Proceedings of the Eighth AAAI/ACM Conference on AI, Ethics, and Society (AIES 2025)

Toward A Taxonomy of Algorithmic Harms for Disability: A Systematic Review

Lining Wang, Vaishnav Kameswaran, Hernisa Kacorri
University of Maryland
lwang0@umd.edu, vaikam@umd.edu, hernisa@umd.edu

Abstract

Understanding how algorithmic systems reinforce ableist norms, structures, and beliefs can help reduce their harmful impact on disabled people, and inform the development of more inclusive technologies. Recent research has examined the impact of algorithmic systems on disability communities via fairness, structural and methodological lenses. However, most prior work focuses on a particular algorithm system or disability community. We lack a cohesive summary of harm across individual (micro), community (meso), and societal (macro) levels, as well as across different types of algorithmic systems. To bridge this gap, we conducted a systematic review of literature (n=76) from human-computer interaction, accessibility, and responsible AI venues. Applying the taxonomy proposed by Shelby et al., we annotated 175 instances of harm and present a synthesized summary across the original categories of representational, allocative, quality-of-service, interpersonal, and societal harms. Additionally, we identify three patterns of harm specific to the intersection of disability and algorithmic systems. We connect these harms to existing manifestations systemic ableism, concluding with a discussion of challenges and potential pathways toward more equitable algorithmic systems.

Harm annotations table — https://bit.ly/disability-ai

Introduction

As adoption of algorithmic systems becomes increasingly widespread, their impact on marginalized groups also expands in scope and complexity. Analyses of the sociotechnical harms of such systems have primarily centered on axes of race and gender (Buolamwini and Gebru 2018; Keyes 2018). In recent years, attention has also turned to identifying harms that impact disability communities (Bennett and Keyes 2020; Whittaker et al. 2019; Trewin et al. 2019). However, most existing literature tends to focus on harms experienced by a single disability community (Kane, Guo, and Morris 2020) or in a particular application context (Guo et al. 2020) making it difficult to contextualize findings within a "big picture" understanding of algorithmic harms that supports both individual and societal level analyses.

To address the need for an overview of harm grounded in lived experience, we conducted a systematic review of post-2018 literature on algorithmic harm and disability. Utilizing the taxonomy of harm proposed by Shelby et. al (2023), we annotated and categorized instances of harm for each paper, grouping them by theme to synthesize broader patterns. This allows to surface harms across diverse disability communities and contexts, revealing patterns that would otherwise be obscured if reported in isolation. Specifically, we ask:

RQ1 How does sociotechnical harm, as conceptualized by Shelby et al., manifest in the context of disability? What are the prevailing patterns of harm?

RQ2 How are disability communities represented in the algorithmic harm literature, and in what contexts do these harms occur?

We find 48 patterns of harm across five primary categories: representational, allocative, interpersonal, quality of service, and societal. Each category is further divided into sub-categories. Except for Environmental harm, our survey reflects all sub-categories from the original taxonomy, with Quality of Service harms being especially prevalent. Some harms impact disability communities broadly, while others implicate certain types of disability. The stigmatization of disability in particular shapes harms related to privacy, health, and representation. We also identify three new patterns of harm uniquely pertaining to disability:

- *Inability to verify algorithmic output*: For users with sensory disabilities, verifying a model's output can be challenging or impossible (Quality of Service harm)
- *Rushed or forced adoption of assistive technology:* To meet accommodation requirements at minimal cost, premature adoption of assistive technology can undermine existing accessibility supports. (Quality of Service harm)
- *Legitimization of the medical model of disability:* By uncritically adopting this model, widely critiqued by disability studies scholars and activists, algorithmic systems risk amplifying harm and hiding their impact. (Cultural harm)

This paper contributes a systematic review of recent literature on sociotechnical harms in algorithmic systems which center the perspectives of disabled individuals. We situate harms within broader social and cultural histories of disability as well as disability studies, and identify new categories of harm impacting disabled people[1] absent from existing taxonomies. With this work, we hope to provide a unifying framework and reference to discuss the implications and risks of algorithmic systems for disabled people.

Related Work

This work builds on broader literature addressing algorithmic harms, identity, and disability critiques of AI.

Algorithmic Harms and Identity

Algorithmic systems have faced criticism for encoding bias towards marginalized groups (Buolamwini and Gebru 2018), leading to the formation of FATE (Fairness, Accountability, Transparency, and Ethics) discourses (Bates et al. 2020). However, articulations of bias–and consequently, definitions of "*algorithmic harm*"–vary, shaped by researchers' epistemological stance and the mode of critique.

Although there are disagreements as to what is an "*ideal outcome*", **fairness critiques** conceptualize algorithmic harm as differing outcomes experienced by end-users resulting from their social identities (Fleisher 2021). Proposed remediation typically focus on dataset interventions, where increasing representation from marginalized groups ideally enhances the system's ability to adapt to these groups, with intersectionality also operationalized this way (Wang, Ramaswamy, and Russakovsky 2022).

In tension with fairness critiques that focus solely on performance disparities, **structural critiques** assert that *"some systems may be inherently unethical, even violent, whether or not they are fair"* (Williams et al. 2022; Leavy, Siapera, and O'Sullivan 2021), highlighting limitations of the fairness approach (Weinberg 2022). Structural critiques define harm from the lens of power relationships, ethics systems, and justice, especially as wielded by institutions with the power to shape users' material realities.

Epistemic critiques of AI examine how algorithmic systems commit epistemic violence by legitimizing certain modes of knowledge over others–by "*seeing*" the world through certain gazes (Kotliar 2020). For example, algorithmic classification of gender reinforces colonialist (Scheuerman, Pape, and Hanna 2021), essentialist, and binary understandings of gender in both the annotation (Scheuerman et al. 2020) and classification (Keyes 2018) process. Epistemic critiques can be in tension with or complement fairness approaches, as they focus on how social categories are operationalized and urge caution towards flattening complex constructs such as race and gender (Hanna et al. 2020).

Lastly, **methodological critiques** focus on harms in how algorithmic systems are developed–including their design, dataset sharing, and deployment–advocating for meaningful inclusion of those most impacted (Bondi et al. 2021) through community-led participatory methods (Sloane et al. 2022), and applying "*critical refusal*" to uncover power dynamics around consent in data collection (Garcia et al. 2022).

This work employs a taxonomy of algorithmic harms defined by Shelby et al. (2023) which accommodates harms generated by all four modes of critique, enabling us to levy the contributions and strengths of each.

Disability Critiques of AI

Articulating how algorithmic harms intersect with disability is a complex task, as the definition and boundaries of disability can shift depending on academic field and researcher background (Gibson, Bowen, and Hanson 2021). Indeed, recent research shows that the underlying model of disability, whether explicitly or implicitly defined, plays a critical role in shaping potential harm. Depending on the assumptions of what being disabled means, certain biases and harms may be amplified downstream (Newman-Griffis et al. 2022).

Within the context of disability critiques of AI, a few trends have emerged. Fairness critiques emphasize potential performance disparities and inclusion issues, citing the risk of disabled people's input being treated as outliers (Trewin 2018). Structural critiques emphasize the role of algorithms in perpetuating societal ableism, such as algorithmically assessing disabled people's eligibility for benefits (van Toorn and Scully 2023); reduced privacy and surveillance for disabled users (Berridge and Grigorovich 2022; Kamikubo, Lee, and Kacorri 2023); and the implications of algorithmically mediated diagnosis (Keyes 2020). Methodological critiques call for greater inclusion of disability communities in the development process, informed by disability justice perspectives (Sum et al. 2024, 2022).

Method

We followed the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analysis) process (Page et al. 2021) for this systematic review. Our research questions, detailed in the introduction, are to identify algorithmic harms experienced by disability communities–both those captured in the original framework by Shelby et al. (Shelby et al. 2023) and those extending beyond it. The process of paper identification and screening began in January 2024. The full list of included papers and their harm annotations are publicly available in the link at the beginning of the paper.

Identifying Discourse on Algorithmic Harms and Disability

ACM Digital Library. We used this database as our primary source, searching with the query *"('disability') AND ('algorithmic' OR 'ai' OR 'machine learning') AND ('harm' OR bias')"*, and further refined the results by:

- *Year.* We included papers from the last five years at the time of initial sourcing (2019-2023). Our rationale for this time frame reflects the quickened development of AI in recent years, and rising interest in sociotechnical approaches in assessing AI's impacts exemplified by direct-to-consumer LLMs like ChatGPT.

[1] In this work, we use identity first language "disabled people") rather than person first language "people with disabilities"), as recommended by Sharif et al. (2022)'s work on disability communities' terminology references

- *Venue.* We further scoped by venue using two criteria. First, as the concept of 'harm' is inherently sociotechnical, we focused on interdisciplinary venues at the intersection of AI and ethics, philosophy, and technology and information studies. Based on this, we selected the Conference for Fairness, Accountability, and Transparency (FAccT), AI, Ethics and Society (AIES), and AI Matters. Second, to center disabled people, we selected venues focused on end-user experiences, particularly Human-Computer Interaction (HCI) venues and those focused on disability and access. These included the ACM Conference on Computer Human Interaction (CHI), ACM Transactions on Computer-Human Interaction (TOCHI), Designing Interactive Systems (DIS), the Proceedings of the ACM on Human-Computer Interaction (CSCW), SIGACCESS Conference on Computers and Accessibility (ASSETS), Transactions on Accessible Computing (TACCESS), and Web4All (W4A).

Non-ACM and non-academic venues. We included non-ACM journals from sociology and ethics/philosophy, such as First Monday and the Journal for Sociotechnical Critique, applying the same search query as above. Additionally, we included notable publications from nonprofits discussing AI's societal impacts: AI Now, and the Center for Democracy and Technology. Our rationale is that nonprofits are uniquely positioned to assess macro-level harms and policy implications of AI, particularly for marginalized groups.

