approach problems in very different ways.

Vasudha points out that, just because technically focused workers might prioritize metrics over social impact, they are *"not malicious or necessarily bad people."* She highlights that technically focused workers, like machine learning researchers, are often unable to see issues of implicit bias. Such workers are simply approaching their work from a very different positional vantage point, one which prioritizes model performance over identity biases.

Participants made their own positionalities clearer through their disagreement with the perspectives of their colleagues. In many cases, participants seemed to distinguish between the workers who "care" about identity and the workers who do not. In others, like with Elliot and Callia, their opinions about how to represent identity revealed an underlying, otherwise implicit positionality. These comparisons highlight distinctly different approaches to identity work in computer vision, driven by implicit and ingrained positional perspectives.

4.3.3 Negotiating Positional Perspectives with Others

As demonstrated by how different workers disagree with each other's outlooks, designing identity in computer vision is a team endeavor. Just as decisions about identity were not made *"in a vacuum,"* they were also not made by individuals. Beyond the role each individual's positionality played in motivating their interest in specific work and guiding the decisions they make in conducting their work, participants regularly had to contend with the positionalities of their colleagues—both those they agreed and disagreed with.

Most often, workers were collaborating with those inside their own teams. Generally, workers seemed to share values and perspectives about identity with those colleagues on their direct teams. Elliot compared their own team with how other more product-focused teams approach their work: *"My team is probably ... the best team in terms of thinking a little bit more critically about machine learning systems as sociotechnical systems, as opposed to just ... algorithms where data comes in and data comes out."* Rather than being surrounded by pragmatic engineers focused on ensuring the best possible product, Elliot is surrounded by fellow researchers who care most about ethics. Given that all members of the team focus on ethical issues, it likely shapes their worldview in how they approach their work. Given the way that Elliot discussed their team as a collective (using language like *"we're interested"*), Elliot sees themselves as part of this larger mission with an ethical focus. They agree with their team members and their approach within the company.

Of course, in industrial contexts, people often collaborate with others with differing perspectives. Kaleigh described how different types of teams often work together to influence the outcome of a product, like both research and product teams. In these cases of collaboration, team members actively learn from one another's perspectives and expertise to shape identity outcomes in product. While positional tensions could cause stress, infighting, and even retaliation, it could also push teams to think outside of their comfort zones. Nitesh described how he *"support[s] diversity"* because *"you get different approaches to a problem."* He felt that *"constructive conflicts"* are reflective of the *"complex world"* that products are meant to serve. In particular, he highlighted that teams with diverse training—such as from diverse academic backgrounds—are beneficial. Workers also brought up how conducting user research and bringing in outside consultants could shift conceptions of identity in products.

However, people often had to work with those they disagreed with. Tensions seemed to occur often when working with colleagues who had had very different roles and goals. Once more, the notion of technical versus non-technical became a point of contention. As Callia explained, from her perspective as an accessibility researcher, *"there's a significant amount of negotiation to orient slash reorient [computer vision engineers] in a way that accounts for the human experience."* Callia viewed her position, and other human-centered researchers, as oppositional to technical researchers, who are mostly focused on solving technical issues and not accounting for human experience. Elliot similarly criticized machine learning colleagues about their approach to racial identity as attribute-based rather than something sociopolitical.

Perspectives on how to implement identity in product was often so personal to workers that negotiations became, as Kaleigh explained, "emotional" for those on the conceding side:

> *Some [product teams] are really, really good about recognizing that things have changed and we need to conceptualize things differently. But for others, I was in a meeting recently, where it was quite emotional for them to let go of a feature that doesn't align with our responsible AI principles and values anymore, [because] it's a feature that they've been working on for years.*

Kaleigh explained that though her team had the power to step in and make executive decisions about changing identity in computer vision, they did their best to work with teams to avoid causing internal conflict. She said that she *"[tries] to avoid the sort of a mallet approach"* as a *"last resort"* for cases where *"thinking through it together hasn't helped."*

Additionally, issues with how to implement identity can occur when interacting with superiors or people with more institutional power within the company; workers in lower positions of power face silencing and even retaliation. Retaliations from those higher on the "chain of command" indicate how some workers might use their power over others to quash specific positional perspectives. For example, Madison described how the company and those in it were not supportive of addressing identity-based biases in product, and so it was an uphill battle with those in more powerful positions who did not value her perspective. She described facing skepticism and pushback, which she tried to navigate by appealing to other peoples' values. Instead of focusing on harms, she would focus on the business benefits of attending to identity bias. Madison explicitly attributed difficulties implementing more fairness ideas in computer vision to "politics" around the identity groups that dominate tech. She felt that, as a woman, her ideas were taken less seriously than those of men:

> *If you have a group of women saying how they think the technical aspects should go, and it's what you're not used to hearing, there's like, no way they'll be taken [seriously] . . . But then if you have people who are white or Asian men, maybe more fitting the traditional personalities of who tech people are, putting forward ideas, it does get a lot more traction.*

Madison described how she has received retaliation from colleagues through bad performance reviews, which impacted her career trajectory. She believes that such perspectives are so deeply ingrained in people's approaches to their work that they are not even necessarily aware they are being biased: *"I guess, it's just that the desire to maintain a white, Asian, cis, male view of the world is so strong, and people don't even realize they're doing it."* To ensure her perspectives are listened to, Madison describes relying on white and Asian cis male allies to communicate her ideas for her.

Further, since many companies provide computer vision solutions for clients, workers also had to negotiate different perspectives with client representatives. Beyond different perspectives on how to approach computer vision tasks broadly (e.g., for or against identity classifications), regional differences were often a source of positional differences, as tech companies generally operate at an international scale. For example, Lydia commented on a specific instance where she and colleagues made a decision to include bias mitigation practices despite a client being unconcerned by bias:

> *Different countries have different laws around bias and fairness . . . like, Japanese data is usually really sexist. And like, they don't care about that in their country. So, like, the customer doesn't care, whereas we would want to mitigate [gender bias] . . . So, I guess for me, it's just kind of having that conversation where there's cultural differences of where that balance should be.*

In this statement, Lydia is also expressing her own view of Japanese culture from the positional vantage point of a United States citizen. She is packaging both the views of her client and the views of another culture together in her assessment of how to approach gender in the product (*"they don't care about that in their country"*). Nicholas, on the other hand, questioned when it was ethical to *"inject"* U.S. *"attitudes"* into products meant to be deployed in other countries: *"We have the right to not do business with them, but do we have the right to inject our beliefs and change algorithms for them?"* Balancing ethics about when to include certain cultural beliefs about identity impacts the outcome of how identity is deployed in specific countries.

Determining how a product should be designed requires negotiation between many different actors. Furthermore, interfacing with others acted to mutually construct how individual participants interpreted identity. As workers encountered different viewpoints from their own, their own viewpoints grew and sometimes changed—what Nitesh described as *"constructive conflict."* Representations of identity in computer vision are therefore not the perspective of a single person or even a single team, but a multitude of actors within (and outside of) the development context.

