



Considerations for AI Fairness for People with Disabilities

Shari Trewin (IBM; trewin@us.ibm.com)

Sara Basson (Google Inc.; basson@google.com)

Michael Muller (IBM Research; michael_muller@us.ibm.com)

Stacy Branham (University of California, Irvine; sbranham@uci.edu)

Jutta Treviranus (OCAD University; jtreviranus@ocadu.ca)

Daniel Gruen (IBM Research; daniel_gruen@us.ibm.com)

Daniel Hebert (IBM; dhebert@us.ibm.com)

Natalia Lyckowski (IBM; nladuca@us.ibm.com)

Erich Manser (IBM; emanser@us.ibm.com)

DOI: [10.1145/3362077.3362086](https://doi.org/10.1145/3362077.3362086)

Abstract

In society today, people experiencing disability can face discrimination. As artificial intelligence solutions take on increasingly important roles in decision-making and interaction, they have the potential to impact fair treatment of people with disabilities in society both positively and negatively. We describe some of the opportunities and risks across four emerging AI application areas: employment, education, public safety, and healthcare, identified in a workshop with participants experiencing a range of disabilities. In many existing situations, non-AI solutions are already discriminatory, and introducing AI runs the risk of simply perpetuating and replicating these flaws. We next discuss strategies for supporting fairness in the context of disability throughout the AI development lifecycle. AI systems should be reviewed for potential impact on the user in their broader context of use. They should offer opportunities to redress errors, and for users and those impacted to raise fairness concerns. People with disabilities should be included when sourcing data to build models, and in testing, to create a more inclusive and robust system. Finally, we offer pointers into an established body of literature on human-centered design processes and philosophies that may assist AI and ML engineers in innovating algorithms that reduce harm and ultimately enhance the lives of people with disabilities.

Introduction

Systems that leverage Artificial Intelligence are becoming pervasive across industry sectors (Costello, 2019), as are concerns that these technologies can unintentionally exclude or lead to unfair outcomes for marginalized populations (Bird, Hutchinson, Kenthapadi, Kiciman, & Mitchell, 2019)(Cutler, Pribik, & Humphrey, 2019)(IEEE & Systems, 2019)(Kroll et al., 2016)(Lepri, Oliver, Letouzé, Pentland, & Vinck, 2018). Initiatives to improve AI fairness for people across racial (Hankerson et al., 2016), gender (Hamidi, Scheurman, & Branham, 2018a), and other identities are emerging, but there has been relatively little work focusing on AI fairness for people with disabilities. There are numerous examples of AI that can empower people with disabilities, such as autonomous vehicles (Brewer & Kameswaran, 2018) and voice agents (Pradhan, Mehta, & Findlater, 2018) for people with mobility and vision impairments. However, AI solutions may also result in unfair outcomes, as when Idahoans with cognitive/learning disabilities had their healthcare benefits reduced based on biased AI (*K.W. v. Armstrong, No. 14-35296 (9th Cir. 2015) :: Justia, 2015*). These scenarios suggest that the prospects of AI for people with disabilities are promising yet fraught with challenges that require the sort of upfront attention to ethics in the development process advocated by scholars (Bird et al., 2019) and practitioners (Cutler et al., 2019).

The challenges of ensuring AI fairness in the context of disability emerge from multiple sources. From the very beginning of al-

gorithmic development, in the problem scoping stage, bias can be introduced by lack of awareness of the experiences and use cases of people with disabilities. Since systems are predicated on data, in data sourcing and data pre-processing stages, it is critical to gather data that include people with disabilities and to ensure that these data are not completely subsumed by data from presumed “normative” populations. This leads to a potential conundrum. The data need to be gathered in order to be reflected in the models, but confidentiality and privacy, especially as regards disability status, might make collecting these data difficult (for developers) or dangerous (for subjects) (Faucett, Ringland, Cullen, & Hayes, 2017)(von Schrader, Malzer, & Bruyère, 2014). Another area to address during model training and testing is the potential for model bias. Owing to intended or unintended bias in the data, the model may inadvertently enforce or reinforce discriminatory patterns that work against people with disabilities (Janssen & Kuk, 2016). We advocate for increased awareness of these patterns, so we can avoid replication of past bias into future algorithmic decisions, as has been well-documented in banking (Bruckner, 2018)(Chander, 2017)(Hurley & Adebayo, 2016). Finally, once a trained model is incorporated in an application, it is then critical to test with diverse users, particularly those deemed as outliers. This paper provides a number of recommendations towards overcoming these challenges.

In the remainder of this article, we overview the nascent area of AI Fairness for People with Disabilities as a practical pursuit and an academic discipline. We provide a series of examples that demonstrate the potential for harm to people with disabilities across four emerging AI application areas: employment, education, public safety, and healthcare. Then, we identify strategies of developing AI algorithms that resist reifying systematic societal exclusions at each stage of AI development. Finally, we offer pointers into an established body of literature on human-centered design processes and philosophies that may assist AI and ML engineers in innovating algorithms that reduce harm and – as should be our ideal – ultimately enhance the lives of people with disabilities.

Related Work

The 2019 Gartner CIO survey (Costello, 2019) of 3000 enterprises across major industries reported that 37% have implemented some form of AI solution, an increase of 270% over the last four years. In parallel, there is increasing recognition that intelligent systems should be developed with attention to the ethical aspects of their behavior (Cutler et al., 2019)(IEEE & Systems, 2019), and that fairness should be considered upfront, rather than as an afterthought (Bird et al., 2019). IEEE’s Global Initiative on Ethics of Autonomous and Intelligent Systems is developing a series of international standards for such processes (Koene, Smith, Egawa, Mandalh, & Hatada, 2018), including a process for addressing ethical concerns during design (P7000), and the P7003 Standard for Algorithmic Bias Considerations (Koene, Dowthwaite, & Seth, 2018). There is ongoing concern and discussion about accountability for potentially harmful decisions made by algorithms(Kroll et al., 2016)(Lepri et al., 2018), with some new academic initiatives – like one at Georgetown’s Institute for Tech Law & Policy (Givens, 2019), and a workshop at the ASSETS 2019 conference(Trewin et al., 2019) – focusing specifically on AI and Fairness for People with Disabilities.

Any algorithmic decision-process can be biased, and the FATE/ML community is actively developing approaches for detection and remediation of bias (Kanellopoulos, 2018)(Lohia et al., 2019). Williams, Brooks and Shmargad show how racial discrimination can arise in employment and education even without having social category information, and how the lack of category information makes such biases harder to detect (Williams, Brooks, & Shmargad, 2018). Although they argue for inclusion of social category information in algorithmic decision-making, they also acknowledge the potential harm that can be caused to an individual by revealing sensitive social data such as immigration status. Selbst et al. argue that purely algorithmic approaches are not sufficient, and the full social context of deployment must be considered if fair outcomes are to be achieved (Selbst, Boyd, Friedler, Venkatasubramanian, & Vertesi, 2019).

Some concerns about AI fairness in the

context of individuals with disabilities or neurological or sensory differences are now being raised (Fruchterman & Mellea, 2018)(Guo, Kamar, Vaughan, Wallach, & Morris, 2019)(Lewis, 2019)(Treviranus, 2019)(Trewin, 2018a), but research in this area is sparse. Fruchterman and Mellea (Fruchterman & Mellea, 2018) outline the widespread use of AI tools in employment and recruiting, and highlight some potentially serious implications for people with disabilities, including the analysis of facial movements and voice in recruitment, personality tests that disproportionately screen out people with disabilities, and the use of variables that could be discriminatory, such as gaps in employment. “Advocates for people with disabilities should be looking at the proxies and the models used by AI vendors for these “hidden” tools of discrimination” (Fruchterman & Mellea, 2018).

Motivating Examples

In October 2018, a group of 40 disability advocates, individuals with disabilities, AI and accessibility researchers and practitioners from industry and academia convened in a workshop (Trewin, 2018b) to identify and discuss the topic of fairness for people with disabilities in light of the increasing mainstream application of AI solutions in many industries (Costello, 2019). This section describes some of the opportunities and risks identified by the workshop participants in the areas of employment, education, public safety and healthcare.

Employment

People with disabilities are no strangers to discrimination in hiring practices. In one recent field study, disclosing a disability (spinal cord injury or Asperger’s Syndrome) in a job application cover letter resulted in 26% fewer positive responses from employers, even though the disability was not likely to affect productivity for the position (Ameri et al., 2018). When it comes to inclusive hiring, it has been shown that men and those who lack experience with disability tend to have more negative affective reactions to working with individuals with disabilities (Popovich, Scherbaum, Scherbaum, & Polinko, 2003). Exclusion can be unin-

tentional. For example, qualified deaf candidates who speak through an interpreter may be screened out for a position requiring verbal communication skills, even though they could use accommodations to do the job effectively. Additional discriminatory practices are particularly damaging to this population, where employment levels are already low: In 2018, the employment rate for people with disabilities was 19.1%, while the employment percentage for people without disabilities was 65.9% (Bureau of Labor Statistics, 2019).

Employers are increasingly relying on technology in their hiring practices. One of their selling points is the potential to provide a fairer recruitment process, not influenced by an individual recruiter’s bias or lack of knowledge. Machine learning models are being used for candidate screening and matching job-seekers with available positions. There are AI-driven recruitment solutions on the market today that analyze online profiles and resumes, the results of online tests, and video interview data, all of which raise potential concerns for disability discrimination (Fruchterman & Mellea, 2018). While the use of AI in HR and recruitment is an increasing trend (Faggella, 2019), there are already cautionary incidents of discrimination, as when Amazon’s AI recruiting solution “learned” to devalue resumes of women (Dastin, 2018).