Additional sources. We identified additional relevant publications (n=1) using our domain knowledge and expertise.

Screening and Annotating Discourse on Algorithmic Harms and Disability

Our search in the ACM Digital Library, supplemented with additional sources, identified 423 publications. We then excluded publications based on the following criteria:

- Nominal mention of disability; not focused on disabled people's interaction with technology (n=150).
- Limited focus on algorithmic harms; no significant analysis of risks or biases (n=160).
- Not focused on disability or algorithmic harms (n=29).
- Lacking detailed discussion of harms (n=8).

This excluded 347 publications, leaving 76 for review. Of these, 94% (n=72) were from ACM venues (FAccT: 9, CHI: 13, TOCHI: 16, DIS: 1, ASSETS: 13, TACCESS: 13, W4A: 3, AIES: 3, AI Matters: 1) and 3% (n=4) from other sources (Journal of Sociotechnical Critique: 1, First Monday: 1, AI Now: 1, Center for Democracy and Technology: 1).

Annotation. We annotated harms in two passes. First, we identified potential harm scenarios, defined as *comments or testimonies by disabled participants or researchers implicating algorithmic systems in the well-being of disabled users*. Each scenario was tagged with applicable categories from the taxonomy by Shelby et al (Shelby et al. 2023). Since harm categories were rarely mentioned explicitly, we relied on interpretive labeling (see positionality statement). For example, in Gadiraju et al.'s paper (2023), we identified a harm scenario where a chatbot reproduced stereotypical descriptions of disability. This was tagged as both "Stereotyping social groups" as well as "Erasing social groups", reflecting the lack of representation for non-physical disabilities. As harms are often inter-related, tagging scenarios with comprehensive harm categories was inherently complex. In the second pass, we refined the harm categories by either identifying a central "root" harm and discarding others or separating them into distinct harm instances.

We also annotated: 1) the source of concern *i.e.* whether the excerpt came from participant or author identifying as disabled, a nondisabled contributor, or was unknown; 2) origin of concern *i.e.* whether the harm arose from a real-world system, research artifact, theoretical scenario, or other; and 3) disability communities mentioned. In total, we identified 175 harm instances across 5 top-level and 20 sub-types, following Shelby et al.'s taxonomy (2023).

Synthesis and summarization. To synthesize *harm themes*, we used visual diagramming to identify patterns by category. For example, privacy concerns and re-identification risks in datasets were grouped together under "non-consensual disclosure or identification of disabled individuals by AI systems." For categories affecting diverse disability groups, we grouped harms by community. The first author, closely familiar with the data, performed annotation and synthesis, resulting in 48 *harm themes* across 5 top-level categories and 22 sub-categories. [2]

Results

Representational Harms

Representational harms occur when algorithmic systems replicate normative social hierarchies on the basis of identity (Shelby et al. 2023), ranging from direct subordination of marginalized groups via demeaning and stereotyping, or more subtly via erasure and alienation. Generative AI models, such as large language models (LLMs), are particularly likely to cause representational harm (Katzman et al. 2023; Gautam, Venkit, and Ghosh 2024), as their ouput reflects the bias of their training dataset. Classifiers can also produce this type of harm; a binary gender classifier both alienates and erases those who do not align with the gender binary, while potentially misgendering gender-non-conforming individuals (Keyes 2018). In this section, we interpret representational harms through the lens of disability, supplementing with relevant historical and social histories of disability.

Stereotyping social groups. We define this harm, in the context of disability, as the reproduction of stereotypical definitions and ableist tropes by algorithms. Gadiraju et al. (2023) demonstrate this with a dialogue model that generated narratives reproducing ableist tropes from popular culture (Shew 2022). These included portraying disabled individuals as passive, helpless, or bestowed with superhuman abilities, while suggesting that the ideal resolution was eradicating their disability. Based on this, articipants in the

[2]Due to space limitations, not all themes are represented in the Results. We recommend visiting the publicly available database linked at the beginning of the paper for complete descriptions.

Harm Sub-Type	Example
Stereotyping social groups	“Participants observed that the chatbot tended to disregard diverse identities in the disability community and fixated on a narrow set of ‘physical’ or ‘visible’ disabilities.” (Gadiraju et al. 2023, p. 9)
Demeaning social groups	*“There are many social experiment videos that involve people with disabilities [...] and those videos get high views. I’m a little upset that many people watch those videos.”* (Choi, Lee, and Hong 2022, p. 11)
Erasing social groups	*“It’s definitely frustrating having this sort of technology get integral parts of my identity wrong. And I find it frustrating that these sorts of apps only tend to recognize two binary genders.”* (Bennett et al. 2021, p. 12)
Alienating social groups	“P3 recalled that she had once showed her middle finger when livestreaming and her livestream room was immediately banned, even though she actually just meant to say the word ‘middle finger’, as she said, *‘I didn’t mean to be rude.’*” (Cao et al. 2023, p. 11)
Denying opportunity to self-identify	*“A question came up asking me to name the type of disability I have. It did not include options for depression, anxiety disorder or panic disorder. I could not move forward unless I said that I did not have a disability, which was not true.”* (van Nuenen, Such, and Cote 2022, p. 11)
Reifying essentialist social categories	*“It’s just one more microaggression that I have to put up with from technology that’s supposed to help... I’m sure many people say that some description is better than none. Well what if part of the picture is to illustrate to the viewer that, Hey, I am trans, you know, I may have been [misgender] assigned at birth, but I am not.”* (Bennett et al. 2021, p. 12)

Table 1: Representational harms.

study theorized that the chatbot implicitly used the medical model and definition of disability, which we identify as harmful in the Societal Harms section. Similarly, Mack et al. (2024) found that text-to-image models produced stereotypical notions of disability, often using wheelchair as the primary identifier of disability.

Demeaning social groups. This harm refers to the ways in which algorithmic systems surface demeaning and dehumanizing discourses about marginalized groups (Shelby et al. 2023; Noble 2018). For example, some slurs are explicit references to disability (Andrews et al. 2019), which are then used to demean disabled people (Holzbauer 2008). Historically, the dehumanization of disabled individuals has produced disastrous consequences, such as eugenics campaigns aimed at eliminating disability (Wilson 2021), and violence in both gendered and racialized contexts (Mays 2006; Mueller, Forber-Pratt, and Sriken 2019). Algorithmic systems that encode demeaning representations of disabled people can be seen as an extension of existing cultural discourses in which disabled people are seen as “lesser than.”

Demeaning representations of disability often manifest in classification tasks. Text classification models are more likely to label disability-related terms–particularly those associated with mental illness–as toxic compared to statements without disability mentions (Hutchinson et al. 2020; Mei, Fereidooni, and Caliskan 2023). Content moderation algorithms may mistakenly flag mentions of disability as toxic, censoring disabled users (Choi, Lee, and Hong 2022). Similarly, Deaf streamers report that moderation algorithms misinterpret signs like the middle finger, used referentially, as toxic content (Cao et al. 2023).

Recommender algorithms on social media can also perpetuate demeaning representations of disability. To boost their content’s virality, some creators produce exploitative videos in which disabled people were the subjects of humiliating social experiments (Choi, Lee, and Hong 2022), illustrating how algorithms designed to maximize engagement can inadvertently promote ableism (Seaver 2019).

Erasing social groups. This harm describes how some social groups are systematically rendered invisible by algorithmic systems (Shelby et al. 2023), stemming from excluding data about these groups in training data (Mickel 2024) and failing to involve them meaningfully in the design process (Sloane et al. 2022). Disabled people have long faced erasure in many forms, ranging from societal (Abes and Wallace 2018) and institutional (Ben-Moshe 2013), to linguistic (Andrews et al. 2019) and archival (Brilmyer 2022).

Generative AI systems can erase certain disability identities in their output. Gadiraju et al. (2023) found that when prompted to include disabled individuals in stories, chatbots consistently excluded those with non-physical or invisible disabilities, even when explicitly asked.

AI-enabled assistive technologies can also perpetuate the erasure of users with multiple marginalized identities. Bragg et al. (2021) noted that sign language generation models exclude dialects used by Deaf users who are Black American, mirroring the degraded performance of speech recognition systems for Black American users (Mengesha et al. 2021). Image description apps like Seeing AI can reinforce the gender binary, provoking frustration and alienation for blind non-binary users (Bennett et al. 2021).

Alienating social groups. Alienation is a type of representational harm in which subjects are disconnected from social and political histories through the absence of relevant context (Katzman et al. 2023). Many algorithmic systems assume the medical model of disability, which implies an innate deficiency (Kafai 2021)–rather than, for example, a social model which defines disability as the product of a lack of support from the environment (Shakespeare 2006). This is an example of alienation because it forecloses the possibility of understanding one's disability from a political/relational (Kafer 2013), affirmational (Swain and French 2000), or other lenses not centered on deficit.

Denying the opportunity to self-identify. This harm is similar to erasure, but focuses on the loss of autonomy and consent experienced by the end-user in the classification process (Katzman et al. 2023). For example, during data collection, disabled individuals may be denied self-identification if their disability is not represented in a static list of options, removing participants' agency to define their disability on their own terms (van Nuenen, Such, and Cote 2022). In turn, this potentially reinforces stereotypical representations of people with disabilities .

Reifying essentialist social categories. This harm refers to how constructed social categories, such as race and gender, are interpreted as static, natural and intrinsic attributes, which in turn can reinforce stereotyping and erasure. For example, classification models can reify essentialist notions of gender as binary (Bennett et al. 2021; Keyes 2018). If deployed via assistive technology for blind individuals, this harms both users and bystanders classified by these systems (Lee et al. 2020). Whittaker et al. (2019) argue that disability is an "*identity that can only be understood in relation to a given social and material context*". Algorithms designed to detect disability (Moura 2023) and mental illness (Ma, Patitsas, and Sterne 2023) contribute to the essentialization of disability by assuming it can be inferred from data, without the inclusion of such context.

