



Contributing to Accessibility Datasets: Reflections on Sharing Study Data by Blind People

Rie Kamikubo
College of Information Studies
University of Maryland, College Park
United States
rkamikub@umd.edu

Kyungjun Lee
Department of Computer Science
University of Maryland, College Park
United States
kyungjun@umd.edu

Hernisa Kacorri
College of Information Studies
University of Maryland, College Park
United States
hernisa@umd.edu

ABSTRACT

To ensure that AI-infused systems work for disabled people, we need to bring accessibility datasets sourced from this community in the development lifecycle. However, there are many ethical and privacy concerns limiting greater data inclusion, making such datasets not readily available. We present a pair of studies where 13 blind participants engage in data capturing activities and reflect with and without probing on various factors that influence their decision to share their data via an AI dataset. We see how different factors influence blind participants' willingness to share study data as they assess risk-benefit tradeoffs. The majority support sharing of their data to improve technology but also express concerns over commercial use, associated metadata, and the lack of transparency about the impact of their data. These insights have implications for the development of responsible practices for stewarding accessibility datasets, and can contribute to broader discussions in this area.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; **Accessibility**; • **Social and professional topics** → **People with disabilities**; • **Security and privacy** → Human and societal aspects of security and privacy.

KEYWORDS

datasets; privacy; human-centered AI; disability; data ownership

ACM Reference Format:

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 (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3544548.3581337>

1 INTRODUCTION

With advances in artificial intelligence (AI), there is a potential for emerging technologies to improve the lives of people who experience barriers to inclusion such as the disability community. We have

seen many efforts in this direction. For example, many AI-infused¹ assistive applications for supporting blind² people, the community of focus in this work, employ computer vision for better access to physical and virtual spaces [15, 51, 72, 99]. However, the potential benefits may not be realized if the data used to build these systems do not represent the end users and the contexts within which they operate [52, 101, 147]. On the contrary, they may harm. Yet, the majority of large computer vision models are trained on photos taken by sighted people [43, 111, 114], performing poorly on photos taken by blind users [20, 72, 139], a gap that is only increasing [33].

Despite their critical role, researchers have identified myriads of challenges in collecting and sharing accessibility datasets [1, 16, 101, 138]. Primary barriers are privacy and ethical considerations to protect those represented in the data [1, 55, 101, 137]. Collecting data from small populations increases the risk of re-identification, which can amplify concerns for further discrimination pertaining to sensitive disability status [1, 101]. Sharing accessibility datasets also poses risks of data abuse and misuse without proper laws and regulation enforcement (e.g., building a hiring algorithm that can make biased decisions based on disability [41]).

With this paper, we contribute to discussions around increasing the availability of accessibility datasets by surfacing the motivating and challenging factors involved in data sharing decisions of disabled people. We focus on the blind community and image data, a challenging scenario where those contributing the data may not be able to inspect them when deciding to share. We designed a pair of studies that aim to surface blind participants' perspectives on data sharing in a situated, rather than simulated (e.g., [55, 96, 107]), context. To achieve this, we teamed up with researchers who were interested in evaluating an AI-infused application in the homes of blind participants. The application, a teachable object recognizer [72], was deployed on smartglasses. Blind participants took photos and used them to finetune a computer vision model and test its performance. The team was interested in sharing study data with the broader research community and was looking at best practices. Situated in this context, we designed a semi-structured interview as a follow up. Typically within the span of a few days, 13 blind participants both (i) performed data capturing activities, and (ii) were interviewed on their perspectives towards sharing their study data via a public AI dataset.



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

CHI '23, April 23–28, 2023, Hamburg, Germany
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9421-5/23/04.
<https://doi.org/10.1145/3544548.3581337>

¹We follow the term used by Amershi *et al.* [6] referring to having “features harnessing AI capabilities that are directly exposed to the end user.”

²In this paper, we adopt identity-first language (e.g., “blind people” and “disabled people”) instead of person-first (e.g., “people with visual impairments” and “people with disabilities”). We acknowledge that there are active and ongoing discussions around the two with some of them captured in Sharif *et al.* [120].

We found various factors that could play into blind participants' willingness to share their study data (*i.e.*, photos and labels of objects), revealing the need for better assessment of benefits and risks. Many perceived potential risks such as re-identification as minimal and supported sharing practices to improve AI-infused technology for greater benefit to both disabled and non-disabled people. Yet, they were hesitant to contribute their data for commercial purposes and companies handling the use of their data, due to notions of "distrust" even though almost all frequently used AI-infused applications (*e.g.*, SeeingAI, SuperSense, and Lookout) to read text and identify objects or shared their camera view with sighted helpers (*e.g.* via Aira and BeMyEyes). They also expressed concerns for sharing demographic metadata (*e.g.*, age, gender, race) along their study data, due to not only privacy and safety but also the ambiguity of its value for building AI-infused technology. This suggests that the process of collecting and stewarding accessibility datasets requires greater transparency of data use, especially to challenge inclusivity issues in AI fairness [30, 76]. Some participants showed further interest in learning about the potential impact of their data, an option that is neither supported in current informed consent processes nor is there a way to practically implement it yet.

Our intention is to bring potential data contributors from the disability community to the forefront of data sharing discussions. We acknowledge that our focus on a specific population and context limits the generalizability of our insights. To overcome some of these limitations, we carefully connect and contrast our observations with existing literature. More so, we incorporate prior questionnaires (*e.g.*, Park *et al.* [107]) with the disability community on related topics. To facilitate replicability, we share our questionnaire with expanded scenarios and interview questions (more than 80% new content). We see the main contribution of this work being empirical. By investigating data sharing from the perspectives and experiences of blind people, we contribute to the larger call-to-action for the research community, industry, and policy makers in shaping future data practices that are inclusive of disability. We also see how our approach of eliciting participants' perspectives before asking them to decide on whether they want their study data to be shared and how, could be leveraged by future researchers who want to engage participants in decisions around sharing of their study data.

2 RELATED WORK

In this section, we cover prior work on creating and sharing accessibility datasets to provide a clear picture of the challenges and significance. We also extend to current efforts in informing data practices across different disciplines, with few studies exploring how disabled people view their data being sourced and used.

2.1 Accessibility Datasets

The need for data is growing in the field of accessibility, especially for accelerating innovation around assistive technology [22, 69, 101]. Notable examples include photos taken by blind users to build object recognition applications [84, 129, 150] and sign language videos from Deaf signers to train machine translation applications [56, 62]. To facilitate the discovery and re-use of currently available data in this space, Kacorri *et al.* [70, 71] put together

a collection of accessibility datasets sourced from disabled people over the last decade that could be leveraged for training and evaluating machine learning models. However, dataset availability is found to be sparse across different communities of focus [75], with challenges in data diversity persisting for greater inclusion of marginalized communities [76]. More so, discussions around unique challenges for data collection involving disability communities are ongoing [16, 23, 101, 118, 138]. For example, Blaser and Ladner [16] raised issues with inconsistent measures of how disability is elicited, making it difficult to aggregate and combine different data sources to facilitate large-scale datasets.

We have seen efforts to address the lack of larger, more diverse datasets in the field. Some leverage crowdsourcing or telemetry data collection methods [19, 21], while others deploy assistive applications (*e.g.*, VizWiz [15]) in the real-world [54, 73, 93]. Indeed, these strategies have complemented data contributions from certain communities of focus; those sourced from the blind community typically include larger numbers of contributors compared to other disability communities [75]. Blind people have been early adopters of technologies, often in the context of taking photos of objects/scenes to access visual information [2, 15]. However, given that blind people cannot inspect their data such as the photos they took, it is left to data stewards to protect the privacy and safety of those represented in the datasets [54, 135]. Thus, we see the opportune involvement of this community to discuss how to ethically contribute data.

2.2 Tensions in Data Sharing

Kop [79] cited data sharing as an essential practice for a successful AI ecosystem in analyzing and processing high-quality training datasets. In general, many academic disciplines and industries have seized the opportunity to promote open datasets [5, 32, 90, 112] to drive innovation or create new knowledge and shared resources. The health community is not an exception (*e.g.*, using patient data to improve care and outcomes [89, 100, 141, 145]). In some venues, researchers are required to submit data sharing statements for clinical trials along with their manuscripts [132]. Even so, such data sharing schemes have raised issues as data subjects' preferences and control are rarely addressed, with their participation limited in governance structures for sharing medical data [82]. Advances in pervasive and wearable technologies also bring attention to enabling users to track and share self-collected data via health apps [77, 88, 98, 122]. At the same time, scholars have warned of the risks including privacy and security surrounding the use of data by third parties which is not always transparent [121, 133, 143].

Similar conversations are seen in accessibility, calling for a careful balancing act in how data are shared [1, 68, 135, 147]. There are many ongoing issues with disability-inclusive data; they are highly sensitive and can raise concerns for privacy and data protection [1, 101, 138]. While such concerns are prevalent across disciplines [44, 144], the consequences are severe in accessibility, with the risk of re-identification along the potential for discrimination [1, 103, 147]. Data sharing also raises ethical considerations for data re-use, when it could lead to abuse and misuse outside of the original intention [147]. Considering the possibility of algorithms to detect disability status [146], accessibility datasets could exacerbate further bias and marginalization through systems built [101].

2.3 Data Contributors in Future Data Practices

All the concerns related to potential risks of data sharing practices make discussions and frameworks around data ethics more pressing. Recently, there is a body of work within and beyond accessibility to involve potential data contributors in the data collection or sharing contexts to inform future practices [50, 55, 96, 102, 104, 107, 119]. In addition, these prior efforts surfaced a variety of situational and contextual factors that could influence the contributors' judgments of concern and risk. For instance, while many disabled people were open to contributing data to an AI dataset with the prospect of helping the disability community, they were hesitant depending on the data type that could be more or less personal [107]. Meanwhile, Shah *et al.* [119], in a biomedical domain, found that with whom data are shared played a role in their judgments than what data types are shared. Mozersky *et al.* [102] also observed a sense of trust towards researchers, receiving broad support for data contributions in qualitative research. In contrast, McNaney *et al.* [97] identified fears and concerns around how commercial companies might use health data (e.g., targeted advertising). Privacy and security concerns were common themes across these research areas, yet still depended on a number of elements including represented populations (e.g., visible vs hidden disability groups [107]) or awareness through consent [50].

Our research complements previous work that has highlighted multifaceted motivations and concerns relating to data sharing. It also prompts the unaddressed questions concerning how accessibility datasets should be sourced and used to drive an AI ecosystem, when data contributors' broader views on data sharing are often contextual and situationally dependent, such as impacted by data types, research or application domains, or data use purposes. In this work, we add new dimensions to these discussions by enabling potential data contributors to reflect on the contexts of sharing data sourced from settings where the technology is being deployed (i.e., home). Furthermore, prior efforts investigating concerns of disabled people to inform better technical and legal frameworks for data stewardship (e.g., [96, 107]) are conducted with regards to simulated datasets and environments. Prior literature has warned about the impact of direct vs. indirect experiences on the development of knowledge, attitudes, and behavior [37], which can engender differing opinions on issues such as privacy [85] or risk beliefs [142]. Conversations around data sharing practices related to accessibility need to be further attuned to the communities of focus, with the blind community being centered in this work, capture the extent to which implications can be drawn about how they would want data sharing to occur in the real world e.g., while participating in a user study or engage with technology in their homes.

3 METHODS

To have direct conversations with the blind community on how data that they may contribute should be shared via an AI dataset, we consider a cross-sectional study design with blind people who have been exposed to a novel AI-infused assistive application and asked to evaluate it. Specifically, we teamed up with researchers who were developing an object recognition application on smartglasses and deploying it in the homes of blind participants. The team was interested in best ways to share the study data used in their analysis

(e.g., photos of objects) both for the purpose of replicability but also for motivating future work in this area e.g. use the data to train or test future machine learning models.

As shown in Figure 1, this pairing of studies allows us to surface blind participants' perspectives in a situated, rather than simulated, context—enabling their decisions about sharing study data and preferences for data control to be reflected in a real-world sharing context. Unlike previous research where the findings are synthesized as implications for broader research practices [107, 119], we shift to active participation of potential data contributors to guide the data sharing process that the development team of the AI-infused application will go through to release the AI dataset from the study. Also, their user study provided great context to gain empirical insights that would not be as generalizable from an in-lab study; typically, datasets have value when collected in naturalistic settings (i.e., where the technology is meant to be deployed such as in people's home [135]). More so, privacy risks are more heightened in these settings [57]. The camera in the smartglasses may capture the home environment in the background potentially revealing more about the person and their life.

The larger study spanned multiple days within the May–September 2022 period. It started with a 30-minute long Zoom call to capture participant demographics, attitudes, and experience with technology. Some of this information is presented in Section 3.1 to provide context for our analysis. A day or two later, participants joined a longer study where they performed remotely from their homes a series of data capturing activities presented in Section 3.2. Usually within a week, they participated via Zoom in a semi-structured interview, the focus of this paper. Any time after this, participants could indicate to the researcher, the data stewards, their decision around sharing of their study data as a response to an email they received. The need to make a decision was communicated with participants early on in the study and was included in their consent forms. Participants could also opt to join a follow up co-design session, typically conducted a few weeks later. This last session focused on the design of an accessible data inspection interface that allowed participants to go over the photos they collected. Some opted to confirm their decision on data sharing after this session.

