

Disability, Bias, and AI

Meredith Whittaker, AI Now Institute at NYU

Meryl Alper, Northeastern University; **Cynthia L. Bennett**, University of Washington; **Sara Hendren**, Olin College; **Liz Kaziunas**, AI Now Institute at NYU; **Mara Mills**, New York University; **Meredith Ringel Morris**, Microsoft Research; **Joy Rankin**, AI Now Institute at NYU; **Emily Rogers**, New York University; **Marcel Salas**, New York University; **Sarah Myers West**, AI Now Institute at NYU

NOVEMBER 2019

Introduction	2
Key Questions	2
Terms and Concepts from Disability Studies	7
Bias at the Intersection of AI and Disability	7
AI's Version of "Normal"	12
Work, Disability, and AI	16
Representation, Privacy, and Disability	19
Designing with, Not For	22
Challenges and Interventions	26
Further Reading	27
Disability Studies: Core Texts and Concepts	27
Disability and Technology	28
AI Bias and the Politics of Large-Scale Technical Systems	31
Workshop Participants	32

获取更多信息无障碍相关文档：www.aar.org.cn/knowledge



Introduction

On March 28, 2019, the AI Now Institute at New York University (NYU), the NYU Center for Disability Studies, and Microsoft convened disability scholars, AI developers, and computer science and human-computer interaction researchers to discuss the intersection of disability, bias, and AI, and to identify areas where more research and intervention are needed.¹

This report captures and expands on some of the themes that emerged during discussion and debate. In particular, it identifies key questions that a focus on disability raises for the project of understanding the social implications of AI, and for ensuring that AI technologies don't reproduce and extend histories of marginalization.

Key Questions

Over the course of the workshop, the following questions emerged, pointing to productive avenues for research and advocacy around disability, bias, and AI.

- AI is already informing decisions about who is fit for work, who requires care, and who receives the benefits and protections of citizenship. In doing so, it is reshaping the domains of employment, caregiving, and much else. **How can we draw on disability activism and scholarship to ensure that we protect people who fall outside of the “norms” reflected and constructed by AI systems?**
- Disabled people are a heterogeneous population.² Even among those who identify as having the “same” disability, differences in race, class, gender, and other identities result in significantly different lived experiences. However, AI systems may misrecognize, or fail to recognize, these important distinctions. **How can work on “AI bias” attend to these interlocking marginalizations, recognizing that the intersections of race, gender, sexual orientation, and class often mediate how “disability” is defined and understood?**

¹ A sincere thank you to Jutta Treviranus and Kate Crawford for their generous edits and suggestions, and to Meredith Ringel Morris, Megan Lawrence, Kate Crawford, Gretchen Krueger, and Alejandro Calcaño Bertorelli for their work organizing and managing the Disability, Bias, and AI Workshop at AI Now.

² Throughout this report, we take an identity-first approach in talking about disability and people, using the phrase “disabled people” to refer to the constituency group of those who are, or who are labeled, disabled. We occasionally switch to the term “people with disabilities” in service of readability and style. In making this choice, we are following disability scholar and activist Simi Linton, who situates the term historically: “Beginning in the early 90s *disabled people* has been increasingly used in disability studies and disability rights circles when referring to the constituency group. Rather than maintaining disability as a secondary characteristic, *disabled* has become a marker of the identity that the individual and the group wish to highlight and call attention to.” Simi Linton, *Claiming Disability: Knowledge and Identity* (New York and London: New York University Press, 1998), 13. Similarly, the #SayTheWord movement on Twitter, which was led by disability rights activists and people in the disabled community, called on those discussing disability to avoid euphemisms like “differently abled” and use “disability” and “disabled” when referring to disability.

- The concept of “normal,” as well as the tools and techniques for enforcing normalcy, have historically constructed the disabled body and mind as deviant and problematic. How can we better assess the normative models encoded in AI systems, and what the consequences of these norms may be? **What standards of “normal” and “ability” are produced and enforced by specific AI systems, and what are the costs of being understood as an “outlier”?** How might these systems contribute to enforcing and creating fixed categories that further marginalize those who don’t “fit” (or those who do)?
- Given the proliferation of AI systems that claim to diagnose disability and illness, how can we ensure that options are available to challenge such diagnoses and intervene in automated responses to them? **How can we better highlight the (often profound) consequences of being diagnosed and pathologized, whether accurately or not, and provide opportunities to “opt out” before such determinations are made?**
- Major legislation prompted by disability advocacy, such as the Americans with Disabilities Act (ADA), has spurred technology companies to make their products and services more accessible to disabled people. Considering the increasing role of algorithms and AI systems across nearly all social institutions, **how might other anti-bias legal frameworks, such as national housing federation laws against discrimination and Section 508 laws mandating accessible digital infrastructure, provide us with new ways to imagine and fight for accountability in AI, in service of disabled people?**
- Over the twentieth and twenty-first centuries, the disability rights movement has overlapped as well as diverged from other social movements in important ways. **What tactics can those working toward AI accountability learn from the disability rights movement,** such as identifying architectural and structural bias across physical and digital spaces? How can those working on disability rights incorporate issues of AI bias and justice into their strategies?
- **What kinds of design, research, and AI engineering practices could produce more desirable futures for disabled people?** How can we assess the effects of these practices, and what other systemic interventions might be needed, beyond those focused on the technology itself?