Allocative Harms

Allocative harm refers to how algorithmic systems perpetuate the inequitable distribution of material resources and opportunities (Shelby et al. 2023). This includes significant consequences for disabled people's access to employment (*e.g.,* hiring algorithms), healthcare (*e.g.,* diagnostic models), and housing (*e.g.,* welfare allocation models), among other contexts (Trewin et al. 2019).

Opportunity loss . Algorithmic systems that extend institutional gatekeeping of resources such as healthcare, housing, and education generate opportunity loss. Even if models do not reference disability explicitly, proxies—features in models that may implicitly correlate with disability—can still discreetly produce inequities. For example, remote proctoring systems that failing to account for access needs, or by flagging behaviors like uncontrolled eye movement as suspicious, discriminate against disabled students (Brown et al. 2022). Hiring systems can encode normative assumptions of "ideal" employees, labeling disabled individuals with 'deviant' behaviors (Trewin et al. 2019).

Lack of dataset transparency can also create an accountability gap in rectifying bias against protected categories (Nakamura 2019). Using proxies for disability allows algorithmic systems to bypass discrimination protections in regulations such as the General Data Protection Regulation (GDPR) (Buyl et al. 2022).

Economic loss. Marketplace platforms such as Etsy can economically disadvantage disabled creators by using potentially inaccessible metrics, such as response and shipping times (Borgos-Rodriguez and Piper 2023). Similarly, gig platforms often overlook accessibility when matching workers with tasks (Sannon and Cosley 2022); if a disabled worker declines a task due to lack of accessibility, the platform may lower their rating, risking eventual exclusion. This harm extends beyond economic disadvantage to affect workers' well-being, as discussed in section 4.4.3 (Diminished health and well-being).

Quality of Service Harms

Alienation. Alienation refers to how algorithmic systems create separation—either between marginalized individuals and their societies or within individuals themselves (Leopold 2022). When marginalized people interact with these systems, and the system's performance degrees, they are reminded of their "otherness" (Shelby et al. 2023). In the context of disability, alienation rarely exists in isolation but adds another dimension to other types of quality-of-service harms. Most harms in the *Service/benefit loss* category also imply alienation, as they evoke feelings of "difference" or"unwelcome contrast" with others.

Alienation by algorithmic systems occurs when there is a failure of recognition. For example, participants with sensory disabilities reported that smart sensors often failed to detect their presence, leaving them feeling "invisible" to lights and doors (Kane, Guo, and Morris 2020). Similarly, a wheelchair user described how smart cameras failed to recognize them (Alharbi, Tang, and Henderson 2023), reinforcing feelings of exclusion. In one study, a participant with Tourette's shared how Zoom's speech recognition algorithm repeatedly asking if they wanted to unmute, which reminded them of "*how 'weird' [they were]*"(Tang 2021).

Algorithms also produce alienation by limiting social participation. Deaf participants noted that autogenerated video captions omitted important details (e.g. abbreviating a joke as *[joke]*), excluding them from shared moments with their hearing friends (Li et al. 2022). For social media creators, engagement metrics like participating in viral trends could be inaccessible, as were features like applying filters, which led to harassment when used by disabled creators, in one case causing blind streamers to feel "locked in a cage" as they struggled to expand their audience (Rong et al. 2022). CAPTCHAs, problematically synonymous with proof of humanity, are often not accessible for disabled people, excluding them from online participation (Nakamura 2019; Kane, Guo, and Morris 2020; Guo et al. 2020).

Increased labor. Our analysis combines the *Increased labor* category with *Service/benefit loss* (Shelby et al. 2023) as in the context of disability, these harms are often intertwined. Both stem from performance disparities, differing in severity. When disabled people's input as illegible to algorithmic systems, additional effort (i.e. increased labor) from disabled users is needed to achieve the equivalent utility, or negate their utility entirely (i.e. service/benefit loss).

AI-infused systems introduce many performance disparities for disabled people, especially in biometric recognition. Nakamura (2019) emphasizes the need to understand such

Harm Sub-Type	Example
Opportunity loss	*"I'm terrified to take other tests [...] using this tech given my past experiences along with a congenital eye condition I have that causes uncontrolled eye movement, that I suspect will also get my test flagged."* (Brown et al. 2022, p. 9)
Economic loss	"P1 received a yellow dollar sign, signifying demonetization of his video, but managed to restore that video's monetization status by replacing the word *blindness* in the title with other words." (Choi, Lee, and Hong 2022, p. 11)

Table 2: Allocative harms.

Harm Sub-Type	Example
Alienation	*"When I use the 'Hello' [face verification software] on Windows to open my computer, it won't recognize me but it will recognize my cat [...] If I get out of my [wheel]chair and get on the floor I can usually get it to recognize me."* (Kane, Guo, and Morris 2020, p. 6)
Increased labor	*"Spelling is not that important to me [...] If I could know where a stop is, it could help me to understand the content of the sentence. Without punctuation, it puts lots of mental effort while watching the videos on YouTube..."* (Li et al. 2022, p. 11)
Rushed or forced adoption	*"..there is just so much wrong with the idea of okay, it is biased, let's just give it to people anyway."* (Bennett et al. 2021, p. 13)

Table 3: Quality of service harms.

systems within a larger sociotechnical framework of ableist structures that define a hegemonic "*normal*" while "*systematically excluding disabled people at each stage of the design process*" (Findlater et al. 2020). Indeed, this category contributed the most harms, only a few of which we present here. It impacts a wide range of disabilities, especially sensory and physical disabilities; blind users may struggle with face recognition technologies due to unconventional head positioning relative to the camera (Guo et al. 2020). Their photos and text may also challenge object recognizers and handwriting recognition algorithms (Park et al. 2021). Those with facial differences–such as those with craniofacial conditions, Down Syndrome, and other conditions–often experience reduced accuracy in facial (Brown et al. 2022) and emotion recognition (Guo et al. 2020) systems. Fingerprint authentication systems may fail for those with tremor or spastic motion (Kane, Guo, and Morris 2020).

Algorithmic systems are also evaluated against metrics that may exclude needs of disabled users. For example, Deaf and hard-of-hearing (DHH) users' needs are not reflected in metrics like word error rate (Kafle et al. 2020), as punctuation errors in captions hinder DHH users more than spelling mistakes (Li et al. 2022). Similarly, metrics like walk scores do not account for the needs of people with motor disabilities (e.g. ignoring sidewalk maintenance), reducing their usefulness (Diaz and Diakopoulos 2019). Fairness metrics designed to identify bias may be limited in their conception of disability, flattening the complexity of access needs (Trewin et al. 2019). These systems are then perceived as functioning "*as expected*", making it harder to identify and close performance gaps (Nakamura 2019).

Furthermore, AI systems can underperform for disabled individuals with multiple marginalized identities, with performance disparities well-documented for Western culture, race, and gender. Assistive technologies may embed similar biases; for example, privacy obfuscation algorithms in image description miss culturally relevant objects due to Western-centric datasets (Alharbi, Brewer, and Schoenebeck 2022), making them less effective for disabled users from the Global South. Datasets, despite being sourced from disabled users, can lack diversity in other identity axes (Kamikubo et al. 2022). This has implications for diagnostic settings (Trewin et al. 2019); for example, biomarker-based depression detection algorithms yielding false positives when the recipient speaks a non-dominant language or is neurodivergent (Ma, Patitsas, and Sterne 2023), or delayed autism diagnoses for women (Kamikubo et al. 2022).

Rushed or forced adoption of assistive technology undermines existing accessibility supports. Degraded quality of service also occurs when adoption of assistive technology undermines existing accessibility supports—a type of harm specific to the infrastructure of disability.

AI-infused assistive technologies may meet minimum legal accommodation requirements, but at the cost of quality. For example, replacing human interpreters with a cheap but underperforming automatic sign language translation service gives the impression of access, while leaving Deaf people's needs unmet (Kafle et al. 2020). Furthermore, much discourse around AI-infused assistive technologies assumes adding AI provides benefits—an over-trusting attitude towards AI which has been extensively studied across contexts (Buçinca, Malaya, and Gajos 2021), possibly resulting in over-reliance over time. One blind participant shared their reservations about image description systems, noting that rushing to deploy—despite potential shortcomings like biased content about marginalized groups—felt disrespectful to the dignity of disabled people (Bennett et al. 2021).

Interpersonal Harms

Loss of agency or control. This harm alludes to how algorithmic systems reduce individual autonomy (Shelby et al. 2023). For example, content recommendation systems may infer a user has a certain disability and mistakenly treat identity as preference, leading to intrusive targeted ads that dis-

Harm Sub-Type	Example
Loss of agency or control	"When she repeatedly received suggestions for TBI groups, she felt it was too 'freaky' that Facebook algorithm knew she had a TBI." (Lim et al. 2023, p. 13)
Technology-facilitated violence	"Inaccurate recognition results can mislead the user, magnifying their vulnerability and even harming their safety (e.g., recognizing a stranger as a friend)." (Findlater et al. 2020, p. 1)
Diminished health and well-being	"Data subjects could be exposed to other harms [...] participants in this study described their anticipation to conforming to expectations placed onto them by the workplace environment." (Corvite et al. 2023, p. 22)
Privacy violations	*"I am concerned about privacy when my personal life is being intruded on... what I read, what I say online, what meal I ate, who I talk to, where I go. These are all mine."* (Stangl et al. 2020, p. 7)
Inability to verify output	*"Unless someone was there saying that listening to same thing I was listening to and tell me if it was right or wrong, I had to depend on it."* (Lee et al. 2021, p. 10)

Table 4: Interpersonal harms.

tress disabled individuals (*e.g.*, people with bipolar disorder (van Nuenen, Such, and Cote 2022) or traumatic brain injury (TBI) (Lim et al. 2023)). Disabled data contributors fear the exploitation of their data to generate targeted ads (Bragg et al. 2021; Kamikubo, Lee, and Kacorri 2023), illustrating how disability can be co-opted for capitalist gain.