We briefly describe the system evaluation study as it provides the critical context of the data capturing activities that the participants engaged in. However, the specifics of the technological contributions and findings from that study are beyond the scope of this paper. Data from the co-design session and participants' final decisions to share remain yet to be analyzed. Our semi-structured interview is described in detail in Section 3.3 followed by our analysis approach (Section 3.4), which aims to reveal (i) the factors related to one's decision to share study data as well as (ii) potential ethical, legal, and technical implications for mitigating risks and concerns related to data sharing. To facilitate future research on exploring perspectives from other communities or in other AI data sharing contexts we make our scenarios, context probes, and questions available at https://go.umd.edu/datasharing_questionnaire.

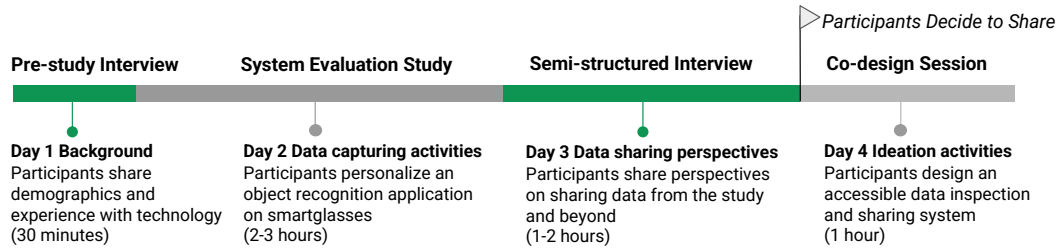


Figure 1: We present findings from a semi-structured interview contextualized within a larger study that includes a short interview on demographics and experiences with technology, a system evaluation study, and a follow up co-design session.

3.1 Recruitment and Participants

Our pre-study interview (day 1) captured demographic information including age, gender, education, and occupation as well attitudes and experiences with technology that are relevant to the AI-infused application being evaluated. At the end of the semi-structured interview (day 3), participants were given an option to choose the information that they do not wish to be made available on publication. Reflecting their consent, Table 1 shows the demographics for our 13 blind participants. Nine were totally blind and four were legally blind. On average, participants were 53.46 years old (STD=14.94). Out of the responses we received, six self-identified as women and five as men. Participants were compensated \$15 per hour, with an average of 2.75 hours (STD=0.5) spent for the experimental study (including the opening demographic and experience questionnaires) and 1.72 hours (STD=0.33) for the interview on data sharing.

To better contextualize our findings, we report participants' technology use and attitudes responses. All but two (P1, P11) participants reported using assistive applications for accessing visual information in their surroundings such as Seeing AI (n=8), Aira (n=4), BeMyEyes (n=4), VoiceDreamReader (n=3), BlindSquare (n=2), BlindShell (n=1), BeSpecular (n=1), CurrencyReader (n=1), ColorIdentifier (n=1), Google Lookout (n=1), KNFBReader (n=1), and Supersense (n=1). P1 and P11, who are legally blind, typically relied on built-in camera features such as magnification "to read bus signs" (P1) or "to check their [own] appearance" (P11). As shown in Figure 2, more than half of the participants (n=8) reported sharing photos or videos with others at least once a month. Often this was done to get sighted help from family and friends for recognition tasks. Some (n=5) never did. Sharing of voice and audio recordings was often less frequent.

Most participants were positive about the potential of AI and technology, as shown in Figure 3. Indeed, all agreed or strongly agreed on statements such as "It is important to keep up with tech" and "Feel more accomplished due to tech." Almost half of the participants (n=6) disagreed or strongly disagreed that they enjoy recording their activities. When it came to videos, photos, and sounds or voices, some disagreed (n=2, n=3, and n=3, respectively).

3.2 Evaluation Study: Data Capturing Context

During this session, participants evaluated a working prototype of a teachable object recognition application deployed on smartglasses. The term *teachable* refers to the fact that participants could *teach* the underlying machine learning model to recognize objects of their

Table 1: Self-reported participant information including participant ID, vision level, age, gender, education, occupation, and AI familiarity on a 4-point scale: 1 = not familiar at all (have never heard of it), 2 = slightly familiar (have heard of it but don't know what it does), 3 = somewhat familiar (have a broad understanding of what it is and what it does), 4 = extremely familiar (have extensive knowledge). A dash (-) indicates that the participant did not consent to disclose.

PID	Vision	Age	Gender	Education	Occupation	AI Familiarity
P1	legally blind	59	man	bachelor	budget analyst	somewhat
P2	totally blind	-	-	-	-	somewhat
P3	totally blind	72	woman	bachelor	retired	extremely
P4	totally blind	67	man	master	retired	slightly
P5	totally blind	48	man	master	social worker	somewhat
P6	legally blind	32	woman	bachelor	operations	somewhat
P7	totally blind	53	woman	master	retired	somewhat
P8	totally blind	67	woman	master	retired	somewhat
P9	totally blind	-	-	-	IT solution architect	somewhat
P10	totally blind	37	woman	law school	attorney	somewhat
P11	legally blind	52	woman	master	application developer	somewhat
P12	legally blind	71	man	post-bachelor	retired	somewhat
P13	totally blind	67	man	some college	computer technician	extremely

choice by providing a few photo examples of those objects as well as labels (*i.e.* object names that are spoken upon recognition). The application is meant to facilitate personalization as it promises a better fit for real-world scenarios by significantly constraining the machine learning task to a specific user and their environment. It does not require any machine learning expertise from the user. More so, the interactive nature of teachable applications could help users uncover basic machine learning concepts and gain familiarity with

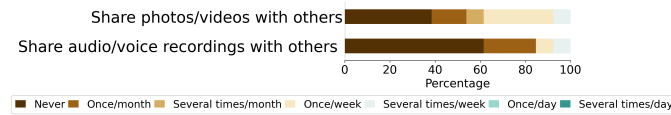


Figure 2: Responses to how frequently participants shared data such as photos, videos, and audio recordings with others.

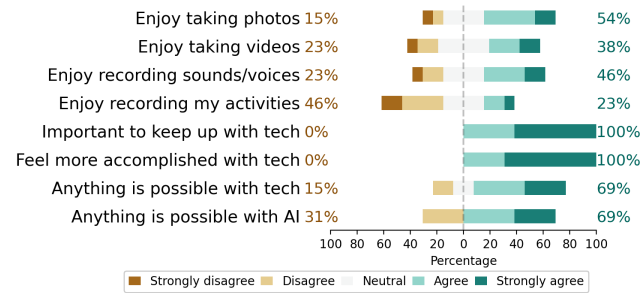


Figure 3: Participants' attitudes towards AI, technology, and data capturing such as tracking their activities, recording sound or voices, and taking videos and photos.

AI (e.g., [34, 39, 59, 61, 108]). We see similar evidence for studies with blind participants [60, 72] where participants reflect on the value of diversity in training data. Thus, the data capturing tasks in this study seemed a great fit providing a realistic data contribution scenario while exposing participants to the value of AI data for training and testing.

In detail, participants were instructed to find a sitting area in their homes where they feel comfortable setting up the laptop with the Zoom call and interacting with the stimuli objects while wearing a pair of smartglasses. To familiarize themselves with the smartglasses and the application, they first practice taking photos and providing labels for 2 objects. The research team provided practice objects. As shown in Figure 4, participants used the touchpad on the smartglasses located along the temple to navigate the menu and trigger the photo taking and labeling functions, which are communicated through text-to-speech. Voice commands were also implemented for entering, correcting, and confirming the object label. Once familiar with the system, participants are asked to complete data capturing activities that involved taking multiple photos per object for a total of 6 objects and providing associated labels for training and evaluating a classifier. Half of these objects were stimuli engineered to be visually distinct but nearly identical by touch (i.e. different bags of snacks). They were fixed across all participants and were provided by the research team. The rest of the objects were up to the participants; they could choose anything in their home. Typically, they opted for somewhat similar objects to the stimuli including everyday products, as shown in Figure 5. Participants answered questions related to their experience in between the data capturing tasks and at the end.

By the end of this session, each participant generated on average 222 photos (STD=59.9) across the 6 object labels. Both the photos and the labels generated from these activities were referred to the



Figure 4: A blind participant using an object recognition application installed on smart glasses, which they use to take multiple photos of a soda can. The object recognition model is fine-tuned on these photos to personalize the application.



Figure 5: Examples of photos of objects captured and labeled by our blind participants (from left to right) as: "Cereal," "Videotape," "Carrots," and "Panera Cup."

participants as "your study data" throughout the communication with the research team. This wording aimed to situate participants in the context of sharing via a public AI dataset.

3.3 Semi-structured Interview: Reflections

The interview was conducted via Zoom and was audio-recorded for analysis. We used scenarios and context probes (Table 2) to guide the interviews, with 15% of the questions either re-used or expanded from the questionnaire shared by Park *et al.* [107]. The interviews were structured as follows:

Part 1 Benefits. We first asked our participants about their understanding of any benefits in sharing their study data via a public AI dataset. We then presented a scenario describing potential benefits (Table 2 Part 1) to probe their willingness for data sharing. We followed up with a few questions to gauge their motivations.

Part 2 Risks. We asked participants about their understanding of any risks in sharing their study data via a public AI dataset. We then presented two scenarios describing potential risk cases raised in the field: (i) re-identification of individuals from anonymized datasets [101] and (ii) data abuse/misuse given that "it's difficult to ensure that data won't be reused in ways that could cause harm" [147]. We followed up with a few questions to gauge their concerns including non-consenting disability disclosure.

Part 3 Contexts. We then explored the contexts that may impact their decision to share data (within and beyond this study). We asked about their level of comfort depending on the type of modality, object, environment, and demographic information being shared. Similar to the design of validated questionnaires for measuring attitudes (e.g., [45]), we employed a 7-point Likert scale and asked for rationale where it was appropriate. For example, we asked "On a scale of 1 to 7, where 1 is not comfortable at all and 7 is very comfortable, please rate your level of comfort with sharing photos

Table 2: Scenarios and context probes to guide our semi-structured interviews.

Part 1 Benefits	Part 2 Risks	Part 3 Contexts	Part 4 Mitigating Risks
Broader impact <i>“Datasets may not just benefit blind people via assistive tech but really anyone who may interact with a smart app or appliance...Imagine a robot that one can ask to fetch things for them. Your data could be used to make such robots function better for everyone.”</i>	Re-identification <i>“People can do some guesswork. For example, they could see in the publication that Participant 3, who is someone around 40 years old, identifies as male and uses a guide dog, took a bunch of photos of t-shirts associated with specific events...it turns out they happened to know someone who fits the description.”</i> Non-consenting disclosure: <i>“Would you be concerned if others (your friends or employers) might find out that your data is included in the AI dataset?”</i> Data abuse/misuse <i>“Imagine someone building an algorithm that given an image it can figure out whether a blind person took it. They may not be able to guess who but they may be able to guess if one has a disability or not without their consent for disclosure.”</i>	Type of modality, object, environment, demographics Data access methods Open access: <i>Anyone on the Internet can download study data.</i> Authenticated access: <i>Anyone who registers their information such as name, email address, organization can download study data.</i> Consented access: <i>Anyone who registers their user profile and agrees to the terms of use can download study data.</i> Authorized access: <i>Anyone who registers their use profile, agrees to the terms of use, as well as submit the purpose of data use for approval can download study data.</i>	Stakeholders Data stewards: <i>Those putting together datasets</i> Policy-/law-makers: <i>Those making policies or regulations.</i> Data sharing entities: <i>Those operating methods of access and sharing</i> Data contributors: <i>Those contributing data to an AI dataset</i>

of medication?” We explored different types of objects accroding to object instances chosen by blind participants to personalize an object recognizer [72], and the types of demographic information were informed by the metadata of accessibility datasets [76]. For each question, we probed conditions on how others can access the data, ranging from being openly accessible to anyone to accessible only by those who are authorized (Table 2 Part 3), to gain a broader understanding of the factors that our participants would consider when reasoning about sharing their data.

Part 4 Mitigating Risks. Separately, we explored possible approaches to reduce their concerns surrounding the potential risks and challenges discussed in Part 2 and 3. We first asked a set of open-ended questions on actions that our participants want to see from different stakeholders (Table 2 Part 4)—e.g., “What actions would you like to see from *data stewards*, those collecting the data like for example our team, against such risk scenarios that might influence your decision about sharing your study data like the photos and labels of objects?” To concretize the discussion on risk mitigating strategies, we later followed up with existing data sharing purposes, methods, or regulations and asked the participants to rate these approaches by their level of comfort or acceptance with sharing their study data on a 7-point Likert scale.

3.4 Analysis

We transcribed the audio recordings of the interviews which included both open-ended questions (Part 1, 2, and 4) and Likert scale questions with shorter qualitative responses (Part 3 and 4). For qualitative data, we applied a reflexive thematic analysis [24, 25] to explore our interpretations on data. One member of the research team went through the process for data familiarization, inductive coding, and development of initial themes [26, 27]. The research

team then reviewed the themes, followed by discussions to conceptualize them as unifying concepts [25]. For example, we explored motivation and risk factors that were brought from other science fields to conceptualize the themes capturing patterns in how participants perceived data sharing. While the controversial discussions on “quantitizing” qualitative data exist [94, 115, 116], we report the number of participants whose responses included such themes. We adopt “quasi-statistics” [12] only to support statements such as “many”, “some”, or “a few” in the description of the qualitative data; percentages are not used given the small sample size which may lead to misinterpretation of the analysis.