Terms and Concepts from Disability Studies

To frame the workshop discussion and provide a shared vocabulary across disciplines and backgrounds, scholars Sara Hendren and Mara Mills began by introducing concepts and terms from disability studies that are especially relevant for a discussion of technological systems and structural discrimination. We summarize their presentation here, to help guide readers in how we are using these terms when researching disability, AI, and bias.

Models of Disability

“Models” of disability are ways of understanding and theorizing disability that shape our shared conception of what disability is and our approaches to it.³

The *medical model* of disability views disability as an impairment, illness, or disorder lodged within the individual. The medical model relies on biomedical standards of “normal” bodies, and classifies what falls outside of these standards as pathological.⁴ It has worked to stigmatize populations that are different, often in ways that justify exploitation.

The *social model* of disability is distinct from the medical model, defining disability as the product of disabling environments and attitudes, not aberrations located in individual bodies. According to this view, bodies and abilities are constructed in relation to these social and material contexts. Prejudice, which can be rooted in the medical model, causes people to classify and evaluate difference as pathological. The social model places the locus of intervention not at the level of the individual—with efforts that would attempt to “fix” disabled people—but at the level of social justice.⁵

Recent research has questioned the binary opposition posed between the medical and social models of disability, pointing to the difficulty of distinguishing between them in practice, and the tendency for social-model orthodoxy to downplay the lived experiences of people with different bodies, along with the very real difficulties of living with, for example, chronic pain.⁶

Moreover, a history of the concept of disability (especially in American contexts) is substantially different from a history of disabled people. As the historian Kim E. Nielsen argues, “When ‘disability’ is considered synonymous with ‘deficiency’ and ‘dependency,’ it contrasts sharply with American ideals of independence and autonomy. Thus disability has served as an effective weapon in contests over power and ideology. For example, at varying times, African Americans, immigrants, gays and lesbians, poor people, and women have been defined categorically as defective citizens incapable of full civic participation.”⁷

Disability as (Minority, Coalitional) Identity

A coalitional identity speaks to collective identities constructed not from a shared list of symptoms, in the case of disability, but rather through recognition of “the social and political

³ Throughout this report, we use the term “model” in two distinct ways: the first refers to theoretical/disciplinary framings of disability, as in this instance; the second refers to a discrete component of AI systems—as in an “AI model”.

⁴ Tobin Siebers, *Disability Theory* (Ann Arbor: University of Michigan Press, 2008), 25, 54.

⁵ Siebers, *Disability Theory*, 73.

⁶ Allison Kafer, *Feminist Queer Crip* (Bloomington and Indianapolis: Indiana University Press, 2013), 6–7.

⁷ Kim E. Nielsen, *A Disability History of the United States* (Boston: Beacon Press, 2012), xii.

circumstances that have forged [disabled people] as a group.”⁸ This works to unite a multiplicity of bodies and experiences and build solidarity between them, instead of producing an exclusionary definition. As scholars Rachel Adams, Benjamin Reiss, and David Serlin put it: “Disability encompasses a broad range of bodily, cognitive, and sensory differences and capacities. It is more fluid than most other forms of identities in that it can potentially happen to anyone at any time, giving rise to the insiders’ acronym for the nondisabled, TAB (or temporarily able-bodied)... Disability can be situational; it can also wax and wane within any particular body. Disability brings together people who may not agree on a common definition.”⁹

Scholar Meryl Alper also notes that “[o]ne billion people, or 15 percent of the world’s population, experience some form of disability, making them one of the largest (though also most heterogeneous) groups facing discrimination worldwide.”¹⁰

Non-disabled

A term that centers disability and places the non-disabled at the margin, used as an alternative to “able-bodied”.