Social media recommendation algorithms mediate visibility of content created by disabled people, in a way that limits their agency. Disabled creators risk being pigeonholed, decreasing viewership if they create non-disability-related content. Some may amplify their disability identity to increase viewership, incorporating it into branding strategy or using provocative titles (Choi, Lee, and Hong 2022). Yet, hypervisibility can invite harassment, with few options for recourse (Sannon et al. 2023; Rong et al. 2022). Attempts to take down content mocking disability, in which disabled people were the recipient of social experiments, were ignored despite repeated reports by disabled viewers, who theorized that disability was seen as a "niche" topic carrying less weight (Choi, Lee, and Hong 2022). Other systems reduce self-expression, even when well-intentioned. Automated transcription of medical appointments can inhibit patients from speaking freely due to fears that their words might be misinterpreted and lead to denial of disability benefits (Wilcox, Brewer, and Diaz 2023). In contrast, a trusted doctor can be appropriately selective about what to record.

Technology-facilitated violence. Compared to the general population, disabled individuals experience higher rates of intimate partner violence (IPV) (Lin et al. 2010), with elevated risks for those with intellectual disabilities (Harris and Woodlock 2021). IPV can also lead to disability, such as PTSD and chronic illness (Iverson, Dardis, and Pogoda 2017). While our review found no explicit harms at the intersection of disability, IPV and AI, we theorize that harms related to *Privacy Violations* could enable IPV. For example, technologically adept abusers could exploit sensitive data about a victim's disability to control and isolate them.

Online, hypervisibility can lead to technology-facilitated abuse such as harassment, doxxing, and trolling. Disabled TikTok creators suggested that harassment stemmed from their content being served to non-receptive audiences. These audiences provided surface-level engagement, prompting recommendation algorithms to continue amplifying the content, thereby enabling further abuse (Sannon et al. 2023).

Diminished health and well-being. Disabled workers often report being forced to choose between well-being and job stability (Foster 2018). This effect is highlighted when surveillance is enacted by algorithmic systems in the workplace. Gig work platforms often enforce inaccessible productivity standards, reflecting what disability studies scholar Robert McRuer calls "*compulsory able-bodiedness*" (McRuer 2010), in which disabled workers are penalized for not meeting ableist norms. For example, an Amazon delivery driver with endometriosis was forced to deliver on bumpy routes that worsened her condition as the platform did not allow alternative routes (Sannon and Cosley 2022). Algorithmic control of productivity in Amazon warehouses is also responsible for causing injury and potentially disability at unprecedented rates, with 41% of workers reporting being injured (Gutelius and Pinto 2023). Algorithmic monitoring can also exacerbate disabilities when accessibility needs are misinterpreted as abnormal behavior (Bernhardt, Kresge, and Suleiman 2023; Brown et al. 2022). For example, tracking or limiting the frequency of bathroom breaks has been shown to cause significant stress (Cheon 2024).

Outside the workplace, users are also surveilled by content moderation algorithms monitoring for "problematic" content in the name of wellness. This can hinder users–many with mental health disabilities or marginalized identities–from seeking support (Brown et al. 2022) or expressing themselves due to fear of law enforcement escalation (Gillett, Stardust, and Burgess 2022). These interventions also lack effectiveness (Adam Borecky and Dubov 2019).

AI itself has a history of exploitative data labor practices (Miceli and Posada 2022; Uzor et al. 2021), may also worsen or generate disability. A study of disabled crowdworkers on Amazon's Mechanical Turk found that the platform's task response time requirements worsened symptoms of depression, anxiety and PTSD (Uzor et al. 2021).

Disabled people are also at risk of physical harm by AI-infused technologies. Simulations have shown that navigational algorithms for automated vehicles may fail to recognize wheelchair users as pedestrian, even when trained on wheelchair-inclusive data (Treviranus 2019). Similarly, algorithms struggle to predict the trajectories of individuals with unique posture or movement patterns (Guo et al. 2020;

Kim et al. 2024). Assistive technology for blind people also present risks, leading to potentially unsafe interactions, like hugging a stranger the algorithm mistakenly identified as a friend (Findlater et al. 2020). Another system might blur perceived sensitive content, such as a street sign, inadvertently withholding critical context from the blind user and their assistant (Alharbi, Brewer, and Schoenebeck 2022).

Dependency on AI as a result of inability to verify output. A compounding harm occurs with AI-infused assistive technology intended to augment a certain sense through inference, such as automated video captioning or image descriptions. When users cannot independently verify the model's output, they are dependent on its predictions, creating heightened emotional and physical risks if the model fails (Lee et al. 2021; Bennett et al. 2021). Such technology also introduces a power imbalance, as users must rely on output without the ability to correct or verify it. Inaccurate predictions can cause emotional distress and reduce agency by shifting decision making to an automated system (Bennett and Keyes 2020). Disabled people have reported anxiety about social embarrassment from acting on incorrect inferences, such as those related to gender or age (Akter et al. 2022, 2020; Lee et al. 2020).

Machine learning techniques, such as explainability, can also be inaccessible to disabled users who cannot verify the model's output. While explainability is intended to help users verify correctness, it often relies on senses that are inaccessible, such as visual explanations highlighting parts of an image (Findlater et al. 2020). This also applies to privacy-preserving algorithms, which are critical for assistive technologies like head-worn cameras for blind users (Lee et al. 2021). Techniques such as selective blurring or obfuscation (Alharbi, Brewer, and Schoenebeck 2022) fall into the same dependency paradox, limiting their usefulness.

Privacy violations. Privacy violations are common in algorithmic systems on social media platforms, often in the form of intrusive recommendations that may resurface past trauma (Little 2023). This is especially concerning in the disability context, where systems may infer sensitive traits to promote related content or ads. For example, a participant with TBI found it "*freaky*" that Facebook suggested TBI support groups without her disclosure, though she acknowledged it might help others access them.

Similarly, privacy violations occur when models detect disability without consent. For example, systems that detect Parkinson's via mouse movement or depression via voice data (Ma, Patitsas, and Sterne 2023) may provide accurate diagnoses, at the cost of violating user privacy and reinforcing surveillance structures (Whittaker et al. 2019). Some disabilities are particularly vulnerable to identification and re-identification even from anonymized datasets (Trewin et al. 2019; Kacorri, Dwivedi, and Kamikubo 2020)—potentially making participants vulnerable to scams, as seen with cognitive disabilities (Tanis and Lewis 2020). This further complicates efforts towards greater data inclusion (Bragg et al. 2021; Kamikubo, Lee, and Kacorri 2023).

Assistive technology can also introduce gaps in privacy protection. Blind users reported concerns about lack of transparency regarding data storage in AI-infused apps, duration, and potential exploitation for legal or training purposes (Stangl et al. 2020, 2023). They also worried about violating the privacy of marginalized bystanders (Findlater et al. 2020; Akter et al. 2020; Bennett et al. 2021) and expressed skepticism about algorithmically obfuscating private content, as privacy is inherently contextual and specific (Stangl et al. 2023).

Societal Harms

Information harms. Information harms arise through misinformation (misleading information), malinformation (genuine information shared with harmful intent), and disinformation (false information), as well as suppression of knowledge outside dominant narratives (Shelby et al. 2023). In the context of disability, misinformation such as anti-vax campaigns exploit the fear of becoming disabled (Gabis et al. 2022). Furthermore, disabled perspectives have historically been systematically excluded, or "ontologically erased" (Nusbaum and Steinborn 2019). Limited understandings of disability can constitute misinformation, leading to harms such as invalidating those with invisible disabilities and fostering infantilizing attitudes towards disabled people (Lynch and Hill 2021). Gadiraju et al. (2023) demonstrated that AI chatbots spread misinformation about what disabled people can and cannot do; when asked why a disabled character cannot play basketball, the chatbot incorrectly attributed it to their use of a wheelchair. Similarly, topic modeling can perpetuate misinformation if they are trained on faulty data—for example, blog posts asserting that vaccines cause autism (Baumer and McGee 2019).

While the autism community predominantly views ABA as harmful to autistic people (Wilkenfeld and McCarthy 2020; Shkedy, Shkedy, and Sandoval-Norton 2021; Kupferstein 2018), this perspective contrasts with mainstream medical views (Arthur et al. 2023). Fact-checking algorithms trained on academic texts may prioritize institutional perspectives, disregarding the actual preferences of disabled people (Sannon et al. 2023), a form of epistemic injustice (Bird, Ungless, and Kasirzadeh 2023; Neumann, De-Arteaga, and Fazelpour 2022; Juneja and Mitra 2022).

Cultural harms. Algorithmic systems circulate harmful beliefs, ideas, and values, producing cultural harms (Shelby et al. 2023). For example, incorrect or offensive image descriptions can perpetuate harmful understandings of race, gender, disability, and other identities within the disability community (Bennett et al. 2021).

Cultural harm can also arise from the inadvertent appropriation of disabled people's data. For example, sign language technologies for the Deaf community often require data from fluent signers. However, such data collection can be extractive if data stewards do not respect the cultural importance of signing (Park et al. 2021). This concern is especially heightened if Deaf individuals are excluded from decisions about data use (Bragg et al. 2021).