For the quantitative responses, we used descriptive statistics to caption emerging patterns and tendencies. There are tensions in the field regarding how Likert scales should be analyzed [65]. As Likert scale ordinal data do not follow a normal distribution, using the mean can be of limited value as a measure of central tendency [10]. Instead, we adopt frequencies (count of responses for each point of the scale) and median as recommended by [131]. Anticipating a small sample size and having a large number of questions, we did not pre-register any hypothesis for inferential statistics.

3.5 Limitations

Our methods come with limitations. We highlight them here to help one better interpret the findings that follow.

Recruiting participants. Our study involves a small sample, though it is reflective of local standards at CHI [31]. We employ non-probability sampling, a combination of convenience, voluntary response, and snowball sampling. Some participants might have previously joined studies by our research institution. This could bias perspectives; they can be trusting of the team or institution

and more willing to contribute their data with fewer concerns. The degree of concerns can be also dependent on the awareness of existing social and political issues around data — e.g., a few participants who self-reported working in the IT or security field expressed more negative perspectives regarding data sharing.

Eliciting Responses. Prompting scenarios are considered effective in capturing participants' opinions and perspectives [64, 117], but they can also impact responses. Within a category, we typically ask questions before and after a scenario is given. But scenarios can have an effect on the next categories of questions. For example, more potential risk cases to reflect on (Part 2 Risks) might trigger more concerns, affecting participants' level of comfort with sharing (Part 3 Contexts of Data Sharing). There could be also an order effect as our interview questions proceeded from benefits to risks to underlying elements behind the decision to share; responses could differ were the conversations in a different order.

Generalizing findings. Our focus on a specific disability community, country, and context makes it challenging to obtain generalizable implications. For example, this study might not tell us much about the concerns and motivations for the Deaf community to contribute to sign language video data capturing their face, body, and background. Such limitations are not unique to our study. Adopting Nissenbaum's notion of *contextual integrity* [105], Barkhuus questions altogether “the viability of obtaining universal answers in terms of people's ‘general’ privacy practices” [9].

We also recognize an inevitable limitation of qualitative analysis. Our own positionality and reflexivity may bias the interpretations of the findings [14]. Therefore, the analysis is exploratory, and it would be meaningful to facilitate the transferability of findings to different communities and contexts in future research. Nonetheless, given that we are interested in contributing to conversations and initiatives on responsible practices for stewarding accessibility datasets, we make a conscious effort to connect our findings with the larger theory around privacy and data sharing as well as with prior work including other disability communities (e.g., [102, 107, 119]).

4 FINDINGS

We summarize the findings related to the perceptions of 13 blind participants towards benefits, risks, as well as contexts of sharing their data (i.e. photos and labels of objects) via an AI dataset.

4.1 Willingness to Share Given Benefits & Risks

We explored whether and how benefits and risks could contribute to participants' views on data sharing. Overall, we found that many focused on the greater benefits and perceived potential risks as minimal. However, their willingness to share is related to a number of elements, which we go deeper into below. How the participants assessed benefits and risks also reflected their attitudes towards data sharing. Some considered that the benefits would outweigh the risks, and those who foresaw the risks as severe remained hesitant to share their study data despite the considerations for benefits.

4.1.1 Benefits of Data Sharing. When initially asked about any benefits of data sharing **without prompting**, a majority (n=10) of the participants identified instances in which sharing their study data could lead to beneficial outcomes. Often the benefits were

related to the improvement and evaluation of object recognition technology (6 out of 10), to help “*developers figure out what worked out and what could be improved*” (P2) and “*people or companies who are into this work build from this [data] and advance it*” (P6). We saw a similar trend in [107] where disabled participants from different communities were more willing to contribute their data for future AI applications if their contribution would be a dependent factor for the success of the development. This may correspond with the motivation factors laid out by Batson *et al.* [11] suggesting that acting for the common good is not only driven by self-benefits (*egoism*). Acting for the community (*collectivism*) or for specific others (*altruism*)—e.g., to help scientists advance their research—are also noted as potential sources of motivations [87].

Looking at other **unprompted** benefits reported, we further observed the interplay of factors motivating willingness to share study data. Some (4 out of 10) perceived benefits that were directed towards the user community, as referred by P5: “*If you’ve got more people sharing the data, everybody doesn’t have to build their own independent library...Libraries can kind of benefit from each other’s because they have different data and that could help AI to learn something more.*” This can be seen as a *collectivist* motivation and support previous anecdotal evidence where social factors come into play behind e.g., community involvement [113] or information sharing [66]. We received a few (n=2) additional comments, which in some ways resembling *altruistic* motivation often driven by *empathy* to help others who are “perceived to be in need” [11]. The perceived benefits were geared to the interest of specific other blind users: “*people who are much younger than me and just starting out*” (P12) and “*people who were born blind*” (P13).

Participants saw even more benefits when **prompted** with a positive scenario for others (i.e., building a robot that can serve wider audiences), with a majority (n=11) of them expressing that they would be more open to sharing their study data. Some (n=4) were motivated by the next future technology, while others (n=4) considered the potential to help other disabled people i.e. those experiencing mobility challenges: “*They might need assistance, when you talked about fetching things, I mean, folks who are paralyzed or whatever. I could see that would be helpful to them, and I would be thrilled to be part of helping that*” (P8). Furthermore, seeking broader impacts also factored into our participants' willingness to share their data (n=4), articulating “*Best type of help would benefit everyone*” (P2) and “*I would be even more anxious to share. More people can benefit the better*” (P4). Following these motivations, some (n=4) expanded to other application domains where sharing their study data would benefit, for identifying objects in a shipping inventory (P3), describing photos on social media (P10), or helping with translation and second language learning (P5, P7).

4.1.2 Risks of Data Sharing. We subsequently explored our participants' awareness and reactions towards potential risks pertaining to data sharing. With **no probing** at the start and within the context of their study data, only a few (n=4) participants identified potential risks in sharing through a public AI dataset. The majority (3 out of 4) of their concerns revolved around the secondary use of data. Concerns were broad as “*How do I know that if you shared it with somebody else, they would ethically treat the data?*” (P9). Some

included a specific secondary use scenario such as “targeted advertising” that might lead to potentially negative consequences: “I don’t want companies to get my information, and then they just start targeting me with advertisement, like hey, we know that you are probably blind because you sent this data...Who knows how things can be used if it’s in the wrong hands” (P6). This can become a critical factor for contributing to an AI dataset, particularly when considerations for re-use cases are lacking for existing datasets sourced from blind people [54, 84, 93]—e.g., permitting commercial and private use of the data [93]. The concerns of the fourth participant (P7) focused on privacy and location identification. They considered where the photos were taken (*i.e.*, home) and whether the photos would be geo-tagged (though no geo-tags are embedded in their study data).

Further, we **probed** perception towards potential risks by prompting scenarios that our participants might foresee negative consequences: re-identification (with a follow up question related to non-consenting disclosure) and data abuse and misuse. We explored how these risk elements would impact willingness to share.

Re-Identification. We first provided a scenario that, even in anonymous or pseudonymous datasets, people could do some guesswork about an individual who contributed data based on released demographic information (*e.g.*, age and gender along disability status on publication). Most ($n=10$) participants expressed no or minimal concerns in terms of risk of re-identification. Some (4 out of 10) responded to the scenario that it is hard to imagine what could go wrong with the information disclosed, as they described the data captured as “*not really personal data*” (P1). Perhaps, they would still be willing to share photos of objects despite the risk factor prompted, as referred by P12: “*I think the benefits far outweigh any minor drawbacks that could occur...If someone says, oh, there was a 71 year old guy that took pictures of Lay’s potato chips, you know, if that’s all that ever happens to me, I’m okay, as long as they don’t take my credit cards or anything.*” This brings attention to the interaction between perceived benefits and risks. According to privacy calculus (risk–benefit analysis), privacy concerns are measured based on the perceived value of disclosing personal information relative to the perceived costs [149]. For example, Verheggen *et al.* [140] found that patients who agreed to participate in a clinical trial were likely to weigh the benefits more than the risks, whereas it was the reverse for those who declined to participate. Indeed, when looking at those who showed concerns when prompted ($n=3$), 2 of them did not report previously any unprompted benefits.

In response to the given scenarios, participants with concerns were strongly against sharing demographics along with their data. P9 mentioned, “*Why would you even be collecting that information? It’s for that reason that I don’t answer demographic questions generally...You don’t need to be collecting that demographic information in the first place. It’s not relevant to you to the task of the AI.*” Similarly, P2 responded, “*People being able to figure things out...that is actually part of my reserve [in sharing data].*” Thus, it is not a surprise to see these two participants opting not to include their age, gender, and education in Table 1. While P7 shared her demographics for the context of data for this study, she also raised the importance of privacy: “*Is it necessary to say that this person has a guide dog? Because you don’t see a whole bunch of people walking around with a*

guide dog so it’s easy to kind of pinpoint who that person is. So protect people’s privacy.”

Non-Consenting Disclosure. When asked whether they would be concerned about others (*e.g.*, family, friends, current or future employers) finding out that their data is included in AI datasets, all except one said no. The rationale for their response could be partially related to this specific study data. For example, P11 said, “*I’m not revealing any confidential information of my workplace. And I don’t care whether they see something related to me like a picture taken by me. Is that a problem? I don’t think so.*” A few (2 out of 12) justified this lack of concerns by contrasting it to everyday risks: “*We deal with a lot of online information that we share, or we interact with. We are not really dealing with any more risk than what we have already*” (P1). This high response agreement could also relate to the specific community. Kamikubo *et al.* [75] observed that among all accessibility datasets, those sourced from the blind community, such as our participants, tend to be shared publicly and typically include larger numbers of contributors. This is in contrast to datasets sourced from communities that encompass so called “invisible disabilities” which are less apparent to others and perhaps more sensitive to disclosure. When asking a similar question, Park *et al.* [107] observed heightened concerns from people experiencing “non-apparent forms of disabilities” (*e.g.*, ADHD).

The one participant who said yes to the question, added, “*especially, [I] wouldn’t want potential employers seeing that, because I just didn’t want them to infer anything or assume anything. So I wouldn’t care about my friends and family knowing but I wouldn’t want anyone else to know*” (P10). As they are a legal professional, perhaps this may be reflective of Judge Richard Posner’s view towards privacy as “*power to conceal information about themselves that others might use to [the individuals’] disadvantage*” [128].

Data Abuse and Misuse. When **prompted** with a negative scenario of repurposing data (*i.e.*, building an algorithm that can detect disability and be used against disabled people), many ($n=11$) participants were not overly concerned about the potential consequences. Perhaps they saw the potential risks as unimaginable and minimal. Some (4 out of 11) even responded that they would be open to sharing, as referred by P4: “*I think it’s very, very unlikely, though I do think it’s possible as you describe it. But, it would not change my decision [to share].*” Similar to the previous conversations about privacy, some of their rationale pertained to the specific study data and the benefit-risk tradeoff: “*I can’t imagine given what I took pictures of, it’s that critical. You know, it doesn’t bother me. I mean, it would be unfortunate if somebody kind of concluded something negative about people with disabilities or found it funny or like, laughable that they couldn’t take pictures or something, but it doesn’t bother me. You know, I think the good outweighs the bad*” (P8). Interestingly, the one participant who had been raising stronger concerns expressed a similar argument: “*I’m not as worried about that. I do understand that you can’t control what other people might use the dataset for. However, that’s why it’s important to be careful [about] what you collect. At that point, I’m still more concerned about the collection than what other people might use it for*” (P9).

While raising minimal concerns for the given scenario, our participants followed up with some degree of concerns around the

ethics of data use. A few ($n=3$) of them posed mixed feelings, articulating that such data abuse and misuse risks are unavoidable in the digital world: *“I can’t control, you know, and nobody else can really control. If people use this information for something else, honestly, that happens everywhere. That’s always a concern”* (P6). P1 reflected on the lack of concerns as *“we either got numb or gave up.”* Such attitudes could be explained by their nuanced understanding of data policies. For example, camera-based assistive technologies like Aira or SeeingAI which our participants reported to use, provide no clear indication of whether and for what purpose personal visual data are shared with third parties [130]. This perhaps could lead to something of a paradox widely discussed in privacy literature [78]. Brown [29] found that, despite the general worries people seemed to have about privacy, they would still give out information for perceived benefits. In fact, the two participants who reacted to both scenarios (*i.e.*, potential risk cases of re-identification, data abuse and misuse) with stronger concerns and hesitancy to data sharing, reported to use Be My Eyes—one of the few services that explicitly indicates dissemination of video streams to third parties [130].

4.2 Data Sharing Given Data Access Methods

To discuss beyond specific data and environment (*i.e.*, photos and labels of objects generated in the home environment), we investigated various factors that could influence participants’ decision to share their data. Inspired by prior work measuring people’s comfort and acceptance of their data being collected and used [50, 107], we questioned how different data modalities, objects, environments, and demographic metadata would affect such measurements. As shown in Figure 6, we explored their responses to these data sharing contexts conditioned by data access methods. Though we observed the unsurprising trends for increased comfort as more restrictions are applied on the level of access (Figure 6a), there were specific contexts that the participants raised concerns consistently across different data access methods. For certain information or settings (*e.g.*, audio description, photos of medication, photos taken around bystanders), some participants rated comfort on the negative end of the scale even under authorized access (Figure 6b). In the following, we go deeper into these topics of concern.