Ableism

Discrimination in favor of the non-disabled or against people with disabilities.

Assistive Technology

“Assistive technology” is shorthand for technology that assists disabled people. While useful in some contexts, this term tends to (1) presume that some technologies offer assistance, while others don’t; and (2) frames the users of “assistive” technologies as passive. Katherine Ott points out that this term is also redundant: “Since all useful technology is assistive, it is peculiar that we stipulate that some devices are assistive while others need no qualification. Besides serving to stigmatize and segregate a benign and inanimate entity—a device or appliance—the term ‘assistive technology’ also needlessly complicates understanding of the devices so designated.”¹¹ Richard Ladner further argues that the term “assistive,” “has the ring of paternalism, a view that disabled people need lots of extra help, are dependent and are not capable human beings.”¹² While Mara Mills points out that, “the phrase advances a technological fix that is unconcerned with education, community support, or social change”¹³

⁸ Linton, *Claiming Disability*, 4.

⁹ Rachel Adams, Benjamin Reiss, and David Serlin, “Disability,” in *Keywords for Disability Studies*, eds. Adams, Reiss, and Serlin (New York and London: New York University Press, 2015), 5–6.

¹⁰ Meryl Alper, “Can Technology Really ‘Give Voice’ to Disabled People?,” *Pacific Standard* (March 6, 2018), <https://psmag.com/social-justice/can-technology-really-give-voice-to-disabled-people>.

¹¹ Katherine Ott, David Serlin, and Stephen Mihm, *Artificial Parts, Practical Lives: Modern Histories of Prosthetics* (New York: New York University Press, 2002), 21. See also Sara Hendren, “All Technology is Assistive,” *Wired* (October 16, 2014), <https://www.wired.com/2014/10/all-technology-is-assistive/>.

¹² Richard E. Ladner, “Accessible Technology and Models of Disability,” in *Design and Use of Assistive Technology: Social, Technical, Ethical, and Economic Challenges*, eds. Meeko Mitsuko K. Oishi, Ian M. Mitchel, and H. F. Machial Van der Loos (New York, Dordrecht, Heidelberg, London: Springer, 2010), 26.

¹³ Mara Mills, “Technology,” in *Keywords in Disability Studies*, 178.

Use and Users

As they work to create technologies that center disabled people, developers are often enthusiastic about involving “users” in codesign processes at early stages. But we need to critically assess what we mean by “use” and take care in how we invoke the design terminology of “the user.” Interactions with technology can be forced, reluctant, or partial; and non-use includes resisters, rejectors, the expelled, and the excluded.¹⁴

“Relationships to technology are generally far more intimate than implied by the term ‘user,’ or by the related term, ‘wearer,’” Mara Mills observes. “Identity formation can occur through technological use, as exemplified by self-advocacy groups for ‘cochlear implant users.’ Group affiliation also results from technical exclusion, as with the category of ‘print disability,’ which unites disparate individuals who cannot read printed materials. Finally, access to assistive devices is always stratified; the devices are not equally available or affordable to those who might benefit from them.”¹⁵

Assistive Pretext

A term coined by Mills to describe scenarios in which technology that is initially created to address a use case specific to disabled people is presented as the precursor—an inspiration or test case—for the development of a technology that is relevant to the general public. A classic example of this is curb cuts in sidewalks. While this is not altogether negative, it can in practice lead to disability being used as an initial pretext for development and funding that is later dropped, abandoning disabled people in favor of introducing technology into larger and more profitable markets.¹⁶ In a similar vein, disability advocate and design strategist Liz Jackson refers to high-tech solutions for the disability community that bypass serious infrastructural and institutional access needs (e.g., a wheelchair that climbs stairs) as “disability dongles.”¹⁷

Software Model of Personhood/the Mechanic View of “Intelligence”

This term describes a view of the body/mind that imagines them to be like an information-processing system, reduced to data, risks, diagnostics, and information to be processed. Such mechanic metaphors threaten to erase the lived experience of bodies, leaving visible only what can be measured and standardized.

¹⁴ Sally Wyatt, “Non-users Also Matter: The Construction of Users and Non-users of the Internet,” in *How Users Matter: the Co-Construction of Users and Technology*, eds. Nelly Oudshoorn and Trevor Pinch (Cambridge, MA: MIT Press, 2005), 67–79.

¹⁵ Mills, “Technology,” in *Keywords for Disability Studies*, 179.

¹⁶ Mara Mills, “Hearing Aids and the History of Electronics Miniaturization,” in *IEEE Annals of the History of Computing* 11 (2011) 24–44, <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5771310>.