The workshop identified several risk scenarios:

- A deaf person may be the first person using sign language interpretation to apply to an organization. Candidate screening models that learn from the current workforce will perpetuate the status quo, and the biases of the past. They will likely exclude candidates with workplace differences, including those who use accommodations to perform their work.
- An applicant taking an online test using assistive technology may take longer to answer questions, especially if the test itself has not been well designed for accessibility. Models that use timing information may end up systematically excluding assistive technology users. Resumes and job applications may not contain explicit information about a person’s disability, but other variables may be impacted, including gaps

in employment, school attended, or time to perform an online task.

- An applicant with low facial affect could be screened out by a selection process that uses video analysis of eye gaze, voice characteristics, or facial movements, even though she is highly skilled. This type of screening poses a barrier for anyone whose appearance, voice or facial expression differs from the average. It could exclude autistic individuals or blind applicants who do not make eye contact, deaf people and others who do not communicate verbally, people with speech disorders or facial paralysis, or people who have survived a stroke, to name a few.

When the available data do not include many people with disabilities, and reflect existing biases, and the deployed systems rely on proxies that are impacted by disability, the risk of unfair treatment in employment is significant. We must seek approaches that do not perpetuate the biases of the past, or introduce new barriers by failing to recognize qualified candidates because they are different, or use accommodations to do their work.

Education

In the United States, people with disabilities have historically been denied access to free public education (Dudley-Marling & Burns, 2014)(Obiakor, Harris, Mutua, Rotatori, & Algozzine, 2012). It was nearly 20 years after the passing of Brown v. Board of Education, which desegregated public schools along racial lines, that the Education for All Handicapped Children Act was passed (Dudley-Marling & Burns, 2014), mandating that all students are entitled to a “free and appropriate public education” in the “least restrictive environment.” Prior to 1975, a mere one in five learners with disabilities had access to public school environments, often in segregated classrooms (Dudley-Marling & Burns, 2014). Despite great strides, some learners with disabilities still cannot access integrated public learning environments (Dudley-Marling & Burns, 2014), K-12 classroom technologies are often inaccessible (Shaheen & Lazar, 2018), postsecondary online learning materials are often inaccessible (Burgstahler, 2015) (Straumsheim, 2017), and e-learning

platforms do not consider the needs of all learners (Cinquin, Guitton, & Sauzéon, 2019).

AI in the education market is being driven by the rapid transition from onsite classroom based education to online learning. Institutions can now expand their online learning initiatives to reach more students in a cost-effective manner. Industry analyst, Global Market Insights (Bhutani & Wadhvani, 2019), predicts that the market will grow to a \$6 Billion dollar industry by 2024. The new generation of online learning platforms are integrated with AI technologies and use them to personalize learning (and testing) for each student, among other applications. Two examples of providers of these systems are from traditional Learning Management System (LMS) vendors like Blackboard; and more recently from the Massive Open Online Course (MOOC) providers like edX.

Personalized learning could provide enormous benefits for learners with disabilities, e.g. (Morris, Kirschbaum, & Picard, 2010). It could be as simple as augmenting existing content with additional illustrations and pictures for students who are classified as visual learners, or as complex as generation of personalized user interfaces (Gajos, Wobbrock, & Weld, 2007). For non-native language speakers, including deaf learners, the system could provide captions for video content so the student can read along with the lecture.

Any system that makes inferences about a student’s knowledge and abilities based on their online interactions runs the risk of misinterpreting and underestimating students with disabilities. Students whose learning capabilities or styles are outside the presumed norm may not receive fair treatment. For example, if there is a rigid time constraint for completing a test or a quiz, a student that has a cognitive disability or test anxiety where they process information more slowly than other students would be assessed as being less capable than they are.

Unlike other areas, in an educational setting, disability information may often be available, and the challenge is to provide differentiated education for a wide range of people, without introducing bias against disability groups.

Public Safety

People with disabilities are much more likely to be victims of violent crime than people without disabilities (Harrell, 2017). Threats come not only from other citizens, but also from law enforcement itself; for example, police officers can misinterpret people with disabilities as being uncooperative or even threatening, and deprive them of access to Miranda warnings (US Department of Justice Civil Rights Division, 2006). Law enforcement's implicit bias and discrimination towards people with disabilities, as well as the potential for technology to address these challenges, are both featured in the 2015 Final Report of the President's Task Force on 21st Century Policing (Policing, 2015).

The application of AI technology to identify threats to public safety and enforce the law is highly controversial (McCarthy, 2019). This includes technology for identifying people, recognizing individuals, and for interpreting behavior (for example, whether someone is acting in a suspicious manner). Aside from the threat to personal privacy, the potential for errors and biased performance is very real. While public discourse and academic attention has so far focused on racial and gender disparities, workshop participants identified serious concerns and also some opportunities for people with disabilities.

Autonomous vehicles must be able to identify people in the environment with great precision. They must reliably recognize individuals who use different kinds of wheelchair and mobility devices, or move in an unusual way. One workshop participant described a wheelchair-using friend who propels themselves backwards with their feet. This is an unusual method of getting around, but the ability to recognize and correctly identify such outlier individuals is a matter of life and death.

Another participant had recently observed, a dishevelled man pacing restlessly in an airport lounge, muttering to himself, clearly in a state of high stress. His behavior could be interpreted by both humans and AI analysis as a potential threat. Instead he may be showing signs of an anxiety disorder, autism, or a strong fear of flying. Deaf signers' strong facial expressions can be misinterpreted (Shaffer & Rogan, 2018), leading to them being wrongly

identified as being angry, and a potential security threat. Someone with an altered gait could be using a prosthesis, not hiding a weapon.

People with cognitive disabilities may be at especially high risk of being misidentified as a potential threat. Combining this with the need to respond quickly to genuine threats creates a dangerous situation and requires careful design of the AI system and its method of deployment.

There may also be opportunities for AI to improve public safety for people with disabilities. For example, AI-based interpretation could be trained to 'understand' a wide range of behaviors including hand flapping, pacing and sign language, as normal. A recent survey and interview study of people who are blind (Branham et al., 2017) suggests that facial and image recognition technologies could better support personal safety for individuals with sensory disabilities. They may support locating police officers and identifying fraudulent actors claiming to be officials. They may allow a person who is blind or deaf to be made aware of a weapon being brandished or discharged. They may support access to facial cues for more cautious and effective interactions with a potential aggressor or a police officer. These technologies may even allow blind individuals to provide more persuasive evidence to catch their perpetrators.

When considering the ethics of proposed projects in this space, the potential risks for individuals with disabilities should also be evaluated and addressed in the overall design. For example, a system could highlight someone with an altered gait, and list possibilities as someone hiding a weapon, or someone using a prosthesis or mobility device. In a situation where facial recognition is being used, a person's profile could help responders to avoid misunderstanding, but again this comes at the cost of sacrificing privacy, and potentially doing harm to other marginalized groups (Hamidi et al., 2018a). An overall balance must be found between using AI as a tool for maintaining public safety while minimizing negative outcomes for vulnerable groups and outlier individuals.

Healthcare

Today there are large disparities in access to healthcare for people with disabilities (Iezzoni, 2011) (Krahn, Walker, & Correa-De-Araujo, 2015), especially those with developmental disabilities (Krahn, Hammond, & Turner, 2006). Patients who are non-verbal or patients with cognitive impairments are often under-served or incorrectly served (Barnett, McKee, Smith, & Pearson, 2011) (Krahn & Fox, 2014) (Krahn et al., 2015). Deaf patients are often misdiagnosed with having a mental illness or a disorder, because of lack of cultural awareness or language barrier (Glickman, 2007) (Pollard, 1994). Another area that is underserved in the current model are patients with rare diseases or genetic disorders, that do not fall within standard protocols (Wastfelt, Fadeel, & Henter, 2006). For older adults with deteriorating health, this may lead to unwanted institutionalization. Promising technological developments, many of which include AI, abound, but need to better incorporate target users in the development process (Haigh & Yanco, 2002).

AI applications in healthcare could help to overcome some of the barriers preventing people getting access to the care and preventative care they need. For example, a non-verbal person may have difficulty communicating a problem they are experiencing. With respect to pain management or prescription delivery, AI can remove the requirement that patients advocate on their own behalf. For complex cases where disabilities or communicative abilities may affect treatment and ability to adhere to a treatment plan, AI could be applied to recognize special needs situations and flag them for extra attention, and build a case for a suitable course of treatment. With respect to rare diseases or genetic disorders, disparate data points can be aggregated such that solution determination and delivery is not contingent on an individual practitioner's know-how.

Unfortunately, there are no standards or regulations to assess the safety and efficacy of these systems. If the datasets don't well-represent the broader population, AI might work less well where data are scarce or difficult to collect. This can negatively impact people with rare medical conditions/disabilities.

For example, if speech pauses are used to diagnose conditions like Alzheimer's disease, a person whose speech is already affected by a disability may be wrongly diagnosed, or their diagnosis may be missed because the system does not work for them. Pauses in speech can be because person is a non-native speaker; and not a marker of disease.

Just as we have seen in the domains of employment, education, and public safety, if healthcare applications are built for the extremes of currently excluded populations, the solution stands to improve fairness in access, instead of locking people out. Across all domains, AI applications pose both risks and opportunities for people with disabilities. The question remains: how and when can fairness for people with disabilities be implemented in the software development process towards minimizing risks and maximizing benefits? In the following section, we address this question for each stage of the AI development process.