4.4 Positional Gaps that Arise During Product Deployment

As demonstrated by the Findings thus far, participants approached their work from their own positional perspectives. Their individual perspectives were negotiated through collaborations with others. Yet, workers situated within tech company contexts are often unable to predict how their own positionalities, as incomplete images of the world, might result in positional gaps in product design. Once products were finished being developed, identity issues sometimes arose. These issues were largely caused by the development team being unable to foresee them, because they did not occupy identity positions that would make such issues obvious.

Many positional gaps are embedded into products because neither the clients nor the workers even realize their perspective could be biased. Irina described a project focused on providing AI-generated insights about video job interviews. She said that the client had originally gone to another annotation company but ran into strong cultural biases in the data annotation: *"They were working with an Indian outsourcing company and people were much more favorable towards Indians."* The client had not expected the annotator to have a deeper understanding or affinity for those applicants which shared the same racial and cultural positionality as them.

Sometimes, issues caused by positional gaps were caught before deployment—particularly for companies with the resources to do internal testing. Madison described a scenario where a product team of primarily men had designed a computer vision wearable. Before deploying the product, they conducted internal testing, commonly referred to as "dogfooding" in the tech industry. She explained that the women who *"dogfooded"* the *"necklace version"* of this wearable quickly *"realized that the camera was like on their breasts ... [because] all the original designs had been developed by flat-chested men."* This example highlights how the positions the men inhabited made it difficult for them to automatically recognize their product was uncomfortable for people with breasts. Further, those with breasts were quickly able to recognize that the product did not work for their body types. As Madison said, "[the women] were able to discover that because that's part of who they are."

To avoid positional gaps, many participants described how fellow workers from marginalized positionalities were a resource for vetting or feedback. Colleagues from marginalized identities could name issues with products that others workers were unable to see due to their limited positional standpoints. For example, Lynn described how the language in their classifier was changed due to the inputs from non-binary colleagues:

> *Instead of saying 'male' or 'female,' we used 'masculine' or 'feminine,' more of a descriptor than a prescripter. There was a big debate about it, the whole company was involved. . . . we had some genderfluid people in the office, we had some trans people in the office, so their opinions were really important to us. And in the end, we delayed the launch.*

Similarly, Elliot described how it was common practice to rely on identity-based "resource groups" at their company. It is common for employees at Maelstrom to do internal testing with these identity-specific groups before deploying products.

In other cases, issues only arose once the product was publicly deployed, and users of the product encountered issues. Often, given the issues were relevant to identity characteristics, users found these issues offensive. As someone who worked on both text and image-based machine translation, he encountered a number of accidents that resulted in major PR problems for the company. For example, he described how religious entities and names were mistranslated. In one specific case, formal Russian names were misgendered as female instead of male. Coleman lamented these mistakes, though he also felt they were a learning experience for him and his colleagues:

> *[It] was a very healthy shock for our ecosystem, because now people have understood, okay, if we do this, we then at least have to give set up some monitoring for the first week . . . maybe have a new release to be able to very quickly un-deploy if something bad comes up.*

Positional gaps were largely inevitable, as teams were unable to fully account for perspectives they did not know were missing.

However, whether pre- or post-deployment, those occupying positions outside of participants' own limited viewpoints could make products more robust and more inclusive—and hopefully fill the gaps product teams did not realize would occur.

5 Discussion

Tech workers all bring their own identities to the table. They operate from their own positional perspectives when conducting identity work for computer vision. These perspectives are influenced and constrained by the industrial contexts that workers are embedded in. For example, workers in smaller companies are often more constrained in how deeply they can engage with identity than those in large companies due to a lack of economic resources and incentives.

Positional perspectives are especially evident in the personal reasons participants expressed for engaging in identity work. Of course, their individual positional perspectives may or may not align with those of their colleagues. Workers often expressed displeasure with their colleagues' worldviews, though they had to work with them and negotiate differing perspectives. Further, given each worker brought their own limited positionalities to their work, teams had gaps in positional worldviews. Such gaps led to unforeseen and undesirable outcomes when it came to product design, such as offensive classifications or hardware that only worked for some bodies. Aware of their own positional gaps, workers were proponents of diversity and would often turn to colleagues from minoritized identities as resources to augment their own limited worldviews. We detail how individual workers' positionalities are mutually constructive and informed by actors across a variety of contexts that influence worker worldviews.

5.1 Positional Approaches in Context

Positionality is the complex, mutually constructive relationship between one's identity and how they view the world around them. Every individual occupies a specific position in the world, impacting how they are viewed and treated; as a result, they view and interact with the world from a specific standpoint [16, 51]. We found that the tech workers who develop computer vision are no different. In implementing identity characteristics for industrial-scale computer vision, workers implicitly rely on their own positionalities. They shape identity in computer vision from their own standpoints. Often, their own interest in identity in computer vision stems from their personal values and their own affinities with identity characteristics. On the other hand, those who do not explicitly value fairness for identity groups or identify strongly with specific identity characteristics express little knowledge or interest in identity issues in computer vision.

However, the process of developing identity in computer vision is not simple and straightforward. We found that individual workers do not make individual decisions about how best to implement identity. Rather, workers operate within a complex environment, informed by many different contexts. We identified four contexts (Company, Development, Outside Development, and Macro Social) in which each individual worker sits (see Figure 1).

5.1.1 *Company Context*

Individual workers are embedded within a specific *Company Context*. Both the economic power and the values of their company impact the approaches workers can take to identity. In some cases, companies also have specific policies governing their approach to developing identity concepts in computer vision. Workers in smaller companies with less economic power often have limited access to resources, such as teams dedicated to assessing the ethics of a project or the ability to collect robust and diverse datasets. Further, such companies are often driven by market demand more so than ethical policies. Small companies like MultiplAI adopt a market-first approach, serving client demands first and foremost and thus deprioritizing more nuanced approaches to identity concepts like gender. Meanwhile, larger companies like Aqueous have core policies driving their approaches to AI, including computer vision. They have teams dedicated to ensuring the fairness of products and whether products align with company policy. They even dedicate resources to overhauling identity concepts as outside perceptions about them change. For example, how Kaleigh was overseeing updating gender in Aqueous' core computer vision product.

The relationship that individual workers have with their company also showcases their positional perspectives. Some workers expressed alignment with the values expressed by their company, indicating that they shared those values. These workers seemed to have a positive perspective of the type of work they could conduct within their company because their company likely valued their approach. On the other hand, some workers disagreed with the values of their company. Workers like Lynn, who disagreed with the values they felt their company prioritized, did not feel empowered to approach development the way they felt it should be done. Lynn disagreed with the company's market-first approach and desired more nuanced, careful, contextual approaches to identity in computer vision. Even while she expressed relatively reductive beliefs about gender herself, she expressed pride in her trans colleagues who pushed back on the initial representation of gender in their demographics model.