¹⁷ S. E. Smith, “Disabled People Don’t Need So Many Fancy New Gadgets. We Just Need More Ramps,” *Vox*, April 30, 2019, <https://www.vox.com/first-person/2019/4/30/18523006/disabled-wheelchair-access-ramps-stair-climbing>.

As Olga Rachello puts it in her review of Mark O’Connell’s *To Be a Machine*, “the metaphor of ‘the mind as a piece of software, an application running on the platform of flesh’ [is] indicative of a [common] conception of the human being [among tech entrepreneurs] It equates the person with the mind, the mind with intelligence, and intelligence with information processing, and argues that large-scale information processing rarely needs to be fully understood by anyone to be effective. Moreover, it tolerates imprecision: to be a machine means in part that we can select certain aspects of our current human state and discard others as irrelevant.”¹⁸

Privatization/Bundling Ethical Choices within AI Consumer Products

Many significant ethical choices are already made by the time AI is delivered to (or used on) consumers, meaning that market forces—not the needs of individuals and communities on whom AI acts—very often determine both which products are offered, and the limits and implications of these products. George Estreich argues that “[b]ecause . . . products need to answer to existing demand, they will reproduce the values of the society in which they are sold. To the extent that they are adopted, they will translate those values into human populations.”¹⁹

Bias at the Intersection of AI and Disability

AI systems are being rapidly integrated into core social domains, making sensitive determinations that shape who receives resources and opportunities and who doesn’t. These systems, often marketed as capable of making smarter, better, and more objective decisions, have been shown repeatedly to produce biased and erroneous outputs, from voice recognition that doesn’t “hear” higher-pitched (i.e., “more feminine”) voices²⁰ to diagnostic systems that work poorly for people with dark skin²¹ to hiring algorithms that downgrade women’s résumés.²² But these concerns go beyond the identification of technical errors. Even when AI “works” as its designers intended, these systems are too often used in ways that serve the interests of those who already possess structural power, at the expense of those who don’t—from providing law enforcement with

¹⁸ Olga Rachello, “Lives of the Immortalists: The Human Stories of People Who Don’t Want to Be Human,” *The New Atlantis*, 54, Winter (2018), <https://www.thenewatlantis.com/publications/lives-of-the-immortalists>.

¹⁹ George Estreich, *Fables and Futures: Biotechnology, Disability, and the Stories We Tell Ourselves*, (Cambridge, MA: MIT Press, 2019), 5.

²⁰ Rachael Tatman, “Gender and Dialect Bias in YouTube’s Automatic Captions,” in *Conference: Proceedings of the First ACL Workshop on Ethics in Natural Language Processing* (January 2017), <http://www.ethicsinnlp.org/workshop/pdf/EthNLP06.pdf>.

²¹ Adewole S. Adamson and Avery Smith, “Machine Learning and Health Care Disparities in Dermatology,” in *JAMA Dermatology*, 154, no. 11 (November, 2018): 1247–1248, <https://jamanetwork.com/journals/jamadermatology/article-abstract/2688587>.

²² Jeffrey Dastin, “Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women,” Reuters, October 9, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

expanded surveillance powers, which target people of color,²³ to Immigration and Customs Enforcement (ICE) using AI systems to track and detain migrant populations.²⁴ Such concerns about AI bias and harm have grown over the past years, in time with a growing body of research and investigative journalism surfacing new and troubling examples.

To date, much of the work examining bias in AI systems has focused on the axes of race and gender, finding that AI bias generally harms women, people of color, gender minorities, and those at the intersections of these identities (although research that takes these intersections into account is less common than research that doesn't).²⁵ Like discourses on diversity and inclusion more generally, disability has been largely omitted from the AI-bias conversation, even as disabled people are affected by these issues in differing ways across axes of identity. Disabled people have been subject to historical and present-day marginalization, much of which has systematically and structurally excluded them from access to power, resources, and opportunity.²⁶ Such patterns of marginalization are imprinted in the data that shapes AI systems, and embed these histories in the logics of AI.²⁷ Recent research demonstrates this, showing that social attitudes casting disability as bad and even violent are encoded in AI systems meant to "detect" hate speech and identify negative/positive sentiment in written text. Researchers found that "a machine-learned model to moderate conversations classifies texts which mention disability as more 'toxic'" while "a machine-learned sentiment analysis model rates texts which mention disability as more negative."²⁸

These findings fit an established pattern. Those who have borne discrimination in the past are most at risk of harm from biased and exclusionary AI in the present. And when these discriminatory logics are reproduced and amplified by AI systems, they are likely to be read as

²³ Drew Harwell, "Oregon Became a Testing Ground for Amazon's Facial-Recognition Policing. But What If Rekognition Gets It Wrong?," *Washington Post*, April 30, 2019, <https://www.washingtonpost.com/technology/2019/04/30/amazons-facial-recognition-technology-is-supercharging-local-police/>.

