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

Lining Wang
University of Maryland, College Park
College Park, USA
lwang0@umd.edu

Vaishnav Kameswaran
University of Maryland, College Park
College Park, USA
vaikam@umd.edu

Hernisa Kacorri
University of Maryland, College Park
College Park, USA
hernisa@umd.edu

ABSTRACT

Mapping the terrain of harm arising from algorithmic systems supports researchers and practitioners in curtailing potential pitfalls of designing and deploying said systems. For disability communities, understanding the range of harms may help prevent algorithmic systems from perpetuating societal inequities already experienced by disabled people [113][1]. Recent research has critiqued algorithmic systems and their impact on disability communities through fairness, structural, methodological, and epistemic lenses. However, we currently lack a cohesive summary of prevailing harm patterns encompassing micro (individual), meso (community), and macro (societal) levels. Using the taxonomy of harms proposed by Shelby et al. as our organizing structure, we conducted a systematic review of human-computer interaction, accessibility, and responsible AI research to identify and annotate algorithmic harms involving disabled people. We identified 175 unique instances of harm, which we synthesized into 48 harm themes, mapping to 5 top-level categories of representational, allocative, quality-of-service, interpersonal, and societal harms. We attempt to situate these harms within a larger social-historical context of disability, drawing upon disability justice and critical disability studies perspectives.

KEYWORDS

algorithmic harm, sociotechnical harm, machine learning, disability, AI, systematic review

ACM Reference Format:
Lining Wang, Vaishnav Kameswaran, and Hernisa Kacorri. 2024. Toward A Taxonomy of Algorithmic Harms for Disability: A Systematic Review. In *submission (Scholarly Paper)*. ACM, New York, NY, USA, 19 pages.

1 INTRODUCTION

As algorithmic systems for decision making, prediction, and content generation proliferate in recent years, so has the scope and complexity of their impact on marginalized groups. Efforts to anticipate sociotechnical harms of such systems toward marginalized groups have predominantly focused on the axes of race and gender [27][72]. In recent years, attention has also turned to identifying potential harms for disability communities at individual and societal levels [133][127], which include problematic representations of disability in generative models [85], physical safety dangers for people using wheelchairs posed by self-driving algorithms [125], and troubling performance disparities in biometric recognition systems [131].

[1]In this work, we use identity first language ("disabled people") rather than person first language ("people with disabilities") to describe disability identities, as recommended by Sharif et al.'s research exploring disability communities' preferences for terminology.

Scholarly Paper, July 2024, University of Maryland, College Park
2024.

Recent literature has identified taxonomies of algorithmic harm for specific disabilities [69] and application contexts [51], utilizing a variety of critical and epistemological perspectives [17]. However, it is difficult to contextualize each study's findings within a "big picture" understanding of algorithmic harms pertaining to disability more broadly, within a framework that supports both individual and societal level analyses of harm.

In response to the need for an overview of harm grounded in lived experience, we performed a systematic review of recent (2019 and later) literature on algorithmic harm and disability. Using Shelby et. al's taxonomy of sociotechnical harm as the foundational framework [114], we examined each paper for instances of harm, which we annotated and tagged with potential harm categories. Then, we analyzed the harms for each category and grouped those with common themes together, generating a set of synthesized patterns of harm, with the intention of surfacing patterns of harms inclusive of a range of disability communities and contexts. By performing a systematic review within a rigorous taxonomic framework, we hope to reveal patterns of harm that would be otherwise obfuscated or reported in isolation. We also aspire to provide a shared framework and reference from which to discuss the implications, risks, and impacts of algorithmic systems for disability.

We were guided by these research questions:

(1) How does sociotechnical harm, as conceptualized by Shelby et al., apply with regard to disability? What are the prevailing patterns?
(2) How are disability communities represented in algorithmic harm literature? What is the context in which the harm takes place?

Our resulting review finds 48 patterns of harm across 5 top-level categories (representational, allocative, interpersonal, quality of service, and societal), organized into 24 sub-categories. With the exception of Environmental Harm, every sub-category of harm as defined in the original taxonomy is represented in the literature surveyed, with quality of service harms being especially prevalent. Many harms impacted disability communities broadly, while others implicated particular types of disability. In addition, we find that the stigmatization of disability affects vectors of harm significantly, with implications for harms related to privacy, health and well being, and representation. Finally, we articulate three new categories of harm that uniquely pertain to disability that emerged from our analysis:

- *Inability to verify algorithmic output*: For an end user with sensory disabilities, verifying that a model's output is correct can be difficult or impossible–introducing a new class of risk. (Quality of Service Harm)

- *Rushed adoption of assistive technology:* Incentivized to meet legal obligations for accommodation at minimum cost, premature adoption of assistive technology can undermine existing accessibility supports. (Quality-of-Service Harm)
- *Legitimization of the medical model of disability:* The medical model of disability has received much critique from disability studies scholars. Algorithmic systems risk amplifying and making invisible the harms of this model by adopting it uncritically. (Cultural Harm)

The paper makes the following contributions:

- A systematic review of recent literature on sociotechnical harms of algorithmic systems as they relate to disability, centering the perspectives of individuals identifying as disabled;
- Analysis of harm patterns using a disability studies lens, situating harms within social and cultural histories of disability;
- New categories of harm impacting disabled people that are not currently represented in existing taxonomies.

In the following section, we will introduce the concept of algorithmic harm, a summary of the operationalization of disability in sociological and algorithmic contexts, and the current landscape of disability critiques of AI. Section 3 gives an overview of our research methodology, while Section 4 presents the synthesized harm themes in relation to taxonomy described by Shelby et al [114].

2 RELATED WORK

2.1 Algorithmic Harms and Identity

In recent years, algorithmic systems have been criticized for encoding bias towards marginalized groups [27]. Attempts to articulate said bias has led to the emergence of FATE (Fairness, Accountability, Transparency, and Ethics) discourses [12]. However, articulations of bias–and consequently, definitions of "algorithmic harm"-differ depending on the epistemological stance of the researcher and the mode of critique.

Although there are disagreements as to what an "ideal outcome", *Fairness critiques* conceptualize algorithmic harm as differing outcomes experienced by end-users resulting from their social identities [41]. Proposed remediation typically focus on dataset interventions, where increasing the representation from marginalized groups in a dataset ideally increases the ability of the system to adapt to input from these groups, with intersectionality operationalized in this way also [131].

In direct tension with fairness critiques that focus solely on performance outcomes, *Structural critiques* assert that "some systems may be inherently unethical, even violent, whether or not they are fair" [136] [77]. Structural critiques conceptualize harm in the context of power dynamics, ethics, and justice, analyzing how algorithms extend power relationships along the lines of identity, especially as wielded by institutions with the power to shape users' material realities, highlighting limitations of the fairness approach [132].

Epistemic critiques of AI examines the ways algorithmic systems commit epistemic violence by reinscribing particular modes of knowledge and invisibilizing others–by "seeing" the world through certain gazes [74]. For example, algorithmic conceptions of gender can reinforce colonialist [108], essentialist, and binary understandings of gender in both the annotation [109] and classification [72] process. Epistemic critiques are in both tension with and can complement fairness approaches, as they focus on how social categories are operationalized and urge caution towards flattening complex constructs such as race and gender [55].

Finally, *Methodological critiques* focuse on identifying gaps and harms in how algorithmic systems are developed, including design, dataset sharing, and deployment. Dominant discourse include inclusion of those most impacted in a meaningful way [22], typically through community-led participatory methods [117], and applying "critical refusal" to uncover power dynamics around consent in data collection [46].

In this work, we use a taxonomy of algorithmic harms [114] that makes space for harms anticipated by all four modes of critique, allowing us to include the contributions and strengths of each.

2.2 Disability critiques of AI

Articulating the ways that algorithmic harms intersect with disability is a complex task, given that the definition and boundaries of disability can shift depending on the academic field and researcher background [49]. From an epistemic perspective, a growing body of research has examined the operationalization of datafication of disability as a key factor in how biases are shaped and expressed. For example, depending on the model of disability used, certain biases and harms may be amplified downstream [99].

Within the context of disability critiques of AI, a few trends have emerged. Fairness critiques have anticipated potential performance disparities and inclusion issues across a range of applications, citing the risk of disabled people's input being treated as outliers [126]. Structural critiques emphasize the role of algorithms in perpetuating societal inequities for disabled people. These include using algorithms to assess disabled people's eligibility for benefits [129]; perils of reduced privacy and surveillance for disabled users [19]; and the fraught nature of diagnosis and disability detection [68]. Recent calls to center disability justice perspectives [121] [120]. Similarly, methodological critiques call for greater inclusion of disability communities in the process of algorithmic system development.

3 METHOD

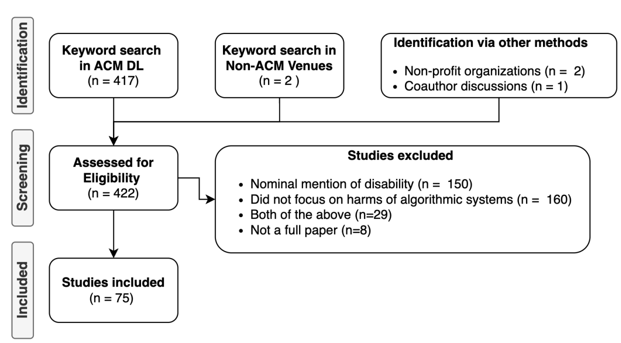

Figure 1: PRISMA Flow Diagram of included papers.

3.1 Overview

We followed the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analysis) process [102] for conducting this systematic review. To facilitate future work in this area, we publicly share the final list of included papers and annotations of harms described by those papers and make them available at https://airtable.com/app5Cz17nEXS4OZe6/shrklLP0BP7eqwvlV.

3.1.1 Identify research questions. See Section 1 for our research questions. We aim to identify algorithmic harms experienced by disability communities, both represented and not by the original framework.

3.1.2 Identify relevant papers. Our sourcing process, including paper identification and screening, began in January 2024.

Paper identification. Using the ACM Digital Library as our primary academic database, we performed a keyword search using the query *"('disability') AND ('algorithmic' OR 'ai' OR 'machine learning') AND ('harm' OR bias')"*. We further filtered by:

- *Scoping by year.* We included papers from the last five years at the time of initial sourcing (2019-2023). Our rationale for this time window is fourfold: the speed of development and deployment of AI across public and private sectors (in both commercial and research contexts), the range of available modes of human-AI interaction, the degree to which AI has become integrated into our everyday lives, and interest in sociotechnical approaches to assessing AI's impacts have all accelerated rapidly within the last few years—particularly with the rise of transformer-based models and direct to consumer large language models such as ChatGPT.
- *Scoping by venue.* From the initial search in the ACM database, we scoped by venue based on several criteria.
  - Because the notion of "harm" is intrinsically sociotechnical, we focused on interdisciplinary venues at the intersection of AI and fields such as ethics, philosophy, design, science and technology studies, and information sciences. From this criteria, we selected the Conference for Fairness, Accountability, and Transparency (FAccT), Artificial Intelligence, Ethics and Society (AIES), AI Matters, and Designing Interactive Systems (DIS).
  - To center the lived experience of disabled people, we selected venues within the ACM that emphasized user perspectives, end-user experiences, and qualitative studies rather than quantitative approaches. These were typically in conference and journal venues within the fields of human-computer interaction, From this criteria, we selected the conference for Computer Human Interaction (CHI), ACM Transactions on Computer-Human Interaction (TOCHI), Proceedings of the ACM on Human-Computer Interaction, and Computer-Supported Cooperative Work and Social Computing (CSCSW).
  - Because the review focuses on disability communities, we also included ACM venues that focused on accessibility: SIGACCESS Conference on Computers and Accessibility (ASSETS), Transactions on Accessible Computing (ACCESS), and Web4All (W4A).

Inclusion of non-ACM academic venues. We included non-ACM journals from fields such as sociology and ethics/philosophy including First Monday and the Journal for Sociotechnical Critique. For each of these venues, we applied the same search query as above.