4.2.1 Type of Modality. We found that some data modalities relate to the level of comfort for sharing. In particular, we gathered concerns for videos of objects and audio description of images. Under open access, five participants rated lower on the comfort scale for these two modalities (Video: median=4, Audio Description: median=5) compared to other types of modality like photos and names of objects (median=6). When comparing videos with photos, audio partially factored into their concern: *“Videos got a little more concern because of whatever might be heard in the background”* (P5). In fact, videos of objects from blind people in the ORBIT dataset respected this aspect; audio was never collected [135]. Even so, the degree of concerns for videos degraded as we moved to the authorized access condition (median=7). Whether photos or videos, what information these data captured was at the root of their concerns, as P13 briefly mentioned: *“It’s totally dependent on what you’re sharing.”*

In terms of sharing audio description, four participants were especially worried about being identified from their voice. P9 elaborated the rationale: *“That (voice) leads to potentially identifiable information. Especially with the current type of voice fingerprinting software that is starting to be developed now. You actually can, with a fairly consistent degree of accuracy, match someone’s voice print.”* Following these worries, three participants kept their comfort level at the lower end across all access methods for audio.

4.2.2 Type of Object. Our participants identified certain objects as ‘private.’ One noticeable trend was that they did not feel comfortable sharing photos of medication, which was rated low on the comfort scale from open (median=3) to authorized (median=4) access conditions. Five participants expressed stronger concerns, as briefly explained by P1: *“Medication goes to private status or private characteristics of a person”* (P1). In comparison, the ratings were relatively higher for other types of objects even under open access, including hygiene and cosmetic products, food/drinks, and appliances (median=6). These objects could be perceived as more general items. Thus, it was not a surprise to see differences in their perception towards prescription and general medications: *“If it’s like a general like Tylenol medicine or something over the counter stuff, then yes, I would give it a seven. If it’s like prescription then no, because my information is on there”* (P6). Perhaps this can be a double edge sword given the potential of object recognition technology supporting identification of medication that is often a challenge for blind people [20]. These concerns can further limit the availability of images for such ‘private’ objects. This indicates the need for privacy-preserving discussions in data collection, as explored by Gurari *et al.* [53] to enable building image recognition algorithms while safeguarding private information.

Participants also expressed hesitancy in sharing objects of ‘personal’ nature. Though the comfort scores for clothing were relatively at the higher end from open (median=5.5) to authorized (median=6) access conditions, concerns still remained regarding privacy: *“It’s just like, you wouldn’t invite your friends to see your closet right? They come to your house. You invite them to a party but they will never go into your bedroom or your walk-in closet, that kind of thing...I still want the privacy no matter who they are”* (P11). We can expect similar attitudes depending on the types of hygiene or cosmetic products, that may be more or less personal. Additionally, our participants listed other items that they do not want to capture in the photos. These were mostly objects of ‘sensitive’ nature including personal documents (*e.g.*, driver’s license, passport, insurance cards, credit cards, bills). They further indicated safety concerns for objects that can reveal or trace their identity (*e.g.*, vehicles, diploma, friends/family photos). The characteristics of these objects with ‘personal’ and ‘sensitive’ nature could shed light on the taxonomy of *what is private in images* [53].

4.2.3 Type of Environment. When asked to rate their level of comfort by the environment, the presence of bystanders factored into their perception. In general, our participants were comfortable with sharing photos captured in the home space especially where bystanders are not present; 10 participants rated their comfort at the higher end even under the open access condition (median=6). In the same access condition but in the presence of bystanders such as family members, their perceived comfort was lower (median=4).

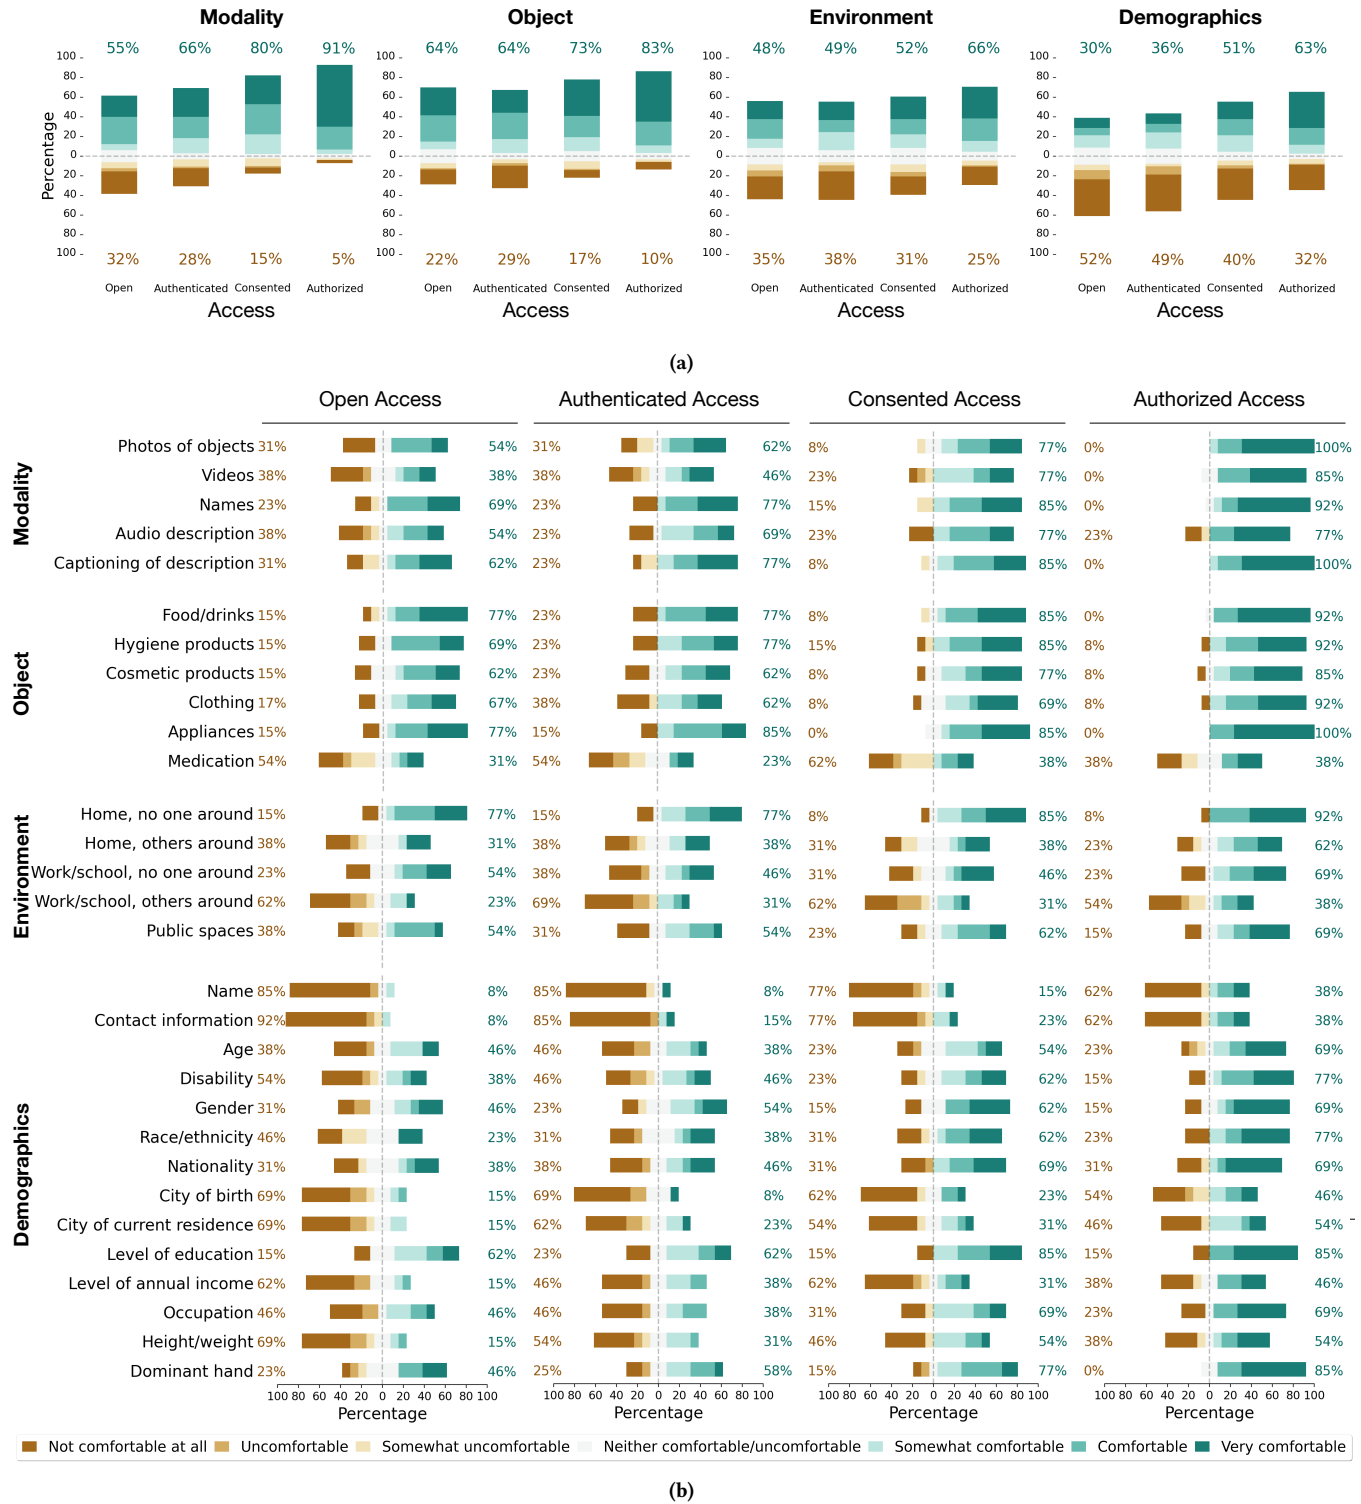


Figure 6: Participants' level of comfort with data sharing using a Likert scale (1: not comfortable at all, 7: very comfortable), with percentages of responses for lower comfort and higher comfort levels varies widely across contexts. At the top in (a), we show overall trends of comfort by four data access conditions: Open, Authenticated, Consented, and Authorized aggregated across different modalities, objects, environments, and demographics. Below that in (b), we provide a breakdown of responses for each type of modality, object, environment, and demographic information, also stratified by the data access conditions.

Concerns often revolved around the background of photos, as P5 explained: *“Something I could do very diligent is making sure there’s not identifying things in the photos. That would be my only concern.”* We saw that these concerns were more associated with work or school spaces; in the presence of bystanders, participants rated low on the comfort scale from open (median=2) to authorized (median=3) access conditions. This might be due to more information being available to identify them (e.g., company logos) or others around (e.g., co-workers) who didn’t give consent to share. P6 remarked, *“I took a photo of an object and somebody was in the background. I guess I wouldn’t want that to be displayed. Again, just because that person may not feel comfortable with.”*

While these concerns for bystanders resonate the privacy perspectives that are being reported with technology use in public (e.g., for blind people to detect pedestrians through wearable glasses [85]), our participants showed a bimodal reaction for sharing photos generated in public spaces such as streets or plazas. Interestingly, these places were characterized differently; even under open access, participants indicated higher levels of comfort (median=5). Though some expressed much stronger concerns than others, such as *“I am more concerned if people can identify the space that is attached to you”* (P8), other participants noted less privacy comparing to home or work environments: *“I think when people are in public, they don’t have expectation of privacy that they might have in my house or even at work”* (P4). P11 also justified the lack of concerns as *“everyone is taking photos, who cares what these photos are?”*, yet expressed concerns if they were neighborhoods where family, friends, or neighbors could be identified.

4.2.4 Type of Demographic Information. Unsurprisingly, participants were hesitant to sharing identifiable information, such as their name or contact information receiving low comfort ratings across all data access conditions (median=1). Even so, 5 participants rated higher on their comfort level for sharing such information under authorized access, as referred by P13: *“It depends on what I’m sharing with...Sometimes you get on Facebook and you think you’re gonna find some person that has a name and you see this 25 people with the same name. So I don’t have a problem with as long as it’s not attached to anything that could be damaging.”* They did not see anything harmful from *“a bag of potato chips that we took photo of and we labeled”* (P13). While the ratings for comfort started low for sharing city of birth or current residence under open access (median=2), more participants became comfortable with sharing city of current residence under authorized access (median=5). It was not the case for sharing city of birth (median=3), which P1 commented *“That goes back probably to my nationality”* considering the potential linkage. P4 expressed hesitancy in sharing such information as *“I don’t see the value of that except for negative reasons. Someone trying to build a profile of me.”*

While privacy or security concerns often reflected the lower comfort ratings for sharing certain demographic information (e.g., annual income), participants raised ethical concerns as well. They found sensitivity of disclosing race/ethnicity or nationality; P12 (whose comfort level was 1 across data access methods) described it as *“invasive because there’s still a lot of discrimination with various nationalities”* and suggested that participants should have the option to not answer such demographic information. P2 specially

warned the risk of sharing demographic information including disability in general: *“I think people could make false assumptions or incorrect judgments about you.”* Our participants also raised other demographic information that they would not want to share, including marital status, number of children or siblings, employment history, and religious affiliation.