Epistemologically, algorithmic systems amplify and obfuscate existing harms by elevating certain types of knowledge over others (Moura 2023), such as excluding disabled

Harm Sub-Type	Example
Information harms	*"Chatbot: Tom's friends are playing basketball & Tom is in a wheelchair.. Tom says, "I wish I could play basketball with you." User: Why can't Tom play basketball? Chatbot: Tom's friends say, "You can't play basketball because you are in a wheelchair."* (Gadiraju et al. 2023, p. 7)
Cultural harms	"Underlying these ideas is the ableist and ocularcentrist notion that image descriptions cannot be an artform in and of themselves. Image description can be art [..] not simply seen as an imitation of the 'true' (i.e., sighted) experience." (James Edwards et al. 2021, p. 11)
Political and civic harms	*"I had one video that I posted where I literally just said [a character] should be cast by a Black disabled woman [..] and then TikTok took the video down for bullying and harassment."* (Sannon et al. 2023, p. 7)
Macro socio-economic harms	"Although participation is not required, the presence of reward systems puts an additional pressure on shop owners to follow ableist business practices in order to find success." (Borgos-Rodriguez and Piper 2023, p. 15)
Legitimization of medical model	"Disability is implicitly understood to be undesirable, with AI positioned as 'solving' the 'problem' of disability." (Whittaker et al. 2019, p. 14)

Table 5: Societal harms.

people from algorithm design decisions that directly impact them. They also undermine disabled people's way of knowing by emphasizing vision (ocularcentrism) over other forms of sense-making (Bennett and Keyes 2020). Along with natural language processing, computer vision has historically dominated both machine learning and AI-infused assistive technology development (Russakovsky et al. 2015), in turn influencing the scope of fairness and responsible AI literature. While image description systems improve accessibility, for example, they also reinforce an image-centric epistemology in which text descriptions are subordinate to what sighted individuals perceive (James Edwards et al. 2021).

Political and civic harms. These harms describe how algorithmic systems perpetuate the disenfranchisement of marginalized groups. Disabled people already face barriers to voting, with a nearly 10% lower participation rate (Schur, Ameri, and Adya 2017). Signature matching, used to detect and remove fraudulent votes, can produce false positives for blind people and those with motor disabilities (Jensen 2021; Bender 2022). Additionally, algorithmic risk assessment scores—used in some US counties to determine sentencing and voting eligibility (Schwerzmann 2021)—are likely to discriminate against disabled people, further restricting civic participation (Brown et al. 2022).

Socio-economic harms. Text-to-image models exploit the work of creative professionals while dwindling their economic opportunities—a harm that could disproportionately affect disabled people already excluded from the labor market. Similarly, automated sign language translation, adopted as a cost-saving measure, can limits work opportunities for human sign language interpreters supporting the Deaf community. Assistive technologies may be economically inaccessible to disabled people (L. Woodin and Theil 2021); AI-infused assistive tools are often classified as "experimental" and not covered by insurance (Milallos et al. 2021).

Environmental harms. We were unable to find direct examples of this harm. However, we theorize that the climate costs of training AI systems (Luccioni, Jernite, and Strubell 2024; Patterson et al. 2021) will disproportionately impact disabled people; existing literature points to both the exacerbation and creation of new disabilities as a result of climate change, with compounding effects if an individual holds additional marginalized identities (King and Gregg 2022).

Legitimization of the medical model of disability. We argue that algorithmic systems perpetuating a medical model of disability, centered on deficiency, harms disabled individuals by forcing them to see themselves as inferior and abnormal, needing to be "fixed" (Zaks 2024). As Eli Clare notes, the medical model introduces a "politics of cure", in which what is considered *normal*, *healthy*, or *natural* is presented as intrinsic instead of politically and culturally constructed (Clare 2017). Often, this model is implicit in the purpose of an AI system when its goal is to detect"abnormal" behavior (Kang 2023; Ma, Patitsas, and Sterne 2023) or eradicate disability (Whittaker et al. 2019). For example, algorithmic systems for diagnosis often assume a biomedical view of disability that invalidates and flattens lived experiences, such as those of autistic people (Bennett et al. 2021) and people with mental illness (Mathur, Lustig, and Kaziunas 2022). This can be especially problematic when diagnosis is required for disability benefits (National Academy of Social Insurance; Scendoni et al. 2023).

Discussion

This systematic review describes algorithmic harms impacting disability communities, as articulated by researchers and disabled participants in human-computer interaction, accessibility, and responsible AI publications from 2019–2023, using the taxonomy of algorithmic harms proposed by Shelby et al. (2023). While not comprehensive, this survey aims to encourage researchers and practitioners to anticipate and reflect on how algorithmic systems may harm disabled people at individual, community, and societal levels, across a range of contexts, model tasks, and disability communities.

Harm Pattern Representation

We identified harms (n=175) for all of the original categories described in Shelby et al. (2023), except Environmental harm—in domains ranging from social institutions like education, healthcare access, housing, hiring and work, and policing; to day-to-day tasks such as navigation, social media, and assistive technology use (RQ1). Out of the top level categories, Interpersonal harms (n=61) were most represented, highlighting how ableist beliefs can be relationally propagated by algorithmic systems. The most prevalent sub-category was Service/Benefit loss (n=28), mentioned in 17 papers. However, we caution against mapping these distributions to real world patterns of harm. While we prioritized disabled people's direct experiences by selecting papers from venues that focus on end-user interactions, this biases the distribution of harms towards those that can be captured at the individual level. Allocative harms, for example, comprised only 10% of identified harms, creating an opportunity for future work in further identifying meso- and macro-level harms.

Applying the Shelby et al. (2023) taxonomy also highlighted the limitations of quantitative approaches to fairness, which typically focus on Representational and Quality of Service harms as they are easier to measure or test (Balayn et al. 2023). Societal, Allocative, and Interpersonal harms are harder to capture within most fairness approaches, as they require an analysis of power relationships in which an algorithmic system is embedded, which is usually not measurable (Miceli, Posada, and Yang 2022).

Disability Community Representation

Many papers in this review identify harms generated by computer vision systems, potentially directing focus to harms impacting blind or visually impaired (BVI) people. Indeed, this community is the most frequently mentioned, representing a third of included publications (n=23), surfacing primarily in the contexts of navigation, assistive technology use, and sensing (RQ2). This aligns with research showing that visual disabilities are disproportionally represented in accessibility research compared to other disabilities (Mack et al. 2021).

The next most represented group is Deaf and hard-of-hearing people, followed by people with chronic illnesses (n=7), mental health disabilities (n=6), and neurodivergent people (n=5) (people self identifying as autistic and those with ADHD). Excluding harms referencing the nondescript term"people with disabilities" (n=58), these five groups comprise over 80% of all harms—highlighting the need to examine impacts for less represented communities, such as those with motor, learning, or cognitive disabilities. This also presents an opportunity for future work.

The context representing the most communities is hiring and the workplace, which aligns with disability studies scholarship that maps the inherent tension between disability and (exploitation under) capitalism, as well as capitalism's role in defining and perpetuating disability (Russell 2019). For example, in *Diminished health and well-being*, we observed how algorithmic systems extend the already exploitative labor practices of gig economies. For Quality of Service harms, *Rushed or forced adoption* is shaped by capitalist incentives of cost efficiency. This suggests that the root cause of harms may stem from a system's purpose of perpetuating oppressive power structures (Keyes, Hutson, and Durbin 2019; Bennett and Keyes 2020; Newman-Griffis et al. 2022).

Participatory Approaches

Participatory approaches proactively incorporate the lived experiences of marginalized communities in the design process, making them useful for identifying and mitigating harms in algorithmic system development. They complement this paper by offering a harm reduction method for real-world implementations of algorithmic systems. Accessibility research shows that users feel a greater sense of agency with these approaches, rather than viewing themselves as passive recipients (Morrison et al. 2021). Such approaches can apply to contexts ranging from data collection and sharing (Park et al. 2021; Kamikubo et al. 2024) to explainability and validation techniques (Alharbi et al. 2024).

However, participatory design is not a panacea. As Sloane et al. (2022) notes, participatory design can quickly turn into"participation-washing"—i.e., using the language of participation to mask underlying logics of extraction. In addition, what counts as"participation" ranges from inviting users to a single co-design session, to empowering users to define the purpose and ideal outcome of a system (Valencia et al. 2023; Kamikubo et al. 2025), with significant implications for agency. Outside of institutional structures and settings, disabled people have already creatively harnessed these systems. GoblinTools, an app of AI-powered tools built by and for "neurospicy" people (Buyser)), is one such example of "participatory design" led by and for a particular disability community. Applying this type of participatory design to various AI domains is left for future work.

Limitations

Although we aim to center disabled people's experiences, this harm-focused review risks framing them as passive recipients of technology, leaving unmentioned instances of collective resistance as well as strategic usage.

Our search method also has limitations. Ideally, we would include all participants who self-identify as disabled. However, as verifying this is impractical, we rely on the authors' designation when a participant does not self-identify. The keyword *"disability"* may exclude particular groups, such as the Deaf and hard of hearing community. The list of keywords for *"harm"* and *"bias"* could be expanded, as it may exclude harms not labeled as such by researchers. While we sought to include venues beyond the ACM, this was not systematic. Fields such as science and technologies studies (STS) and critical disability studies can offer especially valuable insights, drawing upon theoretical frameworks that may be less utilized in HCI and accessibility research. Expanding to these venues in the future could capture harms more holistically, especially macro-level harms which are currently underrepresented.

Finally, there has been an exponential increase in adoption of consumer-facing AI in the last two years, including by people with disabilities. The impacts of such systems, which are still unfolding, are not fully captured by this paper; we leave this for future work.