Considerations for AI Practitioners

In this section, we recommend ways AI practitioners can be aware of, and work towards fairness and inclusion for people with disabilities in their AI-based applications. The section is organized around the typical stages of AI model development: problem scoping, data sourcing, pre-processing, model selection and training, and incorporating AI in an application.

Problem Scoping

Some projects have greater potential to impact human lives than others. To identify areas where special attention may need to be paid to fairness, it can be helpful to apply the BioSS AI Protocol (BioSS, 2019), which recommends asking the following 5 questions about the work being done by AI:

1. Is the work **Advisory**, leaving space for human judgement and decision making?
2. Has the AI been granted any **Authority** over people?
3. Does the AI have **Agency** (the ability to act in a given environment)?
4. What skills and responsibilities are we at risk of **Abdicating**?

5. Are lines of **Accountability** clear, in what are still organizations run by human beings?

AI practitioners can also investigate whether this is an area where people with disabilities have historically experienced discrimination, such as employment, housing, education, and healthcare. If so, can the project improve on the past? Identify what specific outcomes there should be, so these can be checked as the project progresses. Develop a plan for tackling bias in source data to avoid perpetuating previous discriminatory treatment. This could include boosting representation of people with disabilities, adjusting for bias against specific disability groups, or flagging gaps in data coverage so the limits of the resulting model are explicit.

Approaches to developing ethical AI include actively seeking the ongoing involvement of a diverse set of stakeholders (Cutler et al., 2019), and a diversity of data to work with. To extend this approach to people with disabilities, it may be useful to define a set of 'outlier' individuals, and include them in the team, following an inclusive design method as discussed in the following section. These are people whose data may look very different to the average. What defines an outlier depends on the application. Many variables can be impacted by a disability, leading to a potential for bias even where no explicit disability information is available. For example, in speech recognition it could be a person with a stutter or a person with slow, slurred speech. In a healthcare application involving height, this could mean including a person of short stature. Outliers may also include people who belong with one group, but whose data looks more like that of another group. For example, a person who is slow to take a test may not be struggling with the material, but with typing, or accessing the test itself through their assistive technology. By defining outlier individuals up front, the design process can consider at each stage what their needs are, whether there are potential harms that need to be avoided, and how to achieve this.

Related to identifying outliers, developing a measurement plan is also valuable at this stage. If the plan includes expected outcomes for outliers and disability groups, this can impact what data (including people) are

included, and what data and people are left out.

Finally, a word of warning. From a machine learning perspective, an obvious solution to handling a specialized sub-group not typical of the general population might be to develop a specialized model for that group. For example, a specialized speech recognition model tuned to the characteristics of people with slurred speech, or people who stutter. From an ethical perspective, solutions that handle outliers and disability groups by routing them to an alternative service require careful thinking. Individuals may not wish to self-identify as having a disability, and there may be legal protections against requiring self-declaration. Solutions that attempt to infer disability status, or infer a quality that serves as a proxy for disability status also present an ethical minefield. It may be acceptable to detect stuttered speech in order to route a speech sample to a specialized speech recognition model with higher accuracy, but using the same detection system to evaluate a job applicant could be discriminatory, unfair and potentially illegal. Any system that explicitly detects ability-related attributes of an individual may need to make these inferences visible to the user, optional, and able to be challenged when they make wrong inferences. It is crucial to involve members of the affected disability group from the outset. This can prevent wasted time and effort on paths that lead to inappropriate and unfair outcomes, inflexible systems that cannot be effectively deployed at scale, or that will be likely to face legal challenges.

Data Sourcing

When sourcing data to build a model, important considerations are:

1. Does the data include people with disabilities, especially those disabilities identified as being most impacted by this solution? For example, data about the employees of a company with poor diversity may not include anyone who is deaf, or blind. If important groups are missing, or if this information is not known, take steps to find or create such data to supplement the original data source.
2. Might the data embody bias against people with disabilities? Consider whether the

data might capture existing societal biases against people with disabilities. For example, a dataset of housing applications with decisions might reflect a historical reluctance to choose someone with a disability. When this situation is identified, raise the issue.

3. Is disability information explicitly represented in the data? If so, practitioners can use bias detection tests to check for bias, and follow up with mitigation techniques to adjust for bias before training a model (Bellamy et al., 2018).

Sometimes data are constructed by combining several data sources. Depending on the requirements of those sources, some groups of people may not have records in all sources. Be attentive to whether disability groups might fall into this category and be dropped from the combined data set. For example, when combining photograph and fingerprint biometrics, consider what should happen for individuals who do not have fingers, and how they will be represented and handled.

In Europe, GDPR regulations (European Union, 2016) give individuals the right to know what data about them is being kept, and how it is used, and to request that their data be deleted. As organizations move to limit the information they store and the ways it can be used, AI systems may often not have explicit information about disability that can be used to apply established fairness tests and corrections. By being attentive to the potential for bias in the data, documenting the diversity of the data set, and raising issues early, practitioners can avoid building solutions that will perpetuate inequalities, and identify system requirements for accommodating groups that are not represented in the data.

Data Pre-Processing

The process of cleaning and transforming data into a form suitable for machine learning has been estimated to take 80-90% of the effort of a typical data science project (Zhang, Zhang, & Yang, 2003), and the choices made at this stage can have implications for the inclusiveness of the solution.

- **Data cleaning** steps may remove outliers, presumed to be noise or measurement er-

ror, but actually representing non-typical individuals, reducing the diversity in the dataset.

- **Feature selection** may include or exclude features that convey disability status. Besides explicit disability information, other features could be impacted by disability status or the resulting societal disadvantage, providing a proxy for disability status. For example, a preference for large fonts could serve as a proxy for visual impairment, or use of video captions could be correlated with deafness. Household income, educational achievement, and many other variables can also be correlated with disability.
- **Feature engineering** involves deriving new features from the data, either through analysis or combination of existing features. For example, calculating a person's reading level or personality traits based on their writing, or calculating a ratio of days worked to days lived. In both of these examples, the derived feature will be impacted by certain disabilities.

Although accepted practice in many fields is to exclude sensitive features so as not to build a model that uses that feature, this is not necessarily the best approach for algorithmic solutions. The reality is that it can be extremely difficult to avoid including disability status in some way. When possible, including features that explicitly represent disability status allows for testing and mitigation of disability-related bias. Consulting outlier individuals and stakeholder groups identified in the problem scoping stage is valuable to provide a better understanding of the ways that disability can be reflected in the data, and the tradeoffs involved in using or excluding certain features and data values.

Preserving Privacy

People experiencing disabilities may have the most to gain from many smart systems, but are also particularly vulnerable to data abuse and misuse. The current privacy protections do not work for individuals who are outliers or different from the norm. The current response to this data abuse and misuse, by privacy efforts globally, is to de-identify the data. The notion is that if we remove our identity from the data, it can't be traced back to us and it can't

be used against us. The assumption is that we will thereby retain our privacy while contributing our data to making smarter decisions about the design.

While people experiencing disabilities are particularly vulnerable to data abuse and misuse, they are often also the easiest to re-identify. If you are the only person in a neighborhood using a wheelchair, it will be easy to re-identify you. If you are the only person that receives delivery of colostomy bags in your community, it will be very easy to re-identify your purchasing data.

If de-identification is not a reliable approach to maintaining the privacy of individuals that are far from average, but data exclusion means that highly impactful decisions will be made without regard to their needs, what are potential approaches to addressing this dilemma? The primary focus has been on an ill-defined notion of privacy. When we unpack what this means to most people, it is self-determination, ownership of our own narrative, the right to know how our data is being used, and ethical treatment of our story.

To begin to address this dilemma, an International Standards Organization personal data preference standard has been proposed as an instrument for regulators to restore self-determination regarding personal data. The proposal is developed as a response to the all-or-nothing terms of service agreements which ask you to give away your private data rights in exchange for the privilege of using a service. These terms of service agreements are usually couched in legal fine print that most people could not decode even if they had the time to read them. This means that it has become a convention to simply click “I agree” without attending to the terms and the rights we have relinquished. The proposed standard will be part of an existing standard called AccessForAll or ISO/IEC 24751 (ISO/IEC, 2008). The structure of the parent standard enables matching of consumer needs and preferences with resource or service functionality. It provides a common language for describing what you need or prefer in machine-readable terms and a means for service providers or producers to describe the functions their products and services offer. This allows platforms to match diverse unmet consumer needs with the clos-

est product or service offering. Layered on top of the standard are utilities that help consumers explore, discover and refine their understanding of their needs and preferences, for a given context and a given goal. The personal data preference part of this standard will let consumers declare what personal data they are willing to release to whom, for what purpose, what length of time and under what conditions. Services that wish to use the data would declare what data is essential for providing the service and what data is optional. This will enable a platform to support the negotiation of more reasonable terms of service. The data requirements declarations by the service provider would be transparent and auditable. The standard will be augmented with utilities that inform and guide consumers regarding the risks and implications of preference choices. Regulators in Canada and Europe plan to point to this standard when it is completed. This will hopefully wrest back some semblance of self-determination of the use of our data.