The contrast between Lynn's description of gender classification and her description of her colleagues also highlighted her own positional perspectives. To Lynn, biological sex was still always evident on the face. It is possible that her colleagues held very different perspectives on gender. Within a company, workers might have differential views from the others they are working with—both their colleagues and their superiors. Choices are also not made by one person, but by numerous people, with varying positional perspectives. Tech workers need to negotiate their own positional perspectives with those of their colleagues. As such, identity in a product may shift and morph as it takes on numerous perspectives during the development process. For those with very different views, whoever is given the most decision-making power is likely to be most influential. At Aqueous, Kaleigh, in her position as a manager, was given power to veto product team decisions. However, in other cases, product teams or engineering teams might have more power to make final decisions.

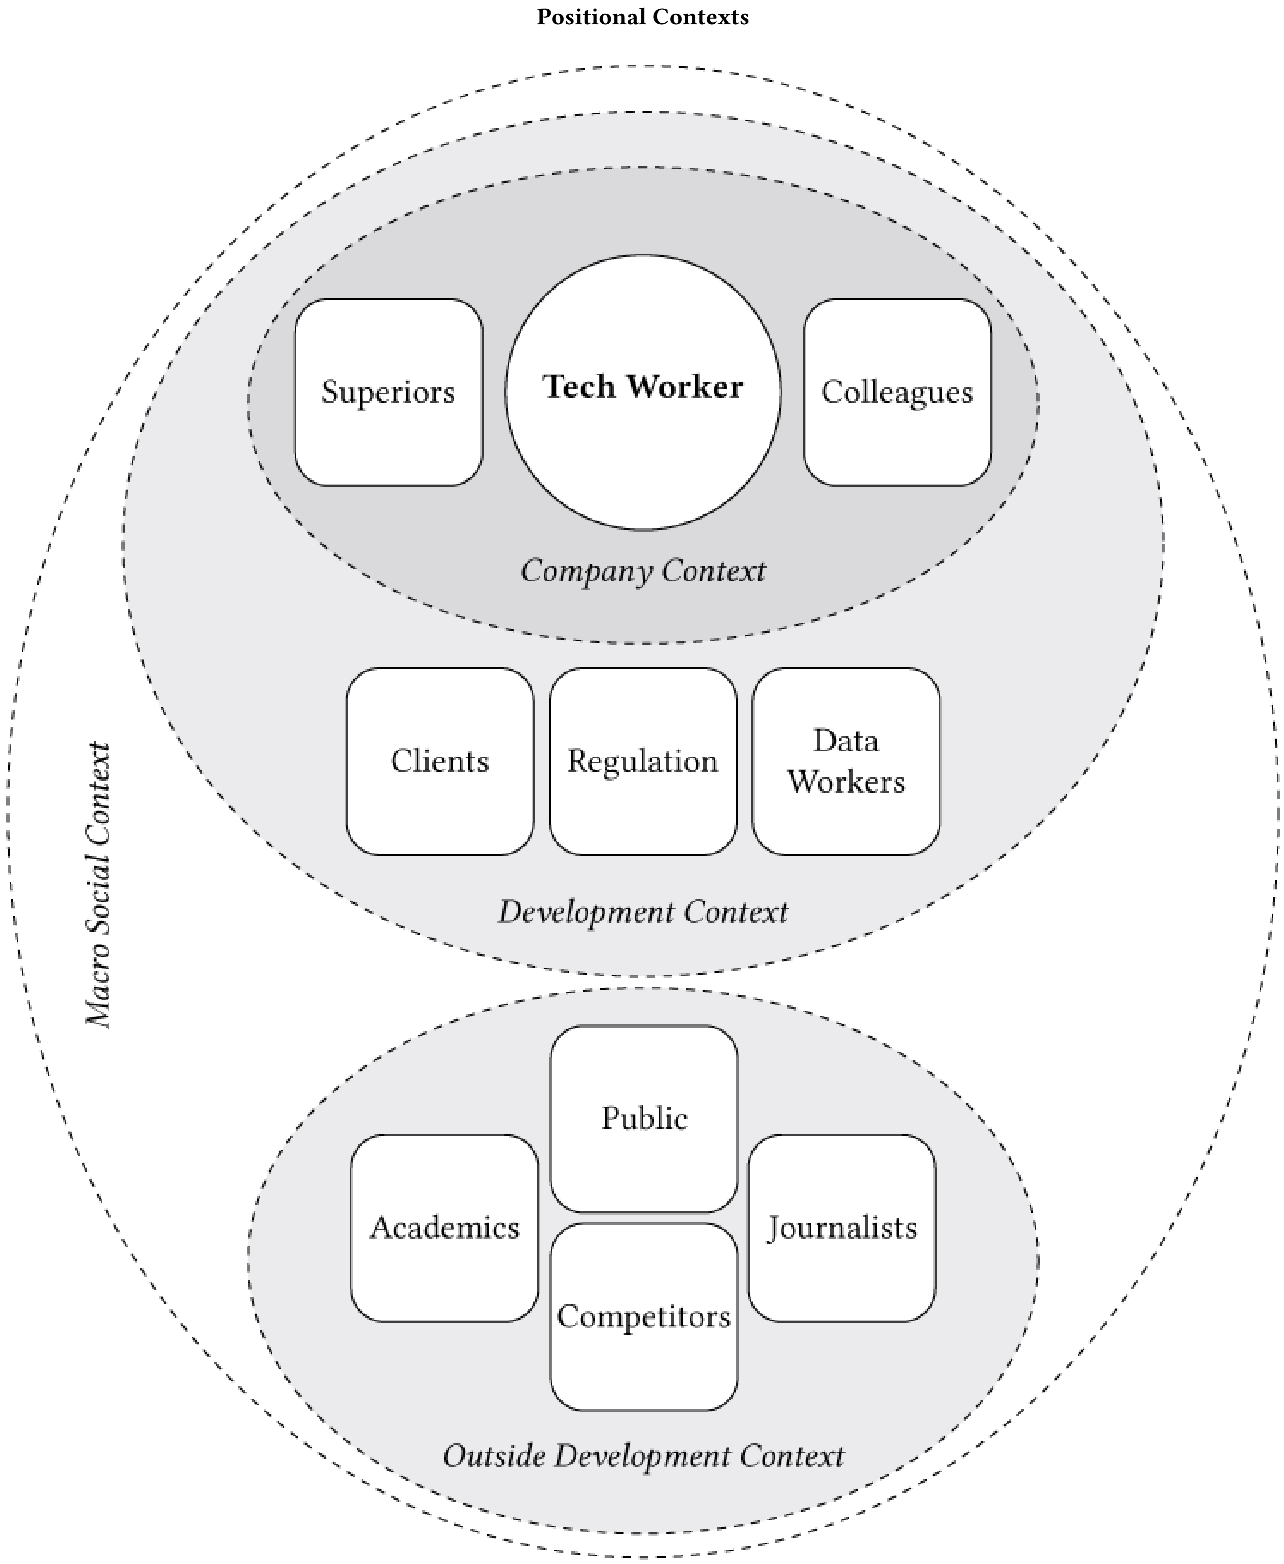

Figure 1: A figure depicting the contexts that different actors who influence the development of developing computer vision are situated in. Each actor is placed in the relevant context they are involved in (e.g., academics are in the Outside Development Context, tech workers are in the Company Context).