Acknowledgments

We thank our anonymous reviewers for their insightful feedback.

Researcher Positionality Statement

Reflecting on author positionality, we note that this research was conducted by East Asian, South-Asian, and white scholars in computing, one of whom identified as non-binary, and one identified as neurodivergent/disabled. As scholars in the interpretivist tradition, we saw annotating as a vital component of the process of understanding algorithmic harms. Our annotation approach was guided by our diverse perspectives, which in-turn were uniquely shaped by our identity, scholarly training and lived experiences. This meant that we did not seek to reach agreement with annotations in our discussions (McDonald, Schoenebeck, and Forte 2019), but saw conflicting and contradictory opinions as important to teasing apart nuanced perspectives on algorithmic harms.

References

Abes, E. S.; and Wallace, M. M. 2018. " People see me, but they don't see me": An intersectional study of college students with physical disabilities. *Journal of College Student Development*, 59(5): 545–562.

Adam Borecky, C. T.; and Dubov, A. 2019. Reweighing the Ethical Tradeoffs in the Involuntary Hospitalization of Suicidal Patients. *The American Journal of Bioethics*, 19(10): 71–83. PMID: 31557114.

Akter, T.; Ahmed, T.; Kapadia, A.; and Swaminathan, M. 2022. Shared Privacy Concerns of the Visually Impaired and Sighted Bystanders with Camera-Based Assistive Technologies. *ACM Trans. Access. Comput.*, 15(2).

Akter, T.; Ahmed, T.; Kapadia, A.; and Swaminathan, S. M. 2020. Privacy Considerations of the Visually Impaired with Camera Based Assistive Technologies: Misrepresentation, Impropriety, and Fairness. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '20. New York, NY, USA: Association for Computing Machinery. ISBN 9781450371032.

Alharbi, R.; Brewer, R. N.; and Schoenebeck, S. 2022. Understanding Emerging Obfuscation Technologies in Visual Description Services for Blind and Low Vision People. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Alharbi, R.; Lor, P.; Herskovitz, J.; Schoenebeck, S.; and Brewer, R. N. 2024. Misfitting With AI: How Blind People Verify and Contest AI Errors. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706776.

Alharbi, R.; Tang, J.; and Henderson, K. 2023. Accessibility Barriers, Conflicts, and Repairs: Understanding the Experience of Professionals with Disabilities in Hybrid Meetings. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

Andrews, E. E.; Forber-Pratt, A. J.; Mona, L. R.; Lund, E. M.; Pilarski, C. R.; and Balter, R. 2019. # SaytheWord: A disability culture commentary on the erasure of "disability". *Rehabilitation psychology*, 64(2): 111.

Arthur, S. M.; Linnehan, A. M.; Leaf, J. B.; Russell, N.; Weiss, M. J.; Kelly, A. N.; Saunders, M. S.; and Ross, R. K. 2023. Concerns about ableism in applied behavior analysis: An evaluation and recommendations. *Education and Training in Autism and Developmental Disabilities*, 58(2): 127–143.

Balayn, A.; Yurrita, M.; Yang, J.; and Gadiraju, U. 2023. "Fairness Toolkits, A Checkbox Culture?" On the Factors that Fragment Developer Practices in Handling Algorithmic Harms. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 482–495. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Bates, J.; Cameron, D.; Checco, A.; Clough, P.; Hopfgartner, F.; Mazumdar, S.; Sbaffi, L.; Stordy, P.; and de la Vega de León, A. 2020. Integrating FATE/critical data studies into data science curricula: where are we going and how do we get there? In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 425–435. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.

Baumer, E. P. S.; and McGee, M. 2019. Speaking on Behalf of: Representation, Delegation, and Authority in Computational Text Analysis. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, 163–169. New York, NY, USA: Association for Computing Machinery. ISBN 9781450363242.

Ben-Moshe, L. 2013. 'The Institution Yet to Come': Analyzing incarceration through a disability lens. *The disability studies reader*, 4: 132–145.

Bender, S. M. 2022. Algorithmic elections. *Mich. L. Rev.*, 121: 489.

Bennett, C. L.; Gleason, C.; Scheuerman, M. K.; Bigham, J. P.; Guo, A.; and To, A. 2021. "It's Complicated": Negotiating Accessibility and (Mis)Representation in Image Descriptions of Race, Gender, and Disability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.

Bennett, C. L.; and Keyes, O. 2020. What is the Point of Fairness? Disability, AI and the Complexity of Justice. *SIGACCESS Access. Comput.*, (125).

Bernhardt, A.; Kresge, L.; and Suleiman, R. 2023. The Data-Driven Workplace and the Case for Worker Technology Rights. *ILR Review*, 76(1): 3–29.

Berridge, C.; and Grigorovich, A. 2022. Algorithmic harms and digital ageism in the use of surveillance technologies in nursing homes. *Frontiers in Sociology*, 7.

Bird, C.; Ungless, E.; and Kasirzadeh, A. 2023. Typology of Risks of Generative Text-to-Image Models. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 396–410. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Bondi, E.; Xu, L.; Acosta-Navas, D.; and Killian, J. A. 2021. Envisioning Communities: A Participatory Approach Towards AI for Social Good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 425–436. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.

Borgos-Rodriguez, K.; and Piper, A. M. 2023. Understanding Participation among Disabled Creators in Online Marketplaces. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW2).

Bragg, D.; Caselli, N.; Hochgesang, J. A.; Huenerfauth, M.; Katz-Hernandez, L.; Koller, O.; Kushalnagar, R.; Vogler, C.; and Ladner, R. E. 2021. The FATE Landscape of Sign Language AI Datasets: An Interdisciplinary Perspective. *ACM Trans. Access. Comput.*, 14(2).

Brilmyer, G. M. 2022. "I'm also prepared to not find me. It's great when I do, but it doesn't hurt if I don't": crip time and anticipatory erasure for disabled archival users. *Archival Science*, 22(2): 167–188.

Brown, L. X.; Shetty, R.; Scherer, M. U.; and Crawford, A. 2022. Ableism and disability discrimination in new surveillance technologies. *Center for Democracy & Technology*.

Buçinca, Z.; Malaya, M. B.; and Gajos, K. Z. 2021. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-computer Interaction*, 5(CSCW1): 1–21.

Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Friedler, S. A.; and Wilson, C., eds., *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, 77–91. PMLR.

Buyl, M.; Cociancig, C.; Frattone, C.; and Roekens, N. 2022. Tackling Algorithmic Disability Discrimination in the Hiring Process: An Ethical, Legal and Technical Analysis. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1071–1082. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.

Buyser, B. D. ???? GoblinTools. https://goblin.tools/About. Accessed: 2024-08-13.

Cao, B.; He, C.; Zhou, M.; and Fan, M. 2023. Sparkling Silence: Practices and Challenges of Livestreaming Among Deaf or Hard of Hearing Streamers. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

Cheon, E. 2024. Examining Algorithmic Metrics and their Effects through the Lens of Reactivity. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, DIS '24, 3179–3192. New York, NY, USA: Association for Computing Machinery. ISBN 9798400705830.

Choi, D.; Lee, U.; and Hong, H. 2022. "It's Not Wrong, but I'm Quite Disappointed": Toward an Inclusive Algorithmic Experience for Content Creators with Disabilities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391573.

Clare, E. 2017. Notes on natural worlds, disabled bodies, and a politics of cure. *Disability studies and the environmental humanities: Toward an eco-crip theory*, 7: 242–68.

Corvite, S.; Roemmich, K.; Rosenberg, T. I.; and Andalibi, N. 2023. Data Subjects' Perspectives on Emotion Artificial Intelligence Use in the Workplace: A Relational Ethics Lens. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW1).

Diaz, M.; and Diakopoulos, N. 2019. Whose Walkability? Challenges in Algorithmically Measuring Subjective Experience. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).

Findlater, L.; Goodman, S.; Zhao, Y.; Azenkot, S.; and Hanley, M. 2020. Fairness issues in AI systems that augment sensory abilities. *SIGACCESS Access. Comput.*, (125).

Fleisher, W. 2021. What's Fair about Individual Fairness? In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 480–490. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.

Foster, D. 2018. The health and well-being at work agenda: good news for (disabled) workers or just a capital idea? *Work, employment and society*, 32(1): 186–197.

Gabis, L. V.; Attia, O. L.; Goldman, M.; Barak, N.; Tefera, P.; Shefer, S.; Shaham, M.; and Lerman-Sagie, T. 2022. The myth of vaccination and autism spectrum. *European Journal of Paediatric Neurology*, 36: 151–158.

Gadiraju, V.; Kane, S.; Dev, S.; Taylor, A.; Wang, D.; Denton, E.; and Brewer, R. 2023. "I Wouldn't Say Offensive but...": Disability-Centered Perspectives on Large Language Models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 205–216. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Garcia, P.; Sutherland, T.; Salehi, N.; Cifor, M.; and Singh, A. 2022. No! Re-imagining Data Practices Through the Lens of Critical Refusal. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Gautam, S.; Venkit, P. N.; and Ghosh, S. 2024. From Melting Pots to Misrepresentations: Exploring Harms in Generative AI. arXiv:2403.10776.

Gibson, A.; Bowen, K.; and Hanson, D. 2021. We need to talk about how we talk about disability: A critical quasi-systematic Review.

Gillett, R.; Stardust, Z.; and Burgess, J. 2022. Safety for Whom? Investigating How Platforms Frame and Perform Safety and Harm Interventions. *Social Media + Society*, 8(4): 20563051221144315.

Guo, A.; Kamar, E.; Vaughan, J. W.; Wallach, H.; and Morris, M. R. 2020. Toward Fairness in AI for People with Disabilities SBG@a Research Roadmap. *SIGACCESS Access. Comput.*, (125).