²⁴ McKenzie Funk, "How ICE Picks Its Targets in the Surveillance Age," *New York Times*, October 2, 2019, <https://www.nytimes.com/2019/10/02/magazine/ice-surveillance-deportation.html>.

²⁵ A notable exception that presents an intersectional analysis, examining AI inaccuracy in the context of both race and gender, is Joy Buolamwini and Timnit Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (New York, 2018), 77–91, <http://proceedings.mlr.press/v81/buolamwini18a.html>.

²⁶ Douglas Baynton, "Disability and the Justification of Inequality in American History," in *The New Disability History: American Perspectives*, eds. Paul K. Longmore and Lauri Umansky (New York: New York University Press, 2001), 33–57. See also Peter Cryle and Elizabeth Stephens, *Normality: A Critical Genealogy* (Chicago: University of Chicago Press, 2017).

²⁷ For an example of the way discriminatory logics and practices constitute the data that informs AI systems, see Rashida Richardson, Jason Schultz, and Kate Crawford, "Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice," *New York University Law Review Online* (February 13, 2019), <https://ssrn.com/abstract=3333423>.

²⁸ Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl, "Unintended Machine Learning Biases as Social Barriers for Persons with Disabilities," in *SIG Access Newsletter*, 125 (October 2019), <http://sigaccess.org/newsletter/2019-10/hutchinson.html>.

authoritative, the product of sophisticated technology. Beyond biased data, additional risks are presented by the significant power asymmetries between those with the resources to design and deploy AI systems, and those who are classified, ranked, and assessed by these systems. The history of disability itself can be read as the story of the ways in which various systems of classification (i.e., medical, scientific, legal) interface with social institutions and their articulations of power and knowledge.

Centering disability in the discussion of AI bias complicates the dominant approaches to addressing these challenges. The disability studies and disability rights communities have long deliberated over issues of discrimination, bias, and the means by which the technology that mediates our world can alleviate or contribute to these harms. Further, the AI-bias debate engages many concerns that affect disabled people in particular and acute ways, including issues of privacy, consent, and the high stakes of (mis)classification in the context of asymmetric power relationships (such as, for instance, between patient and doctor). At its most basic level, excluding consideration of disability (and disability scholarship and activism) from the discourse on AI bias means that efforts to remediate such bias won't include disabled people, and thus they will be more likely to experience harm and marginalization.

AI systems model the world based on what's in the data they're given. If something is missing from the data, say images of people with dark skin, these people will be missing from the AI model, and thus won't be recognized or included. Such exclusions can lead to significant harm. Take, for example, the use of AI in developing autonomous vehicles. If the data used to train a pedestrian recognition system doesn't include representations of people using scooters or wheelchairs, it's likely that such people won't be "recognized" as pedestrians. In 2018, an autonomous Uber in Arizona killed Elaine Herzberg, a pedestrian who was pushing a bicycle when she was killed.²⁹ A recent National Transportation Safety Board investigation found significant problems with Uber's autonomous system, including its shocking failure to "recognize" pedestrians outside of crosswalks. The investigation also found that Uber's system had a hard time classifying Herzberg: "When the car first detected her presence, 5.6 seconds before impact, it classified her as a vehicle. Then it changed its mind to 'other,' then to vehicle again, back to 'other,' then to bicycle, then to 'other' again, and finally back to bicycle."³⁰ Did the system misclassify Herzberg due to the presence of the bicycle? Would it similarly misclassify people on scooters and in wheelchairs? There are countless examples of AI systems known or strongly suspected to function incorrectly for disabled people.³¹ And as disability scholar Karen Nakamura

²⁹ Troy Griggs and Daisuke Wakabayashi, "How a Self-Driving Uber Killed a Pedestrian in Arizona," *New York Times*, March 21, 2018, <https://www.nytimes.com/interactive/2018/03/20/us/self-driving-uber-pedestrian-killed.html>.

³⁰ Aarian Marshall and Alex Davies, "Uber's Self-Driving Car Didn't Know Pedestrians Could Jaywalk," *Wired*, November 5, 2019, <https://www.wired.com/story/ubers-self-driving-car-didnt-know-pedestrians-could-jaywalk/>.