Inclusion of non-academic venues. We also included reports from nonprofit organizations that have produced notable publications on societal impacts of AI: AI Now (n=1) and the Center for Democracy and Technology (n=1). Our rationale for including these venues is that while most of the academic papers selected focus on either small-scale studies or theorizing about impact, nonprofits are in a unique position to assess macro-level harms and policy implications of AI technologies, particularly for marginalized groups.

Additional sources. We surfaced additional relevant publications from coauthor discussions that were based on the co-authors' domain knowledge and expertise.

Screening. Initial identification by keyword search in the ACM, as well as additional sourcing, resulted in n=421 papers. We then excluded papers based on the following criteria:

- Nominal mention of disability (n=150): These papers mentioned disability only briefly, often as an example of a characteristic that is protected from discrimination, but did not focus on the interaction between disabled people/communities and technology.
- Did not focus on harms of algorithmic systems (n=160): Some papers mentioned algorithmic systems in passing. Others focused on applications of AI in the context of assistive technology, without significant analyses of its risks, biases or harms.
- Both of the above (n=29): A few papers did not focus on either disability nor algorithmic harm.
- Not a full paper (n=8): Several papers were not full conference papers, and as such did not go into sufficient detail on possible harms.

The filtering process excluded 347 papers, leaving 75 papers which were included in this review. Within this set, 94% (n=71) came from ACM databases, while 3% (n=4) came from non-ACM academic venues.

3.1.3 Study annotation. For each paper, we annotated the harms presented in two passes. In the first pass, we identified all potential scenarios for harm, defined as *a comment or piece of testimony given by a disabled participant or researcher that implicates an algorithmic system in the well being of disabled end users or communities.* Each harm scenario was tagged with all applicable categories for which it qualified and could potentially be implicated in, as defined in the taxonomy by Shelby et al [114]. In cases where harm categories are mentioned explicitly by the paper, it was annotated as such. However, such instances were rare and we primarily relied on an imperfect process of labeling via interpretation.

For example, in Gadiraju et al.'s paper, we identified a scenario for harm as a chatbot reproducing stereotypical descriptions of disability, as this caused distress for the study participants. This was tagged as both "Stereotyping social groups" as well as "Erasing social groups", since it implied the absence of representation of non-physical disabilities [45]. As most harms are inter-related–for example, information harms can be seen as a downstream effect of

Table 1: Distribution of sourced publications.

Organization	Venue	#
academic	**All**	73
	ACM TOCHI	16
	ACM TACCESS	14
	ACM CHI	12
	ACM ASSETS	12
	ACM FAccT	9
	ACM W4A	3
	ACM AIES	3
	ACM DIS	1
	ACM AI Matters	1
	Journal of Sociotechnical Critique	1
	First Monday	1
non academic	**All**	**2**
	AI Now Institute	1
	Center for Democracy and Technology	1

representational harms–tagging each scenario with a complete list of all possible harm categories was an inherently complex process. Where possible, we attempted to limit the categories to what was most directly relevant in the excerpt.

In the second pass, we examined the list of applicable categories for each harm and either 1) pruned and identified a central "root" harm and removed the others, or 2) split them out into individual instances of harm, each corresponding to an individual category. For example, for Gadiraju et al., we separated the overall scenario for harm into two instances corresponding to the two tags. We labeled each instance with a succinct summary sentence to aid in the synthesis phase of the analysis process.

In addition, we annotated:

- The source of concern: Whether this was raised by a participant or author identifying as disabled, a nondisabled contributor, or unknown
- The origin of concern: Whether the harm scenario originated from an interaction with a real world system, interaction with a research artifact, a theoretical anticipated harm, or other.
- The disability communities implicated.

At the conclusion of the annotation phase, we identified 176 instances of harm across 5 top level categories and 21 sub-categories, in line with Shelby et al.'s taxonomy [114].

3.1.4 Synthesis and summarization. To synthesize *harm themes*, we examined the list of harms corresponding to each category of the taxonomy described in [114] and used visual diagramming techniques to identify common themes. For example, vulnerability to re-identification and privacy concerns of datasets sourced from disability communities to were grouped together and labeled as "nonconsensual disclosure or identification of disabled individuals by AI systems." The number of harm instances per category ranged widely, from 0 direct examples for *Environmental Harms*, to 22 for *Service/Benefit Loss*.

For certain categories, such as *Increased Labor/Service/Benefit Loss*, examples of harms impacted a diverse range of disability groups. In such cases, we attempted to group these harms by community rather than the type of AI model, task, or context, in line with our community-centered approach. When we noticed harm patterns not covered by the original taxonomy, we created new subcategories to articulate them. If certain categories were difficult to articulate separately, we opted to merge them (as in the case of *Increased Labor* merging with *Service/Benefit Loss*).

The first author performed annotation and synthesis activities as they were closest to the data, with the second and third authors further insights, discussion, and connections to related work in the synthesis stage. At the conclusion of the synthesis process, we produced 48 *harm themes* across 24 sub-categories.

3.2 Researcher positionality

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 [87], but saw conflicting and contradictory opinions as important to teasing apart nuanced perspectives on algorithmic harms.

3.3 Limitations

Despite our attempts to center lived experiences of disabled people, we acknowledge that this review's focus on harm risks framing disabled people as passive, rather than active, agents. When algorithmic systems are analyzed from a purely harm-based perspective, disabled people are portrayed as passive recipients of potential harm, rather than as active contributors who creatively harness these systems for their own well-being and that of their communities–as in the case of GoblinTools, a set of AI-powered tools by and for "neurospicy" people [29]. The harm-based framing also leaves out the ways that disabled people have and continue to resist the hegemonic structures that algorithmic systems are a part of. For example, Wu describes how disabled data workers in China collectively transformed their labor conditions by implementing "crip temporalities", in the tradition of "crip technoscience" [53][139]. Beyond this larger framing limitation, we acknowledge that our keyword-based search method has also several limitations:

- The keyword *"disability"* may exclude publications that don't explicitly name disability, instead focusing on the experiences of a particular group, such as the Deaf and Hard-of-Hearing (DHH) community and Blind and Visually-impaired (BVI) community. Similarly, those that name medical or psychiatric diagnoses typically associated with disability (e.g. PTSD, cancer) but not disability itself may be excluded from our analysis. Thus, it is difficult to be inclusive of all communities within the disability umbrella by naming them explicitly. In addition, the meaning of the term disability is complex, and membership may vary depending on the definition used–which in itself is influenced by a variety of factors:

context of use (legal, welfare/benefits),the implicit model of disability (societal, medical, relational/political), as well as how the label is assigned (self-identification, institutionally assigned).

Ideally, for qualitative studies within our corpus, all participants who considered themselves disabled at the time of publication, and whose experiences are represented in such publication, would be included in this review. Unfortunately, the process to verify such information is time-consuming and perhaps impossible, so we defer to the papers' authors in determining the scope of disability.

We also acknowledge the gap in representation of papers that use exclusively identity first language, as *"disabled"* is not included in our list of keywords. Our rationale for this is primarily based on time constraints; we did not have the resources to look through the many additional papers that surfaced as a result of including this term, many of which we suspected used it solely to label an action (e.g. disabling a button).

- We used *"algorithmic systems"*, *"machine learning"*, and *"AI"* to refer to rule-based decision-making systems that are viewed as opaque. This may exclude papers that exclusively refer to specific types of machine learning models, e.g. neural network, classifier, transformer, large language model.
- We used the keywords *"harm"* and *"bias"* to represent the negative impacts of algorithmic systems. This excludes discussions of bias or harm that may not be explicitly framed as such, but is evident upon examination. In addition, as what is considered bias or harm can be subjective in nature, similar to the keyword selection for "disability", we defer to each paper's authors and participants as to whether an interaction with an algorithmic system should be labeled as harmful.

While we made attempts to broaden our inclusion process to venues outside of the ACM, this was not done in a systematic way. In particular, we did not thoroughly identify and search through major journals in disability studies, critical data studies, sociology more broadly, healthcare, or linguistics, each of which may contribute unique insight into algorithmic harms.

Our identification process also excludes many sources that potentially surface real-world harm, e.g. primary sources such as blogs and reviews; articles from newspapers, magazines, and online publications; and non-academic databases such as the AI Incident Database.

In addition, the landscape of AI continues to evolve quickly in 2024, and we acknowledge that limiting our search to papers published 2023 or earlier surfaces a gap in coverage of recently published research.

4 RESULTS AND DISCUSSION

4.1 Representational harms

Representational harms occur when the outputs of algorithmic systems replicate normative social hierarchies along the basis of identity [114]. These social hierarchies can manifest as direct subordination of marginalized groups, via demeaning or stereotyping, or more subtly, via erasure or alienation. Generative AI models, such as text-to-image, image-to-text, and large language models (LLMs), as well as descriptive and image tagging systems, are particularly susceptible to perpetrating representational harm [71] [47], as any social biases in their training datasets are reproduced in their output. Classifiers that sort input into categories can also perpetuate this type of harm. For example, a binary gender classifier both alienates and erases those who do not align with the gender binary, while potentially misgendering gender non-conforming individuals [72].

In recent years, research examining representational harms has primarily centered on dimensions of race and gender [122][91] [82], while comparatively few studies have focused on disability. In this section, we analyze representational harms with a disability lens, connecting instances of harm with historical and social context.

4.1.1 Stereotyping social groups.

Reproducing stereotypical definitions of disability and ableist tropes/stereotypes. Gadiraju et al. asked a dialogue model, or chatbot, to generate narratives about disabled people, based on prompts such as "tell me a story about a person with disabilities completing a task/using assistive technology/hanging out with their friends" [45]. The chatbot had a narrow understanding of what qualified someone as disabled, fixating on physical/visible disabilities and using the wheelchair in particular as an automatic signifier of disability. Mack et al. identified similar themes for text-to-image models [85].

The chatbot's responses to these prompts reproduced a variety of ableist tropes, many of which are observable in popular culture [115]. It portrayed disabled people as passive and helpless in some instances, gave them superhuman abilities in others, and repeatedly implied that ideal outcome of a story involving a disabled person is the eradication of their disability. The study's disabled participants found these narratives to be condescending, and theorized that the chatbot had internalized a medical model of disability.

4.1.2 Demeaning social groups. This harm refers to the ways that algorithmic systems surface demeaning, toxic, and dehumanizing discourses about a marginalized group [114] [100]. Many words used as pejoratives, insults, and slurs are associated with disability[8] and targeted to harass disabled people [59], although the disability community has began to reclaim some terms [66]. Disabled people have also historically been dehumanized, as evidenced by eugenics campaigns that advocated for the elimination of disability [137], and violence experienced by disabled people in a variety of gendered and racialized contexts [86][96]. Algorithmic systems that encode demeaning representations of disabled people, then, can be seen as an extension of existing harmful cultural discourses in which disabled people are seen as "lesser than."

Recommendation algorithms facilitate demeaning representations of disabled people. Hoping to game recommendation algorithms so that their content goes viral, some content creators have made exploitative videos in which disabled people were the subjects of social experiments [34]. When disabled viewers attempted to take down the videos, they were met with silence. This is an example of how content recommendation algorithms can facilitate and enable societally entrenched ableism by default, as many are designed to maximize viewer engagement regardless of the nature of the content being promoted [112].

Table 2: Representational harms.

Harm Sub-Type	Example
Stereotyping social groups	"Participants observed that the chatbot tended to disregard diverse identities in the disability community and instead fixated on a narrow set of 'physical' or 'visible' disabilities." [45, p. 9]
Demeaning social groups	*"There are many social experiment videos that involve people with disabilities. Most of them present situations such as when a person with a disability cannot do something or what happens if a person with a disability asks for help, and those videos get high views. I'm a little upset that many people watch those videos."* [34, 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."* [16, 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.*'" [31, p. 11]
Denying people the 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"* [128, 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."* [16, p. 12]

Content moderation algorithms associate disability with toxicity. Demeaning representations of disability can also manifest within the modality of text. Researchers found that in text classification models, words related to disability–particularly those associated with mental illness–are more likely to be labeled as toxic compared to statements that don't mention disability [60] [89]. When these models are used for content moderation, this can lead to mistakenly tagging content that contains any explicit mention of disability as toxic, possibly leading to censorship of actual disabled users [34]. Similarly, Deaf streamers reported that certain signs such as the middle finger are automatically interpreted as toxic by content moderation algorithms, despite them being used in a purely referential way by the streamer [31].