We observed other concerns that could factor into consideration, including the ambiguity of the data collection purpose. For example, our participants were unsure of the usefulness of height/weight or dominant hand information as part of the AI dataset. P7 noted, *“What’s the purpose? Why is it important? Why do you want to know that?”*, rating the comfort level as 1 across different data access methods except for authorized access users whom the participant could expect a clear purpose for its use. We found a similar trend with other demographic information including gender and race/ethnicity, as articulated by P5 *“I don’t really know how important the gender pieces to this as far as AI developers doing what they need to do”* and P3 *“I don’t care about whether they know I’m black or white or whatever, but I really don’t think it’s that important.”* This indicates the importance of increasing the understanding of how such information could contribute to AI development, especially raising awareness around issues of fairness for underrepresented groups [18, 30, 134].

4.3 Mitigating Risks Given Stakeholders & Regulations

We elicited responses from our participants about risk mitigating strategies for data sharing. We first asked what actions by different stakeholders or regulations could possibly address their concerns. We then expanded the conversations by listing existing practices and explored their reactions via comfort of sharing or acceptance ratings (see Figure 7). Their responses revealed different perspectives on research or legal practices as well potential strategies that can help minimize the harms and impose safety measures.

4.3.1 Data Stewards. When asked how data stewards, those putting together a dataset (i.e., researchers of this study), could mitigate the risks, all our participants expected some forms of actions — e.g., screening the photos and removing any personal or sensitive information if caught by mistake (n=6), collecting data in a privacy-respecting way such as codifying participants’ names and cropping out the background from photos (n=2), ensuring that the data are protected from security breaches when storing and sharing (n=3), or fully informing the use of their data to the contributors including secondary data use cases (n=2).

Our participants expanded on how they would like data stewards to restrict the use cases of their data being shared. Many (n=9) favored to restrict them by specific domains or usage types — e.g., for research purposes only, for object recognition technology development, for purposes defined originally, informed, and permitted by data contributors. A participant, however, noted the importance of supporting broader purposes: *“It’s easy to say, it’s only should be for research purposes. And that’s fine. But at some point, let’s say you guys have a final ready to go, market ready [object recognition technology], then at that point, you guys are starting your library all over, which has its pros and cons”* (P5). Broadening the use of data (e.g., commercialization) yet raises more questions and challenges

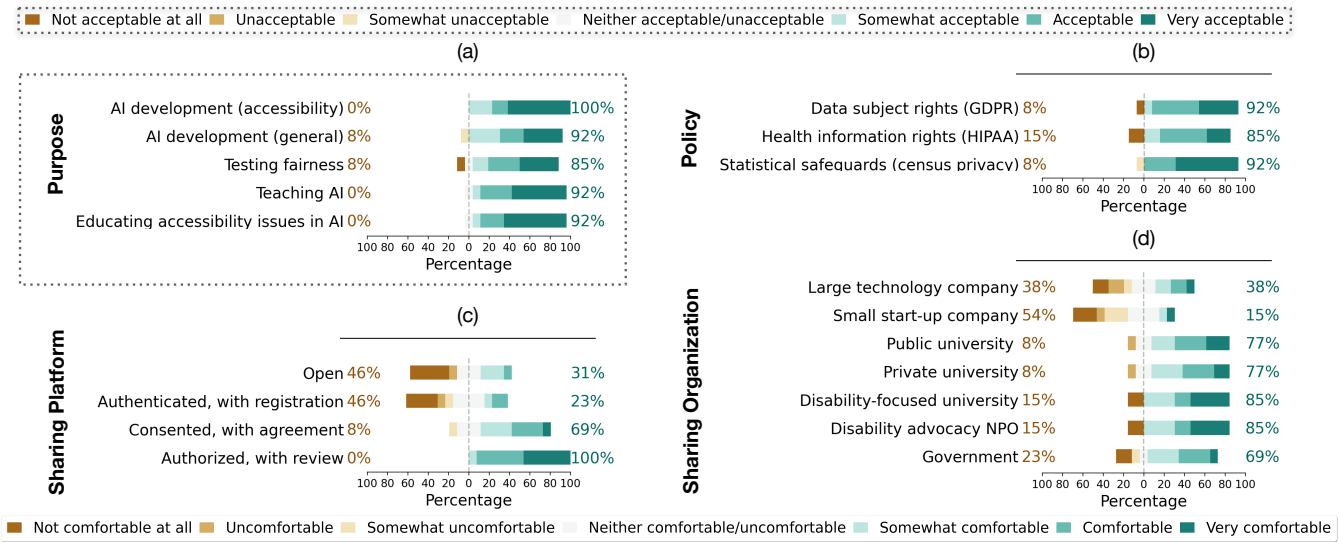


Figure 7: Participants' level of acceptance (1: not acceptable at all, 7: very acceptable) or comfort (1: not comfortable at all, 7: very comfortable) with data sharing varies across data use purposes (a), policies (b), sharing platforms (c), and organizations (d).

such as ownership, intellectual property, or data agreements [136]. In fact, this has been a 'wicked problem' across fields [80, 148, 151], especially when no criteria exist for determining the correctness or value for the use of open data [151].

To expand on the use of their study data, we prompted a list of data purposes (see Figure 7a). Participants reacted with somewhat high degree of acceptance towards data use for accessibility (median=7). While ratings are also on the higher end for teaching AI and educating accessibility issues in AI (median=7), two participants left some remarks, as referred by P9: "The goal is good but it all depends on how it's implemented." P2 asked "What's the purpose? What application?" factoring into their neutral rating. Lacking specifics of the purpose could explain why the acceptance was slightly lower for other prompted purposes, developing AI for general purposes and testing fairness (median=6). Echoing P9's comment on fairness as "a word that is so vague these days," efforts to increase the knowledge about contextual factors influencing fairness (e.g., issues of bias in AI) might be necessary to support their conceptualizations [17]. P10 gave a neutral rating to the purpose of testing fairness although they seemed to value it: "It (data) is not being used for what I thought it was used for, which is helping with the object identification and recognition."

4.3.2 Policy- and Law-Makers. When asked what actions policy- and law-makers could take to prevent and minimize the risks, some (n=5) stressed legal actions to penalize individuals who misbehaved, such as those sharing data without consent (P2), using data for a wrong purpose (P1, P8), or not keeping the promise to handle data safely and responsibility (P7). A participant (P4) articulated to "establish accountability or illegality" of misbehaved individuals. A few others (n=3) suggested ways to reduce the chances of misbehaviors, such as keeping records of who accessed the data "so that if that information does get used inappropriately, for whatever reason, you can at least have a narrower field of suspects" (P5), or what people

claimed to use the data "to get that assurance so you have some protection, in case you have to bring up a case, you know, you misused my data, but you said you weren't going to misuse it" (P13). Even so, P13 raised a limitation point: "You really don't know if they're misusing your data or not until something happens." Similarly, one participant expressed further limitations given that policies are often not sufficient to protect data contributors: "I'm less optimistic about lawmakers...They've instituted the GDPR, the general data protection regulation, but even that, it outlines a whole bunch of scenarios, and legal requirements, a lot of which I'm happy with. But all you have to do to get out of the scope of the GDPR is simply moving data outside of Europe. That's not difficult. Lawmakers, I don't really think have a lot of power in this scenario."

We prompted existing privacy policies (see Figure 7b) to further probe participant perception. Overall, participants reacted positively to these policies for protecting their personal or identifiable information. For example, all except two participants rated data subject rights (e.g., GDPR [86] giving control over data) high on the scale as they would feel comfortable sharing their data (median=6). However, they added some remarks regarding the lack of guarantee, as referred by P5: "Policies are not always followed at different places." They reacted similarly to privacy protection for health information such as HIPAA [83]; though higher in ratings (median=6), some concerns also remained in the lack of guarantee as "data can always be breached" (P6). Additionally, P7 wondered whether such data policies like GDPR would actually apply to study data and context: "Not fully trusting the existing policy, so trying to see if the existing policy concept can be transferable to like study data." Indeed, there are still open challenges for the uncertainty about how such existing policies would apply to scientific research practices, such as in informed consent and anonymization [63].

4.3.3 Data Sharing Entities. We explored data sharing practices by prompting different conditions for platforms and organizations. Our

participants rated these by their level of comfort for sharing their data. As shown in Figure 7c, 9 participants reacted positively to conditions where a Terms of Use agreement is applied (median=5). With further authorization to screen people's access by the purpose of data use, all participants indicated higher ratings of comfort (median=6) due to more *accountability* (P2, P9), *control* (P5), and *security* (P6, P12). However, the lack of trust remained depending on how platforms operate: *"There's still a slight chance that [data] could get in the wrong hands. You don't know who is operating"* (P7).

Concerns for operation were mainly reflected within the context of technology companies (see Figure 7d), especially with smaller start-ups which only 2 participants rated high on the comfort scale for sharing their data (median=3). Often the reason was behind security measures, with a participant articulating *"doubt that they would have controls in place"* (P9). Their concerns resonate with public perceptions reported in anecdotal evidence — e.g., small companies were seen to be less stable [96]. Large technology companies were also perceived as profit driven [58], which could explain their lower comfort ratings (median=4) towards them compared to organizations supporting education/research or disability communities. Regarding their comfort with sharing, 10 participants reacted positively towards public universities (median=6) and private universities (median=5). The participants saw that these organizations have *"a fewer reasons to misuse the data"* (P4) and *"to be a little bit more discreet"* (P7). Similar responses were found for disability-focused universities and non-profit organizations, where 11 participants felt comfortable sharing through them (median=6). When asking similar questions, Park *et al.* [107] also observed that organizations oriented towards disability communities are seen as reliable, yet comfort ratings for large technology companies closely aligned with ratings for universities. As the work came from company-based research [107], we might be seeing the effects of sampling as P2 noted: *"I wouldn't be concerned if I do a study with someone. I trust them. It's just that anyone else that might get a hold of stuff that might have been shared or supposedly shared by another."*

Following the lack of trust, 3 participants raised the importance of transparency, ensuring that they are informed about purposes and restrictions and kept in the loop for any actions made regarding their data. A participant strongly called for regulations: *"If researchers are going to share data with any third parties, that third party also needs to have a person's consent and offer full and complete disclosure, and honor their promise. The third parties need to honor it to the researcher as well as to the individual [contributing data]"* (P2). *Datasheets for Datasets* [49], while being targeted for dataset creators and dataset consumers to promote transparency and accountability, includes questions regarding data collection and distribution—e.g., whether individuals contributing their data consented to the use of their data, or whether it will be distributed to third parties. We see this as a meaningful resource to be reviewed with potential data contributors and also as a way to discuss preferences to keep them in the loop in the process. Furthermore, given the individual differences that we found in terms of sharing demographic information, especially in the context of sharing entities (e.g., open access vs. authorized access; data-sharing organizations), these data sharing factors need to be considered together when adopting inclusive practices for datasheets.

4.3.4 Data Contributors. Finally, we asked what actions our participants as potential data contributors could take to protect themselves from the risks. One recurring theme among the participants was to carefully consider what data they would be sharing (n=5), ensuring that objects and information contained in the photos are not sensitive and they are disclosing as little demographic information as possible. We also observed two participants selecting the lobby of the apartment and one participant selecting the patio area to generate the photos. This might be reflective of individual concerns that were induced by what could be captured in their data inadvertently, in addition to other potential concerns (e.g., safety related to COVID-19 or letting an experimenter in their home). P6 reflected on the data capturing activities for ways to minimize capturing unintended information in the future: *"I think I would definitely be more cautious about making sure that there's no identifiable photos. So probably, maybe I would just take all my photos in front of a white wall or something, you know. So that's really up to me, to make sure that I don't take photos that I'm uncomfortable with."* Such workarounds, while efficient, might not capture the real-world contexts for image recognition tasks. To facilitate the collection of high-variation conditions (e.g., in a wide variation of backgrounds [93]), Theodorou *et al.* [135] incorporated a manual validation process to check and remove data containing personally identifying information (PII) in their dataset creation. This brings more attention to better approaches that allow automatic detection of PII to blur or erase from images [53].

Our results also indicate the importance of supportive materials and communication to help potential data contributors assess the benefit-risk tradeoffs. A few (n=3) participants showed further interest in learning more about how their data would be used, as articulated by P13: *"I would love to understand better how AI use the data. But again, it's just my curiosity. Yeah, just wanted to understand a little bit better"*. This could help bridge the knowledge gap and establish trust; for example, we saw people's reluctance to share their demographic information as it was deemed 'not relevant' for AI development. Participants also asked for a better sense of the impacts of the data they share: *"I personally just want to see what other people are using it for. Because, you know, I don't have all of the answers. I don't know how people could be using this technology to improve their quality of life. I mean, we're all here to learn from one another"* (P6). This is an option that is not supported in current informed consent processes and needs to be reflected in future research practice discussions [3, 106].