³¹ Anhong Guo, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach, and Meredith Ringel Morris, "Toward Fairness in AI for People with Disabilities: A Research Roadmap," in ACM ASSETS 2019 Workshop on AI Fairness for People with Disabilities, Issue 125 (August 19, 2019), <http://www.sigaccess.org/newsletter/2019-10/guo.html>.

points out: “Wheelchair users are constantly being run over by human car drivers that do not recognize them as humans (the ‘I didn’t even see you’ excuse that bicyclists are also familiar with), yet the datasets being used to train automobile vision systems also embed similar limitations not only due to the lack of wheelchairs and scooters in training datasets, but the trainers themselves may be misrecognizing them.”³² If such AI systems aren’t built and implemented in ways that put disabled people’s safety first, such histories and harms will be encoded, automated, and amplified.³³

But the issues go beyond this. The category of “disability” complicates past classifications, and thus perturbs calls to simply include disabled people in datasets, which are constructed around rigid models of categorization, however many categories they might include. Disability encompasses a vast and fluid number of physical and mental health conditions (such as asthma, depression, and post-traumatic stress disorder) which can come and go throughout a person’s lifetime (or even in the course of a single day).³⁴ Each has its own histories and specificities, some of which are uncommon in relation to the average of a general population, and many of which aren’t visible. Simply expanding a dataset’s parameters to include new categories, in an attempt to account for “disability,” won’t work to ensure disabled people are represented (and may require increased surveillance and invasion of privacy in the process). **Indeed, the way in which “disability” resists fitting into neat arrangements points to bigger questions about how other identity categories, such as race, sexual orientation, and gender, are (mis)treated as essential, fixed classifications in the logics of AI systems, and in much of the research examining AI and bias.**

Integrating disability into the AI bias conversation helps illuminate the tension between AI systems’ reliance on data as the primary means of representing the world, and the fluidity of identity and lived experience. Especially given that the boundaries of disability (not unlike those of race and gender) have continually shifted in relation to unstable and culturally specific notions of “ability,” something that has been constructed and reconstructed in relationship to the needs of industrial capitalism, and the shifting nature of work.³⁵ As scholars Susan Burch and Ian Sutherland note: “Disability is often less about physical or mental impairment than it is about how society responds to impairments.”³⁶

³² Karen Nakamura, “My Algorithms Have Determined You’re Not Human: AI-ML, Reverse Turing-Tests, and the Disability Experience,” in the 21st International ACM SIGACCESS Conference (October 2019): 1–2, <https://doi.org/10.1145/3308561.3353812>.

³³ John D. Kraemer and Connor S. Benton, “Disparities in Road Crash Mortality among Pedestrians Using Wheelchairs in the USA: Results of a Capture–Recapture Analysis,” in *BMJ Open* (October 19, 2015), <https://doi.org/10.1136/bmjopen-2015-008396>.

³⁴ Shari Trewin, “AI Fairness for People with Disabilities: Point of View,” Cornell University, Computer Science, Artificial Intelligence, November 26, 2018, <https://arxiv.org/abs/1811.10670>.

³⁵ Sarah F. Rose, *No Right to Be Idle: The Invention of Disability, 1840s–1930s* (Chapel Hill: University of North Carolina Press, 2017), 2.

³⁶ Susan Burch and Ian Sutherland, “Who’s Not Here Yet? American Disability History,” *Radical History Review* 94 (2006), 127–47.

Disability is also mediated by other identity categories, including race, gender, sexual orientation, class, and age, which often inform what “counts” as disability. Indeed, discrimination against people of color, women, and other historically marginalized groups has often been justified by representing these groups as disabled, meaning that the very definition of (dis)ability is gendered, racialized, and based on class status. As scholar Douglas Baynton lays out, “the *concept* of disability has been used to justify discrimination against other groups by attributing disability to them.”³⁷ Thus disability is entwined with, and serves to justify, practices of marginalization. As such, **examination of AI bias cannot simply “add” disability as one more stand-alone axis of analysis, but must pay critical attention to interlocking structures of marginalization, and the way in which these mediate how “disability” itself is defined and who has the power to affix these definitions.**³⁸