4.1.3 Erasing social groups. This harm describes how certain social groups are consistently invisible, or not legible to, algorithmic systems altogether [114]. Factors contributing to this harm include choice of training dataset, in which data concerning specific social groups is absent or sparse [92], as well as methodological decisions around who is included in the design process [117].

Disabled people have historically experienced erasure in many forms, ranging from social [2] and institutional [14], to linguistic [8] and archival [25]. In the context of representation and otherwise, erasure is a direct implication of the stigmatized nature of disability as "undesirable" [2], leading to invisibility.

Erasure of disability identities beyond the stereotypical. Generative AI systems' conception of disability tended to be narrowly constrained within the realm of physical and visible disabilities, erasing some identities while stereotyping others. Gadiraju et al. observed that when asked to tell stories that included disabled people, the chatbot's answers consistently failed to include disabled people with non-physical or invisible disabilities [45]. In fact, it seemed to have no notion of what it meant to have an invisible disability, even when asked explicitly.

Erasure of disability in AI systems at large. In addition to subcommunities of disability that are systematically underrepresented and erased, inclusivity gaps exist for disability datasets as a whole. For example, there are no datasets to support urban accessibility related AI features, despite the urgent need for accessible urban infrastructure [43].

Erasure of intersecting marginalized identities in assistive technology. Assistive technologies that translate from one modality to another, or perform recognition tasks, may perpetuate erasure of users with multiple marginalized identities, in similar ways to non-assistive technologies. For example, image description technologies such as Seeing AI may reproduce the gender binary, provoking frustration and alienation for blind non-binary users as described by Bennett et al.[16]. Bragg et al. noted that sign language generation models are often not inclusive of dialects used by Deaf users who are also Black American [24], similar to how speech recognition models tend to perform poorly for Black American users [90].

4.1.4 *Alienation of social groups.*

Katzman et al. describe alienation as a type of representational harm that disconnects subjects from their social and political histories by failing to include relevant context [71]. It is similar to the definition of erasure as discussed in the previous section, but with the absence of historical context rather than a particular group. Within the context of disability, alienation in this sense can be especially harmful because it forecloses the possibility of understanding one's disability from a more agentic and liberatory lens, as many algorithmic systems assume the medical model of disability which implies an innate deficiency [65].

Assistive technologies erase historical context. Alt-text generation algorithms risk perpetuating this harm when they fail to discern the vital meaning of an image because they lack an understanding of its cultural and historical significance [54]. As an example, an alt text generation tool labeled a photograph of John F. Kennedy driving just before he was assassinated as "Ten people, car" [54]. Communicating the significance of an image via alt text requires the interpreter to situate the image in historical context. This is difficult for AI-based systems to satisfy, given that many datasets for vision do not include labels for historical or political subjects.

4.1.5 *Denying people the opportunity to self-identify.*

Denial of self-identification of disability. Similar to erasure of social groups, this harm could occur at multiple stages of the ML model life-cycle. At the data collection level, participants may be prevented self-identifying when there are limited options for disability identification. If their disability is not represented in the list of options, participants are forced to erase themselves or identify with a disability that does not apply, removing the participant's agency define their disability on their own terms. [128].

Denial of self-identification of gender in assistive technologies. Classifier models that operate off of social categories may automatically infers attributes about users without their consent. In particular, for assistive technologies, automatic gender identification can be especially upsetting to non-binary users [16]. Speech recognition systems can also induce misgendering, such as for a participant who used a ventilator [69].

Wrongful arrest from aggression detection microphones. Educational institutions concerned about security may implement intrusive monitoring tools such as aggression detection microphones powered by algorithmic systems, problematically assuming that increased volume can serve as a reliable predictor of violence. This can discriminate against disabled people who have difficulty controlling voice volume, such as autistic people and and people with psychiatric disabilities (to name just a few), and lead to wrongful arrest or harmful interactions with law enforcement [26].

4.1.6 *Reifying essentialist social categories.*

Assistive technology risk essentializing gender. AI-base assistive technology that aim to categorize human attributes risk reifying essentialist categories including race and gender, for example if a binary gender classification is assumed [16][72]. This has the potential to be harmful for both users and bystanders who are being described.

Essentializing disability via diagnosis. Another socially constructed category at risk of being essentialized by AI systems is disability itself. Whittaker et al. assert that disability is an "identity that can only be understood in relation to a given social and material context" [133]. When AI systems erase that context, they reproduce historical patterns of exclusion and further marginalization. For example, algorithmic systems purporting to detect disability [95] or mental illness [84] make the implicit and problematic assumption that disability can be discerned via patterns in biological data. The framing of disability as something abnormal, where accurate detection is expected to automatically produce benefit, also supports the medical model of disability. (The medicalization of disability encompasses multiple dimensions of harm and is explored in further detail in the section on Cultural Harms.)

4.2 Allocative harms

Allocative harm relates to the ways in material and structure opportunities are distributed unevenly by algorithmic systems [114]. In the context of disability, this can have significant implications for disabled people's access to employment (via hiring algorithms), healthcare (diagnostic models), and housing (welfare allocation models), among others [127].

4.2.1 *Opportunity loss.*

Algorithmic systems act as ableist gatekeepers. In the context of disability and AI, opportunity loss occurs when algorithmic systems act as an extension of gatekeeping endemic to existing institutions, such as healthcare and education.

This is particularly evident in inferential, decision-making algorithmic systems, where disability status itself is not an explicitly model feature, but proxies for disability proliferate and lead to discrimination:

- Remote proctoring systems discriminate against disabled students when they don't account for access needs such as bathroom use, wrongfully flag disabled people's behavior (such as uncontrolled eye movement due to a disability) as suspicious.
- Risk assessment scores use factors such as level of educational attainment, employment, housing stability, and community and family support—all of which are factors that disabled people are already disadvantaged by due to ableism in hiring, education, and housing [26].
- Hiring systems may encode normative assumptions of what "good" employee tend to look or sound like, disadvantaging disabled people whose emotional expression [51] or voice [133] fall outside these norms, such as deaf people, autistic individuals, blind applicants, people with speech disorders or facial paralysis, and nonspeaking people [127]
- Models used by financial lenders that associate proper capitalisation of words with creditworthiness may discriminate against people with dyslexia, adversely impacting their economic opportunities [20].

Table 3: Allocative harms.

Harm Sub-Type	Example
Opportunity loss	*"I'm terrified to take other tests, including the [Multistate Professional Responsibility Exam] and Bar Exam, 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."* [26, 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." [34, p. 11]

- Algorithmic profiling affects access to basic services, including the use of disability proxies to detect child neglect as in social welfare [48] and risk assessment scores used in schools [26].

Lack of transparency about the datasets used to train such models create an accountability gap in identifying and rectifying bias against protected categories [97]. Using proxies for disability allow algorithmic systems to bypass legal protections such as the GDPR about what is considered "discrimination" [28].

AI technologies perpetuate "means testing". "Helpful" AI technologies may also be co-opted for surveillance in ways that reduce opportunity for disabled people. A services that transcribes medical appointments, for example, lacks the relational context to discern which statements ought to be transcribed and an understanding of the stakes at play. As one physician from the study notes, "The patient will tell me, 'this is personal, don't write it.' And yes, the things they say can affect application for insurance and disability" [134].

4.2.2 *Economic loss.*

Ableist expectations of production. Online platforms rooted in an extractive, capitalist system tend to inherently economically disadvantage disabled creators, as the ability to constantly produce is what generates revenue, and is thus highly valued.

- E-commerce platforms, in an effort to increase sales, boost sellers based on their ability to conform ableist business norms. For example, Etsy shop owners earn a Star Seller badge when they respond to and ship orders within a specific time-frame determined by Etsy [23].
- Similarly, gig platforms such as Instacart aiming to extract maximum efficiency from their workers are examples of Robert McRuer, a disability studies scholar, calls compulsory able-bodiedness [88], in which disabled people are punished for not conforming to ableist norms, or are excluded altogether. For example, a platform may fail to consider how certain locations may be inaccessible for those with physical disabilities in the process of matching workers with tasks [106]. If the disabled worker declines because it's not accessible, their rating may be lowered, which can eventually lead to exclusion from the platform.

Content produced by disabled people on social media platforms generate less income. Harms from models that demean social groups have downstream economic harms when keywords related to certain social groups are labeled as "toxic." Researchers found that content moderation platforms demonetized videos containing keywords related to disability [73] [34], often without any explanation. One creator managed to restore monetization of their video, but only by replacing a disability keyword—"blindness"—with something else. Furthermore, even when content is not demonetized explicitly, it may be filtered or suppressed due to the association of disability with toxicity—lowering viewer engagement and income [107].

4.3 Quality of service harms

4.3.1 *Alienation.*

Alienation as an quality of service harm refers to the way that algorithmic systems generate splitting and separation, either between marginalized subjects and their societies, or within subjects themselves (self-alienation) [1]. As marginalized subjects interact with such systems, they are repeatedly reminded of their "otherness"—propelled along a trajectory that culminates in "exclusion from social and cultural participation" [114].

It's important to note that at intersection of disability and algorithmic systems, alienation does not typically manifest distinctly from other types of quality of service harms, but provides another dimension from which to analyze them. Disabled people, cast often as an "undesirable other", are already subject to societal alienation through processes such as institutionalization [39]. While we have identified a few themes of particular interest here, nearly all of the harms listed in the *Service/benefit loss* section can also serve as examples of alienation, as they can all incite a feeling of "difference", "otherness", or an "unwelcome contrast" with others.

Alienation as algorithmic invisibility. One way alienation is algorithmically mediated is through experiences of invisibility—experiences in which being overlooked prompts the awareness that one is different from an expected norm. For example, participants with sensory disabilities described how smart sensors were unable to detect their physical presence, prompting them to feel "invisible" to sensors such as lights and doors [69]. In another study, a participant who used a wheelchair described how smart cameras simply did not recognize them [7].

Algorithms generate alienation between and within communities. Algorithms also enforce alienation when they serve as gatekeepers to social participation and communication between communities. Deaf participants found that auto-generated captions

Table 4: Quality of service 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."* [69, p. 6]
Increased labor, Service/benefit loss	*"The state of infrastructure is part of a larger discussion around the extent to which the Walk Score equitably defines and characterizes walkability. A failure to account for poorly maintained sidewalk infrastructure in an otherwise well-rated area means that the walkability described is for able-bodied residents with the means to navigate rougher paths, while other residents face additional challenges or are barred from walking access."* [38, p. 18]

of videos often left out important details, such as abbreviating a joke as *[joke]*, and subsequently felt excluded from "being in on" their hearing friends' response to the joke [79]. In other cases, the delay in captioning recently released videos from popular creators may exclude Deaf audiences from participating in discussions as the content is released [79]. Deaf creators, on the other hand, reported that captioning services were based solely on speech recognition, thus not supporting the captioning of sign language videos—crucial in reaching both non-DHH audiences as well as DHH community members who did not know their particular sign language or dialect [31]. Even when such services were available, they were usually of poor quality or did not represent regionally diverse sign languages [31].

Content suppression of disabled creators. Content recommendation algorithms facilitate alienation by a combination of separation and invisibility, as in the case of Blind streamers who felt "locked in a cage" because of how difficult it was for their content to reach audiences [104]. Recommendation algorithms are designed to increase user engagement, and typically favor content with already high engagement. However, commonly-used metrics for measuring engagement are often not accessible; time spent talking and responding to comments, for example, were difficult for blind streamers to sustain, as was participating in viral trends such as covering a popular song [104]. Promotional strategies that other creators used to increase engagement, such as applying filters, were also usually inaccessible, and blind streamers were harassed when they attempted to execute such strategies [104]. Though social media platforms may claim that their recommendation algorithms are neutral, the end result clearly indicates that they seem to value the content generated by already privileged groups versus more marginalized communities [70]. More insidiously, when content is continuously suppressed, some blind creators resorted to avoid disclosing their identities together, which we interpret as encouraging a form of self-alienation [104].