5.1.2 Development Context

The *Development Context* was the context in which products were being developed, encompassing the company but also actors involved in development outside of a singular company. In some cases, clients, regulations, and data workers were involved in the process of developing a product. Clients bring their own positions to the table through their demand for specific features when they hire a company to develop a product. Once more, this was evident in MultiplAI's early clients demanding gender classification for marketing purposes; this led MultilplAI's model having gender embedded as a feature from the company's infancy. In other cases, a product that was developed years ago was still relied upon by clients. Even though Aqueous was interested in updating gender in their models, they also had clients that had been using it for years. Some product teams would utilize client reliance as a reason to cling to older models of identity classification that they were attached to as initial developers.

While computer vision is subject to little federal regulation, some use cases of computer vision products necessitate compliance with federal laws. For example, Resoom uses computer vision for assessing job candidate interviews. Thus, the model is subject to the U.S. Equal Employment Opportunity Commission's regulations. Regulatory requirements govern the types of identity groups workers must attend to in the design of computer vision. It necessitates that workers design for specific categories and ensure some level of fairness for those categories. For example, the EEOC provides a list of the "minimum" categories that must be attended to for "race/ethnicity": "White; Black or African American; Hispanic or Latino; American Indian or Alaska Native; Asian; and Native Hawaiian or Other Pacific Islander," as determined by the Census Bureau [26]. While workers might push to attend to race beyond these six categories, they are not required to do so. Meanwhile, the gender categories the EEOC required at the time of this study were only "male" and "female," which were the only categories Resoom attended to in their model. However, the EEOC plans to add "non-binary" to its list of gender categories, likely forcing Resoom to include "non-binary" gender options in their fairness mitigation strategies [1]. Regulation can force workers to attend to identity in ways that otherwise would not, but also has the potential to limit whether workers attend to categories beyond necessity.

Further, many companies will hire data workers for their data needs, prior to being able to train and evaluate models. The positionalities of those who curate and label the data for training models add another layer of complexity to defining identity. EnVision Data is an example of a company that provides data services for computer vision. While tech workers define identity concepts early on during the development process and constrain data worker positionalities via instruction guidelines, data workers also introduce their own positional perspectives on identity while conducting data work. Tech workers act to control data worker positionalities, attempting to maintain their own perspectives. Often, these interactions further expose how workers impose their positionalities on identity concepts. As they come into contact with data workers who view identity differently, especially when they live in a different social context, tech workers attempt to reorient data workers to their positional worldviews.

5.1.3 Outside Development Context

Beyond the Development Context itself, many tech workers are attuned to those in the *Outside Development Context*. Those in the Outside Development Context context are still directly engaging with AI and the tech industry, but are not directly involved in product development or internal company contexts. They were aware of ongoing conversations about identity in computer vision among the public, the press, academics, and their competitors. The press and the public were often viewed as sources of contention. In some cases, participants viewed public outcry or poor press coverage as a lesson for what not to do. For example, Lynn described learning from critical PR coverage of big tech when she was otherwise unsure how to proceed. In other cases, workers seemed to denounce the perspectives of the public. For example, Coleman spoke sarcastically about public outcry towards model mistakes, even if they were opportunities to identify issues otherwise unseen. Academia was also viewed as a resource for workers who did not want to rely solely on their own intuition but wanted to implement best practices.

5.1.4 Macro Social Context

Finally, overarching the development of computer vision itself, workers are always influenced by the *Macro Social Context* in which they are embedded. Workers brought to the table opinions and beliefs shaped by the institutions they had learned from and the society in which they were embedded. Many participants expressed beliefs that were learned from their education, before they came to industry. For example, Coleman expressed having an internal desire to *"toss [data] into the data grinder and build models."* This approach stems from how he learned to conduct research. Many disagreements seemed to stem from different disciplinary training which led workers to fulfill different roles. Workers in more social scientist roles expressed a dislike for approaches that prioritized engineering goals.

Further, participants acknowledged that identity concepts were culturally contingent; they could differ across cultures. Lydia felt that her clients in Japan did not care about sexist data, but it is also possible that views on sexism simply differ between the U.S., where Lydia is based, and Japan. As all participants came from Western countries in the Global North, their approach to identity was informed by what was culturally familiar to them through their socialization.

5.2 Attending to Positionality

Given that humans are involved in the design process of identity characteristics, they will always bring their own perspectives to the table—perspectives which are influenced by the contexts in which they are situated. Like a snake eating its own tail, workers are constantly influenced by context as they themselves influence those contexts. Rather than attempt to "solve" positionality, viewing it as a subjectivity which should be stripped for the sake of objectivity, practitioners and researchers alike can explicitly attend to it. The goal would be to identify and attend to gaps before they become unforeseen outcomes. We propose attending to positionality from two perspectives: (1) attending to contexts and (2) attending to actors within those contexts. Both perspectives are not mutually exclusive; they can be attended to in tandem or in relationship with one another.

5.2.1 Attending to Contexts

Given that contexts mutually influence one another, one can imagine starting at the highest-level context—the Macro Social Context—or at the one most constrained to the creation of the product—the Company Context. Understanding how worker positionalities are shaped and constrained by context can ground research on AI development in industrial context in the social, cultural, and material conditions of work. Companies might also consider adopting contextually informed approaches to improve identity practices and enable their workers to more explicitly contend with positionality, which has been otherwise implicit.

Understanding the **Macro Social Context** in which development is embedded opens up opportunities to understand more about how social categories of identity are structured before attempting to define them for technical systems like computer vision. Further, examining the social context of development might illuminate how categorical histories influence the way that workers approach identity problems in computer vision. They might be taking identity categories for granted, treating them as given, while the categories themselves are socially contingent. For example, workers in the United States might be making decisions about identity that are untenable in other contexts, like India. Tech workers attempting to define identity for more constrained environments, like medical contexts, might fail to account for how domain experts use identity information. If workers first contend with the social context governing identity, their perspectives may become more grounded in the specific context of use rather than personal experience.