Another approach to self-determination and data that is being explored with the Platform Co-op Consortium ([Platform Cooperativism Consortium, 2019](#)) is the formation of data co-ops. In a data co-op, the data producers would both govern and share in the profit (knowledge and funds) arising from their own data. This approach is especially helpful in amassing data in previously ignored domains, such as rare illnesses, niche consumer needs or specialized hobbies. In smart cities, for example, there could be a multiplicity of data domains that could have associated data co-ops. Examples include wayfinding and traffic information, utility usage, waste management, recreation, consumer demands, to name just a few. This multiplicity of data co-ops would then collaborate to provide input into more general urban planning decisions.

Model Training and Testing

When developing a model, there are many bias testing methods available to researchers. However, when applying these techniques to fairness for people with disabilities, some limitations become evident. Firstly, group-based methods require large enough numbers of individuals in each group to allow for statistical comparison of outcomes, and secondly, they

often rely on binary in-group/out-group comparisons, which can be difficult to apply to disability. This section will expand on each of these points and suggest ways to address these limitations in model testing.

When examining fairness for people with disabilities, there may be few examples in the training data. As defined by the United Nations Convention on the Rights of People with Disabilities ([UN General Assembly, 2007](#)), disability “results from the interaction between persons with impairments and attitudinal and environmental barriers that hinders their full and effective participation in society.” As such, disability depends on context and comes in many forms, including physical barriers, sensory barriers, and communication barriers. One important consequence of experiencing a disability is that it can lead us to do things in a unique way, or to look or act differently. As a result, disabled people may be outliers in the data, or align with one of many very different sub-groups.

Today’s methods for bias testing tend to split individuals into members of a protected group, and ‘others’. However, disability status has many dimensions, varies in intensity and impact, and often changes over time. Furthermore, people are often reluctant to reveal a disability. Binarized approaches that combine many different people into a broad ‘disabled’ category will fail to detect patterns of bias that apply, differently and distinctively, against specific groups within the disability umbrella. For example, an inaccessible online testing Web site will not disadvantage wheelchair users, but those who rely on assistive technologies or keyboard-only control methods to access the Web may be unable to complete the tests. More sensitive analysis methods not based on binary classifications are needed to address these challenges. Other protected attributes like race and gender that have traditionally been examined with binary fairness metrics are also far more complex and nuanced in reality ([Keyes, 2018](#)) ([Hamidi, Scheuerman, & Branham, 2018b](#)), and new approaches suitable for examining disability-related bias would support a more sophisticated examination of these attributes too.

To test for fairness, an audit based on identified test cases and expected outcomes is

valuable. These can be developed with the outlier individuals and key stakeholder groups identified at the outset. For models that interpret humans (speech, language, gesture, facial expression) where algorithmic fairness is important, the goal is to develop a method that works well for as many groups as possible (e.g. speech recognition for deaf speakers), and to document the limitations of the model. For models that allocate people to groups (e.g. job candidate, loan applicant), allocative fairness is important. In selecting a measure for allocative fairness, we argue that an individual fairness approach rather than a group fairness approach is preferable. While group fairness seeks to equalize a measure across groups, individual fairness aims for ‘similar’ individuals to receive similar outcomes. For example, in deciding whether to grant loan applications, even if the presence of a disability is statistically correlated with unemployment over the whole dataset, it would still be unfair to treat an employed person with a disability the same as the unemployed group, simply because of their disability. Individual fairness aligns better with the societal notion of fairness, and legal mandates against discrimination.

Deployment in Real Applications

In this stage, the trained model is incorporated into an application, typically with an interface for people to use, or an API to connect to. Testing with diverse users, especially outliers, is essential. Understanding how different people will perceive and use an AI-based system is also important, for example, to see if certain people are more likely than others to ascribe trust to an AI-based system, or be more likely to feel insulted by an AI system’s terse replies or lack of context.

As a matter of course, quality assurance should include testing by people with disabilities, covering as broad a set of disability groups as possible. This would include both testing the user interface of the system itself to ensure it is accessible, and the system’s performance on diverse data inputs. The test process should deliberately include outlier individuals to test the limits of the system and the mechanisms for addressing system failure.

Because disability manifests in such diverse ways, there will be situations where the ap-

plication is presented with an individual quite unlike those in the training data. For example an automated telephone help system may have difficulty interpreting an individual with a speech impairment, especially if they are not speaking in their native language. Developers can avoid discrimination by providing an alternative to using AI, for example by supporting typed input in addition to speech. Users should have the ability to opt out of AI-based interpretation.

Disability may also impact the accuracy of inputs to an AI-based system. An example is a video-based personality analysis concluding that an autistic interviewee is untrustworthy because they did not make eye contact with the interviewer, and then feeding that into an applicant selection model. For people with disabilities, it is essential to have the opportunity to inspect and correct the data used to make decisions about them.

Equally important is the ability to query and challenge AI decisions, receiving some form of explanation of the factors that most impacted the decision. If these factors were affected by a person's disability, the decision may be discriminatory. Any AI-based system that makes decisions affecting people should include both an opportunity to dispute a decision, and provide a manual override for outlier individuals where the model is unreliable.

As many AI systems by their very nature learn and thus modify their behavior over time, ongoing mechanisms to monitor for fairness to people with disabilities should be incorporated. These can include ongoing auditing and reviews of performance as well as periodic explicit testing to verify that changes in the system's operation aimed at improving performance do not introduce disparities in how decisions are made for specific sub-populations. This is crucial to ensure that as a system gets better for people overall it doesn't unfairly get worse for some.

Design Approaches

In several of the stages of AI development described above, and particularly in the problem scoping and testing/deployment phases, we have encouraged AI/ML engineers to seek the ongoing involvement of people with disabili-

ties. For AI/ML practitioners, this may seem like a daunting task, with questions ranging from how to find diverse users, how to ethically and respectfully engage them, and by what methods one can reliably incorporate their feedback to improve systems.

The field of Human-Computer Interaction has long contemplated these questions, and has developed a number of design philosophies, methodologies, and methods to guide the practice. Some of these pertain specifically to engaging people with disabilities, including Universal Design (Story, Mueller, & Mace, 1998), Ability-Based Design (Wobbrock, Kane, Gajos, Harada, & Froehlich, 2011), Design for User Empowerment (Ladner, 2015), and Design for Social Accessibility (Shinohara, Wobbrock, & Pratt, 2018). In this section, we endeavor to provide a brief overview of just three potential approaches that AI/ML developers might choose as they seek to integrate users into their process. Our purpose, then, is not to provide a comprehensive review of all or even a few methodologies, but rather to offer links into the literature for those who want to learn more about these approaches, or perhaps seek out a collaborator with such expertise.

Below, we overview three distinct approaches to human-centered design: Inclusive Design, Participatory Design, and Value-Sensitive Design. Each of these has developed from different intellectual traditions, and therefore varies in the degree to which they explicitly include people with disabilities in their theoretical frameworks. People with disabilities are often excluded from design processes, and designs rarely anticipate end-users' needs to appropriate and adapt designs (Derboven, Geerts, & De Grooff, 2016). We therefore will begin here by introducing some basic rationale for why it is important to include people with disabilities directly in software development efforts.

Firstly, the opportunity to fully participate in society is every person's right, and digital inclusion is fundamental to those opportunities today. All of us will very likely experience disability at some stage in our lives, and our technology must be robust enough to accommodate the diversity of the human experience, especially if it is used in critical decision-making

areas. This cannot happen unless this diversity is considered from the outset, hence the disability rights movement's mantra of "nothing about us, without us."

Including people with disabilities may bring compelling new design ideas and ultimately expand the potential user base of the product. People with disabilities (including seniors) have been described as the original life hackers and personal innovators ([Harley & Fitzpatrick, 2012](#)) ([Storni, 2010](#)), as they often have to find creative workarounds and innovate new technologies in a world that is not built to their specifications. Second, people with disabilities can be seen as presenting valuable diverse cases which should be brought from the "edge" and into the "center" of our design thinking. In her foundational paper on Feminism in Human Computer Interaction, Bardzell advocated to study not only the conceptual "center" of a distribution of users, but also the edge cases ([Bardzell, 2010](#)). She argued that design often functions with a default "user" in the designers' minds, and that default user is often male, white, educated, and non-disabled. In Bardzell's analysis, accommodating the edge cases is a way both to broaden the audience (or market) for a design, and to strengthen the design against unanticipated changes in users, usage, or contexts of use. These ideas are echoed by other researchers ([Krischkowsky et al., 2015](#)) ([Muller et al., 2016](#)) ([Tscheligi et al., 2014](#)). One canonical example of how designs made with and for people with disabilities can actually improve the user experience of everyone (an example of "universal design"), is the curb cut. Curb cuts – the ramps that allow people in wheelchairs to transition from public sidewalks to cross the street – also serve parents with prams, workers with heavy wheeled loads, and pedestrians on scooters. Both Downey and Jacobs ([Downey, 2008](#)) ([Jacobs, 1999](#)) have advocated for electronic curb cuts; one example of such a feature is the zooming capability of the browser, which supports easier reading for people with low vision or people who are far away from the screen.

In practice, the best way of centering marginalized perspectives will require that we include people with disabilities in our core design practices. Fortunately, there is a rich history of work on design that centers users at

the margins. In particular, we believe the theoretical approaches of Inclusive Design (specifically as it evolved in Canada), Participatory Design and Value Based Design are particularly valuable when designing for – and with – people with disabilities.