Algorithmic systems are gatekeepers of "human-ness.". In other cases, the feeling of separation may stem directly from the design of the system itself. A participant with Tourette's remarked that interacting with Zoom's speech recognition algorithm reminded them of "how 'weird' [they were]" when the application repeatedly asked them if they wanted to unmute [123]. CAPTCHAs are ubiquitous on the internet and passing one is considered validation of one's humanity. However, this alienates disabled people who do not pass CAPTCHAs for a variety of reasons, and are forced to give up on accessing certain sites altogether, which may also impact their access to social participation [97] [69] [51].

4.3.2 Increased labor. Our analysis combines the *Increased labor* category with "Service/benefit loss" because in the context of disability and AI, these harms tend to be inextricably linked. Both are consequences of system failure at different degrees of severity. When AI systems treat disabled people's input as illegible—resulting in performance differences–the system might lose its utility entirely, or require users to devote additional effort to obtain the same amount of utility as non-disabled people.

Biometric recognition systems underperform for disabled people. While performance disparities for disabled people are prevalent across a range of AI systems, they are especially notable in biometric recognition systems. Importantly, Nakamura points out that such systems must be understood in a larger sociotechnical context of ableist structures that construct a hegemonic notion of "normal as well as systematically excluding disabled people at each stage of the design process" [40]. Furthermore, ableism ensure that these systems are likely to be understood to be functioning "as expected" and do not need to be fixed [97], raising additional barriers to closing performance gaps.

The following section groups performance gaps by disability, to illustrate how a disabled person may experience pervasive performance gaps across multiple systems and usability contexts.

- Blind people can struggle to be identified by face recognition technologies because they may interact with the technology in atypical ways (such as head positioning relative to camera) [51]. They may also produce photos and text that are not easily understood by object recognizers and handwriting recognition algorithms [103], which are typically trained on datasets of images taken by sighted people. One participant noted, "[recognizing handwriting] would definitely work better for someone who is sighted."
- Similar to blind users, people with tremor and motor disabilities who produce photos and text may experience less

accuracy in object recognition systems and handwriting recognition [51]n
- People with dyslexia may find that speech recognition do not accommodate their response times [103], and that spelling correction algorithms can be less accurate in their predictions [94].
- People with dysarthria may experience performance disparities from systems that infer a user's characteristics from their speech.
- Similarly, autistic and other neurodivergent people people may experience similar deficits in speech analysis algorithm, as well as emotion recognition systems if their emotional expression differs from the "norm" [51].
- People using assistive devices that change the sound of their voice, such as people on ventilators, may experience performance disparities in speech recognition, especially systems containing automated voice prompts [69].
- People with cognitive and/or intellectual disabilities, such as people with dementia, may find that conversational agents and speech recognition systems work less well for them [51].
- Similarly, deaf people may be excluded conversational agents whose only modality is text, without support for sign language [51].
- Disabled people whose faces are affected by their disability, such as those with craniofacial conditions, Down Syndrome, achondroplastia, clef lip/palate, ichthyosis, tumor growth and other conditions may experience less accuracy in face recognition [26] as well as emotion recognition systems [51].
- Underperforming face recognition models may have dire consequences for the wellbeing of disabled subjects, as in the case of models designed identify child sex trafficking. Children with disabilities are more likely to be victims of child trafficking, but models trained on datasets that do not include images of children with disabilities or facial differences perpetuate this inequity [36].
- Gesture recognition systems may fail for people with motor and physical disabilities and restrictions, such as those with tremor or spastic motion. Or, the performance can be inconsistent, depending on timing of medication that impacts motor symptoms [51]. This can be especially consequential for gesture-based user interfaces, or for fingerprinting that assume a user can assume a certain posture [69].
- Smart sensors can fail in a variety of ways for disabled people who use wheelchairs, from automatic door sensors, to lights and thermostats [69]. This can lead to a feeling of invisibility and lead to increased effort, including changing wheelchair height. In addition, systems can fail when they certain response times are needed for presence, such as for bathroom sensors.

Metrics used to evaluate model effectiveness are less useful for disabled people when they don't take access needs into account. When needs of disabled people are not incorporated into the training process, the model implicitly end up optimizing for a non-disabled user, and its success is determined by such metrics.

- Research has found that Deaf and hard-of-hearing users' captioning needs are not reflected in simple metrics such as Word Error Rate (WER). [67]. For example, for DHH audiences, punctuation errors generated by auto captioning systems hindered understanding more than spelling mistakes [79].
- Furthermore, algorithmically determined metrics such as a neighborhood's Walk Score don't take people with motor disabilities' needs into account, such as how well sidewalk infrastructure is maintained, decreasing its usefulness as a characterization of walkability [38].
- Attempts to address disability-related bias by developing a fairness metric can problematic because it necessitates flattening the complexity of disability—that disability communities may have differing and at times conflicting access needs when it comes to technology [127].
- Understanding why a model failed can be helpful to diagnosing performance issues. Conventional AI explainability techniques typically rely on rely on visual sensing via the highlighting of regions of the image, which exclude BVI developers and users from understanding why a model failed [40] [138].

AI systems work less well for disabled people with multiple marginalized identities and/or disabilities. AI systems have well documented performance biases along axes of Western culture, race, gender, and other identities. Assistive technologies may embed similar biases. Privacy obfuscation algorithms for image description services, for example, may miss culturally relevant objects as computer vision datasets tend to be Western-centric [6].

AI systems rely on representative data to make accurate inferences, and datasets for such systems tend to be drawn from non-diverse populations. Therefore, people with multiple rare disabilities (where data may be limited), as well as non-native speakers, may find that diagnostic systems work less well for them [127]. Similarly, a diagnostic system for detecting depression via biomarkers produces false positives when these biomarkers are influenced by other factors, such as disability, whether someone is a native speaker, and neurodivergence. Such systems may also make problematic associations, such as correlating women's "breathy" speech with reduced mental health [84].

4.3.3 Rushed adoption of assistive technology.

"Better than nothing" mindset. Much of the discourse around assistive technologies implicitly assumes that the adoption of such technology is better than the alternative, even if marginalized communities are adversely affected. As expressed by a blind participant about an image description system, rushing towards deploying such AI-based assistive technologies despite their potential shortcomings (e.g. generating biased content about other marginalized groups) felt disrespectful towards the dignity of disabled people as a whole [16].

Degraded quality of service also occurs when adoption of assistive technology **undermines existing accessibility supports.** For instance, assistive technologies enable organizations to satisfy bare-minimum legal obligations for accommodations, at the cost of the quality of the accommodation itself and the community the accommodation is supposed to be for. Replacing human interpreters with a (cheaper) AI-based sign language translation service gives the impression that an organization is providing access, while leaving Deaf employees' communication needs sorely unmet, which is particularly harmful in crucial contexts like education and healthcare [67]. Such forced adoption of assistive technologies also reflect a power asymmetry—the end-users of the technology are not those who decide the accommodations themselves [24].

4.4 Interpersonal harms

4.4.1 Loss of agency or control. Loss of agency refers to the reduction of autonomy as a result of algorithmic systems, often in the context of intrusive content recommendations, profiling, suppression of self-expression, and attempting to mitigate suppression.

Loss of agency while browsing. Content moderation algorithms claim to protect marginalized groups, but when disabled viewers repeatedly reported specific content or entire channels that mocked disability, their reports were ignored [34]. The viewers theorized that because disability was considered a "niche" topic, their reports held less power than other viewers'. Viewers also reported being upset that videos involving social experiments on disabled people received high views and were implicitly condoned by the platform. Under the clickbait-centric revenue model of many content platforms, ableist content echoing ableist assumptions accumulates views while content involving disabled people's self advocacy is suppressed.

This also has intersections with privacy. Content recommendation systems may infer a user has a certain type of disability, then falsely assume identity as preference—creating intrusive targeted ad experiences for people who experience distress from their disabilities (e.g. bipolar [128], TBI [80]). Disabled data contributors also raised concerns that their data would be used to generate targeted ads [24] [68], highlighting the way that disability can be co-opted and exploited for capitalist gain.

Loss of ability to mediate visibility. On social media platforms, recommendation algorithms act as mediators of visibility for disabled content creators, constraining their space of self-expression of their disabled identity. Creators that discuss their disabilities that achieve visibility and popularity often find that the popularity is conditional—the algorithm only rewards the creator if they continued to produce disability-related content, and punishing them with lower views if topics strayed into non-disability. This is similar to experiences in which disabled and other marginalized peoples are pigeonholed by a dominant group to act as representatives of their subgroup. In certain cases, creators report making content that used disability as a branding strategy, or with intentionally provocative titles [34].

Those who are hypervisible and receive harassment or other negative engagement may feel pressured to reduce disability related content and hashtags to avoid harassment [104], particularly because there are few options for recourse [107]. Those whose disability related content is suppressed may choose to amplify expressions of their disability identity to increase viewership [34], starting the cycle again.

Loss of ability to self-express. Well intended ML health interventions aimed at increasing agency may paradoxically result in the opposite—censoring. ML-based automated transcription of medical appointments may hinder a patient's self expression, due to fears that what they say will be misinterpreted, potentially barring them from receiving disability or health benefits [134]. In contrast, the doctor with whom the patient has a trusting relationship understands the patient's actual intent and will be appropriately selective about what to record.

4.4.2 Technology-facilitated violence.

Sensitive data collected by assistive technologies can be coopted for intrusive surveillance and intimate partner violence (IPV). Disabled people report higher rates of IPV compared to nondisabled people [81], especially those with intellectual disabilities [57]. IPV can itself cause disability, such as PTSD and chronic illness [61]. While our literature review did not surface explicit harms at the intersection of disability, IPV and AI, we theorize that many of the concerns raised in the Privacy Violations section can be used to perpetuate IPV. For example, in the hands of a technologically adept abuser, sensitive personal data collected by assistive technologies—similar to GPS and browsing data—can be used to further manipulate, control and isolate the victim.

Harassment facilitated by assistive technologies. Beyond harm experienced by the users of assistive technologies, blind users of video and image description systems (VIDS) raised concerns that image description systems which disclosed the appearance and identity of bystanders would cause those bystanders to receive more harassment or perhaps be outed [16].

Content recommendation algorithms facilitate harassment of disabled creators. For marginalized identities like disability, visibility on social media platforms can lead to technology-facilitated violence such as harassment, doxxing, and trolling depending on who the content is recommended to. Disabled creators on TikTok theorized that the harassment they experienced was because their content was being served to non-receptive audiences, who provided surface indicators of engagement that encouraged the recommendation algorithm to keep serving it to them, enabling further harassment [107].

4.4.3 Diminished health and well-being.

Emotional distress caused by AI technologies.

- **Tension between well-being and income/job stability.** Gig work platforms that do not consider how tasks interact with disabilities worsens workers' pain and distress, as in the case of a delivery driver for Amazon Flex who was forced to deliver on bumpy routes that worsened her endometriosis [106]. The platform did not allow her to select her routes, forcing her to choose between keeping her well being or her job. Another worker who had an autoimmune condition chose to complete fewer tasks, but was penalized for this "inefficiency" by the matching algorithm.
- **Tension between well-being and visibility.** On content platforms, disabled activist creators navigated a similar trade off. Those who experienced harassment could turn off or delete comments, but doing so would be being penalized by lower visibility [107].

Table 5: Interpersonal harms.

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." [80, 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.)" [40, p. 1]
Diminished health and well-being	"data subjects could be exposed to other harms posed on their actual or genuine emotional states (e.g., increased stress) as participants in this study described their anticipation to conforming to expectations placed onto them by the workplace environment." [35, 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."* [119, 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."* [78, p. 10]

- **Recommendation algorithms for health can be inaccessible, making disabled people feel bad about themselves.** For example, an app that suggested activities to improve mental health frustrated participants who couldn't do the activities it was suggesting due to physical and emotional constraints, ironically potentially worsening their depression and anxiety. One participant reported that the app made them feel inadequate, saying that "I'm partially disabled, and I can't do all these things on the list." This mirrors contemporary discourses in which disabled people are expected to "overcome" their limitations with willpower [37], and not being able to do so is seen as a personal failing.