Each of these identity categories also has a social and political history attached to it. Examining the history of a category can reveal normative and prejudicial assumptions about the people grouped under a category. Examining social categories can also reveal which types of identity categories are perceived as rigid and which are not—for example, in some cultures, gender is viewed as rigid and binary, while in other cultures, gender is viewed as fluid and non-binary [21, 48, 59, 68]. Further, given identity is always evolving, turning to developing conversations about identity categories in the Macro Social Context is beneficial to staying engaged in contextual politics and local communities. After examining the history of institutions like the U.S. Census, workers might choose to explore more community-grounded approaches to categories like gender, race, or ethnicity. They might assess what identity categories they have chosen to be rigid and why. Given critiques of computer vision reflecting narrow perspectives on identity (e.g., [6, 28, 57]), contending with the Macro Social Context of identity before development can expand the narrow positional worldviews of tech workers. It can benefit companies trying to understand the broader conversations around identity so that they can initiate the development of identity categories in knowledge grounded in current community-centered language rather than outdated practices (e.g., gender as a binary). They might instead consider other ways of viewing the world.

One might also take a step down, to examine the **Development Context** of a specific computer vision product. Attending to the Development Context might mean examining how the role of regulation influences approaches to identity in technical artifacts, much like how the identity categories outlined by the EEOC influenced approaches at Resoom. It might also mean understanding how the current landscape of B2B businesses in the computer vision space has constructed status quo approaches to identity in the computer vision industry, potentially shaping the way individual workers are primed to think about identity problems in the field.

Similarly, examining the **Company Context** would mean grounding understanding identity in the larger industrial context shaping the project. The companies that individuals work in heavily influence how workers can approach identity concepts. Economic conditions, internal policies, and company values influenced whether workers could approach identity in the way they desired. Focusing on each of these three factors can reveal how company context shapes the positions workers occupy. For example, understanding the economic conditions of the company reveals the resources that are and are not available to workers. Workers in smaller companies are often unable to access the same resources, like having multiple researchers focused on developing best practices. Similarly, internal policies established by companies might provide justifications for certain approaches to work, while denying other potential approaches. Kaleigh, for example, would often rely on her company's policies to justify pushing for better approaches to gender in their computer vision model. Other companies might adopt policies that do the opposite, prioritizing, perhaps, the technical over the social. A lack of company policies might also mean workers rely more on management, intuition, or market incentives. Company values might invite certain types of workers to succeed in their approaches more than others. For example, Lynn felt her company did not value the same things she valued and left the company. Meanwhile, Kenny seemed to embrace the same company's values of a market-driven approach in his role as the vice president of business; he prioritized bringing in clients over developing nuanced approaches to identity like Lynn. Not only does understanding the company context benefit research—and critique—by grounding it more acutely in the realities constraining and enabling certain workers, but companies can also benefit from understanding how their own company culture shapes development. Companies might also consider shifting priorities to better allocate resources or develop policies for scoping requirements for identity concepts, especially given changing legal landscapes and public perceptions around AI.

Finally, understanding the **Outside Development Context** can reveal how those uninvolved in the direct development process of computer vision can still influence worker approaches. Many workers were aware of how the public, journalists, academics, and corporate competitors perceived identity in computer vision. Workers could often use these outside perceptions to influence their colleagues or managers. Further, companies often responded to these outside influences, creating or updating policies, reallocating budgets, and developing new company identities. Assessing how identity is being discussed by actors in the Outside Development Context can benefit researchers attempting to understand current practices in industry. For example, it can reveal the mistakes their competitors have made and how the public responded to those mistakes, as well as how their competitors are adjusting to public discourse. One could imagine how the role of competitors, journalists, activists, educators, and others could have significant impact on how these types of products are built and used.

5.2.2 Attending to Actors

A more precise method for attending to positionality in computer vision is to attend to the different actors involved in the development process. Attending to specific groups of workers is certainly not a new approach in HCI (e.g., [17, 24, 30, 42, 47]). Others have uncovered that, for example, even programmers with a grasp on gender beyond the binary have a tendency not to consider how best to implement gender or when it is relevant to users [8]. The work at hand took this approach—examining how tech workers expressed their positionalities during the development of computer vision to understand how positionality shaped how identity is embedded into computer vision artifacts. There is still further opportunity to engage with the ways worker positionality influences identity practices in technology development for both researchers and practitioners. Attending to different positional actors can further reveal how perspectives shape identity outcomes in AI.

Much like this study, one might consider grounding understanding in worker positionalities. Beyond broadly understanding the role of positionality in identity development, there are still many opportunities to create better practices for documenting and attending to positionality. For example, one might center the positional perspectives of workers to better develop policies for explicitly engaging with worker positionality during development. In this study, the process of defining identity for computer vision projects was not explicitly part of development approaches. Workers did not explicitly engage with how their own positionalities influenced the way they perceived or implemented identity in computer vision. Given different types of workers have different perspectives, it could be fruitful to focus in depth on how specific types of workers reason about identity. One might compare researcher approaches with that of engineers, for example. One might also choose to examine the role of management in defining identity, to determine whether and how often identity comes from a bottom up or a top-down perspective in industrial contexts. Much like Wang et al. found when examining the working relationships between tech workers and data workers [69], the positional values of tech workers in positions of institutional power may have more impact on identity outcomes in product. Understanding the ways managers imagine identity characteristics can reveal (mis)alignments between management and other actors in the pipeline, like users and product developers. It could also reveal points of intervention for reimagining identity characteristics which might otherwise be lost when focusing only on non-management level workers.

Another opportunity is to understand how other actors influence the perspectives of core workers. To do this, one might examine the perspectives of actors within the Development Context but outside the core company context—like clients, data workers, and regulators.

Clients, those who request computer vision services from other companies, have their own expectations about identity. Understanding the positionalities of clients can also reveal why identity in computer vision products is designed the way it is and may reveal points of intervention for shifting design practices. For example, Kenny explained that marketing clients drove the use of discrete identity categories like binary gender early on at his company. Talking with client representatives in marketing contexts can reveal what worldviews drive their desires for such discrete categories.

Much like clients, there is also further opportunity to understand the positionalities of data workers who provide data services for computer vision. Data workers, as underpaid and largely invisible in the development of AI [25, 44, 45], are still crucial to its development. As scholars increasingly examine the ways data workers are disempowered in the development of AI, they might also examine the ways data workers make decisions about identity in their work. Data workers, as they interface with tech workers, might influence the way workers consider identity in data.

Finally, understanding the role of regulation and its implications for identity in AI can reveal alignments and gaps between those directly within the company context and those outside it. Understanding the perspectives of policymakers and their interactions with tech companies in designing policy can reveal their values, experiences, and perspectives. For example, do council members at the EEOC even consider how their decisions influence identity categories in AI? When creating facial recognition laws, what positional perspectives are policymakers bringing to the table? Given regulation influences and constrains the way tech workers engage with identity categories, expanding understanding of identity development beyond solely product can paint a richer picture of the many positional worldviews influencing identity in computer vision.

One might also consider engaging the positionalities of those actors in an Outside Development Context of computer vision altogether. Given products are deployed and impact actors outside of their development, understanding the role of the public, journalists, and academics might reveal different positional perspectives on identity than what workers exhibit. Further, examining the relationships between outside actors and how they communicate with those in a Development Context can further ground knowledge on how tech workers perceive and are shaped by outside actors.