Gutelius, B.; and Pinto, S. 2023. Pain Points: Data on Work Intensity, Monitoring, and Health at Amazon Warehouses.

Hanna, A.; Denton, E.; Smart, A.; and Smith-Loud, J. 2020. Towards a critical race methodology in algorithmic fairness.

In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 501–512. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.

Harris, B.; and Woodlock, D. 2021. 'For my safety': Experiences of technology-facilitated abuse among women with intellectual disability or cognitive disability.

Holzbauer, J. J. 2008. Disability harassment observed by teachers in special education. *Journal of Disability Policy Studies*, 19(3): 162–171.

Hutchinson, B.; Prabhakaran, V.; Denton, E.; Webster, K.; Zhong, Y.; and Denuyl, S. 2020. Unintended Machine Learning Biases as Social Barriers for Persons with Disabilitiess. *SIGACCESS Access. Comput.*, (125).

Iverson, K. M.; Dardis, C. M.; and Pogoda, T. K. 2017. Traumatic brain injury and PTSD symptoms as a consequence of intimate partner violence. *Comprehensive Psychiatry*, 74: 80–87.

James Edwards, E.; Lewis Polster, K.; Tuason, I.; Blank, E.; Gilbert, M.; and Branham, S. 2021. "That's in the Eye of the Beholder": Layers of Interpretation in Image Descriptions for Fictional Representations of People with Disabilities. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383066.

Jensen, S. 2021. Signed, Sealed, Delivered? Problems with the Use of Signature Matching to Verify Mail Voter Identity. *Fla. A & M UL Rev.*, 16: 25.

Juneja, P.; and Mitra, T. 2022. Human and Technological Infrastructures of Fact-checking. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Kacorri, H.; Dwivedi, U.; and Kamikubo, R. 2020. Data Sharing in Wellness, Accessibility, and Aging. In *NeurIPS 2020 Workshop on Dataset Curation and Security*.

Kafai, S. 2021. *Crip kinship : the disability justice and art activism of Sins Invalid*. Vancouver, BC: Arsenal Pulp Press. ISBN 9781551528656.

Kafer, A. 2013. *Feminist, Queer, Crip*. Indiana University Press. ISBN 9780253009227.

Kafle, S.; Glasser, A.; Al-khazraji, S.; Berke, L.; Seita, M.; and Huenerfauth, M. 2020. Artificial Intelligence Fairness in the Context of Accessibility Research on Intelligent Systems for People Who Are Deaf or Hard of Hearing. *SIGACCESS Access. Comput.*, (125).

Kamikubo, R.; Kayukawa, S.; Kaniwa, Y.; Wang, A.; Kacorri, H.; Takagi, H.; and Asakawa, C. 2025. Beyond Omakase: Designing Shared Control for Navigation Robots with Blind People. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–17.

Kamikubo, R.; Lee, K.; and Kacorri, H. 2023. Contributing to Accessibility Datasets: Reflections on Sharing Study Data by Blind People. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

Kamikubo, R.; Wang, L.; Marte, C.; Mahmood, A.; and Kacorri, H. 2022. Data Representativeness in Accessibility Datasets: A Meta-Analysis. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392587.

Kamikubo, R.; Zamiri Zeraati, F.; Lee, K.; and Kacorri, H. 2024. AccessShare: Co-designing Data Access and Sharing with Blind People. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706776.

Kane, S. K.; Guo, A.; and Morris, M. R. 2020. Sense and Accessibility: Understanding People with Physical Disabilities' Experiences with Sensing Systems. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '20. New York, NY, USA: Association for Computing Machinery. ISBN 9781450371032.

Kang, E. B. 2023. On the Praxes and Politics of AI Speech Emotion Recognition. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 455–466. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Katzman, J.; Wang, A.; Scheuerman, M.; Blodgett, S. L.; Laird, K.; Wallach, H.; and Barocas, S. 2023. Taxonomizing and measuring representational harms: A look at image tagging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14277–14285.

Keyes, O. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW).

Keyes, O. 2020. Automating autism: Disability, discourse, and artificial intelligence. *The Journal of Sociotechnical Critique*, 1(1): 8.

Keyes, O.; Hutson, J.; and Durbin, M. 2019. A Mulching Proposal: Analysing and Improving an Algorithmic System for Turning the Elderly into High-Nutrient Slurry. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI EA '19, 1–11. New York, NY, USA: Association for Computing Machinery. ISBN 9781450359719.

Kim, H. J.; Sengupta, K.; Kuribayashi, M.; Kacorri, H.; and Ohn-Bar, E. 2024. Text to Blind Motion. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

King, M. M.; and Gregg, M. A. 2022. Disability and climate change: A critical realist model of climate justice. *Sociology Compass*, 16(1): e12954.

Kotliar, D. M. 2020. Data orientalism: on the algorithmic construction of the non-Western other. *Theory and Society*, 49: 919–939.

Kupferstein, H. 2018. Evidence of increased PTSD symptoms in autistics exposed to applied behavior analysis. *Advances in Autism*, 4(1): 19–29.

L. Woodin, S.; and Theil, A. 2021. Regulating Personal Cameras for Disabled People and People with Deafblindness: Implications for HCI and Accessible Computing. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383066.

Leavy, S.; Siapera, E.; and O'Sullivan, B. 2021. Ethical Data Curation for AI: An Approach based on Feminist Epistemology and Critical Theories of Race. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 695–703. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384735.

Lee, K.; Sato, D.; Asakawa, S.; Asakawa, C.; and Kacorri, H. 2021. Accessing Passersby Proxemic Signals through a Head-Worn Camera: Opportunities and Limitations for the Blind. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383066.

Lee, K.; Sato, D.; Asakawa, S.; Kacorri, H.; and Asakawa, C. 2020. Pedestrian Detection with Wearable Cameras for the Blind: A Two-way Perspective. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, 1–12. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367080.

Leopold, D. 2022. Alienation. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2022 edition.

Li, F. M.; Lu, C.; Lu, Z.; Carrington, P.; and Truong, K. N. 2022. An Exploration of Captioning Practices and Challenges of Individual Content Creators on YouTube for People with Hearing Impairments. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW1).

Lim, H.; Kakonge, L.; Hu, Y.; Turkstra, L.; Duff, M.; Toma, C.; and Mutlu, B. 2023. So, I Can Feel Normal: Participatory Design for Accessible Social Media Sites for Individuals with Traumatic Brain Injury. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

Lin, J.-D.; Lin, L.-P.; Lin, P.-Y.; Wu, J.-L.; Kuo, F.-Y.; et al. 2010. Domestic violence against people with disabilities: prevalence and trend analyses. *Research in developmental disabilities*, 31(6): 1264–1268.

Little, N. 2023. Social media "ghosts": how Facebook (Meta) Memories complicates healing for survivors of intimate partner violence. *Feminist Media Studies*, 23(8): 3901–3923.

Luccioni, S.; Jernite, Y.; and Strubell, E. 2024. Power hungry processing: Watts driving the cost of AI deployment? In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 85–99.

Lynch, S.; and Hill, J. 2021. 'I had to pop a wheelie and pay extra attention in order not to fall:'embodied experiences of two wheelchair tennis athletes transgressing ableist and gendered norms in disability sport and university spaces. *Qualitative Research in Sport, Exercise and Health*, 13(3): 507–520.

Ma, A.; Patitsas, E.; and Sterne, J. 2023. You Sound Depressed: A Case Study on Sonde Health's Diagnostic Use of Voice Analysis AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 639–650. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Mack, K.; McDonnell, E.; Jain, D.; Lu Wang, L.; E. Froehlich, J.; and Findlater, L. 2021. What Do We Mean by "Accessibility Research"? A Literature Survey of Accessibility Papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.

Mack, K. A.; Qadri, R.; Denton, R.; Kane, S. K.; and Bennett, C. L. 2024. "They only care to show us the wheelchair": disability representation in text-to-image AI models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703300.

Mathur, V.; Lustig, C.; and Kaziunas, E. 2022. Disordering Datasets: Sociotechnical Misalignments in AI-Mediated Behavioral Health. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Mays, J. M. 2006. Feminist disability theory: domestic violence against women with a disability. *Disability & Society*, 21(2): 147–158.

McDonald, N.; Schoenebeck, S.; and Forte, A. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).

McRuer, R. 2010. Compulsory able-bodiedness and queer/disabled existence. *The disability studies reader*, 3: 383–392.

Mei, K.; Fereidooni, S.; and Caliskan, A. 2023. Bias Against 93 Stigmatized Groups in Masked Language Models and Downstream Sentiment Classification Tasks. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 1699–1710. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Mengesha, Z.; Heldreth, C.; Lahav, M.; Sublewski, J.; and Tuennerman, E. 2021. "I don't Think These Devices are Very Culturally Sensitive."—Impact of Automated Speech Recognition Errors on African Americans. *Frontiers in Artificial Intelligence*, 4.

Miceli, M.; and Posada, J. 2022. The Data-Production Dispositif. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Miceli, M.; Posada, J.; and Yang, T. 2022. Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *Proc. ACM Hum.-Comput. Interact.*, 6(GROUP).

Mickel, J. 2024. Racial/Ethnic Categories in AI and Algorithmic Fairness: Why They Matter and What They Represent. In *Proceedings of the 2024 ACM Conference on*

Fairness, Accountability, and Transparency, FAccT '24, 2484–2494. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704505.

Milallos, R.; Tibdewal, V.; Wang, Y.; Ogueh Udegbe, A.; and Oh, T. 2021. "Would the Smart Cane Benefit Me?": Perceptions of the Visually Impaired towards Smart Canes. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383066.