Inclusive Design

The practice and theoretical framing of inclusive design, that emerged and evolved in Canada and with global partners since the emergence of the Web, takes advantage of the affordances or characteristics of digital systems ([Pullin, Treviranus, Patel, & Higginbotham, 2017](#)) ([Treviranus, 2016](#)). In contrast to related universal design theories, that emerged from architectural and industrial design fields, the Canadian inclusive design practice aims to use the mutability and connectivity of networked digital systems to achieve one-size-fits-one designs within an integrated system, thereby increasing the adaptability and longevity of the system as a whole ([Lewis & Treviranus, 2013](#)). Rather than specifying design criteria, or accessibility checklists, the theory specifies a process or mindset, called the "three dimensions of inclusive design."

1. Recognize that everyone is unique, and strive for a design that is able to match this uniqueness in an integrated system. Support self-awareness of this uniqueness (use data to make people 'smarter' about themselves, not just machines smarter).
2. Create an inclusive co-design process. The most valuable co-designers are individuals that can't use or have difficulty using the current designs. Continuously ask whose perspective is missing from the decision making "table" and how can they help make the "table" more inclusive.
3. Recognize that all design operates within a complex adaptive system of systems. Be cognizant of the entangled impact and friction points. Strive for designs that are beneficial to this larger system of systems.

This practice of inclusive design critiques and counters reliance on probability and population based statistics, pointing out the risks of basing critical decisions solely on the majority or statistical average ([Treviranus, 2014](#)). With

respect to fairness in AI, people with disabilities are most disadvantaged by population-based AI decisions. The only common defining characteristic of disability is difference from the norm. If you are not like the norm, probability predictions based on population data are wrong. Even if there is full representation and all human bias is removed, statistically-based decisions and predictions will be biased against small minorities and outliers. Current automated decisions focus on the good of the majority, causing greater disparity between the majority and smaller minorities. Inclusive design practitioners are investigating learning models that do not give advantage to being like the majority, causing learning models to attend to the full diversity of requirements (Treviranus, 2019). This is hypothesized to also support better context shifting, adaptability and detection of weak signals.

Given that data is from the past, optimization using past data will not achieve the culture shift inclusive design practitioners hope to achieve. Hence inclusive design practitioners are combining existing data, data from alternative scenarios, and modelled or simulated data to assist in decision making. Storytelling, bottom-up personalized data, or small (n=1), thick (in context) data is also employed to overcome the bias toward numerical data (Clark, 2018) (Pullin et al., 2017).

Participatory Design

Participatory Design (PD) emphasizes the active role of people who will be affected by a technology, as co-designers of that technology. There are many methods and rationales for these approaches, which can be found in (Muller & Druin, 2012), among other references. Some PD methods were originally proposed as “equal opportunity” practices, e.g., (Kuhn & Muller, 1993), because the methods involved low-technology or “lo-fi” prototyping practices that did not require extensive computer knowledge to contribute to the design. However, the needs of people with disabilities were generally not considered during the early days of PD.

This omission has now been partially rectified. Börjesson and colleagues (Börjesson, Barendregt, Eriksson, & Torgersson, 2015) recently

published an overview of theory and methods for work with developmentally diverse children. Katan et al. (2015) used interactive machine learning in participatory workshops with people with disabilities (Katan, Grierson, & Fiebrink, 2015).

Some of these methods amount to merely asking people about their needs, e.g., (Holbø, Bøthun, & Dahl, 2013) (Krishnaswamy, 2017). However, other approaches involve bringing people with disabilities (and sometimes their informal carers) into the design process as active co-designers, e.g., (Gomez Torres, Parmar, Aggarwal, Mansur, & Guthrie, 2019) (Hamidi et al., 2016) (Lee & Riek, 2018) (McGrenere et al., 2003) (Sitbon & Farhin, 2017) (Wilde & Marti, 2018) (Williams et al., 2018). In general, the methods that include direct participation by people with disabilities in design activities are more powerful, and tend to include deeper understandings than are possible through the less engaged survey methods.

We note that this is an active research area, with discussion of needs that are not yet met, e.g., (Holone & Herstad, 2013) (Oswal, 2014), and many opportunities to improve and innovate the participatory methods. For people who are new to PD, we suggest beginning with an orientation to the diversity of well-tested methods, e.g., (Muller & Druin, 2012), followed by a “deeper dive” into methods that have been used with particular populations and/or with particular challenges.

Value-Sensitive Design

Participatory design originated primarily from the workplace democracy movement in Scandinavia, e.g., (Bjerknes, Ehn, & Kyng, 1987), and then developed in many directions. One of the core assumptions of the workplace applications was a division of labor among workers and managers. In that context, PD methods were seen as ways to reduce power differences during the two-party design process, by facilitating the voice of the workers in relation to management. These assumptions have tended to carry through into work with people with disabilities, in which the two parties are reconceived as people with disabilities and designers, or people with disabilities and providers, with designers as mediators.

Value Sensitive Design (VSD) offers a broader perspective regarding the stakeholders in design (Friedman, Hendry, & Borning, 2017). For this paper, VSD crucially expands concepts of stakeholders into direct stakeholders (people who have contact with a design or a technology, including users, designers, providers) and indirect stakeholders (people who are affected by the design or technology, even if they do not have direct contact with it). For people with disabilities, there are often multiple stakeholders in complex relationships, e.g., (Zolyomi, Ross, Bhattacharya, Milne, & Munson, 2018).

While there are disagreements about the details of a value-centric approach, e.g., (Borning & Muller, 2012) (Le Dantec, Poole, & Wyche, 2009) (Muller & Liao, 2017), there is consensus that values matter, and that values can be formative for designs. There may be values differences among people with disabilities, carers, and medical professionals, e.g., (Draper et al., 2014) (Felzmann, Beyan, Ryan, & Beyan, 2016), and therefore an explicit and focused values inquiry may be helpful in “satisficing” these complex assumptions, views, and needs (Cheon & Su, 2016). While VSD has tended to have fewer well-defined methods, Friedman et al. (2017) recently published a survey of values-centric methods (Friedman et al., 2017).

Conclusion

We have outlined a number of situations in which AI solutions could be disadvantageous for people with disabilities if researchers and practitioners fail to take necessary steps. In many existing situations, non-AI solutions are already discriminatory, and introducing AI runs the risk of simply perpetuating and replicating these flaws. For example, people with disabilities may already face discrimination in hiring opportunities. With AI-driven hiring systems, models that recognize good candidates by matching to the existing workforce will perpetuate that status quo. In education, an AI system that draws inferences based on a student’s online interactions might misinterpret speed for competency, if the student is using assistive technologies. In public safety, AI systems might misinterpret a person with a cog-

nitive disability as a potential threat. In AI systems for healthcare, where speech characteristics can be used to diagnose cognitive impairments, a person with a speech impediment can be misdiagnosed.

To avoid such erroneous conclusions and potentially damaging outcomes, a number of steps are proposed. AI systems should be prioritized for fairness review and ongoing monitoring, based on their potential impact on the user in their broader context of use. They should offer opportunities to redress errors, and for users and those impacted to raise fairness concerns. People with disabilities should be included when sourcing data to build models. Such “outlier” data - the edge cases - will create a more inclusive and robust system. From the perspective of people with disabilities, there can be privacy concerns with self-identification, but there can be risk of exclusion from the data models if users opt not to participate or disclose. There are methods provided to increase participation while protecting user privacy, such as the personal data preferences standard. In deploying the AI application, it is critical to test the UI and system preferences with outlier individuals. Users should be able to pursue workarounds, and ultimately override the system where models may be unreliable.

AI has been shown to help improve the lives of people with disabilities in a number of different environments whether it be navigating a city, re-ordering prescriptions at a local pharmacy through a telephone or text service, or facilitating public safety. Almost everyone in the greater community is directly connected with someone with a disability whether it be a family member, a colleague, a friend, or a neighbor. While AI technology has significantly improved the lives of those in the disabled community, there are always ways in which we can continue to advocate for fairness and equality and challenge the status quo.

When creating AI that is designed to help the community, we must take into consideration a disabled person user approach. The AI designed should consider disabled people as a focal point. To borrow upon the concept driven by Eric Ries in *The Lean Startup* (Ries, 2011), we need to deploy minimum viable products (MVPs) that are not perfected but rather im-

proved upon by the user involved.

In a sense, what is needed is an incremental and algorithmic approach that continues to challenge the status quo and strives to improve the standardization of fairness and equality. This should be from a multi-industrial approach with key players in different industries.

This paper suggests challenging every day practices that may prove to inhibit people with disabilities, and is a starting point to bring awareness for the need for equality. It is important to remember that promoting this goal is a process. Success will require a series of incremental steps of further learning, thought provoking peer discussion, and changes at the local and municipal level. Only when these incremental changes are met will we drive sustainable outcomes for people with disabilities using AI systems.

Acknowledgments

The authors thank all participants of the IBM-MIT AI Week 2018 Workshop on AI Fairness for People with Disabilities for their contributions to the discussions leading to this article. Special thanks to Peter Fay, Phill Jenkins, Audrey Hasegawa and Morgan Foreman.