Attending to actors within different contexts provides both researchers and practitioners with a number of opportunities to better understand the role of human positionality in computer vision development. It can help to identify positional gaps, before they become unforeseen and undesirable outcomes when products are deployed. Such knowledge can lend to hiring decisions within companies, better practices for attending to positionality and documenting identity decisions during development, and more contextually informed research that extends beyond simple but untenable recommendations focused on improving industry practice.

5.3 Considerations and Future Work

Above, we presented a framework for identifying different contexts that influence positionality and different actor perspectives from which to examine and attend to positionality. We believe that attending to positionality is only a first step in creating more ethical AI futures. In this section, we want to explicitly acknowledge that actually attending to positionality is extremely complicated, as it is often a nuanced, implicit, and intersectional phenomenon that is neither static nor generalizable. While some of the points in this section might be seen as limitations, we instead encourage readers to view them as *considerations* for research and design—open questions that complicate attending to positionality, which may result in tradeoffs or creative workarounds, and which present opportunities for deeper engagement.

The first consideration researchers and designers should attend to is **the difficulty unearthing positionality with relevant actors.** In the work at hand, something that was difficult to analyze were instances where, as external researchers, otherwise visible or expressed identities (e.g., race, gender pronouns) did not seem to be relevant to discuss from participants' perspectives. For example, we noticed that men participants seemed less likely to discuss the role of their own gender in product development than women and non-binary participants; we also saw participants who were both white and non-white never bring up racial or ethnic identities as relevant to their work. We characterize these as gaps between participants' acknowledgments of their own positionalities and the characteristics we as researchers observed as informed by our positionalities. Such gaps brought up methodological questions about how best to get at positionalities, which might feel so naturalized to participants that they are perceived as inconsequential. We must also consider how those of marginalized positionalities may be overburdened by positionally-informed approaches to AI, especially given those with more privileged positions may not view their identities are relevant to conducting identity work.

The second consideration is that **positional perspectives carry specific values and decisions must be made for which values will be implemented.** Not all values should be weighted equally. While HCI has a long history of embracing value-sensitive design [23], many researchers have also actively questioned: whose values should we be sensitive to (e.g., [3, 36, 40])? For example, while Lynn seemed to support her genderqueer colleagues in pushing back against gender classification, she also expressed perspectives that supported the belief that gender characteristics are evident from facial structures. Her perspective sits in opposition with more critical work on gender classification in computer vision [33, 55–57]. Such values are also not static and may become more salient to product decisions as cultural discourse and political climates shift. Emerging values for more ethical AI practices may become outdated and viewed as harmful in the years to come. Figuring out which values to represent and whose to exclude is only the first piece of the puzzle. We must also figure out how to design AI systems that are flexible to shifting values. Further, we must determine ways to decide who is responsible for making such value decisions among the many actors involved in researching, developing, using, and regulating AI.

The third and final consideration, which is highly consequential, and also difficult to disentangle, is **the role of *intersectionality* [15] in shaping and constraining positionality.** In this work, we approached positionality from the perspective of individuals, which were otherwise shaped by outside relationships, contexts, and institutions. Intersectionality, as a theory aimed at understanding how structures of power marginalize and dominate people with overlapping identity characteristics (e.g., race and gender), is integral to actually attending to positionality meaningfully in development contexts. Collins specifically realigns feminist standpoint theories with the reality of overlapping systems of domination [13]. In reality, not all workers have the same power to shape and implement identity in product, and this is compounded when multiple positions are disempowered (e.g., a Black woman researcher will likely be able to exercise less power than a White man engineer). In our own future work, we plan to attend explicitly to the role *positional power* plays in product development.

6 Conclusion

All individuals have their own positionality, the perspective that they hold as a result of their own identities and interactions with others and the world around them. In the development of computer vision, positionality is critical to how workers approach defining and implementing identity. In this paper, we showed how the practices of tech workers—like researchers and engineers—reflect positionality. Not only do their practices reflect their own personal values and experiences, but they also show when workers have differential worldviews and must negotiate them with their colleagues. Further, we demonstrated how the contexts in which workers are embedded shape their positional approaches to computer vision, sometimes enabling and sometimes constraining their perspectives. Finally, we saw how positional gaps within tech workforces can lead to unforeseen outcomes around identity issues in products. Workers, acknowledging their own limitations, advocate for more diverse workforces that they can use as resources for improving identity approaches within their companies. Positionality, as a subjective and value-laden reality, is not an issue to be solved. Instead, it offers opportunities for explicit critical engagement so that researchers and practitioners can attend to positional gaps before they become undesirable and even offensive outcomes. We thus proposed implications for more deeply engaging with positionality across contexts and actors in the field of computer vision development.

Acknowledgments

We would like to thank the reviewers for their time and dedication to improving this paper. We would also like to thank Mary L. Gray, Allison Woodruff, Casey Fiesler, and Robin Burke for their continuing feedback during the course of this project. We would also like to thank Kenneth Holstein for guidance in the early stages of this project. Finally, we would like to thank Samantha Dalal for her feedback in the process of writing this paper. We would like to thank This work was completed while the first author was funded by a Microsoft PhD Research Fellowship.

References

[1] 2022. EEOC to Add Non-Binary Gender Option to Discrimination Charge Intake Process. https://www.eeoc.gov/newsroom/eeoc-add-non-binary-gender-option-discrimination-charge-intake-process.

[2] Carmen Aguado. 2011. Facebook or Face Bank. *Loyola of Los Angeles Entertainment Law Review* 32, 2 (2011), 187–228.

[3] Tamara Alsheikh, Jennifer A. Rode, and Siân E. Lindley. 2011. (Whose) Value-Sensitive Design: A Study of Long- Distance Relationships in an Arabic Cultural Context. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. Association for Computing Machinery, New York, NY, USA, 75–84. https://doi.org/10.1145/1958824.1958836

[4] Fabio Bacchini and Ludovica Lorusso. 2019. Race, Again: How Face Recognition Technology Reinforces Racial Discrimination. *Journal of Information, Communication and Ethics in Society* 17, 3 (Aug. 2019), 321–335. https://doi.org/10.1108/JICES-05-2018-0050

[5] Alistair Barr. 2015. Google Mistakenly Tags Black People as 'Gorillas,' Showing Limits of Algorithms. *Wall Street Journal* (2015).

[6] Cynthia L. Bennett, Cole Gleason, Morgan Klaus Scheuerman, Jeffrey P. Bigham, Anhong Guo, and Alexandra To. 2021. “It’s Complicated”: Negotiating Accessibility and (Mis)Representation in Image Descriptions of Race, Gender, and Disability. In *CHI Conference on Human Factors in Computing Systems (CHI ’21)*. ACM.