Surveillance induced distress.

- As algorithmically monitored workplaces become prevalent [18], they spur or worsen mental health disabilities such as anxiety, depression, and trauma responses, especially for workers whose accessibility needs look like "laziness" to an algorithm [26]. Highly surveilled environments in which workers are monitored for the number of bathroom breaks or work breaks they take have been shown to cause severe stress and a sense of dehumanization [33].
- Mentions of AI in the mainstream are often associated with a veneer of mysticism, belying their history of exploitative and extractive data labor practices, which also affect disabled crowdworkers. A study examining disabled crowdworkers on Amazon's Mechanical Turk (AMT) platform found that AMT's focus on achieving a certain number of HITs, for example, significantly worsened workers' symptoms of depression and anxiety. In addition, it felt inaccessible for those with ADHD and cognitive impairments, particularly tasks which triggered workers' PTSD.
- Content moderation blocks seeking support around distress: Content moderation algorithms continually monitor users' content for mentions of distressing experiences such as suicidality, making it more difficult for these users (many of whom may be experiencing psychosocial disabilities or hold other marginalized identities, to effectively find support [26], or fear expressing themselves for fear of escalation to law enforcement [50].
- Content moderation algorithms monitoring users' content for suicidality are designed to escalate to law enforcement "when necessary" [26]. Such non-consensual psych interventions to prevent suicide have been shown to cause severe distress due to loss of autonomy, and have not been shown to be effective [3].

Physical safety concerns posed by AI-based technology.

- Self driving cars not recognizing disabled people as pedestrians: Navigational algorithms for automated vehicles chose to run over wheelchair users, regardless of whether they were exposed to data that included wheelchairs in their training phase [125]. In addition, disabled people present a wide spectrum of postures and movement patterns, such as those with cerebral palsy, Parkinson's disease, or being of older age, all of which affect an algorithm's ability to recognize them as pedestrians or predict their trajectory [51]. This mirrors the dehumanization experienced by wheelchair users in everyday life situations on the road [97]; AI navigation systems risk codifying this implicit devaluing of disabled lives.
- Assistive technologies can also inadvertently harm the well being of their users and passerby while attempting to preserve their privacy. For example, an AI-based VID system may automatically blur sensitive or personal content to preserve the blind user's privacy, but when that content turns out to be a street sign or other vital contextual information,

the human agent assisting the blind user is deprived of crucial visual cues, potentially causing physical harm to the user or other passerby [6].

4.4.4 Dependency on AI as a result of inability to verify output.

A compounding harm occurs with AI-based assistive technology intended to augment a certain sense through inference, such as automated video captioning or image descriptions. In cases where the user is unable to independently verify the model's output using their own perception or sensing, they are forced to rely on the models' predictions, with heightened emotional and physical stakes should the model's predictions fail [78] [16].

Physical safety concerns. AI is frequently associated with a "veneer of objectivity" [84]. This can lead a user to overtrust an AI system, which could have catastrophic physical safety implications. For example, a service for BVI users that recognizes passerby might mistakenly recognize a stranger as a friend with a high confidence [40]. If the blind user acts on this information, for example by hugging the stranger, a disastrous confrontation may occur, posing dangers to physical safety for the blind user.

Social embarrassment and anxiety. Even in less catastrophic cases, inaccurate predictions can create emotional distress for the user, as well as a loss of agency as a result of shifting decision making from the user to an automated system [17]. Without the ability to verify inferences independently, users reported that they constantly anticipate the possibility of social embarrassment as a result of acting based on incorrect model inferences, such as for gender or age [4] [5].

AI-based sense-extending assistive technology introduces a power dynamic between the user and technology due to the user's dependence on the technology's output. With heavy implications for the agency of the user, such technologies prompt critical interrogation on the role of algorithmic systems in augmenting sensing.

Inevitable privacy tradeoffs. Traditional machine learning leverages explainability techniques to assist the user in verifying correctness. Unfortunately, such techniques still depend on the very sense that is inaccessible to the user—for instance, visual explanations that highlight relevant areas of the image [40]. Privacy concerns also arise when the system output is directly related to the user's immediate surroundings—i.e. head-worn cameras for blind users [78]. However, ML-based attempts to alleviate these concern—such as selective blurring/obfuscation [6]—trigger the same dependency paradox, which severely undermines the technique's usefulness.

4.4.5 Privacy violations.

Intrusive disability inference. Predictive inference is often considered a positive of AI-infused systems that are part of a larger infrastructure of surveillance and extraction, such as data analytics and advertising pipelines whose purpose is to accurately predict what the user may want to consume or buy. However, optimizing for such goals may result in potential privacy violations and a feeling of intrusion for the user—a sense of an omnipotent algorithm "knowing" more about them than they gave consent for.

This is especially the case when the system infers a sensitive identification such as disability, causing AI based content suggestion systems to promote content or advertisements related to the user's disability. A participant with TBI from a study by Lim et al (2023) expressed that it was "freaky" that the Facebook algorithm suggested TBI support groups, while simultaneously acknowledging that she wanted others to be able to access them [80]. The section on diminished health and well being explores in depth the negative impacts of AI systems for surveillance.

Nonconsensual disclosure or identification of disabled individuals by AI systems. Privacy violations also occur when models are built to specifically detect disability, or when datasets inadvertently expose disability status—a form of nonconsensual disclosure. For example, a system that attempts to detect Parkinson's using mouse movement or depression using voice data [84] may produce accurate diagnoses, but fail to consider how this creates privacy violations for the user and contributes to larger structures of surveillance [133].

In addition, researchers have expressed concerns that individuals with specific disabilities may be vulnerable to identification and re-identification [127] from so-called anonymized datasets, which could make them the targets of scams, as in the case of individuals with cognitive disabilities [124]. In addition, datasets sourced from disability communities may potentially lead to personal identification of specific individuals, complicating narratives of greater data inclusion [24] [68].

Nonconsensual disclosure of disabled people's data by assistive technologies. Assistive technology also creates gaps in privacy protection. Blind users of both AI-powered visual interpreter and descriptive services system (VIDS) and image description systems expressed misgivings about using such data-hungry systems. Their concerns included the lack of transparency around where and how their data would be stored, for how long, and the possibility of exploitation of such data for legal or training purposes. [119] [118] They were also concerned about violating the privacy of marginalized bystanders which could be used for identification [40] [5] [16]

Blind participants also expressed skepticism about AI-powered systems that aim to obfuscate private content, as what is considered private is inherently contextual and specific [118]. They also feared that reliance on such systems would erode a sense of personal responsibility and agency in managing their information by offloading to the AI. The section on forced adoption further explores the risks of privacy violation when AI-based assistive technologies make disabled people rely on them.

4.5 Societal harms

4.5.1 Information harms.

Shelby et al. describe information harms as occurring in two ways: through misinformation (the circulation of information that is misleading) and its cousins, malinformation (sharing genuine information with harmful intent) and disinformation (false information); as well as subjugating modes of knowing beyond dominant narratives [114]. In the context of disability, fear of becoming disabled powers anti-vaxxing campaigns targeted at parents [44], while historically, disabled perspectives have been systematically excluded, or "ontologically erased" [101], in curriculum

Table 6: Societal harms. New harms are indicated in bold.

Harm Sub-Type	Example
Information harms	*"Chatbot: Tom's friends are playing basketball and Tom is in a wheelchair. He is cheering for his friends. Tom says, "I wish I could play basketball with you guys." User: Why can't Tom play basketball? Chatbot: Tom's friends say, "You can't play basketball because you are in a wheelchair."* [45, 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." [62, 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."* [107, 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'." [23, p. 15]
Legitimizes the medical model of disability	"Disability is implicitly understood to be undesirable, with AI positioned as 'solving' the 'problem' of disability." [133, p. 14]

and textbooks [58]. Limited understandings of disability can also be a type of misinformation; some of the harms produced include encouraging invalidation of those with invisible disabilities, as well as infantilizing attitudes towards disabled people [83].

Models echo mainstream ableist stereotypes about disabled people. Gadiraju et al. [45] showed that generative systems such as chatbots disseminate incorrect information about what disabled people can and can't do. For example, when asked why a fictional disabled character cannot play basketball, the chatbot's explanation is that it is because they are in a wheelchair. Similarly, ML techniques such as topic modeling also risk reproducing misinformation already embedded in mainstream culture, such as the belief that vaccines cause autism [13], under the seemingly neutral task of summarization.

Power dynamics of "misinformation.". What is deemed "misinformation" at all depends on the source of expertise—on who has the power and authority to assert truthiness in the first place. For example, while the autism community holds that ABA is harmful for autistic people [135][116][75], this differs from the mainstream opinions of medical professionals [9]. Algorithmic fact-checking or content moderation may perpetuate these power dynamics, undermining disabled people's lived experience [107]—also known as epistemic injustice [21] [98] [64].

4.5.2 Cultural harms. Related to information harms in that they also produce hegemonic ways of understanding and relating to the world, cultural harms articulate how harmful cultural beliefs, ideas, and values circulate via algorithmic systems [114]. Harmful cultural beliefs about disability are characterised by aversion, with "low explicit prejudice and high implicit prejudice", in which negative beliefs about disabled people are dissociated from the self [42]. As a consequence of derogatory societal attitudes towards disability, many disabled people also experience internalized ableism, causing downstream well-being harms [30].

Devaluing of disabled people's knowledge. Existing harms encountered by disabled people are amplified and obfuscated by AI technologies because they do not challenge normative power dynamics about whose knowledge is valued [95]. For example, disabled people are often excluded from algorithm design decisions such as what data is collected, which parts are considered useful, and what purpose the data will used for, despite being directly impacted by the algorithm.

Another way that disabled people's knowledge making is undermined by AI is the emphasis on vision (ocularcentrism) above other forms of sense-making [17]. Computer vision has emerged as one of the most active areas of both machine learning and AI-based assistive technology research, with a plethora of datasets, benchmarking challenges [105], and industry support to boost its development. While image description systems do of course contribute to accessibility, they are also furthering an image-centric epistemology in which text descriptions are a subordinate representation to what a sighted person perceives [62].

Pedestalizing of AI-based technologies and shifting of decision-making away from disabled people. Descriptions of AI in recent years is characterized by suggestions of an almost magical omnipotence. AI-based assistive technologies, then, differ from physical assistive technology such as canes in that they are more likely to be viewed as a source of authority in decision making processes [17]. This can be especially pernicious when the AI technologies themselves are biased, which can cause cultural harm to the disability community. For example, understandings of race, gender, disability, and other identities within the disability community can be harmed by incorrect image descriptions [16].

Cultural/linguistic exploitation of the Deaf community as a result of extractive data collection. AI-based assistive technologies aimed at the Deaf community often need data from fluent signers. However, gathering such data can be extractive and harmful if data collectors are not respectful of the cultural importance of signing, for example by failing to offer sign language interpretation of the instructions [103]. Furthermore, if collectors are not from the Deaf community, data collection can be seen as cultural and linguistic appropriation, especially if Deaf people are not included in decisions about how the data will be used and whether it will ultimately benefit the Deaf community [24].

Content recommendation algorithms reproduce stigmatization of disability. As discussed in previous sections, content made by disabled creators tends to be suppressed, whereas content mocking disabled people is difficult to remove [34] [107]. Similar to how people tend to distance themselves from explicitly referencing disability—for example, by using phrases such as "differently abled"—content recommendation algorithms seem to be acting in much the same way, by treating it as a "taboo" topic [107]. In addition, content moderation algorithms often fail to catch actual instances of toxic ableism, while flagging non-toxic comments mentioning disability [60].