References

- Ameri, M., Schur, L., Adya, M., Bentley, F. S., McKay, P., & Kruse, D. (2018, mar). The Disability Employment Puzzle: A Field Experiment on Employer Hiring Behavior. *ILR Review*, 71(2), 329–364. Retrieved from <http://journals.sagepub.com/doi/10.1177/0019793917717474> doi: 10.1177/0019793917717474
- Bardzell, S. (2010). Feminist HCI: Taking Stock and Outlining an Agenda for Design. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1301–1310). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/1753326.1753521> doi: 10.1145/1753326.1753521
- Barnett, S., McKee, M., Smith, S. R., & Pearson, T. A. (2011, mar). Deaf sign language users, health inequities,

and public health: opportunity for social justice. *Preventing chronic disease*, 8(2), A45. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/21324259><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3073438>

- Bellamy, R. K. E., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., ... Zhang, Y. (2018, oct). AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias. Retrieved from <http://arxiv.org/abs/1810.01943>
- Bhutani, A., & Wadhvani, P. (2019). *Artificial Intelligence (AI) in Education Market Size worth \$6bn by 2024*. Global Market Insights. Retrieved from <https://www.gminsights.com/pressrelease/artificial-intelligence-ai-in-education-market>
- Bioss. (2019). *The Bioss AI Protocol*. Retrieved 2019-08-30, from <http://www.bioss.com/ai/>
- Bird, S., Hutchinson, B., Kenthapadi, K., Kiciman, E., & Mitchell, M. (2019). Fairness-Aware Machine Learning: Practical Challenges and Lessons Learned. In *Companion proceedings of the 2019 world wide web conference on - www '19* (pp. 1297–1298). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3308560.3320086> doi: 10.1145/3308560.3320086
- Bjerknes, G., Ehn, P., & Kyng, M. (1987). *Computers and Democracy - A Scandinavian Challenge*. Avebury.
- Börjesson, P., Barendregt, W., Eriksson, E., & Torgersson, O. (2015). Designing Technology for and with Developmentally Diverse Children: A Systematic Literature Review. In *Proceedings of the 14th international conference on interaction design and children* (pp. 79–88). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2771839.2771848> doi: 10.1145/2771839.2771848
- Borning, A., & Muller, M. (2012). Next Steps for Value Sensitive Design. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1125–

- 1134). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2207676.2208560> doi: 10.1145/2207676.2208560
- Branham, S. M., Abdolrahmani, A., Easley, W., Scheuerman, M., Ronquillo, E., & Hurst, A. (2017). "Is Someone There? Do They Have a Gun". In *Proceedings of the 19th international acm sigaccess conference on computers and accessibility - assets '17* (pp. 260–269). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3132525.3132534> doi: 10.1145/3132525.3132534
- Brewer, R. N., & Kameswaran, V. (2018). Understanding the Power of Control in Autonomous Vehicles for People with Vision Impairment. In *Proceedings of the 20th international acm sigaccess conference on computers and accessibility - assets '18* (pp. 185–197). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3234695.3236347> doi: 10.1145/3234695.3236347
- Bruckner, M. A. (2018). The Promise and Perils of Algorithmic Lenders' Use of Big Data. *Chicago-Kent Law Review*, 93. Retrieved from [https://heinonline.org/HOL/Page?handle=hein.journals/chknt93{id=15{&div={&collection=](https://heinonline.org/HOL/Page?handle=hein.journals/chknt93{id=15{&div={&collection=mlr115{id=1081{id=}}div={&collection=}})
- Bureau of Labor Statistics, U. D. o. L. (2019). *PERSONS WITH A DISABILITY: LABOR FORCE CHARACTERISTICS — 2018*. Retrieved from <https://www.bls.gov/news.release/pdf/disabl.pdf>
- Burgstahler, S. (2015, dec). Opening Doors or Slamming Them Shut? Online Learning Practices and Students with Disabilities. *Social Inclusion*, 3(6), 69. Retrieved from <http://www.cogitatiopress.com/ojs/index.php/socialinclusion/article/view/420> doi: 10.17645/si.v3i6.420
- Chander, A. (2017). The Racist Algorithm? *Michigan Law Review*, 115(6), 1023. Retrieved from [https://heinonline.org/HOL/Page?handle=hein.journals/](https://heinonline.org/HOL/Page?handle=hein.journals/mlr115{id=1081{id=}}div={&collection=}})
- Cheon, E., & Su, N. M. (2016). Integrating roboticist values into a Value Sensitive Design framework for humanoid robots. *11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. doi: 10.1109/HRI.2016.7451775
- Cinquin, P.-A., Guitton, P., & Sauzéon, H. (2019, mar). Online e-learning and cognitive disabilities: A systematic review. *Computers & Education*, 130, 152–167. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0360131518303178> doi: 10.1016/J.COMPEDU.2018.12.004
- Clark, C. B. (2018, jun). *Many more stories: Co-design and creative communities*. OCAD University. Retrieved from <http://openresearch.ocadu.ca/id/eprint/2079/>
- Costello, K. (2019). *Gartner Survey Shows 37 Percent of Organizations Have Implemented AI in Some Form*. Stamford, Conn.: Gartner. Retrieved from <https://www.gartner.com/en/newsroom/press-releases/2019-01-21-gartner-survey-shows-37-percent-of-organizations-have>
- Cutler, A., Pribik, M., & Humphrey, L. (2019). *Everyday Ethics for Artificial Intelligence* (Tech. Rep.). IBM. Retrieved from <https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf>
- Dastin, J. (2018). *Amazon scraps secret AI recruiting tool that showed bias against women - Reuters*. Retrieved 2019-08-30, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- Derboven, J., Geerts, D., & De Grooff, D. (2016). The Tactics of Everyday Practice: A Semiotic Approach to Appropriation. *Interaction Design and Architecture(s)*, 29(29), 99–120. Retrieved from <https://lirias.kuleuven.be/389602?limo=0>
- Downey, G. (2008). *Closed captioning: Sub-*

- titling, stenography and the digital convergence of text with television.* Johns Hopkins University Press.
- Draper, H., Sorell, T., Bedaf, S., Syrdal, D. S., Gutierrez-Ruiz, C., Duclos, A., & Amirabdollahian, F. (2014, oct). Ethical Dimensions of Human-Robot Interactions in the Care of Older People: Insights from 21 Focus Groups Convened in the UK, France and the Netherlands. In *International conference on social robotics* (LNCS 8755 ed., pp. 135–145). Springer. Retrieved from http://link.springer.com/10.1007/978-3-319-11973-1_{_}14 doi: 10.1007/978-3-319-11973-1_14
- Dudley-Marling, C., & Burns, M. B. (2014). Two Perspectives on Inclusion in the United States, *Global Education Review*, 2014. *Global Education Review*, 1(1), 14–31. Retrieved from <https://eric.ed.gov/?id=EJ1055208>
- European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Da.*
- Faggella, D. (2019, apr). *Machine Learning in Human Resources – Applications and Trends.*
- Faucett, H. A., Ringland, K. E., Cullen, A. L. L., & Hayes, G. R. (2017, oct). (In)Visibility in Disability and Assistive Technology. *ACM Trans. Access. Comput.*, 10(4), 14:1—14:17. Retrieved from <http://doi.acm.org/10.1145/3132040> doi: 10.1145/3132040
- Felzmann, H., Beyan, T., Ryan, M., & Beyan, O. (2016, jan). Implementing an Ethical Approach to Big Data Analytics in Assistive Robotics for Elderly with Dementia. *SIGCAS Comput. Soc.*, 45(3), 280–286. Retrieved from <http://doi.acm.org/10.1145/2874239.2874279> doi: 10.1145/2874239.2874279
- Friedman, B., Hendry, D. G., & Borning, A. (2017). A Survey of Value Sensitive Design Methods. *Foundations and Trends^R in Human–Computer Interaction*, 11(2), 63–125. Retrieved from <http://dx.doi.org/10.1561/1100000015> doi: 10.1561/1100000015
- Fruchterman, J., & Mellea, J. (2018). *Expanding Employment Success for People with Disabilities — Benetech — Software for Social Good* (Tech. Rep.). benetech. Retrieved from <https://benetech.org/about/resources/expanding-employment-success-for-people-with-disabilities-2/>
- Gajos, K. Z., Wobbrock, J. O., & Weld, D. S. (2007). Automatically Generating User Interfaces Adapted to Users' Motor and Vision Capabilities. In *Proceedings of the 20th annual acm symposium on user interface software and technology* (pp. 231–240). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/1294211.1294253> doi: 10.1145/1294211.1294253
- Givens, A. R. (2019). *Institute Announces New Project on Algorithmic Fairness for People with Disabilities.* Institute for Technology Law and Policy, Georgetown Law.
- Glickman, N. (2007). *Do You Hear Voices? Problems in Assessment of Mental Status in Deaf Persons With Severe Language Deprivation* (Vol. 12). Oxford University Press. Retrieved from <https://www.jstor.org/stable/42658866> doi: 10.2307/42658866
- Gomez Torres, I., Parmar, G., Aggarwal, S., Mansur, N., & Guthrie, A. (2019). Affordable Smart Wheelchair. In *Extended abstracts of the 2019 chi conference on human factors in computing systems* (pp. SRC07:1—SRC07:6). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3290607.3308463> doi: 10.1145/3290607.3308463
- Guo, A., Kamar, E., Vaughan, J. W., Wallach, H., & Morris, M. R. (2019, jul). *Toward Fairness in AI for People with Disabilities: A Research Roadmap.* Retrieved from <http://arxiv.org/abs/1907.02227>
- Haigh, K. Z., & Yanco, H. A. (2002). *Automation as Caregiver: A Survey of Issues*