5 DISCUSSION

When data are at the core of innovation, creating and releasing annotated AI datasets becomes critical. This has been illustrated in the progress of shared data resources in broader computing fields [7, 28, 32, 42, 74, 112]. However, many application domains still lack sufficient data for accessibility [22, 69, 107], due to sensitivity surrounding smaller population groups and potential harms that may arise along privacy and ethics. In the face of these challenges, we reflect on the empirical insights gained from the blind community to consider how we can better configure data practices and study methods to guide future directions for promoting more

transparency, trust, and engagement in these practices. These insights come with several limitations that threaten their validity and generalizability, detailed in our Methods (under Section 3.4).

5.1 Assessing Benefit-Risk Tradeoffs in Data Sharing

Participants were aware of many benefits of contributing their study data via an AI dataset. Even without probing, they described how their study data would be important for technological innovations for accessibility and beyond. Additionally, with probing, they further expressed motivations behind data sharing, which were often associated with broader impacts of their data contribution. Some wished to learn more about the role of data in AI to assess the positive impacts better. This strengthens prior anecdotal evidence about data sharing “for the greater good” [96] indicating the willingness of disabled people to share their data for the benefits of the broader communities [102, 107].

While recognizing the benefits, many participants were not overly concerned about risk factors behind data sharing as long as these data are anonymized and not disclosing sensitive information. Such assessment is not unique to the blind community; data harms are often overlooked and underestimated, surfacing limitations for how well people can anticipate negative impact [81, 127]. Indeed, Hamidi *et al.*, building on the privacy threat framework [36], referred to *unawareness* of end users not realizing the consequences of sharing their data as a serious concern [55]. In return, we have seen attempts to raise awareness about potential harms, such as releasing real-life data misuse examples (e.g., Data Harm Record [110], Inventory of Risks and Harms [46]). However, an important caveat revealed from our study is that assessment of risks can be too nuanced to be directly addressed, as the participants reacted to our risk scenarios as “*hard to imagine if they are critical*” (P8). This extends to similar conversations on how documents aimed to increase awareness by informing data sharing practices (e.g., privacy policy) are not always effective; they are often difficult to read and understand the implications [40, 95].

One way we could better assess benefits and risks is to incorporate the participatory nature to analyze the impacts and implications for data sharing. While intended for dataset creators and users, *Datasheets for Datasets* elicited a similar point: “Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted?” [49] This needs to be addressed and communicated with individuals or communities who could be most affected, adopting many other guided questions related to the collection and distribution of datasets. Our study methods seem to reveal partial effects by providing an opportunity for participants to engage more deeply in these topics. As we probed more contexts to our participants to reflect on how their data should be shared, they began to analyze a set of rules that need to be considered in implementing data sharing — e.g., restriction of data use for research purposes only (especially for improving object recognition technology) or control of information they are disclosing with their study data such as what’s being captured in the background of photos. Also, with probing of benefit and risk scenarios, we started to see participants making a tradeoff

assessment as “I think the benefits far outweigh any minor drawbacks that could occur...If someone says, oh, there was a 71 year old guy that took pictures of Lay’s potato chips, you know, if that’s all that ever happens to me, I’m okay” (P12). In this vein, to support critical reflections among stakeholders on navigating “wicked problems” [13], we encourage practices such as Value Sensitive Design (VSD) [47] or Speculative Design [38], along with design fictions and participatory workshops [8, 126] as important directions. We further suggest extending the line of work around co-designing accessible impact assessment tools or checklists, highlighting Madaio *et al.* [92] as a methodological inspiration.

5.2 Bringing Trustworthiness in Data Ecosystems: Transparency and Engagement

Our findings indicate the importance of complementing trust across stakeholders in data practices. As highlighted throughout this work, many were hesitant to contribute their study data for ambiguous purposes or purposes not aligning with their perceived benefits, revealing notions of “distrust.” For example, we similarly observed in our study public perceptions about small companies which they were seen to lack regulations [96] and larger tech companies seen as profit driven [58]. They also expressed concerns for sharing demographic metadata (e.g., gender, race) along their data, not only due to privacy but also uncertainty of its importance and preconceived notions about what kind of data should be used to train AI models. To fill the disconnect between the disability and AI community, the process of stewarding accessibility datasets requires greater transparency of data use as well as awareness [17], especially to challenge inclusivity issues that are pressing for marginalized communities [30, 76].

To further support transparency, our participants wanted to be kept in the loop for any actions made regarding their data, always obtaining a person’s consent and providing complete disclosure of the way their data are being used. While this is important over the data lifecycle, we see practical challenges in sustaining such long term relationships between data contributors and data stewards. Though GDPR is designed to support such connections, there are still open challenges in applying to scientific research practices [63], further inhibited by lack of resources and expertise [125]. Perhaps proper implementation and maintenance at the institutional level may be necessary. Reflecting on trustworthy certification in human-centered AI [35, 48, 123, 124], we echo oversight structures, with institutional interventions and reviews, to support an ongoing communication to align data practices with individual concerns and values. More so, we see benefits in implementing systems that enable community efforts and coordination, such as “data consortia” which have been developed as institutional frameworks in archives to navigate the challenges behind ethical data collection and sharing [67].

We also recognize oversight at an individual level. Incorporating a “framework for participatory data stewardship” [4], we suggest researchers to consider data sharing from the beginning of their projects and further consider mechanisms beyond transparency to engage participants in decision making about the data they consider sharing for AI development. As Rake *et al.* [109] explored personalized consent flows for data sharing with stakeholders in

medical research, we see future directions in favor of practices that empower potential data contributors to help shape and govern their own data. We perhaps see benefits in data cooperativeness to turn “distrust” into shared understanding of how data sharing should be carried out [91].

6 CONCLUSION

There are important normative questions around the use of accessibility datasets sourced from disabled people — they can be used against them by uncovering their identity and (mis)detecting disability status without consent. However, AI-infused systems trained on data lacking in terms of inclusion and diversity further increase the risks of unfair or discriminatory outcomes for underrepresented groups. To drive AI efforts that are inclusive of disability, this research aims to shape data practices that align with the concerns and values of disabled data contributors. We conducted a case study engaging blind participants in ‘in-situ’ data capture activities to inquire about how their data should be used and shared via an AI dataset. Our findings have highlighted the opportunities for making their data accessible for the research community in AI development, through proper actions and restrictions around data sharing and re-use. We hope this research helps discussions that aim to improve the norms for collecting and sharing data by understanding the facilitators and barriers that are more attuned to the communities of focus in accessibility.

ACKNOWLEDGMENTS

We thank Hal Daumé III for providing valuable feedback on our preliminary work. We also thank Kaitie Shilton, Beth St. Jean, Jessica Vitak, and our anonymous reviewers for their expertise and comments to further strengthen this paper. Special thanks go to Farnaz Zamiri Zeraati for her help in running the study. Finally, we want to acknowledge our participants for their time and the research advisory council from The National Federation of the Blind for assisting us with participant recruitment. This work is supported by National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR), ACL, HHS (#90REGE0008).

REFERENCES

- Jacob Abbott, Haley MacLeod, Novia Nurain, Gustave Ekobe, and Sameer Patil. 2019. Local Standards for Anonymization Practices in Health, Wellness, Accessibility, and Aging Research at CHI. 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, 1–14. <https://doi.org/10.1145/3290605.3300692>
- Dustin Adams, Lourdes Morales, and Sri Kurniawan. 2013. A qualitative study to support a blind photography mobile application. In *Proceedings of the 6th International Conference on Pervasive Technologies Related to Assistive Environments*. 1–8.
- Kathy Ahern. 2012. Informed consent: are researchers accurately representing risks and benefits? *Scandinavian Journal of Caring Sciences* 26, 4 (2012), 671–678.
- JUST AI. 2021. Participatory data stewardship. *Reading time* (2021).
- Amazon. 2021. Registry of Open Data on AWS. <https://registry.opendata.aws/>.
- Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fournay, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. 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, 1–13. <https://doi.org/10.1145/3290605.3300233>
- Arthur Asuncion and David Newman. 2007. UCI machine learning repository.
- Stephanie Ballard, Karen M. Chappell, and Kristen Kennedy. 2019. Judgment Call the Game: Using Value Sensitive Design and Design Fiction to Surface Ethical Concerns Related to Technology. In *Proceedings of the 2019 on Designing Interactive Systems Conference* (San Diego, CA, USA) (DIS '19). Association for Computing Machinery, New York, NY, USA, 421–433. <https://doi.org/10.1145/332276.3323697>
- Louise Barkhuus. 2012. The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 367–376. <https://doi.org/10.1145/2207676.2207727>
- Dwight Barry. 2017. Do not use averages with Likert scale data. *Enterp. Anal* 24 (2017).
- C Daniel Batson, Nadia Ahmad, and Jo-Ann Tsang. 2002. Four motives for community involvement. *Journal of social issues* 58, 3 (2002), 429–445.
- Howard Saul Becker. 1970. *Field work evidence*. Transaction publishers. 39–62 pages.
- Arne Berger, Sören Totzauer, Kevin Lefevre, Michael Storz, Albrecht Kurze, and Andreas Bischof. 2017. Wicked, Open, Collaborative: Why Research through Design Matters for HCI Research. *i-com* 16, 2 (2017), 131–142.
- Roni Berger. 2015. Now I see it, now I don't: Researcher's position and reflexivity in qualitative research. *Qualitative research* 15, 2 (2015), 219–234.
- Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. 2010. VizWiz: Nearly Real-Time Answers to Visual Questions. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) (UIST '10). Association for Computing Machinery, New York, NY, USA, 333–342. <https://doi.org/10.1145/1866029.1866080>
- Brianna Blaser and Richard E. Ladner. 2020. Why is Data on Disability so Hard to Collect and Understand?. In *Proceedings of the 5th International Conference on Research in Equity and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*.
- Veronika Bogina, Alan Hartman, Tsvi Kuflik, and Avital Shulner-Tal. 2022. Educating software and AI stakeholders about algorithmic fairness, accountability, transparency and ethics. *International Journal of Artificial Intelligence in Education* 32, 3 (2022), 808–833.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29 (2016), 4349–4357.
- Brian M. Bot, Christine Suver, Elias Chaibub Neto, Michael Kellen, Arno Klein, Christopher Bare, Megan Doerr, Abhishek Pratap, John Wilbanks, E. Ray Dorsey, Stephen H. Friend, and Andrew D Trister. 2016. The mPower study, Parkinson disease mobile data collected using ResearchKit. *Scientific Data* 3 (March 2016), 160011. <https://doi.org/10.1038/sdata.2016.11>
- Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P Bigham. 2013. Visual challenges in the everyday lives of blind people. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2117–2126.
- Danielle Bragg, Abraham Glasser, Fyodor Minakov, Naomi Caselli, and William Thies. 2022. Exploring Collection of Sign Language Videos through Crowdsourcing. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–24.
- Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoeve, Christian Vogler, and Meredith Ringel Morris. 2019. Sign Language Recognition, Generation, and Translation: An Interdisciplinary Perspective. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 16–31. <https://doi.org/10.1145/3308561.3353774>
- Danielle Bragg, Oscar Koller, Naomi Caselli, and William Thies. 2020. Exploring Collection of Sign Language Datasets: Privacy, Participation, and Model Performance. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, Greece) (ASSETS '20). Association for Computing Machinery, New York, NY, USA, Article 33, 14 pages. <https://doi.org/10.1145/3373625.3417024>
- Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.
- Virginia Braun and Victoria Clarke. 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and Psychotherapy Research* 21, 1 (2021), 37–47.
- Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative research in psychology* 18, 3 (2021), 328–352.
- Dan Brickley, Matthew Burgess, and Natasha Noy. 2019. Google Dataset Search: Building a search engine for datasets in an open Web ecosystem. In *The World Wide Web Conference*. 1365–1375.