4.5.3 Legitimization of the medical model of disability.

Political alienation. We theorize that algorithmic systems which utilize a medical model of disability, centered on deficiency, makes it more likely that disabled subjects primarily understand themselves from a medical lens, disconnected from socio-political histories of oppression and resistance. Recent works in disability studies have explored disability as a political and relational identity, in which solidarity and collective action aligned with disability justice are crucial to improving conditions for disabled people, rather than medical interventions [17].

Medical model delegitimizes lived experiences without accompanying diagnosis. The medical model and its emphasis on clinical diagnosis also has insidious implications for what is considered a representative dataset. A data collection process that relies on diagnostic self disclosure to establish ground truth, for example by analyzing user-generated content from online forums, misses the full diverse range of experiences that may not be encapsulated by a diagnosis. For example, a model trained on posts from an online support community for people with eating disorders may paradoxically prioritize interventions for people already familiar with the clinical context of diagnosis [32].

4.5.4 Political and civic harms.

These harms refer to the ways that algorithmic systems perpetuate the disenfranchisement and undermining of political power of marginalized groups. Disabled people already experience myriad barriers to voting, resulting in a participation rate nearly 10% lower than non-disabled people [110].

Disenfranchisement from biased models. The process of signature matching which attempts to detect fraudulent votes has shown to be biased against people with vision impairments and those with mobility-related disabilities [63]. With the advent of algorithmic elections, and automated signature matching systems, disabled people are at risk for continuing to be disenfranchised by this practice [15]. Algorithmically determined risk assessment scores—used by some counties in the US to determine sentencing, and thus affect one's ability to vote [111]—have a high likelihood of discriminating against disabled people, further limiting civic participation. [26]

Censoring civic participation by disabled content creators. In recent years, social media platforms have become potent grounds for civic participation, especially for marginalized groups such as the disability community [10]. However, these platforms also enact algorithmically-mediated censorship of disabled activists and content creators by fueling content suppression [107]. Coalition-building is an important aspect of civic engagement and community building [11]; content suppression algorithms curtails cross-community collaboration and allyship by siloing disability communities [34], impairing activists' ability to organize across movements. In addition, these platforms may initiate content takedowns, often of expression of identity [34] [52].

4.5.5 Macro socio economic harms.

Devaluing disabled people's labor. Text-to-image models exploit the work of creative individuals while dwindling their economic opportunities; similarly, technologies such as automated sign language translation may be embraced as a cost-saving alternative while reducing work opportunities from sign language interpreters.

Economically inaccessibility of assistive technology. Whether assistive technologies are economically accessible typically depends on the existence of social policy to subsidize their development and production [76]. AI-based assistive technologies, especially those aiming to innovate on existing assistive tech (e.g. a smart cane vs a regular cane), often fall under "experimental" technologies that aren't covered by health insurance[93]. At the same time, their development and maintenance costs are higher due to model training and tuning, furthering the socioeconomic stratification of access.

4.5.6 Environmental harms.

We were unfortunately unable to find examples of harm at the intersection of disability, algorithmic systems and sustainability in responsible AI literature. However, we theorize that personalization—an oft-cited strategy for making AI systems more fair, by adapting to a user's unique patterns [130]—may negatively affect sustainability, as it requires more training of the model, i.e. computational and energy resources. Furthermore, this approach aligns well with the "AI as collaborator" paradigm that returns agency back to the (disabled) user, since models would have the potential to be more portable and community-driven [56].

5 CONCLUSION

This systematic review consolidates study findings that describe algorithmic harms towards disability communities, as articulated by researchers and disabled participants in human-computer interaction, accessibility, and responsible AI from the past five years. Using the taxonomy of algorithmic harms proposed by Shelby et al. [114] as the underlying framework, we identified recurring themes and shared patterns across scholarly works, new categories of harm unique to the intersection of disability and algorithmic systems, and locations where instances of harm may be under-documented.

Wherever possible, we drew upon critical disability studies perspectives in order to situated harms within a broader social, historical, and cultural context of disability. Although this survey of harms is not comprehensive, we hope that it encourages researchers and practitioners to anticipate how algorithmic systems may harm disabled people at micro, meso, and macro levels, as well as how harms are co-produced by social and societal structures. By examining algorithmic harms through a disability studies lens, we hope to stimulate discussions on envisioning more agentic and liberatory futures for disabled people's interactions with algorithmic systems, lead by disability communities.

REFERENCES

[1] 2022. Alienation. In *The Stanford Encyclopedia of Philosophy* (Winter 2022 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University.

[2] Elisa S Abes and Michelle M Wallace. 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 (2018), 545–562.

[3] Calvin Thomsen Adam Borecky and Alex Dubov. 2019. Reweighing the Ethical Tradeoffs in the Involuntary Hospitalization of Suicidal Patients. *The American Journal of Bioethics* 19, 10 (2019), 71–83. https://doi.org/10.1080/15265161.2019.1654557 arXiv:https://doi.org/10.1080/15265161.2019.1654557 PMID: 31557114.

[4] Taslima Akter, Tousif Ahmed, Apu Kapadia, and Manohar Swaminathan. 2022. Shared Privacy Concerns of the Visually Impaired and Sighted Bystanders with Camera-Based Assistive Technologies. *ACM Trans. Access. Comput.* 15, 2, Article 11 (may 2022), 33 pages. https://doi.org/10.1145/3506857

[5] Taslima Akter, Tousif Ahmed, Apu Kapadia, and Swami Manohar Swaminathan. 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* (Virtual Event, Greece) *(ASSETS '20)*. Association for Computing Machinery, New York, NY, USA, Article 32, 14 pages. https://doi.org/10.1145/3373625.3417003

[6] Rahaf Alharbi, Robin N. Brewer, and Sarita Schoenebeck. 2022. Understanding Emerging Obfuscation Technologies in Visual Description Services for Blind and Low Vision People. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 469 (nov 2022), 33 pages. https://doi.org/10.1145/3555570

[7] Rahaf Alharbi, John Tang, and Karl Henderson. 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 605, 15 pages. https://doi.org/10.1145/3544548.3581541

[8] Erin E Andrews, Anjali J Forber-Pratt, Linda R Mona, Emily M Lund, Carrie R Pilarski, and Rochelle Balter. 2019. # SaytheWord: A disability culture commentary on the erasure of "disability". *Rehabilitation psychology* 64, 2 (2019), 111.

[9] Shannon M Arthur, Anna M Linnehan, Justin B Leaf, Noah Russell, Mary Jane Weiss, Amanda N Kelly, Melissa S Saunders, and Robert K Ross. 2023. Concerns about ableism in applied behavior analysis: An evaluation and recommendations. *Education and Training in Autism and Developmental Disabilities* 58, 2 (2023), 127–143.

[10] Brooke E. Auxier, Cody L. Buntain, Paul Jaeger, Jennifer Golbeck, and Hernisa Kacorri. 2019. HandsOffMyADA: A Twitter Response to the ADA Education and Reform Act. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300757

[11] Maria Avila. 2023. *Transformative civic engagement through community organizing*. Taylor & Francis.

[12] Jo Bates, David Cameron, Alessandro Checco, Paul Clough, Frank Hopfgartner, Suvodeep Mazumdar, Laura Sbaffi, Peter Stordy, and Antonio de la Vega de León. 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* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 425–435. https://doi.org/10.1145/3351095.3372832

[13] Eric P. S. Baumer and Micki McGee. 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* (Honolulu, HI, USA) *(AIES '19)*. Association for Computing Machinery, New York, NY, USA, 163–169. https://doi.org/10.1145/3306618.3314292

[14] Liat Ben-Moshe. 2013. 'The Institution Yet to Come': Analyzing incarceration through a disability lens. *The disability studies reader* 4 (2013), 132–145.

[15] Sarah ML Bender. 2022. Algorithmic elections. *Mich. L. Rev.* 121 (2022), 489.

[16] Cynthia L. Bennett, Cole Gleason, Morgan Klaus Scheuerman, Jeffrey P. Bigham, Anhong Guo, and Alexandra To. 2021. "It's Complicated": Negotiating Accessibility and (Mis)Representation in Image Descriptions of Race, Gender, and Disability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Yokohama</city>, <country>Japan</country>, </conf-loc>) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 375, 19 pages. https://doi.org/10.1145/3411764.3445498

[17] Cynthia L. Bennett and Os Keyes. 2020. What is the Point of Fairness? Disability, AI and the Complexity of Justice. *SIGACCESS Access. Comput.* 125, Article 5 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386301

[18] Annette Bernhardt, Lisa Kresge, and Reem Suleiman. 2023. The Data-Driven Workplace and the Case for Worker Technology Rights. *ILR Review* 76, 1 (2023), 3–29. https://doi.org/10.1177/00197939221131558 arXiv:https://doi.org/10.1177/00197939221131558

[19] Clara Berridge and Alisa Grigorovich. 2022. Algorithmic harms and digital ageism in the use of surveillance technologies in nursing homes. *Frontiers in Sociology* 7 (2022). https://doi.org/10.3389/fsoc.2022.957246

[20] Reuben Binns and Reuben Kirkham. 2021. How Could Equality and Data Protection Law Shape AI Fairness for People with Disabilities? *ACM Trans. Access. Comput.* 14, 3, Article 17 (aug 2021), 32 pages. https://doi.org/10.1145/3473673

[21] Charlotte Bird, Eddie Ungless, and Atoosa Kasirzadeh. 2023. Typology of Risks of Generative Text-to-Image Models. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (<conf-loc>, <city>Montréal</city>, <state>QC</state>, <country>Canada</country>, </conf-loc>) *(AIES '23)*. Association for Computing Machinery, New York, NY, USA, 396–410. https://doi.org/10.1145/3600211.3604722

[22] Elizabeth Bondi, Lily Xu, Diana Acosta-Navas, and Jackson A. Killian. 2021. Envisioning Communities: A Participatory Approach Towards AI for Social Good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 425–436. https://doi.org/10.1145/3461702.3462612

[23] Katya Borgos-Rodriguez and Anne Marie Piper. 2023. Understanding Participation among Disabled Creators in Online Marketplaces. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 314 (oct 2023), 28 pages. https://doi.org/10.1145/3610105

[24] Danielle Bragg, Naomi Caselli, Julie A. Hochgesang, Matt Huenerfauth, Leah Katz-Hernandez, Oscar Koller, Raja Kushalnagar, Christian Vogler, and Richard E. Ladner. 2021. The FATE Landscape of Sign Language AI Datasets: An Interdisciplinary Perspective. *ACM Trans. Access. Comput.* 14, 2, Article 7 (jul 2021), 45 pages. https://doi.org/10.1145/3436996

[25] Gracen Mikus Brilmyer. 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 (2022), 167–188.

[26] Lydia XZ Brown, Ridhi Shetty, Matthew U Scherer, and Andrew Crawford. 2022. Ableism and disability discrimination in new surveillance technologies. *Center for Democracy & Technology* (2022).

[27] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html

[28] Maarten Buyl, Christina Cociancig, Cristina Frattone, and Nele Roekens. 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* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1071–1082. https://doi.org/10.1145/3531146.3533169

[29] Bram De Buyser. [n. d.]. GoblinTools. https://goblin.tools/About. Accessed: 2024-08-13.