- and Technologies* (Tech. Rep.). AAAI.
- Hamidi, F., Müller, C., Baljko, M., Schorch, M., Lewkowicz, M., & Stangl, A. (2016). Engaging with Users and Stakeholders: The Emotional and the Personal. In *Proceedings of the 19th international conference on supporting group work* (pp. 453–456). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2957276.2996292> doi: 10.1145/2957276.2996292
- Hamidi, F., Scheuerman, M. K., & Branham, S. M. (2018a). Gender Recognition or Gender Reductionism? In *Proceedings of the 2018 chi conference on human factors in computing systems - chi '18* (pp. 1–13). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3173574.3173582> doi: 10.1145/3173574.3173582
- Hamidi, F., Scheuerman, M. K., & Branham, S. M. (2018b). Gender Recognition or Gender Reductionism?: The Social Implications of Embedded Gender Recognition Systems. In *Proceedings of the 2018 chi conference on human factors in computing systems* (pp. 8:1–8:13). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3173574.3173582> doi: 10.1145/3173574.3173582
- Hankerson, D., Marshall, A. R., Booker, J., El Mimouni, H., Walker, I., & Rode, J. A. (2016). Does Technology Have Race? In *Proceedings of the 2016 chi conference extended abstracts on human factors in computing systems - chi ea '16* (pp. 473–486). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=2851581.2892578> doi: 10.1145/2851581.2892578
- Harley, D., & Fitzpatrick, G. (2012). Appropriation of social networking by older people: two case studies. In *Ecsw 2011 workshop on fostering social interactions in the ageing society: Artefacts - methodologies - research paradigms*. Aarhus, Denmark.
- Harrell, E. (2017). *Crime Against Persons with Disabilities, 2009-2015 - Statistical Tables* (Tech. Rep.). Bureau of Justice Statistics. Retrieved from <http://www.bjs.gov/index.cfm?ty=pbdetail&iid=5986>
- Holbø, K., Bøthun, S., & Dahl, Y. (2013). Safe Walking Technology for People with Dementia: What Do They Want? In *Proceedings of the 15th international acm sigaccess conference on computers and accessibility* (pp. 21:1–21:8). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2513383.2513434> doi: 10.1145/2513383.2513434
- Holone, H., & Herstad, J. (2013). Three Tensions in Participatory Design for Inclusion. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 2903–2906). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2470654.2481401> doi: 10.1145/2470654.2481401
- Hurley, M., & Adebayo, J. (2016). Credit Scoring in the Era of Big Data. *Yale Journal of Law and Technology*, 18. Retrieved from <https://heinonline.org/HOL/Page?handle=hein.journals/yjolt18&id=148&div=&collection=>
- IEEE, G. I. o. E. o. A., & Systems, I. (2019). *Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems* (First edit ed.). IEEE. Retrieved from <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/%0Aautonomous-systems.html>
- lezzoni, L. I. (2011, oct). Eliminating Health And Health Care Disparities Among The Growing Population Of People With Disabilities. *Health Affairs*, 30(10), 1947–1954. Retrieved from <http://www.healthaffairs.org/doi/10.1377/hlthaff.2011.0613> doi: 10.1377/hlthaff.2011.0613
- ISO/IEC. (2008). *Information technology — Individualized adaptability and accessibility in e-learning, education and training — Part 1: Framework and reference model*. Author.
- Jacobs, S. (1999). *Section 255 of the Telecommunications Act of 1996: Fueling the Creation of New Elec-*

- Kuhn, S., & Muller, M. (1993). Participatory Design. *Communications of the ACM*, 36(6).
- K.W. v. Armstrong, No. 14-35296 (9th Cir. 2015) :: Justia.* (2015). Retrieved from <https://law.justia.com/cases/federal/appellate-courts/ca9/14-35296/14-35296-2015-06-05.html>
- Ladner, R. (2015). Design for user empowerment. *Interactions*, XXII(2). Retrieved from <http://interactions.acm.org/archive/view/march-april-2015/design-for-user-empowerment>
- Le Dantec, C. A., Poole, E. S., & Wyche, S. P. (2009). Values As Lived Experience: Evolving Value Sensitive Design in Support of Value Discovery. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1141–1150). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/1518701.1518875> doi: 10.1145/1518701.1518875
- Lee, H. R., & Riek, L. D. (2018, may). Reframing Assistive Robots to Promote Successful Aging. *ACM Trans. Hum.-Robot Interact.*, 7(1), 11:1—11:23. Retrieved from <http://doi.acm.org/10.1145/3203303> doi: 10.1145/3203303
- Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P. (2018, dec). Fair, Transparent, and Accountable Algorithmic Decision-making Processes. *Philosophy & Technology*, 31(4), 611–627. Retrieved from <http://link.springer.com/10.1007/s13347-017-0279-x> doi: 10.1007/s13347-017-0279-x
- Lewis, C. (2019, jun). Implications of Developments in Machine Learning for People with Cognitive Disabilities. *SIGACCESS Newsletter, Issue 124*. Retrieved from <http://www.sigaccess.org/newsletter/2019-06/lewis.html>
- Lewis, C., & Treviranus, J. (2013, sep). Public Policy and the Global Public Inclusive Infrastructure Project. *interactions*, 20(5), 62–66. Retrieved from <http://doi.acm.org/10.1145/2510123> doi: 10.1145/2510123
- Lohia, P. K., Natesan Ramamurthy, K., Bhide, M., Saha, D., Varshney, K. R., & Puri, R. (2019, may). Bias Mitigation Post-processing for Individual and Group Fairness. In *Icassp 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 2847–2851). IEEE. Retrieved from <https://ieeexplore.ieee.org/document/8682620> doi: 10.1109/ICASSP.2019.8682620
- McCarthy, O. J. (2019). *AI & Global Governance: Turning the Tide on Crime with Predictive Policing - United Nations University Centre for Policy Research*. United Nations University Centre for Policy Research. Retrieved from <https://cpr.unu.edu/ai-global-governance-turning-the-tide-on-crime-with-predictive-policing.html>
- McGrenere, J., Davies, R., Findlater, L., Graf, P., Klawe, M., Moffatt, K., ... Yang, S. (2003). Insights from the Aphasia Project: Designing Technology for and with People Who Have Aphasia. In *Proceedings of the 2003 conference on universal usability* (pp. 112–118). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/957205.957225> doi: 10.1145/957205.957225
- Morris, R. R., Kirschbaum, C. R., & Picard, R. W. (2010). Broadening Accessibility Through Special Interests: A New Approach for Software Customization. In *Proceedings of the 12th international ACM SIGACCESS conference on computers and accessibility* (pp. 171–178). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/1878803.1878834> doi: 10.1145/1878803.1878834
- Muller, M., & Druin, A. (2012). Participatory design: The third space of HCI. In J. Jacko (Ed.), *Human computer interaction handbook* (3rd ed., pp. 1125–1154). CRC Press.
- Muller, M., & Liao, V. (2017). *Using participatory design fictions to explore ethics and values for robots and agents*.
- Muller, M., Neureiter, K., Krischkowsky, A., Tscheligi, M., Al Zubaidi-Polli, A. M., & Tscheligi, M. (2016). Collaborative Appropriation: How Couples, Teams,

- Groups and Communities Adapt and Adopt Technologies. In *Proceedings of the 19th acm conference on computer supported cooperative work and social computing companion - cscw '16 companion* (pp. 473–480). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=2818052.2855508> doi: 10.1145/2818052.2855508
- Obiakor, F. E., Harris, M., Mutua, K., Rotatori, A., & Algozzine, B. (2012). Making Inclusion Work in General Education Classrooms. *Education and Treatment of Children*, 35(3), 477–490. doi: 10.1353/etc.2012.0020
- Oswal, S. K. (2014, may). Participatory Design: Barriers and Possibilities. *Commun. Des. Q. Rev.*, 2(3), 14–19. Retrieved from <http://doi.acm.org/10.1145/2644448.2644452> doi: 10.1145/2644448.2644452
- Platform Cooperativism Consortium. (2019). *What Is a Platform Co-op?* Retrieved 2019-08-30, from <https://platform.coop/>
- Policing, P. T. F. o. s. C. (2015). *Final Report of the President's Task Force on 21st Century Policing*. Washington, DC: Office of Community Oriented Policing Services. Retrieved from <https://ric-zai-inc.com/Publications/cops-p311-pub.pdf>
- Pollard, R. (1994). Public mental health service and diagnostic trends regarding individuals who are deaf or hard of hearing. *Rehabilitation Psychology*, 39(3), 147–160.
- Popovich, P. M., Scherbaum, C. A., Scherbaum, K. L., & Polinko, N. (2003, mar). The Assessment of Attitudes Toward Individuals With Disabilities in the Workplace. *The Journal of Psychology*, 137(2), 163–177. Retrieved from <http://www.tandfonline.com/doi/abs/10.1080/00223980309600606> doi: 10.1080/00223980309600606
- Pradhan, A., Mehta, K., & Findlater, L. (2018). "Accessibility Came by Accident". In *Proceedings of the 2018 chi conference on human factors in computing systems - chi '18* (pp. 1–13). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3173574.3174033> doi: 10.1145/3173574.3174033
- Pullin, G., Treviranus, J., Patel, R., & Higginbotham, J. (2017). Designing interaction, voice, and inclusion in AAC research. *Augmentative and Alternative Communication*, 33(3), 139–148. Retrieved from <https://doi.org/10.1080/07434618.2017.1342690> doi: 10.1080/07434618.2017.1342690
- Ries, E. (2011). *The lean startup : how today's entrepreneurs use continuous innovation to create radically successful businesses*. Retrieved from https://books.google.com/books?hl=en&lr=&id=r9x-OXdzpPcC&oi=fnd&pg=PA15&dq=+The+Lean+Startup&ots=0s-bDboHfV&sig=jVn-Rw46-A2D4H_{_}9HnK7gTi5MCg{#}v=onepage&q=TheLeanStartup&f=false
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the conference on fairness, accountability, and transparency - fat* '19* (pp. 59–68). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3287560.3287598> doi: 10.1145/3287560.3287598
- Shaffer, I. R., & Rogan, I. (2018). Exploring the Performance of Facial Expression Recognition Technologies on Deaf Adults and Their Children. In *Proceedings of the 20th international acm sigaccess conference on computers and accessibility - assets '18* (pp. 474–476). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3234695.3240986> doi: 10.1145/3234695.3240986
- Shaheen, N. L., & Lazar, J. (2018). K–12 Technology Accessibility: The Message from State Government. *Journal of Special Education Technology*, 33(2), 83–97.
- Shinohara, K., Wobbrock, J. O., & Pratt, W. (2018). Incorporating Social Factors in