The social model of disability, which cautions against an understanding of disability solely as an individual attribute or pathology, also complicates dominant approaches to “solving” AI bias. Such technical approaches tend to classify people by a single (usually demographic) variable, such as race or gender, and then apply a variety of methods to test or modify a given AI system to ensure that it functions similarly across all “types” of people. For example, such a test might assess whether faces with lighter or darker skin pigmentation are recognized at equal rates by facial-recognition systems.³⁹ These approaches have already received deserved criticism for their limited and essentialist reading of fluid and socially constructed categories of identity, like race and gender, in addition to criticism that they promote technical fixes to social problems.⁴⁰ **But with respect to disability, their limitations become clearer, especially if we understand disability as a product of disabling environments, and thus as an identity that can only be understood in relation to a given social and material context.** Meryl Alper explains: “Efforts to better include individuals with disabilities within society through primarily technological interventions rarely take into account all the other ways in which culture, law, policy, and even technology itself can also marginalize and exclude.”⁴¹

Including concepts from disability studies can help us refocus and refine our approach to AI bias, moving away from an emphasis on the technology alone, and toward an approach that accounts for the context in which such technology is produced and situated, the politics of classification, and the ways in which fluid identities are (mis)reflected and calcified through such technology.

³⁷ Douglas Baynton, “Disability and the Justification of Inequality in American History,” in *The New Disability History: American Perspectives*, 18.

³⁸ Alison M. Parker, “Intersecting Histories of Gender, Race, and Disability,” *Journal of Women’s History* 27, no.1 (Spring 2015): 178–186.

³⁹ Patrick Grother, Mei Ngan, and Kayee Hanaoka, “Ongoing Face Recognition Vendor Test (FRVT) Part 1: Verification,” National Institute of Standards and Technology, July 2019, https://www.nist.gov/system/files/documents/2019/10/16/frvt_report_2019_10_16.pdf

⁴⁰ See, for example, Sarah Myers West, Meredith Whittaker, and Kate Crawford, “Discriminating Systems: Gender, Race, and Power in AI,” April 2019, <https://ainowinstitute.org/discriminatingystems.pdf>.

⁴¹ Alper, “Can Technology Really ‘Give Voice’ to Disabled People?”

AI's Version of "Normal"

In modeling the world through data, AI systems necessarily produce and reflect a normative vision of the world. These systems function to classify new information based on a given model: Does a novel input comport with the model, or constitute an outlier? Is it normal, or aberrant?

When considering AI and disability, It's important to acknowledge this logic as core to how AI systems function. Despite the fact that there are over a billion disabled people worldwide, most social structures and institutions treat disabled people as deficient or dependent, as "atypical" or "abnormal." Indeed, as Simi Linton points out, "*normal* and *abnormal* . . . are often used to distinguish between people with and without disabilities."⁴² Scholar Lennard J. Davis focuses on the construction of normalcy itself, and whom such a construction excludes: "To understand the disabled body, one must return to the concept of the norm, the normal body. . . . the 'problem' is not the person with disabilities; the problem is the way that normalcy is constructed to create the 'problem' of the disabled person."⁴³ Of course, such standards are socially constructed, and highly contingent: normal in 2019 looks much different than normal in 1919, and normal in the United States cannot be generalized and mapped onto the rest of our heterogeneous and dynamic world.

Versions of normalcy reflected in the cultures and logics of corporate and academic tech environments are encoded in data and design and amplified through AI systems. And these systems are being applied in ways that allocate resources and opportunities.

Examining which standards and expectations are encoded in these systems, and what their consequences may be, is an urgent undertaking, especially when considering disability. **What "norms" are produced and enforced by AI systems? And, given that these systems are being used in high-stakes domains, where presenting as an "outlier" may have significant life consequences, how might AI contribute to enforcing and creating fixed categories that further marginalize those who don't "fit"?**

Scholar Jutta Treviranus recounts testing an AI model designed to guide autonomous vehicles, hoping to understand how it would perform when it encountered people who fell outside the norm and "did things unexpectedly." To do this, she exposed the model to footage of a friend of hers who often propels herself backward in a wheelchair. Treviranus recounts, "When I presented a capture of my friend to the learning models, they all chose to run her over. . . . I was told that the learning models were immature models that were not yet smart enough to recognize people in wheelchairs. . . . When I came back to test out the smarter models they ran her over with greater confidence. I can only presume that they decided, based on the average behavior of wheelchairs, that wheelchairs go in the opposite direction." In reflecting on this experience, Treviranus states, "I should have expected this. . . . When we process the scatter-plot of data to guide our decisions

⁴² Linton, *Claiming Disability*, 22.