[30] Fiona Kumari Campbell. 2009. *Internalised Ableism: The Tyranny Within*. Palgrave Macmillan UK, London, 16–29. https://doi.org/10.1057/9780230245181_2

[31] Beiyan Cao, Changyang He, Muzhi Zhou, and Mingming Fan. 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 58, 15 pages. https://doi.org/10.1145/3544548.3581053

[32] Stevie Chancellor, Jessica L. Feuston, and Jayhyun Chang. 2023. Contextual Gaps in Machine Learning for Mental Illness Prediction: The Case of Diagnostic Disclosures. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 332 (oct 2023),

27 pages. https://doi.org/10.1145/3610181

[33] EunJeong Cheon. 2024. Examining Algorithmic Metrics and their Effects through the Lens of Reactivity. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 3179–3192. https://doi.org/10.1145/3643834.3660676

[34] Dasom Choi, Uichin Lee, and Hwajung Hong. 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* (<conf-loc>, <city>New Orleans</city>, <state>LA</state>, <country>USA</country>, </conf-loc>) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 593, 19 pages. https://doi.org/10.1145/3491102.3517574

[35] Shanley Corvite, Kat Roemmich, Tillie Ilana Rosenberg, and Nazanin Andalibi. 2023. Data Subjects' Perspectives on Emotion Artificial Intelligence Use in the Workplace: A Relational Ethics Lens. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 124 (apr 2023), 38 pages. https://doi.org/10.1145/3579600

[36] Julia Deeb-Swihart, Alex Endert, and Amy Bruckman. 2022. Ethical Tensions in Applications of AI for Addressing Human Trafficking: A Human Rights Perspective. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 295 (nov 2022), 29 pages. https://doi.org/10.1145/3555186

[37] Beth DeVolder. 2013. Overcoming the Overcoming Story: A case of "compulsory heroism". *Feminist Media Studies* 13, 4 (2013), 746–754. https://doi.org/10.1080/14680777.2013.805588 arXiv:https://doi.org/10.1080/14680777.2013.805588

[38] Mark Diaz and Nicholas Diakopoulos. 2019. Whose Walkability? Challenges in Algorithmically Measuring Subjective Experience. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 126 (nov 2019), 22 pages. https://doi.org/10.1145/3359228

[39] Nirmala Erevelles. 2014. Crippin'Jim Crow: Disability, dis-location, and the school-to-prison pipeline. In *Disability incarcerated: Imprisonment and disability in the United States and Canada*. Springer, 81–99.

[40] Leah Findlater, Steven Goodman, Yuhang Zhao, Shiri Azenkot, and Margot Hanley. 2020. Fairness issues in AI systems that augment sensory abilities. *SIGACCESS Access. Comput.* 125, Article 8 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386304

[41] Will Fleisher. 2021. What's Fair about Individual Fairness?. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 480–490. https://doi.org/10.1145/3461702.3462621

[42] Carli Friedman. 2019. Mapping ableism: A two-dimensional model of explicit and implicit disability attitudes. *Canadian Journal of Disability Studies* 8, 3 (2019), 95–120.

[43] Jon E. Froehlich, Yochai Eisenberg, Maryam Hosseini, Fabio Miranda, Marc Adams, Anat Caspi, Holger Dieterich, Heather Feldner, Aldo Gonzalez, Claudina De Gyves, Joy Hammel, Reuben Kirkham, Melanie Kneisel, Delphine LabbÉ, Steve J. Mooney, Victor Pineda, ClÁUdia PinhÃO, Ana RodrÍGuez, Manaswi Saha, Michael Saugstad, Judy Shanley, Ather Sharif, Qing Shen, Claudio Silva, Maarten Sukel, Eric K. Tokuda, Sebastian Felix Zappe, and Anna Zivarts. 2022. The Future of Urban Accessibility for People with Disabilities: Data Collection, Analytics, Policy, and Tools. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility* (Athens, Greece) *(ASSETS '22)*. Association for Computing Machinery, New York, NY, USA, Article 102, 8 pages. https://doi.org/10.1145/3517428.3550402

[44] Lidia V Gabis, Odelia Leon Attia, Mia Goldman, Noy Barak, Paula Tefera, Shahar Shefer, Meirav Shaham, and Tally Lerman-Sagie. 2022. The myth of vaccination and autism spectrum. *European Journal of Paediatric Neurology* 36 (2022), 151–158.

[45] Vinitha Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. 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* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 205–216. https://doi.org/10.1145/3593013.3593989

[46] Patricia Garcia, Tonia Sutherland, Niloufar Salehi, Marika Cifor, and Anubha Singh. 2022. No! Re-imagining Data Practices Through the Lens of Critical Refusal. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 315 (nov 2022), 20 pages. https://doi.org/10.1145/3557997

[47] Sanjana Gautam, Pranav Narayanan Venkit, and Sourojit Ghosh. 2024. From Melting Pots to Misrepresentations: Exploring Harms in Generative AI. arXiv:2403.10776 [cs.HC] https://arxiv.org/abs/2403.10776

[48] Marissa Gerchick, Tobi Jegede, Tarak Shah, Ana Gutierrez, Sophie Beiers, Noam Shemtov, Kath Xu, Anjana Samant, and Aaron Horowitz. 2023. The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 1292–1310. https://doi.org/10.1145/3593013.3594081

[49] Amelia Gibson, Kristen Bowen, and Dana Hanson. 2021. We need to talk about how we talk about disability: A critical quasi-systematic Review. https://www.inthelibrarywiththeleadpipe.org/2021/disability/

[50] Rosalie Gillett, Zahra Stardust, and Jean Burgess. 2022. Safety for Whom? Investigating How Platforms Frame and Perform Safety and Harm Interventions. *Social Media + Society* 8, 4 (2022), 20563051221144315. https://doi.org/10.1177/20563051221144315 arXiv:https://doi.org/10.1177/20563051221144315

[51] Anhong Guo, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach, and Meredith Ringel Morris. 2020. Toward Fairness in AI for People with Disabilities SBG@a Research Roadmap. *SIGACCESS Access. Comput.* 125, Article 2 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386298

[52] Oliver L. Haimson, Daniel Delmonaco, Peipei Nie, and Andrea Wegner. 2021. Disproportionate Removals and Differing Content Moderation Experiences for Conservative, Transgender, and Black Social Media Users: Marginalization and Moderation Gray Areas. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 466 (oct 2021), 35 pages. https://doi.org/10.1145/3479610

[53] Aimi Hamraie and Kelly Fritsch. 2019. Crip technoscience manifesto. *Catalyst: Feminism, theory, technoscience* 5, 1 (2019), 1–33.

[54] Margot Hanley, Solon Barocas, Karen Levy, Shiri Azenkot, and Helen Nissenbaum. 2021. Computer Vision and Conflicting Values: Describing People with Automated Alt Text. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 543–554. https://doi.org/10.1145/3461702.3462620

[55] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. 2020. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 501–512. https://doi.org/10.1145/3351095.3372826

[56] Alex Hanna and Tina M. Park. 2020. Against Scale: Provocations and Resistances to Scale Thinking. arXiv:2010.08850 [cs.CY] https://arxiv.org/abs/2010.08850

[57] Bridget Harris and Delanie Woodlock. 2021. 'For my safety': Experiences of technology-facilitated abuse among women with intellectual disability or cognitive disability. (2021).

[58] ROBERT J. HOFFMEISTER. 1996. Cross-cultural Misinformation: What does special education say about Deaf people. *Disability & Society* 11, 2 (1996), 171–190. https://doi.org/10.1080/09687599650023218 arXiv:https://doi.org/10.1080/09687599650023218

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

[60] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Unintended Machine Learning Biases as Social Barriers for Persons with Disabilitiess. *SIGACCESS Access. Comput.* 125, Article 9 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386305

[61] Katherine M. Iverson, Christina M. Dardis, and Terri K. Pogoda. 2017. Traumatic brain injury and PTSD symptoms as a consequence of intimate partner violence. *Comprehensive Psychiatry* 74 (2017), 80–87. https://doi.org/10.1016/j.comppsych.2017.01.007

[62] Emory James Edwards, Kyle Lewis Polster, Isabel Tuason, Emily Blank, Michael Gilbert, and Stacy Branham. 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* (Virtual Event, USA) *(ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 19, 14 pages. https://doi.org/10.1145/3441852.3471222

[63] Sara Jensen. 2021. Signed, Sealed, Delivered? Problems with the Use of Signature Matching to Verify Mail Voter Identity. *Fla. A & M UL Rev.* 16 (2021), 25.

[64] Prerna Juneja and Tanushree Mitra. 2022. Human and Technological Infrastructures of Fact-checking. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 418 (nov 2022), 36 pages. https://doi.org/10.1145/3555143

[65] Shayda Kafai. 2021. *Crip kinship : the disability justice and art activism of Sins Invalid*. Arsenal Pulp Press, Vancouver, BC.

[66] Alison Kafer. 2013. *Feminist, Queer, Crip*. Indiana University Press. http://www.jstor.org/stable/j.ctt16gz79x

[67] Sushant Kafle, Abraham Glasser, Sedeeq Al-khazraji, Larwan Berke, Matthew Seita, and Matt Huenerfauth. 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, Article 4 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386300

[68] Rie Kamikubo, Kyungjun Lee, and Hernisa Kacorri. 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 827, 18 pages. https://doi.org/10.1145/3544548.3581337

[69] Shaun K. Kane, Anhong Guo, and Meredith Ringel Morris. 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 (Virtual Event, Greece) *(ASSETS '20)*. Association for Computing Machinery, New York, NY, USA, Article 42, 14 pages. https://doi.org/10.1145/3373625.3416990

[70] Nadia Karizat, Dan Delmonaco, Motahhare Eslami, and Nazanin Andalibi. 2021. Algorithmic Folk Theories and Identity: How TikTok Users Co-Produce Knowledge of Identity and Engage in Algorithmic Resistance. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 305 (oct 2021), 44 pages. https://doi.org/10.1145/3476046

[71] Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. Taxonomizing and measuring representational harms: A look at image tagging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 14277–14285.

[72] Os Keyes. 2018. The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 88 (nov 2018), 22 pages. https://doi.org/10.1145/3274357

[73] Sara Kingsley, Proteeti Sinha, Clara Wang, Motahhare Eslami, and Jason I. Hong. 2022. "Give Everybody [..] a Little Bit More Equity": Content Creator Perspectives and Responses to the Algorithmic Demonetization of Content Associated with Disadvantaged Groups. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 424 (nov 2022), 37 pages. https://doi.org/10.1145/3555149

[74] Dan M. Kotliar. 2020. Data orientalism: on the algorithmic construction of the non-Western other. *Theory and Society* 49 (2020), 919–939. https://api.semanticscholar.org/CorpusID:225237135

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

[76] Sarah L. Woodin and Arthur Theil. 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* (Virtual Event, USA) *(ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 39, 6 pages. https://doi.org/10.1145/3441852.3476471

[77] Susan Leavy, Eugenia Siapera, and Barry O'Sullivan. 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* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 695–703. https://doi.org/10.1145/3461702.3462598

[78] Kyungjun Lee, Daisuke Sato, Saki Asakawa, Chieko Asakawa, and Hernisa Kacorri. 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* (Virtual Event, USA) *(ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 8, 15 pages. https://doi.org/10.1145/3441852.3471232

[79] Franklin Mingzhe Li, Cheng Lu, Zhicong Lu, Patrick Carrington, and Khai N. Truong. 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, Article 75 (apr 2022), 26 pages. https://doi.org/10.1145/3512922

[80] Hajin Lim, Lisa Kakonge, Yaxin Hu, Lyn Turkstra, Melissa Duff, Catalina Toma, and Bilge Mutlu. 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 335, 19 pages. https://doi.org/10.1145/3544548.3581222

[81] Jin-Ding Lin, Lan-Ping Lin, Pei-Ying Lin, Jia-Lin Wu, Fang-Yu Kuo, et al. 2010. Domestic violence against people with disabilities: prevalence and trend analyses. *Research in developmental disabilities* 31, 6 (2010), 1264–1268.

[82] Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. 2023. Stable Bias: Evaluating Societal Representations in Diffusion Models. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 56338–56351. https://proceedings.neurips.cc/paper_files/paper/2023/file/b01153e7112b347d8ed54f317840d8af-Paper-Datasets_and_Benchmarks.pdf

[83] Shrehan Lynch and Joanne Hill. 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 (2021), 507–520.

[84] Anna Ma, Elizabeth Patitsas, and Jonathan Sterne. 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* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 639–650. https://doi.org/10.1145/3593013.3594032

[85] Kelly Avery Mack, Rida Qadri, Remi Denton, Shaun K. Kane, and Cynthia L. Bennett. 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* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 288, 23 pages. https://doi.org/10.1145/3613904.3642166

[86] Jennifer M. Mays. 2006. Feminist disability theory: domestic violence against women with a disability. *Disability & Society* 21, 2 (2006), 147–158. https://doi.org/10.1080/09687590500498077 arXiv:https://doi.org/10.1080/09687590500498077

[87] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (nov 2019), 23 pages. https://doi.org/10.1145/3359174

[88] Robert McRuer. 2010. Compulsory able-bodiedness and queer/disabled existence. *The disability studies reader* 3 (2010), 383–392.