- Accessible Design. In *Proceedings of the 20th international acm sigaccess conference on computers and accessibility - assets '18* (pp. 149–160). New York, New York, USA: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?doid=3234695.3236346> doi: 10.1145/3234695.3236346
- Sitbon, L., & Farhin, S. (2017). Co-designing Interactive Applications with Adults with Intellectual Disability: A Case Study. In *Proceedings of the 29th australian conference on computer-human interaction* (pp. 487–491). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3152771.3156163> doi: 10.1145/3152771.3156163
- Storni, C. (2010). Multiple Forms of Appropriation in Self-Monitoring Technology: Reflections on the Role of Evaluation in Future Self-Care. *International Journal of Human-Computer Interaction*, 26(5), 537–561. Retrieved from <https://doi.org/10.1080/10447311003720001> doi: 10.1080/10447311003720001
- Story, M. F., Mueller, J. L., & Mace, R. L. (1998). *The Universal Design File: Designing for People of All Ages and Abilities. Revised Edition*. Center for Universal Design, NC State University, Box 8613, Raleigh, NC 27695-8613 (\$24). Tel: 800-647-6777 (Toll Free) Web site: <http://www.design.ncsu.edu>. Retrieved from <https://eric.ed.gov/?id=ED460554>
- Straumsheim, C. (2017, mar). Berkeley Will Delete Online Content. *Inside Higher Ed*. Retrieved from <https://www.insidehighered.com/news/2017/03/06/u-california-berkeley-delete-publicly-available-educational-content>
- Treviranus, J. (2014). *The Value of the Statistically Insignificant*. Educause Review.
- Treviranus, J. (2016). Life-long Learning on the Inclusive Web. In *Proceedings of the 13th web for all conference* (pp. 1:1—1:8). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2899475.2899476> doi: 10.1145/2899475.2899476
- Treviranus, J. (2019). The Value of Being Different. In *Proceedings of the 16th web for all 2019 personalization - personalizing the web* (pp. 1:1—1:7). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3315002.3332429> doi: 10.1145/3315002.3332429
- Trewin, S. (2018a, nov). AI Fairness for People with Disabilities: Point of View. Retrieved from <http://arxiv.org/abs/1811.10670>
- Trewin, S. (2018b, nov). Will AI Methods Treat People with Disabilities Fairly? *Medium, MIT-IBM Watson AI Lab*. Retrieved from medium.com/@MITIBMLab/will-ai-methods-treat-people-with-disabilities-fairly-7626b38f9cb5
- Trewin, S., Morris, M. R., Azenkot, S., Branham, S., Bleuel, N., Jenkins, P., ... Lasecki, W. (2019). *AI Fairness for People with Disabilities*. Retrieved from https://assets19.sigaccess.org/ai_fairness_workshop_program.html
- Tscheligi, M., Krischkowsky, A., Neureiter, K., Inkpen, K., Muller, M., & Stevens, G. (2014). Potentials of the "Unexpected": Technology Appropriation Practices and Communication Needs. In *Proceedings of the 18th international conference on supporting group work* (pp. 313–316). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2660398.2660427> doi: 10.1145/2660398.2660427
- UN General Assembly. (2007). *Convention on the Rights of Persons with Disabilities : resolution / adopted by the General Assembly, 24 January 2007, A/RES/61/106*. Retrieved from <https://www.un.org/development/desa/disabilities/convention-on-the-rights-of-persons-with-disabilities/convention-on-the-rights-of-persons-with-disabilities-2.html>
- US Department of Justice Civil Rights Division. (2006). *COMMONLY ASKED QUESTIONS ABOUT THE AMERICANS WITH DISABILITIES ACT AND LAW ENFORCEMENT*. Retrieved 2019-

- 08-30, from <https://www.ada.gov/q{&}a{ }law.htm>
- von Schrader, S., Malzer, V., & Bruyère, S. (2014, dec). Perspectives on Disability Disclosure: The Importance of Employer Practices and Workplace Climate. *Employee Responsibilities and Rights Journal*, 26(4), 237–255. Retrieved from <https://doi.org/10.1007/s10672-013-9227-9> doi: 10.1007/s10672-013-9227-9
- Wastfelt, M., Fadeel, B., & Henter, J.-I. (2006, jul). A journey of hope: lessons learned from studies on rare diseases and orphan drugs. *Journal of Internal Medicine*, 260(1), 1–10. Retrieved from <http://doi.wiley.com/10.1111/j.1365-2796.2006.01666.x> doi: 10.1111/j.1365-2796.2006.01666.x
- Wilde, D., & Marti, P. (2018). Exploring Aesthetic Enhancement of Wearable Technologies for Deaf Women. In *Proceedings of the 2018 designing interactive systems conference* (pp. 201–213). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3196709.3196777> doi: 10.1145/3196709.3196777
- Williams, B. A., Brooks, C. F., & Shmargad, Y. (2018). How Algorithms Discriminate Based on Data they Lack: Challenges, Solutions, and Policy Implications. *Journal of Information Policy*, 8, 78. Retrieved from <https://www.jstor.org/stable/10.5325/jinfopoli.8.2018.0078> doi: 10.5325/jinfopoli.8.2018.0078
- Wobbrock, J. O., Kane, S. K., Gajos, K. Z., Harada, S., & Froehlich, J. (2011, apr). Ability-Based Design. *ACM Transactions on Accessible Computing*, 3(3), 1–27. Retrieved from <http://portal.acm.org/citation.cfm?doid=1952383.1952384> doi: 10.1145/1952383.1952384
- Zhang, S., Zhang, C., & Yang, Q. (2003, may). Data preparation for data mining. *Applied Artificial Intelligence*, 17(5-6), 375–381. Retrieved from <http://www.tandfonline.com/doi/abs/10.1080/713827180> doi: 10.1080/713827180
- Zolyomi, A., Ross, A. S., Bhattacharya, A., Milne, L., & Munson, S. A. (2018). Values, Identity, and Social Translucence: Neurodiverse Student Teams in Higher Education. In *Proceedings of the 2018 chi conference on human factors in computing systems* (pp. 499:1—499:13). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/3173574.3174073> doi: 10.1145/3173574.3174073



Shari Trewin is a technical leader in IBM's Accessibility Leadership Team, an ACM Distinguished Scientist, and Chair of ACM SIGACCESS. She has a Ph.D. in Computer Science and Artificial Intelligence from

the University of Edinburgh, with 17 patents and over 60 publications in technologies to remove barriers for people experiencing disabilities.



Sara Basson works at Google as "Accessibility Evangelist," with the goal of making the experience of Googlers more accessible and usable. She previously worked at IBM Research on Speech Technology, Accessibility, and Education Transformation. Sara holds a Ph.D. in Speech, Hearing, and Language Sciences

from The Graduate Center of CUNY.



Michael Muller is an ACM Distinguished Scientist, and a member of the SIGCHI Academy with expertise in participatory design and participatory analysis. In the AI Interactions group of IBM Research AI, he focuses on

the human aspects of data science; ethics and values in applications of AI to human issues.



domains of personal mobility and parenting with disability.

Stacy Branham is an Assistant Professor of Informatics at the University of California, Irvine. Her research explores the role of technology in social integration for people with disabilities, with active lines of inquiry in the



Dr. Treviranus has co-designed numerous policies related to digital inclusion.

Jutta Treviranus is the director and founder of the Inclusive Design Research Centre (IDRC), established in 1993. The IDRC leads many international research networks to proactively model more inclusive socio-technical practices, including in the area of data science (e.g., Project We Count).



in User Experience, Visualization, Social Software, and related areas.

Dan Gruen is a Cognitive Scientist in IBM Research's Center for Computational Health focusing on AI explainability, clinical reasoning, and complex medical cases. Dan has a PhD in Cognitive Science from UCSD and holds over 50 US and international patents



and is a community advocate for persons with disabilities and involved with the local independent living movement.

Daniel Hebert has his computer science degree from Rensselaer Polytechnic Institute and has been with IBM since 2006 as a QA, engineer, and developer. He enjoys working on smart home and accessibility projects,



box and remain ever flexible.

Natalia Lyckowski is an Application Developer with IBM Global Financing. She is also a passionate volunteer educator on inclusion and diversity awareness. Most proudly, she is the mother of a young autistic man who has taught her



workplace, Erich is an active mentor, disability rights advocate, and fundraiser. He is also a nationally recognized blind athlete in marathon and triathlon racing.

Erich Manser is an IBM Accessibility evangelist exploring new approaches to accessible technology, new applications of emerging technologies, and contributing to accessibility standards. Beyond the