[7] Abeba Birhane and Vinay Uday Prabhu. 2021. Large Image Datasets: A Pyrrhic Win for Computer Vision?. In *Proceedings - 2021 IEEE Winter Conference on Applications of Computer Vision, WACV 2021*. Institute of Electrical and Electronics Engineers Inc., 1536–1546. https://doi.org/10.1109/WACV48630.2021.00158 arXiv:2006.16923

[8] Elijah Bouma-Sims and Yasemin Acar. 2023. Beyond the Boolean: How Programmers Ask About, Use, and Discuss Gender. *Proceedings of the ACM on Human-Computer Interaction* 7, 1 CSCW (April 2023). https://doi.org/10.1145/3579461

[9] Joy Buolamwini and Timnit Gebru. 2018. *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification **. Technical Report. 1–15 pages.

[10] Scott Allen Cambo and Darren Gergle. 2022. Model Positionality and Computational Reflexivity: Promoting Reflexivity in Data Science. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3491102.3501998

[11] Le Chen, Ruijun Ma, Anikó Hannák, and Christo Wilson. 2018. Investigating the Impact of Gender on Rank in Resume Search Engines. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA.

[12] Patricia Hill Collins. 1990. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. Vol. 21. https://doi.org/10.2307/2074808

[13] Patricia Hill Collins. 1998. Some Group Matters: Intersectionality, Situated Standpoints, and Black Feminist Thought. In *Fighting Words Black Women and the Search for Justice*. University of Minnesota Press.

[14] Kate Crawford and Jason Schultz. 2019. AI Systems as State Actors. https://columbialawreview.org/content/ai-systems-as-state-actors/.

[15] Kimberle Crenshaw. 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence Against Women of Color. *Source: Stanford Law Review* 43, 6 (1991), 1241–1299. arXiv:1229039

[16] Sofia Marques da Silva and Joan Parker Webster. 2018. *Positionality and Standpoint*. https://doi.org/10.1002/9781118933732.ch22 https://doi.org/10.1002/9781118933732.ch22.

[17] Rachel Daigle. 2003. Student Workers: The Heart of the Help Desk. In *Proceedings of the 31st Annual ACM SIGUCCS Fall Conference*. Association for Computing Machinery, New York, NY, USA, 193–195. https://doi.org/10.1145/947469.947520

[18] Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2021. Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations. *Transactions of the Association for Computational Linguistics* 10 (Oct. 2021), 92–110. https://doi.org/10.1162/tacl_a_00449 arXiv:2110.05719

[19] Hannah Davis. 2020. A Dataset Is a Worldview. *Medium* (2020).

[20] Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. On the Genealogy of Machine Learning Datasets: A Critical History of ImageNet. *Big Data and Society* 8, 2 (Sept. 2021), 205395172110359. https://doi.org/10.1177/20539517211035955

[21] Qwo-Li Driskill. 2016. *Asegi Stories: Cherokee Queer and Two-Spirit Memory.*

[22] E. Fosch-Villaronga, A. Poulsen, R. A. Søraa, and B. H.M. Custers. 2021. A Little Bird Told Me Your Gender: Gender Inferences in Social Media. *Information Processing and Management* 58, 3 (May 2021). https://doi.org/10.1016/J.IPM.2021.102541

[23] Batya Friedman. 1996. Value-Sensitive Design. *Interactions* 3, 6 (Dec. 1996), 16–23. https://doi.org/10.1145/242485.242493

[24] Anushri Ghode. 2019. Data Work by Frontline Health Workers in Pregnancy Care. In *Proceedings of the 10th Indian Conference on Human-Computer Interaction*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3364183.3364201

[25] Mary L Gray and Suri Siddharth. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Vol. 53. arXiv:1011.1669v3

[26] Elizabeth M. Grieco and Rachel C. Cassidy. 2001. *Overview of Race and Hispanic Origin: Census 2000 Brief.* Technical Report.

[27] Margot Hanley, Solon Barocas, Karen Levy, Shiri Azenkot, and Helen Nissenbaum. 2021. Computer Vision and Conflicting Values: Describing People with Automated Alt Text. In *AIES 2021 - Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, Inc, 543–554. https://doi.org/10.1145/3461702.3462620 arXiv:2105.12754

[28] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2019. Towards a Critical Race Methodology in Algorithmic Fairness. *FAT** (Dec. 2019). https://doi.org/10.1145/3351095.3372826 arXiv:1912.03593

[29] Amy Hawkins. 2017. KFC China Is Using Facial Recognition Tech to Serve Customers - but Are They Buying It?

[30] Päivi Heikkilä, Anita Honka, and Eija Kaasinen. 2018. Quantified Factory Worker: Designing a Worker Feedback Dashboard. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*. Association for Computing Machinery, New York, NY, USA, 515–523. https://doi.org/10.1145/3240167.3240187

[31] Kenneth Holstein, Hal Daumé III, Miroslav Dudík, Hanna Wallach, and Jennifer Wortman Vaughan. 2019. *Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?* Technical Report. ACM. 16 pages. https://doi.org/10.1145/3290605.3300830 arXiv:1812.05239v1

[32] Lilly Irani, Niloufar Salehi, Joyojeet Pal, Andrés Monroy-Hernández, Elizabeth Churchill, and Sneha Narayan. 2019. Patron or Poison? Industry Funding of HCI Research. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*. Association for Computing Machinery, 111–115. https://doi.org/10.1145/3311957.3358610

[33] Os Keyes. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 1–22. https://doi.org/10.1145/3274357

[34] S L Kogan and M J Muller. 2006. Ethnographic Study of Collaborative Knowledge Work. *IBM Systems Journal* 45, 4 (2006), 759–771. https://doi.org/10.1147/sj.454.0759

[35] Kiely Kuligowski. 2019. Facial Recognition Advertising Targets Customers.

[36] Christopher A. Le Dantec, Erika Shehan Poole, and Susan P. Wyche. 2009. Values as Lived Experience: Evolving Value Sensitive Design in Support of Value Discovery. In *Conference on Human Factors in Computing Systems - Proceedings*. ACM Press, New York, New York, USA, 1141–1150. https://doi.org/10.1145/1518701.1518875

[37] Lan Li, Tina Lassiter, Joohee Oh, and Min Kyung Lee. 2021. Algorithmic Hiring in Practice: Recruiter and HR Professional’s Perspectives on AI Use in Hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES ’21)*. Association for Computing Machinery, New York, NY, USA, 166–176. https://doi.org/10.1145/3461702.3462531

[38] Calvin A. Liang, Sean A. Munson, and Julie A. Kientz. 2021. Embracing Four Tensions in Human-Computer Interaction Research with Marginalized People. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 2 (April 2021), 14. https://doi.org/10.1145/3443686

[39] Wendy E MacKay. 1999. Is Paper Safer? The Role of Paper Flight Strips in Air Traffic Control. *ACM Trans. Comput.-Hum. Interact.* 6, 4 (Dec. 1999), 311–340. https://doi.org/10.1145/331490.331491