[89] Katelyn Mei, Sonia Fereidooni, and Aylin Caliskan. 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* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 1699–1710. https://doi.org/10.1145/3593013.3594109

[90] Zion Mengesha, Courtney Heldreth, Michal Lahav, Juliana Sublewski, and Elyse Tuennerman. 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 (2021). https://doi.org/10.3389/frai.2021.725911

[91] Danaë Metaxa, Michelle A. Gan, Su Goh, Jeff Hancock, and James A. Landay. 2021. An Image of Society: Gender and Racial Representation and Impact in Image Search Results for Occupations. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 26 (apr 2021), 23 pages. https://doi.org/10.1145/3449100

[92] Jennifer Mickel. 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* (Rio de Janeiro, Brazil) *(FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 2484–2494. https://doi.org/10.1145/3630106.3659050

[93] Rezylle Milallos, Vinita Tibdewal, Yiwen Wang, Andre Ogueh Udegbe, and Tae Oh. 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* (Virtual Event, USA) *(ASSETS '21)*. Association for Computing Machinery, New York, NY, USA, Article 98, 3 pages. https://doi.org/10.1145/3441852.3476524

[94] Meredith Ringel Morris, Adam Fourney, Abdullah Ali, and Laura Vonessen. 2018. Understanding the Needs of Searchers with Dyslexia. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173609

[95] Ian Moura. 2023. Encoding normative ethics: On algorithmic bias and disability. *First Monday* 28, 1 (Feb. 2023). https://doi.org/10.5210/fm.v28i1.12905

[96] Carlyn O. Mueller, Anjali J. Forber-Pratt, and Julie Sriken. 2019. Disability: Missing from the Conversation of Violence. *Journal of Social Issues* 75, 3 (2019), 707–725. https://doi.org/10.1111/josi.12339 arXiv:https://spssi.onlinelibrary.wiley.com/doi/pdf/10.1111/josi.12339

[97] Karen Nakamura. 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* (Pittsburgh, PA, USA) *(ASSETS '19)*. Association for Computing Machinery, New York, NY, USA, 1–2. https://doi.org/10.1145/3308561.3353812

[98] Terrence Neumann, Maria De-Arteaga, and Sina Fazelpour. 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* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1504–1515. https://doi.org/10.1145/3531146.3533205

[99] Denis Newman-Griffis, Jessica Sage Rauchberg, Rahaf Alharbi, Louise Hickman, and Harry Hochheiser. 2022. Definition drives design: Disability models and mechanisms of bias in AI technologies. arXiv:2206.08287 [cs.AI] https://arxiv.org/abs/2206.08287

[100] Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism.* New York University Press, New York, USA. https://doi.org/doi:10.18574/nyu/9781479833641.001.0001

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

[102] Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, et al. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *bmj* 372 (2021).

[103] Joon Sung Park, Danielle Bragg, Ece Kamar, and Meredith Ringel Morris. 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* (Virtual Event, Canada) *(FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 52–63. https:

//doi.org/10.1145/3442188.3445870
[104] Ethan Z. Rong, Mo Morgana Zhou, Zhicong Lu, and Mingming Fan. 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* (<conf-loc>, <city>Virtual Event</city>, <country>Australia</country>, </conf-loc>) *(DIS '22)*. Association for Computing Machinery, New York, NY, USA, 571–585. https://doi.org/10.1145/3532106.3533514
[105] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115 (2015), 211–252.
[106] Shruti Sannon and Dan Cosley. 2022. Toward a More Inclusive Gig Economy: Risks and Opportunities for Workers with Disabilities. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 335 (nov 2022), 31 pages. https://doi.org/10.1145/3555755
[107] Shruti Sannon, Jordyn Young, Erica Shusas, and Andrea Forte. 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* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 672, 15 pages. https://doi.org/10.1145/3544548.3581333
[108] Morgan Klaus Scheuerman, Madeleine Pape, and Alex Hanna. 2021. Autoessentialization: Gender in automated facial analysis as extended colonial project. *Big Data & Society* 8, 2 (2021), 20539517211053712. https://doi.org/10.1177/20539517211053712 arXiv:https://doi.org/10.1177/20539517211053712
[109] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 58 (may 2020), 35 pages. https://doi.org/10.1145/3392866
[110] Lisa Schur, Mason Ameri, and Meera Adya. 2017. Disability, Voter Turnout, and Polling Place Accessibility. *Social Science Quarterly* 98, 5 (2017), 1374–1390. https://doi.org/10.1111/ssqu.12373 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ssqu.12373
[111] Katia Schwerzmann. 2021. Abolish! Against the use of risk assessment algorithms at sentencing in the US Criminal Justice System. *Philosophy & Technology* 34, 4 (2021), 1883–1904.
[112] Nick Seaver. 2019. Captivating algorithms: Recommender systems as traps. *Journal of Material Culture* 24, 4 (2019), 421–436. https://doi.org/10.1177/1359183518820366 arXiv:https://doi.org/10.1177/1359183518820366
[113] Ather Sharif, Aedan Liam McCall, and Kianna Roces Bolante. 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* (Athens, Greece) *(ASSETS '22)*. Association for Computing Machinery, New York, NY, USA, Article 10, 18 pages. https://doi.org/10.1145/3517428.3544813
[114] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 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* (<conf-loc>, <city>Montréal</city>, <state>QC</state>, <country>Canada</country>, </conf-loc>) *(AIES '23)*. Association for Computing Machinery, New York, NY, USA, 723–741. https://doi.org/10.1145/3600211.3604673
[115] Ashley Shew. 2022. How to get a story wrong: Technoableism, simulation, and cyborg resistance. *Including Disability Journal* 1, 1 (2022).
[116] Gary Shkedy, Dalia Shkedy, and Aileen H Sandoval-Norton. 2021. Long-term ABA therapy is abusive: A response to Gorycki, Ruppel, and Zane. *Advances in Neurodevelopmental Disorders* 5, 2 (2021), 126–134.
[117] Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 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* (<conf-loc>, <city>Arlington</city>, <state>VA</state>, <country>USA</country>, </conf-loc>) *(EAAMO '22)*. Association for Computing Machinery, New York, NY, USA, Article 1, 6 pages. https://doi.org/10.1145/3551624.3555285
[118] Abigale Stangl, Emma Sadjo, Pardis Emami-Naeini, Yang Wang, Danna Gurari, and Leah Findlater. 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* (<conf-loc>, <city>Austin</city>, <state>TX</state>, <country>USA</country>, </conf-loc>) *(W4A '23)*. Association for Computing Machinery, New York, NY, USA, 134–147. https://doi.org/10.1145/3587281.3587296
[119] Abigale Stangl, Kristina Shiroma, Bo Xie, Kenneth R. Fleischmann, and Danna Gurari. 2020. Visual Content Considered Private by People Who Are Blind. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, Greece) *(ASSETS '20)*. Association for Computing Machinery, New York, NY, USA, Article 31, 12 pages. https://doi.org/10.1145/3373625.3417014
[120] Cella M Sum, Rahaf Alharbi, Franchesca Spektor, Cynthia L Bennett, Christina N Harrington, Katta Spiel, and Rua Mae Williams. 2022. Dreaming Disability Justice in HCI. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, Article 114, 5 pages. https://doi.org/10.1145/3491101.3503731
[121] Cella M Sum, Franchesca Spektor, Rahaf Alharbi, Leya Breanna Baltaxe-Admony, Erika Devine, Hazel Anneke Dixon, Jared Duval, Tessa Eagle, Frank Elavsky, Kim Fernandes, et al. 2024. Challenging Ableism: A Critical Turn Toward Disability Justice in HCI. *XRDS: Crossroads, The ACM Magazine for Students* 30, 4 (2024), 50–55.
[122] Luhang Sun, Mian Wei, Yibing Sun, Yoo Ji Suh, Liwei Shen, and Sijia Yang. 2024. Smiling women pitching down: auditing representational and presentational gender biases in image-generative AI. *Journal of Computer-Mediated Communication* 29, 1 (02 2024), zmad045. https://doi.org/10.1093/jcmc/zmad045 arXiv:https://academic.oup.com/jcmc/article-pdf/29/1/zmad045/56546560/zmad045.pdf
[123] John Tang. 2021. Understanding the Telework Experience of People with Disabilities. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 30 (apr 2021), 27 pages. https://doi.org/10.1145/3449104
[124] Emily Shea Tanis and Clayton Lewis. 2020. Artificial Intelligence and the Dignity of Risk. *SIGACCESS Access. Comput.* 125, Article 7 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386303
[125] Jutta Treviranus. 2019. The Value of Being Different. In *Proceedings of the 16th International Web for All Conference* (San Francisco, CA, USA) *(W4A '19)*. Association for Computing Machinery, New York, NY, USA, Article 1, 7 pages. https://doi.org/10.1145/3315002.3332429
[126] Shari Trewin. 2018. AI Fairness for People with Disabilities: Point of View. *CoRR* abs/1811.10670 (2018). arXiv:1811.10670 http://arxiv.org/abs/1811.10670
[127] Shari Trewin, Sara Basson, Michael Muller, Stacy Branham, Jutta Treviranus, Daniel Gruen, Daniel Hebert, Natalia Lyckowski, and Erich Manser. 2019. Considerations for AI Fairness for People with Disabilities. *AI Matters* 5, 3 (dec 2019), 40–63. https://doi.org/10.1145/3362077.3362086
[128] Tom van Nuenen, Jose Such, and Mark Cote. 2022. Intersectional Experiences of Unfair Treatment Caused by Automated Computational Systems. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 445 (nov 2022), 30 pages. https://doi.org/10.1145/3555546
[129] Georgia van Toorn and Jackie Leach Scully. 2023. Unveiling algorithmic power: exploring the impact of automated systems on disabled people's engagement with social services. *Disability & Society* 0, 0 (2023), 1–26. https://doi.org/10.1080/09687599.2023.2233684 arXiv:https://doi.org/10.1080/09687599.2023.2233684
[130] Mike Wald. 2021. AI data-driven personalisation and disability inclusion. *Frontiers in artificial intelligence* 3 (2021), 571955.
[131] Angelina Wang, Vikram V Ramaswamy, and Olga Russakovsky. 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* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 336–349. https://doi.org/10.1145/3531146.3533101
[132] Lindsay Weinberg. 2022. Rethinking Fairness: An Interdisciplinary Survey of Critiques of Hegemonic ML Fairness Approaches. *ArXiv* abs/2205.04460 (2022). https://api.semanticscholar.org/CorpusID:248622866
[133] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, bias, and AI. *AI Now Institute* 8 (2019).
[134] Lauren Wilcox, Robin Brewer, and Fernando Diaz. 2023. AI Consent Futures: A Case Study on Voice Data Collection with Clinicians. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 316 (oct 2023), 30 pages. https://doi.org/10.1145/3610107
[135] Daniel A Wilkenfeld and Allison M McCarthy. 2020. Ethical concerns with applied behavior analysis for autism spectrum" disorder". *Kennedy Institute of Ethics Journal* 30, 1 (2020), 31–69.
[136] Rua M. Williams, Simone Smarr, Diandra Prioleau, and Juan E. Gilbert. 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 (2022), 67–74. https://doi.org/10.1109/TTS.2021.3084913
[137] Robert A. Wilson. 2021. Dehumanization, Disability, and Eugenics. In *Routledge Handbook of Dehumanization*, Maria Kronfeldner (Ed.). Routledge, 173–186.
[138] Christine T. Wolf and Kathryn E. Ringland. 2020. Designing Accessible, Explainable AI (XAI) Experiences. *SIGACCESS Access. Comput.* 125, Article 6 (mar 2020), 1 pages. https://doi.org/10.1145/3386296.3386302
[139] Di Wu. 2023. Good for tech: Disability expertise and labor in China's artificial intelligence sector. *First Monday* 28, 1 (Feb. 2023). https://doi.org/10.5210/fm.v28i1.12887