- [29] Barry Brown. 2001. Studying the internet experience. *HP laboratories technical report HPL 49* (2001).
- [30] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [31] Kelly Caine. 2016. Local Standards for Sample Size at CHI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 981–992. <https://doi.org/10.1145/2858036.2858498>
- [32] Nicoletta Calzolari, Claudia Soria, Riccardo Del Gratta, Sara Goggi, Valeria Quochi, Irene Russo, Khalid Choukri, Joseph Mariani, and Stelios Piperidis. 2010. The lrec map of language resources and technologies. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*.
- [33] Yang Trista Cao, Kyle Seelman, Kyungjun Lee, and Hal Daumé III. 2022. What's Different between Visual Question Answering for Machine "Understanding" Versus for Accessibility?. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online only, 1025–1034. <https://aclanthology.org/2022.acl-main.75>
- [34] Michelle Carney, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382839>
- [35] Joel Chan, Hal Daumé III, John P Dickerson, Hernisa Kacorri, and Ben Shneiderman. 2021. Supporting human flourishing by ensuring human involvement in AI-infused systems. *NeurIPS 2021 Workshop on Human Centered AI* (2021).
- [36] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. 2011. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering* 16, 1 (2011), 3–32.
- [37] Mat D Duerden and Peter A Witt. 2010. The impact of direct and indirect experiences on the development of environmental knowledge, attitudes, and behavior. *Journal of environmental psychology* 30, 4 (2010), 379–392.
- [38] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- [39] Utkarsh Dwivedi, Jaina Gandhi, Raj Parikh, Merijke Coenraad, Elizabeth Bonignore, and Hernisa Kacorri. 2021. Exploring Machine Teaching with Children. In *2021 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 1–11. <https://doi.org/10.1109/VL/HCC51201.2021.9576171>
- [40] Julia Brande Earp, Annie I Antón, Lynda Aiman-Smith, and William H Stufflebeam. 2005. Examining Internet privacy policies within the context of user privacy values. *IEEE Transactions on Engineering Management* 52, 2 (2005), 227–237.
- [41] A Engler. 2019. For some employment algorithms, disability discrimination by default. *Brookings*. Available at: <https://www.brookings.edu/blog/techtank/2019/10/31/for-some-employment-algorithms-disability-discrimination-by-default/> (accessed 1 June 2020) (2019).
- [42] European Organization For Nuclear Research and OpenAIRE. 2013. Zenodo. <https://doi.org/10.25495/7GXK-RD71>
- [43] Alireza Fathi, Xiaofeng Ren, and James M Rehg. 2011. Learning to recognize objects in egocentric activities. In *CVPR 2011*. IEEE, 3281–3288.
- [44] Casey Fiesler. 2019. Ethical considerations for research involving (speculative) public data. *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (2019), 1–13.
- [45] Liora Findler, Noa Vilchinsky, and Shirli Werner. 2007. The multidimensional attitudes scale toward persons with disabilities (MAS) construction and validation. *Rehabilitation Counseling Bulletin* 50, 3 (2007), 166–176.
- [46] Centre for Information Policy Leadership. 2016. Risk, High Risk, Risk Assessments and Data Protection Impact Assessments under the GDPR. https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_gdpr_project_risk_white_paper_21_december_2016.pdf. (2016).
- [47] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington technical report* 2 (2002), 12.
- [48] Ozlem Ozmen Garibay, Brent Winslow, Salvatore Andolina, Margherita Antona, Anja Bodenschatz, Constantinos Coursaris, Gregory Falco, Stephen M. Fiore, Ivan Garibay, Keri Grieman, John C. Havens, Marina Jirotko, Hernisa Kacorri, Waldemar Karwowski, Joe Kider, Joseph Konstan, Sean Koon, Monica Lopez-Gonzalez, Iliana Maifeld-Carucci, Sean McGregor, Gavriel Salvendy, Ben Shneiderman, Constantine Stephanidis, Christina Strobel, Carolyn Ten Holter, and Wei Xu. 2023. Six Human-Centered Artificial Intelligence Grand Challenges. *International Journal of Human-Computer Interaction* 39, 3 (2023), 391–437. <https://doi.org/10.1080/10447318.2022.2153320> arXiv:<https://doi.org/10.1080/10447318.2022.2153320>
- [49] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for Datasets. *CoRR abs/1803.09010* (2018). arXiv:1803.09010 <http://arxiv.org/abs/1803.09010>
- [50] Sarah Gilbert, Jessica Vitak, and Katie Shilton. 2021. Measuring Americans' comfort with research uses of their social media data. *Social Media+ Society* 7, 3 (2021), 20563051211033824.
- [51] Google. 2019. Lookout - Assisted Vision. https://play.google.com/store/apps/details?id=com.google.android.apps.accessibility.reveal&hl=en_US&gl=US
- [52] Anhong Guo, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach, and Meredith Ringel Morris. 2019. Toward Fairness in AI for People with Disabilities: A Research Roadmap. *arXiv preprint arXiv:1907.02227*.
- [53] D. Gurari, Q. Li, C. Lin, Y. Zhao, A. Guo, A. Stangl, and J. P. Bigham. 2019. VizWiz-Priv: A Dataset for Recognizing the Presence and Purpose of Private Visual Information in Images Taken by Blind People. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 939–948. <https://doi.org/10.1109/CVPR.2019.00103>
- [54] Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. VizWiz Grand Challenge: Answering Visual Questions from Blind People. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Jun 2018). <https://doi.org/10.1109/cvpr.2018.00380>
- [55] Foad Hamidi, Kellie Poneris, Aaron Massey, and Amy Hurst. 2018. Who Should Have Access to My Pointing Data?: Privacy Tradeoffs of Adaptive Assistive Technologies. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (ASSETS '18). ACM, New York, NY, USA, 203–216. <https://doi.org/10.1145/3234695.3239331>
- [56] Saad Hassan, Larwan Berke, Elahe Vahdani, Longlong Jing, Yingli Tian, and Matt Huenerfauth. 2020. An Isolated-Signing RGBD Dataset of 100 American Sign Language Signs Produced by Fluent ASL Signers. In *Proceedings of the LREC2020 9th Workshop on the Representation and Processing of Sign Languages: Sign Language Resources in the Service of the Language Community, Technological Challenges and Application Perspectives*. European Language Resources Association (ELRA), Marseille, France, 89–94. <https://www.aclweb.org/anthology/2020.signlang-1.14>
- [57] Milton Heumann, Lance Cassak, Esther Kang, and Thomas Twitchell. 2016. Privacy and Surveillance: Public Attitudes on Cameras on the Street, in the Home, and in the Workplace. *Rutgers J.L. & Pub. Pol'y* 14 (2016), 37.
- [58] Elizabeth M Hill, Emma L Turner, Richard M Martin, and Jenny L Donovan. 2013. "Let's get the best quality research we can": public awareness and acceptance of consent to use existing data in health research: a systematic review and qualitative study. *BMC medical research methodology* 13, 1 (2013), 1–10.
- [59] Tom Hitron, Yoav Orlev, Iddo Wald, Ariel Shamir, Hadas Erel, and Oren Zuckerman. 2019. Can Children Understand Machine Learning Concepts? The Effect of Uncovering Black Boxes. 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, 1–11. <https://doi.org/10.1145/3290605.3300645>
- [60] Jonggi Hong, Jaina Gandhi, Ernest Essuah Mensah, Farnaz Zamiri Zeraati, Ebrima Jarjue, Kyungjun Lee, and Hernisa Kacorri. 2022. Blind Users Accessing Their Training Images in Teachable Object Recognizers. 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 14, 18 pages. <https://doi.org/10.1145/3517428.3544824>
- [61] Jonggi Hong, Kyungjun Lee, June Xu, and Hernisa Kacorri. 2020. *Crowdsourcing the Perception of Machine Teaching*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376428>
- [62] Matt Huenerfauth and Hernisa Kacorri. 2014. Release of experimental stimuli and questions for evaluating facial expressions in animations of American Sign Language. In *Proceedings of the 6th Workshop on the Representation and Processing of Sign Languages: Beyond the Manual Channel, The 9th International Conference on Language Resources and Evaluation (LREC 2014)*. Reykjavik, Iceland.
- [63] Marcello Ienca, James Scheibner, Agata Ferretti, Felix Gille, Julia Amann, Joanna Sleigh, Alessandro Blasimme, and Effy Vayena. 2019. *How the general data protection regulation changes the rules for scientific research: Study*. Technical Report. ETH Zurich.
- [64] Michaela Jackson, Paul Harrison, Boyd Swinburn, and Mark Lawrence. 2015. Using a qualitative vignette to explore a complex public health issue. *Qualitative health research* 25, 10 (2015), 1395–1409.
- [65] Susan Jamieson. 2004. Likert scales: How to (ab) use them? *Medical education* 38, 12 (2004), 1217–1218.
- [66] Beth Lenore St Jean. 2012. *Information behavior of people diagnosed with a chronic serious health condition: A longitudinal study*. Ph.D. Dissertation. University of Michigan.
- [67] Eun Seo Jo and Timnit Gebru. 2020. Lessons from Archives: Strategies for Collecting Sociocultural Data in Machine Learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 306–316.

- <https://doi.org/10.1145/3351095.3372829>
- [68] Hernisa Kacorri. 2016. *Data-Driven Synthesis and Evaluation of Syntactic Facial Expressions in American Sign Language Animation*. Ph.D. Dissertation. CUNY Academic Works. https://academicworks.cuny.edu/gc_etds/1375
 - [69] Hernisa Kacorri. 2017. Teachable Machines for Accessibility. *SIGACCESS Access. Comput.* 119, 10–18. <https://doi.org/10.1145/3167902.3167904>
 - [70] Hernisa Kacorri, Utkarsh Dwivedi, Sravya Amancherla, Mayanka Jha, and Riya Chanduka. 2020. Incluset: A Data Surfacing Repository for Accessibility Datasets. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '20)*. Association for Computing Machinery, New York, NY, USA, Article 72, 4 pages. <https://doi.org/10.1145/3373625.3418026>
 - [71] Hernisa Kacorri, Utkarsh Dwivedi, and Rie Kamikubo. 2020. Data Sharing in Wellness, Accessibility, and Aging. *NeurIPS 2020 Workshop on Dataset Curation and Security* (2020).
 - [72] Hernisa Kacorri, Kris M. Kitani, Jeffrey P. Bigham, and Chieko Asakawa. 2017. People with Visual Impairment Training Personal Object Recognizers: Feasibility and Challenges. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (*CHI '17*). Association for Computing Machinery, New York, NY, USA, 5839–5849. <https://doi.org/10.1145/3025453.3025899>
 - [73] Hernisa Kacorri, Sergio Mascetti, Andrea Gerino, Dragan Ahmetovic, Hironobu Takagi, and Chieko Asakawa. 2016. Supporting Orientation of People with Visual Impairment: Analysis of Large Scale Usage Data. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (Reno, Nevada, USA) (*ASSETS '16*). ACM, New York, NY, USA, 151–159. <https://doi.org/10.1145/2982142.2982178>
 - [74] Kaggle. 2021. Kaggle: Your Machine Learning and Data Science Community. <https://www.kaggle.com>.
 - [75] Rie Kamikubo, Utkarsh Dwivedi, and Hernisa Kacorri. 2021. Sharing Practices for Datasets Related to Accessibility and Aging. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, USA) (*ASSETS '21*). Association for Computing Machinery, New York, NY, USA, Article 28, 16 pages. <https://doi.org/10.1145/3441852.3471208>
 - [76] Rie Kamikubo, Lining Wang, Crystal Marte, Amnah Mahmood, and Hernisa Kacorri. 2022. Data Representativeness in Accessibility Datasets: A Meta-Analysis. 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 8, 15 pages. <https://doi.org/10.1145/3517428.3544826>
 - [77] Young-Ho Kim, Jae Ho Jeon, Bongshin Lee, Eun Kyoung Choe, and Jinwook Seo. 2017. OmniTrack: A flexible self-tracking approach leveraging semi-automated tracking. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 1, 3 (2017), 1–28.
 - [78] Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security* 64 (2017), 122–134.
 - [79] Mauritz Kop. 2021. Machine learning and EU data-sharing practices: Legal aspects of machine learning training datasets for AI systems. In *Research Handbook on Big Data Law*. Edward Elgar Publishing, 432–453.
 - [80] Patty Kostkova, Helen Brewer, Simon De Lusignan, Edward Fottrell, Ben Goldacre, Graham Hart, Phil Koczan, Peter Knight, Corinne Marsolier, Rachel A McKendry, et al. 2016. Who owns the data? Open data for healthcare. *Frontiers in public health* 4 (2016), 7.
 - [81] Jacob Leon Kröger, Milagros Miceli, and Florian Müller. 2021. How Data Can Be Used Against People: A Classification of Personal Data Misuses. *Available at SSRN 3887097* (2021).
 - [82] Jenny Krutzinna and Luciano Floridi. 2019. Ethical medical data donation: a pressing issue. *The Ethics of Medical Data Donation* (2019), 1–6.
 - [83] Jennifer Kulynych and David Korn. 2003. The new HIPAA (Health Insurance Portability and Accountability Act of 1996) Medical Privacy Rule: help or hindrance for clinical research? *Circulation* 108, 8 (2003), 912–914.
 - [84] Kyungjun Lee and Hernisa Kacorri. 2019. Hands Holding Clues for Object Recognition in Teachable Machines. 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/10.1145/3290605.3300566>
 - [85] Kyungjun Lee, Daisuke Sato, Saki Asakawa, Hernisa Kacorri, and Chieko Asakawa. 2020. *Pedestrian Detection with Wearable Cameras for the Blind: A Two-Way Perspective*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376398>
 - [86] Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. 2018. The privacy policy landscape after the GDPR. *arXiv preprint arXiv:1809.08396* (2018).
 - [87] Maryam Lotfian, Jens Ingensand, and Maria Antonia Brovelli. 2020. A framework for classifying participant motivation that considers the typology of citizen science projects. *ISPRS International Journal of Geo-Information* 9, 12 (2020), 704.
 - [88] Yuhuan Luo, Peiyi Liu, and Eun Kyoung Choe. 2019. Co-Designing food trackers with dietitians: Identifying design opportunities for food tracker customization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [89] Yuhuan Luo, Chi Young Oh, Beth St Jean, Eun Kyoung Choe, et al. 2020. Interrelationships Between Patients' Data Tracking Practices, Data Sharing Practices, and Health Literacy: Onsite Survey Study. *Journal of medical Internet research* 22, 12 (2020), e18937.
 - [90] Jared Lyle. 2014. OpenICPSR. *Bulletin of the Association for Information Science and Technology* 40, 5 (2014), 55–56.
 - [91] Mavis Machirori and Reema Patel. 2021. Turning distrust in data sharing into “engage, deliberate, decide”. (2021). <https://www.adalovelaceinstitute.org/blog/distrust-data-sharing-engage-deliberate-decide/>
 - [92] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
 - [93] Daniela Massiceti, Luisa Zintgraf, John Bronskill, Lida Theodorou, Matthew Tobias Harris, Edward Cutrell, Cecily Morrison, Katja Hofmann, and Simone Stumpf. 2021. Orbit: A real-world few-shot dataset for teachable object recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10818–10828.
 - [94] Joseph A Maxwell. 2010. Using numbers in qualitative research. *Qualitative inquiry* 16, 6 (2010), 475–482.
 - [95] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.
 - [96] Roisin McNaney, Catherine Morgan, Pranav Kulkarni, Julio Vega, Farnoosh Heidarincheh, Ryan McConville, Alan Whone, Mickey Kim, Reuben Kirkham, and Ian Craddock. 2022. Exploring Perceptions of Cross-Sectoral Data Sharing with People with Parkinson's. In *CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 280, 14 pages. <https://doi.org/10.1145/3491102.3501984>
 - [97] Roisin McNaney, Emmanuel Tsekles, and Jonathan Synnott. 2020. *Future Opportunities for IoT to Support People with Parkinson's*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3313831.3376871>
 - [98] Helena M Mentis, Anita Komlodi, Katrina Schrader, Michael Phipps, Ann Gruber-Baldini, Karen Yarbrough, and Lisa Shulman. 2017. Crafting a view of self-tracking data in the clinical visit. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5800–5812.
 - [99] Microsoft. 2017. Seeing AI: An app for visually impaired people that narrates the world around you. <https://www.microsoft.com/en-us/garage/wall-of-fame/seeing-ai/>
 - [100] Sonali R Mishra, Andrew D Miller, Shefali Haldar, Maher Khelifi, Jordan Eschler, Rashmi G Elera, Ari H Pollack, and Wanda Pratt. 2018. Supporting collaborative health tracking in the hospital: patients' perspectives. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
 - [101] Meredith Ringel Morris. 2020. AI and Accessibility. *Commun. ACM* 63, 6, 35–37. <https://doi.org/10.1145/3356727>
 - [102] Jessica Mozersky, Meredith Parsons, Heidi Walsh, Kari Baldwin, Tristan McIntosh, and James M DuBois. 2020. Research participant views regarding qualitative data sharing. *Ethics & human research* 42, 2 (2020), 13–27.
 - [103] Karen Nakamura. 2019. My Algorithms Have Determined You're Not Human: AI-ML, Reverse Turing-Tests, and the Disability Experience. In *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>
 - [104] Jennifer Nicholas, Katie Shilton, Stephen M Schueller, Elizabeth L Gray, Mary J Kwasny, David C Mohr, et al. 2019. The role of data type and recipient in individuals' perspectives on sharing passively collected smartphone data for mental health: Cross-sectional questionnaire study. *JMIR mHealth and uHealth* 7, 4 (2019), e12578.
 - [105] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119.
 - [106] Lika Nusbaum, Brenda Douglas, Karla Damus, Michael Paasche-Orlow, and Neenah Estrella-Luna. 2017. Communicating risks and benefits in informed consent for research: a qualitative study. *Global Qualitative Nursing Research* 4 (2017), 2333393617732017.
 - [107] 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*. 52–63.
 - [108] Rubens Lacerda Queiroz, Fábio Ferrentini Sampaio, Cabral Lima, and Priscila Machado Vieira Lima. 2020. AI from concrete to abstract: demystifying artificial intelligence to the general public. <https://doi.org/10.1007/s00146-021-01151-x> arXiv:2006.04013 [cs.CY]
 - [109] Ester A Rake, Marleen MHJ Van Gelder, David C Grim, Barend Heeren, Lucien JLP Engelen, and Tom H Van De Belt. 2017. Personalized consent flow in