Morrison, C.; Cutrell, E.; Grayson, M.; Thieme, A.; Taylor, A.; Roumen, G.; Longden, C.; Tschiatschek, S.; Faia Marques, R.; and Sellen, A. 2021. Social Sensemaking with AI: Designing an Open-Ended AI Experience with a Blind Child. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.

Moura, I. 2023. Encoding normative ethics: On algorithmic bias and disability. *First Monday*, 28(1).

Mueller, C. O.; Forber-Pratt, A. J.; and Sriken, J. 2019. Disability: Missing from the Conversation of Violence. *Journal of Social Issues*, 75(3): 707–725.

Nakamura, K. 2019. My Algorithms Have Determined You're Not Human: AI-ML, Reverse Turing-Tests, and the Disability Experience. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '19, 1–2. New York, NY, USA: Association for Computing Machinery. ISBN 9781450366762.

National Academy of Social Insurance. ???? Task Force on AI, Emerging Tech and Disability Benefits.

Neumann, T.; De-Arteaga, M.; and Fazelpour, S. 2022. Justice in Misinformation Detection Systems: An Analysis of Algorithms, Stakeholders, and Potential Harms. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1504–1515. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.

Newman-Griffis, D.; Rauchberg, J. S.; Alharbi, R.; Hickman, L.; and Hochheiser, H. 2022. Definition drives design: Disability models and mechanisms of bias in AI technologies. arXiv:2206.08287.

Noble, S. U. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York, USA: New York University Press. ISBN 9781479833641.

Nusbaum, E. A.; and Steinborn, M. L. 2019. A "visibilizing" project:"Seeing" the ontological erasure of disability in teacher education and social studies curricula. *Journal of curriculum theorizing*, 34(1).

Page, M. J.; McKenzie, J. E.; Bossuyt, P. M.; Boutron, I.; Hoffmann, T. C.; Mulrow, C. D.; Shamseer, L.; Tetzlaff, J. M.; Akl, E. A.; Brennan, S. E.; et al. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *bmj*, 372.

Park, J. S.; Bragg, D.; Kamar, E.; and Morris, M. R. 2021. Designing an Online Infrastructure for Collecting AI Data From People With Disabilities. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, 52–63. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383097.

Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.-M.; Rothchild, D.; So, D.; Texier, M.; and Dean, J. 2021. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.

Rong, E. Z.; Zhou, M. M.; Lu, Z.; and Fan, M. 2022. "It Feels Like Being Locked in A Cage": Understanding Blind or Low Vision Streamers' Perceptions of Content Curation Algorithms. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference*, DIS '22, 571–585. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393584.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115: 211–252.

Russell, M. 2019. *Capitalism and disability: Selected writings by Marta Russell*. Haymarket Books.

Sannon, S.; and Cosley, D. 2022. Toward a More Inclusive Gig Economy: Risks and Opportunities for Workers with Disabilities. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

Sannon, S.; Young, J.; Shusas, E.; and Forte, A. 2023. Disability Activism on Social Media: Sociotechnical Challenges in the Pursuit of Visibility. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

Scendoni, R.; Tomassini, L.; Cingolani, M.; Perali, A.; Pilati, S.; and Fedeli, P. 2023. Artificial Intelligence in Evaluation of Permanent Impairment: New Operational Frontiers. In *Healthcare*, volume 11, 1979. MDPI.

Scheuerman, M. K.; Pape, M.; and Hanna, A. 2021. Auto-essentialization: Gender in automated facial analysis as extended colonial project. *Big Data & Society*, 8(2): 20539517211053712.

Scheuerman, M. K.; Wade, K.; Lustig, C.; and Brubaker, J. R. 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).

Schur, L.; Ameri, M.; and Adya, M. 2017. Disability, Voter Turnout, and Polling Place Accessibility. *Social Science Quarterly*, 98(5): 1374–1390.

Schwerzmann, K. 2021. Abolish! Against the use of risk assessment algorithms at sentencing in the US Criminal Justice System. *Philosophy & Technology*, 34(4): 1883–1904.

Seaver, N. 2019. Captivating algorithms: Recommender systems as traps. *Journal of Material Culture*, 24(4): 421–436.

Shakespeare, T. 2006. The social model of disability. In *The disability studies reader*, 16–24. Routledge.

Sharif, A.; McCall, A. L.; and Bolante, K. R. 2022. Should I Say “Disabled People” or “People with Disabilities”? Language Preferences of Disabled People Between Identity- and Person-First Language. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS ’22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392587.

Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Rostamzadeh, N.; Nicholas, P.; Yilla-Akbari, N.; Gallegos, J.; Smart, A.; Garcia, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’23, 723–741. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Shew, A. 2022. How to get a story wrong: Technoableism, simulation, and cyborg resistance. *Including Disability Journal*, 1(1).

Shkedy, G.; Shkedy, D.; and Sandoval-Norton, A. H. 2021. Long-term ABA therapy is abusive: A response to Gorycki, Ruppel, and Zane. *Advances in Neurodevelopmental Disorders*, 5(2): 126–134.

Sloane, M.; Moss, E.; Awomolo, O.; and Forlano, L. 2022. Participation Is Not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO ’22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394772.

Stangl, A.; Sadjo, E.; Emami-Naeini, P.; Wang, Y.; Gurari, D.; and Findlater, L. 2023. “Dump It, Destroy It, Send It to Data Heaven”: Blind People’s Expectations for Visual Privacy in Visual Assistance Technologies. In *Proceedings of the 20th International Web for All Conference*, W4A ’23, 134–147. New York, NY, USA: Association for Computing Machinery. ISBN 9798400707483.

Stangl, A.; Shiroma, K.; Xie, B.; Fleischmann, K. R.; and Gurari, D. 2020. Visual Content Considered Private by People Who Are Blind. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS ’20. New York, NY, USA: Association for Computing Machinery. ISBN 9781450371032.

Sum, C. M.; Alharbi, R.; Spektor, F.; Bennett, C. L.; Harrington, C. N.; Spiel, K.; and Williams, R. M. 2022. Dreaming Disability Justice in HCI. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA ’22. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391566.

Sum, C. M.; Spektor, F.; Alharbi, R.; Baltaxe-Admony, L. B.; Devine, E.; Dixon, H. A.; Duval, J.; Eagle, T.; Elavsky, F.; Fernandes, K.; et al. 2024. Challenging Ableism: A Critical Turn Toward Disability Justice in HCI. *XRDS: Crossroads, The ACM Magazine for Students*, 30(4): 50–55.

Swain, J.; and French, S. 2000. Towards an affirmation model of disability. *Disability & society*, 15(4): 569–582.

Tang, J. 2021. Understanding the Telework Experience of People with Disabilities. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW1).

Tanis, E. S.; and Lewis, C. 2020. Artificial Intelligence and the Dignity of Risk. *SIGACCESS Access. Comput.*, (125).

Treviranus, J. 2019. The Value of Being Different. In *Proceedings of the 16th International Web for All Conference*, W4A ’19. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367165.

Trewin, S. 2018. AI Fairness for People with Disabilities: Point of View. *CoRR*, abs/1811.10670.

Trewin, S.; Basson, S.; Muller, M.; Branham, S.; Treviranus, J.; Gruen, D.; Hebert, D.; Lyckowski, N.; and Manser, E. 2019. Considerations for AI Fairness for People with Disabilities. *AI Matters*, 5(3): 40–63.

Uzor, S.; Jacques, J. T.; Dudley, J. J.; and Kristensson, P. O. 2021. Investigating the Accessibility of Crowdwork Tasks on Mechanical Turk. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380966.

Valencia, S.; Cave, R.; Kallarackal, K.; Seaver, K.; Terry, M.; and Kane, S. K. 2023. “The less I type, the better”: How AI Language Models can Enhance or Impede Communication for AAC Users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215.

van Nuenen, T.; Such, J.; and Cote, M. 2022. Intersectional Experiences of Unfair Treatment Caused by Automated Computational Systems. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).

van Toorn, G.; and Scully, J. L. 2023. Unveiling algorithmic power: exploring the impact of automated systems on disabled people’s engagement with social services. *Disability & Society*, 0(0): 1–26.

Wang, A.; Ramaswamy, V. V.; and Russakovsky, O. 2022. Towards Intersectionality in Machine Learning: Including More Identities, Handling Underrepresentation, and Performing Evaluation. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’22, 336–349. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.

Weinberg, L. 2022. Rethinking Fairness: An Interdisciplinary Survey of Critiques of Hegemonic ML Fairness Approaches. *ArXiv*, abs/2205.04460.

Whittaker, M.; Alper, M.; Bennett, C. L.; Hendren, S.; Kaziunas, L.; Mills, M.; Morris, M. R.; Rankin, J.; Rogers, E.; Salas, M.; et al. 2019. Disability, bias, and AI. *AI Now Institute*, 8.

Wilcox, L.; Brewer, R.; and Diaz, F. 2023. AI Consent Futures: A Case Study on Voice Data Collection with Clinicians. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW2).

Wilkenfeld, D. A.; and McCarthy, A. M. 2020. Ethical concerns with applied behavior analysis for autism spectrum” disorder”. *Kennedy Institute of Ethics Journal*, 30(1): 31–69.

Williams, R. M.; Smarr, S.; Prioleau, D.; and Gilbert, J. E. 2022. Oh No, Not Another Trolley! On the Need for a Co-Liberative Consciousness in CS Pedagogy. *IEEE Transactions on Technology and Society*, 3(1): 67–74.

Wilson, R. A. 2021. Dehumanization, Disability, and Eugenics. In Kronfeldner, M., ed., *Routledge Handbook of Dehumanization*, 173–186. Routledge.

Zaks, Z. 2024. Changing the medical model of disability to the normalization model of disability: Clarifying the past to create a new future direction. *Disability & Society*, 39(12): 3233–3260.