[40] Noëmi Manders-Huits. 2011. What Values in Design? The Challenge of Incorporating Moral Values into Design. *Science and Engineering Ethics* 17, 2 (2011), 271–287. https://doi.org/10.1007/s11948-010-9198-2

[41] Marian Mazzone and Ahmed Elgammal. 2019. Art, Creativity, and the Potential of Artificial Intelligence. *Arts* 8, 1 (March 2019), 26. https://doi.org/10.3390/arts8010026

[42] Janis Lena Meissner, Nicolas Pretterhofer, Nadja Bergmann, and Edeltraud Haselsteiner. 2022. The Hidden Technological Labour of Service Workers in Health and Beauty Shops. In *2022 Symposium on Human-Computer Interaction for Work*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3533406.3533413

[43] Jacob Metcalf, Emanuel Moss, and Danah Boyd. 2019. Owning Ethics: Corporate Logics, Silicon Valley, and the Institutionalization of Ethics. *Social Research: An International Quarterly* 86, 2 (2019), 449–476. https://doi.org/10.1353/SOR.2019.0022

[44] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. (May 2022). https://doi.org/10.48550/arxiv.2205.11963 arXiv:2205.11963

[45] 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), 25. https://doi.org/10.1145/3415186

[46] T Moses, J. & Knutsen. 2019. Ways of Knowing: Competing Methodologies in Social and Political Research.

[47] 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*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3290605.3300356

[48] PBS. 2015. A Map of Gender-Diverse Cultures. http://www.pbs.org/independentlens/content/two-spirits_map-html/.

[49] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where Responsible AI Meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1 (April 2021). https://doi.org/10.1145/3449081

[50] Meredith Roaten. 2022. Report: U.S., China in Artificial Intelligence Arms Race. *National Defense Magazine* (2022).

[51] Kristina Rolin. 2009. Standpoint Theory as a Methodology for the Study of Power Relations. *Hypatia* 24, 4 (2009), 218–226. https://doi.org/10.1111/j.1527-2001.2009.01070.x

[52] Johnny Saldana. 2013. The Coding Manual for Qualitative Researchers Instruction. *Sage* (2013), 299.

[53] Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A Smith. 2022. Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics:*

Human Language Technologies. Association for Computational Linguistics, Seattle, United States, 5884–5906. https://doi.org/10.18653/v1/2022.naacl-main.431

[54] Morgan Klaus Scheuerman, Alex Hanna, and Emily Denton. 2021. Do Datasets Have Politics? Disciplinary Values in Computer Vision Dataset Development. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (Oct. 2021). https://doi.org/10.1145/3476058

[55] Morgan Klaus Scheuerman, Madeleine Pape, and Alex Hanna. 2021. Auto-Essentialization: Gender in Automated Facial Analysis as Extended Colonial Project:. *https://doi.org/10.1177/20539517211053712* 8, 2 (Dec. 2021), 205395172110537. https://doi.org/10.1177/20539517211053712

[56] Morgan Klaus Scheuerman, Jacob M Paul, and Jed R Brubaker. 2019. How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis and Image Labeling Services. 144 (2019), 33. https://doi.org/10.1145/3359246

[57] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1 (2020).

[58] Irving Seidman. 1998. *Interviewing as Qualitative Research: A Guide for Researchers in Education and the Social Sciences*. Teachers College Press.

[59] Phoenix A Singer. 2020. Colonialism, Two-Spirit Identity, and the Logics of White Supremacy. (2020), 1–7.

[60] Mona Sloane, Emanuel Moss, and Rumman Chowdhury. 2022. A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability. *Patterns* 3, 2 (Feb. 2022), 100425. https://doi.org/10.1016/j.patter.2021.100425 arXiv:2106.12403

[61] Stephen C. Slota, Kenneth R. Fleischmann, Sherri Greenberg, Nitin Verma, Brenna Cummings, Lan Li, and Chris Shenefiel. 2020. Good Systems, Bad Data?: Interpretations of AI Hype and Failures. *Proceedings of the Association for Information Science and Technology* 57, 1 (Oct. 2020). https://doi.org/10.1002/pra2.275

[62] Eliza Strickland. 2019. IBM Watson, Heal Thyself: How IBM Overpromised and Underdelivered on AI Health Care. *IEEE Spectrum* 56, 4 (April 2019), 24–31. https://doi.org/10.1109/MSPEC.2019.8678513

[63] Lucy Suchman. 1993. Do Categories Have Politics? The Language/Action Perspective Reconsidered. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW '93*. Springer Netherlands, 1–14. https://doi.org/10.1007/978-94-011-2094-4_1

[64] Hamed Taherdoost. 2018. Sampling Methods in Research Methodology; How to Choose a Sampling Technique for Research. *SSRN Electronic Journal* (April 2018). https://doi.org/10.2139/ssrn.3205035

[65] Eddie L. Ungless, Björn Ross, and Anne Lauscher. 2023. Stereotypes and Smut: The (Mis)Representation of Non-cisgender Identities by Text-to-Image Models. (May 2023). arXiv:2305.17072

[66] Jennifer Valentino-DeVries. 2020. How the Police Use Facial Recognition, and Where It Falls Short. *The New York Times* (2020).

[67] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018). https://doi.org/10.1145/3173574

[68] Ben Vincent and Ana Manzano. 2017. History and Cultural Diversity. In *Genderqueer and Non-Binary Genders*. Palgrave Macmillan UK, 11–30. https://doi.org/10.1057/978-1-137-51053-2_2

[69] Ding Wang, Shantanu Prabhat, and Nithya Sambasivan. 2022. Whose AI Dream? In Search of the Aspiration in Data Annotation.. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery. https://doi.org/10.1145/3491102.3502121

[70] Fei Wang and Anita Preininger. 2019. AI in Health: State of the Art, Challenges, and Future Directions. *Yearbook of Medical Informatics* 28, 1 (Aug. 2019), 16–26. https://doi.org/10.1055/s-0039-1677908

[71] Tom Warren. 2023. Microsoft and Google Are about to Open an AI Battle. *The Verge* (2023).

[72] Rosa Wevers. 2018. Unmasking Biometrics ' Biases: Facing Gender, Race , Class and Ability in Biometric Data Collection. *Tijdschrift voor Mediageschiedenis* 21, 2 (Nov. 2018), 89–105. https://doi.org/10.18146/tmg21368

[73] Steve Woolgar and Lucy Alice. Suchman. 1989. Plans and Situated Actions: The Problem of Human Machine Communication. *Contemporary Sociology* 18, 3 (1989), 414. https://doi.org/10.2307/2073874

[74] Haonan Zhong, Jiamin Chang, Ziyue Yang, Tingmin Wu, Pathum Chamikara Mahawaga Arachchige, Chehara Pathmabandu, and Minhui Xue. 2023. Copyright Protection and Accountability of Generative AI: Attack, Watermarking and Attribution. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 94–98. https://doi.org/10.1145/3543873.3587321

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009