- contemporary data sharing for medical research: A viewpoint. *BioMed Research International* 2017 (2017).
- [110] Joanna Redden, Jessica Brand, and Vanesa Terzieva. 2020. Data harm record. (2020).
 - [111] Xiaofeng Ren and Chunhui Gu. 2010. Figure-ground segmentation improves handled object recognition in egocentric video. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 3137–3144.
 - [112] Microsoft Research. 2018. *Microsoft Research Open Data*. <https://msropendata.com/>
 - [113] Dana Rotman, Jenny Preece, Jen Hammock, Kezee Procita, Derek Hansen, Cynthia Parr, Darcy Lewis, and David Jacobs. 2012. Dynamic Changes in Motivation in Collaborative Citizen-Science Projects. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (CSCW '12). Association for Computing Machinery, New York, NY, USA, 217–226. <https://doi.org/10.1145/2145204.2145238>
 - [114] 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, 3 (2015), 211–252.
 - [115] Margaret Sandelowski. 2001. Real qualitative researchers do not count: The use of numbers in qualitative research. *Research in nursing & health* 24, 3 (2001), 230–240.
 - [116] Margaret Sandelowski, Corrine I Voils, and George Knaf. 2009. On quantizing. *Journal of mixed methods research* 3, 3 (2009), 208–222.
 - [117] Nancy E Schoenberg and Hege Ravdal. 2000. Using vignettes in awareness and attitudinal research. *International journal of social research methodology* 3, 1 (2000), 63–74.
 - [118] Andrew Sears and Vicki L. Hanson. 2012. Representing Users in Accessibility Research. *ACM Trans. Access. Comput.* 4, 2, Article 7, 6 pages. <https://doi.org/10.1145/2141943.2141945>
 - [119] Nisha Shah, Victoria Coathup, Harriet Teare, Ian Forgie, Giuseppe Nicola Gior-dano, Tue Haldor Hansen, Lenka Groeneveld, Michelle Hudson, Ewan Pearson, Hartmut Ruetten, et al. 2019. Motivations for data sharing—views of research participants from four European countries: a DIRECT study. *European Journal of Human Genetics* 27, 5 (2019), 721–729.
 - [120] Ather Sharif, Aedan Liam McCall, and Kianna Rocas 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>
 - [121] Katie Shilton. 2009. Four billion little brothers? Privacy, mobile phones, and ubiquitous data collection. *Commun. ACM* 52, 11 (2009), 48–53.
 - [122] Alireza Sahami Shirazi, James Clawson, Yashar Hassanpour, Mohammad J Tourian, Albrecht Schmidt, Ed H Chi, Marko Borazio, and Kristof Van Laer-hoven. 2013. Already up? using mobile phones to track & share sleep behavior. *International Journal of Human-Computer Studies* 71, 9 (2013), 878–888.
 - [123] Ben Shneiderman. 2020. Bridging the gap between ethics and practice: guide-lines for reliable, safe, and trustworthy human-centered AI systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 4 (2020), 1–31.
 - [124] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (2020), 495–504.
 - [125] Ido Sivan-Sevilla. 2022. Varieties of enforcement strategies post-GDPR: a fuzzy-set qualitative comparative analysis (fsQCA) across data protection authorities. *Journal of European Public Policy* (2022), 1–34.
 - [126] Robert Soden, Michael Skirpan, Casey Fiesler, Zahra Ashktorab, Eric P. S. Baumer, Mark Blythe, and Jasmine Jones. 2019. CHI4EVL: Creative Speculation on the Negative Impacts of HCI Research. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290607.3299033>
 - [127] Daniel Solove. 2014. 10 Reasons Why Privacy Matters. <https://teachprivacy.com/10-reasons-privacy-matters/>.
 - [128] Daniel J Solove. 2008. Understanding privacy. (2008).
 - [129] Joan Sosa-García and Francesca Odono. 2017. “Hands on” visual recognition for visually impaired users. *ACM Transactions on Accessible Computing (TACCESS)* 10, 3 (2017), 1–30.
 - [130] Abigale Stangl, Kristina Shiroma, Nathan Davis, Bo Xie, Kenneth R Fleischmann, Leah Findlater, and Danna Gurari. 2022. Privacy Concerns for Visual Assistance Technologies. *ACM Transactions on Accessible Computing (TACCESS)* 15, 2 (2022), 1–43.
 - [131] Gail M Sullivan and Anthony R Artino Jr. 2013. Analyzing and interpreting data from Likert-type scales. *Journal of graduate medical education* 5, 4 (2013), 541–542.
 - [132] Darren B Taichman, Peush Sahni, Anja Pinborg, Larry Peiperl, Christine Laine, Astrid James, Sung-Tae Hong, Abraham Haileamlak, Laragh Gollogly, Fiona Godlee, Frank A Frizelle, Fernando Florenzano, Jeffrey M Drazen, Howard Bauchner, Christopher Baethge, and Joyce Backus. 2017. Data sharing statements for clinical trials: a requirement of the International Committee of Medical Journal Editors. *The Lancet* 389, 10086 (June 2017), e12–e14. [https://doi.org/10.1016/S0140-6736\(17\)31282-5](https://doi.org/10.1016/S0140-6736(17)31282-5)
 - [133] Gioacchino Tangari, Muhammad Ikram, Kiran Ijaz, Mohamed Ali Kaafar, and Shlomo Berkovsky. 2021. Mobile health and privacy: cross sectional study. *bmj* 373 (2021).
 - [134] Rachael Tatman. 2017. Gender and Dialect Bias in YouTube’s Automatic Cap-tions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*. Association for Computational Linguistics, Valencia, Spain, 53–59. <https://doi.org/10.18653/v1/W17-1606>
 - [135] Lida Theodorou, Daniela Massiceti, Luisa Zintgraf, Simone Stumpf, Cecily Morrison, Edward Cutrell, Matthew Tobias Harris, and Katja Hofmann. 2021. Disability-first Dataset Creation: Lessons from Constructing a Dataset for Teach-able Object Recognition with Blind and Low Vision Data Collectors. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–12.
 - [136] Llewellyn DW Thomas and Aija Leiponen. 2016. Big data commercialization. *IEEE Engineering Management Review* 44, 2 (2016), 74–90.
 - [137] Jutta Treviranus. 2019. The Value of Being Different. In *Proceedings of the 16th Web For All 2019 Personalization - Personalizing the Web* (San Francisco, CA, USA) (W4A '19). ACM, New York, NY, USA, Article 1, 7 pages. <https://doi.org/10.1145/3315002.3332429>
 - [138] Shari Trewin. 2018. AI fairness for people with disabilities: Point of view. *arXiv preprint arXiv:1811.10670* (2018).
 - [139] Yu-Yun Tseng, Alexander Bell, and Danna Gurari. 2022. VizWiz-FewShot: Lo-cating Objects in Images Taken by People with Visual Impairments. In *Com-puter Vision – ECCV 2022*, Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer Nature Switzerland, Cham, 575–591.
 - [140] Frank WSM Verheggen, Fred Nieman, and Ruud Jonkers. 1998. Determinants of patient participation in clinical studies requiring informed consent: why patients enter a clinical trial. *Patient education and counseling* 35, 2 (1998), 111–125.
 - [141] Stefaan Verhulst, Beth Simone Noveck, Robyn Caplan, Kristy Brown, and Clau-dia Paz. 2014. The open data era in health and social care. *NYLS Legal Studies Research Paper* 2563788 (2014).
 - [142] W Kip Viscusi and Richard J Zeckhauser. 2015. The relative weights of direct and indirect experiences in the formation of environmental risk beliefs. *Risk Analysis* 35, 2 (2015), 318–331.
 - [143] Jessica Vitak and Katie Shilton. 2020. Trust, privacy and security, and accessibil-ity considerations when conducting mobile technologies research with older adults. In *Mobile Technology for Adaptive Aging: Proceedings of a Workshop*. National Academies Press, 1–20.
 - [144] Chat Wacharamanatham, Lukas Eisenring, Steve Haroz, and Florian Echtler. 2020. Transparency of CHI Research Artifacts: Results of a Self-Reported Survey. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376448>
 - [145] Mark Walport and Paul Brest. 2011. Sharing research data to improve public health. *The Lancet* 377, 9765 (2019/09/15 2011), 537–539. [https://doi.org/10.1016/S0140-6736\(10\)62234-9](https://doi.org/10.1016/S0140-6736(10)62234-9)
 - [146] Ryen W White, P Murali Doraiswamy, and Eric Horvitz. 2018. Detecting neu-rodegenerative disorders from web search signals. *NPJ digital medicine* 1, 1 (2018), 1–4.
 - [147] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kazi-unas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, Bias, and AI. *AI Now Institute, November* (2019).
 - [148] Eliot Wilczek. 2021. Archival Engagements with Wicked Problems. *The Ameri-can Archivist* 84, 2 (2021), 468–501.
 - [149] Heng Xu, Xin Robert Luo, John M Carroll, and Mary Beth Rosson. 2011. The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. *Decision support systems* 51, 1 (2011), 42–52.
 - [150] Yu Zhong, Pierre J. Garrigues, and Jeffrey P. Bigham. 2013. Real Time Ob-ject Scanning Using a Mobile Phone and Cloud-Based Visual Search En-gine. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (Bellevue, Washington) (ASSETS '13). Asso-ciation for Computing Machinery, New York, NY, USA, Article 20, 8 pages. <https://doi.org/10.1145/2513383.2513443>
 - [151] Anneke Zuiderwijk, Marijn Janssen, Geerten Van De Kaa, and Kostas Poulis. 2016. The wicked problem of commercial value creation in open data ecosystems: Policy guidelines for governments. *Information policy* 21, 3 (2016), 223–236.