⁴³ Lennard J. Davis, "Introduction: Disability, Normality, and Power," in *The Disability Studies Reader*, ed. Lennard J. Davis (New York: Routledge, 2016), 1–2.

using standard data analysis, our decisions will be determined by the middle which is the norm and represents the majority.”⁴⁴ Adding more data to an AI system doesn’t address the issue—it simply reinforces the normative model at the core of a given system’s calculations, meaning that those who fall outside of this norm become increasingly remote “outliers.” No amount of “additional data” can solve this problem.

It is noteworthy that much of the AI targeted at disabled people implicitly promises to make them more like non-disabled people, based on an implied understanding that non-disabled is the norm. For example, an AI-enabled app called Ava promises to allow Deaf people to take part in spoken conversations. Those engaged in the conversation open the app on their phones, and Ava converts spoken word into text, allowing people to follow along.⁴⁵ Ava’s premise is that deafness is a hindrance to communication. But, as disability activist and scholar Eli Clare points out, “many deaf people claim themselves not as disabled, but as a linguistic minority. They locate the trouble they experience not in their inability to hear but in the non-deaf world’s unwillingness to learn and use sign language.” Thus, the Deaf community doesn’t universally welcome technologies that “bring with them the non-deaf world’s hope of eradicating both deafness as a medical condition and being Deaf as an identity.”⁴⁶

AI-enabled technologies are also targeting autistic people, promising to teach them to act more neurotypically.⁴⁷ Such systems coach people clinically diagnosed with autism spectrum disorder (ASD) to make eye contact, to display emotions through facial expression, and to learn how to read emotions on someone else’s face. AI-based “emotion recognition” then assesses their performance and provides an evaluation of their progress based on one or another model of non-disabled expression and interaction.⁴⁸ Katta Spiel and their coauthors observe the way in which these “technologies embody normative expectations of a neurotypical society, which predominantly views autism as a medical deficit in need of ‘correction’.”⁴⁹ In both this case and

⁴⁴ Jutta Treviranus, “Sidewalk Toronto and Why Smarter Is Not Better,” October 30, 2018, <https://medium.com/datadriveninvestor/sidewalk-toronto-and-why-smarter-is-not-better-b233058d01c8>.

⁴⁵ “Ava,” “Live Subtitles for Any Conversation,” accessed October 31, 2019, <https://www.ava.me>.

⁴⁶ Eli Clare, *Brilliant Imperfection: Grappling With Cure* (Durham, NC and London: Duke University Press, 2017), 91. See also Leila Monaghan, Constanze Schmaling, Karen Nakamura, and Graham H. Turner, eds., *Many Ways to Be Deaf: International Variation in Deaf Communities* (Washington, DC: Gallaudet University Press, 2013).

⁴⁷ Brian Scassellati, Laura Boccanfuso, Chien-Ming Huang, Marilena Mademtzi, Meiying Qin, Nicole Salomons, Pamela Ventola, and Frederick Shic, “Improving Social Skills in Shildren with ASD Using a Long-Term, In-Home Social Robot” in *Science Robotics* 3, no. 21 (August 22, 2018), <https://doi.org/10.1126/scirobotics.aat7544>.

⁴⁸ Note that AI-based emotion recognition is based on no solid scientific evidence, as psychologist Lisa Feldman Barrett documents in a recent survey paper reviewing all of the available literature. See Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, and Seth D. Pollak, “Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements,” *Psychological Science in the Public Interest* 20 (July 2019): 1–68, <https://doi.org/10.1177/1529100619832930>.

⁴⁹ Katta Spiel, Christopher Frauenberger, Os Keyes, and Geraldine Fitzpatrick, “Agency of Autistic Children in Technology Research—A Critical Literature Review,” *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, no. 6, Article 38 (November 2019): 40 pages, <https://doi.org/10.1145/3344919>

that of the Ava app, disability is implicitly understood to be undesirable, with AI positioned as “solving the ‘problem’ of disability.”⁵⁰

The increasing use of AI systems in high-stakes decision-making and resource distribution may force individuals to make themselves legible to these systems, even as the “norms” they enforce may not be desirable or comfortable. This phenomenon is known as a *reverse Turing Test*. In the AI research field, the Turing Test is a benchmark that deems an AI system “intelligent” if it can convince a person that they’re talking to another human, and not to a machine. A reverse Turing Test refers to a pattern whereby humans are required to prove *their* humanity to AI systems, adjusting their behavior or appearance to fit the system’s normative, predetermined categories. Disability scholar Karen Nakamura observes that reverse Turing Tests may stigmatize and discriminate against disabled people by judging them as nonhuman.⁵¹

We already encounter such reverse Turing Tests in mundane interactions, such as filling out a CAPTCHA before we submit a form, or clicking a box to prove we are not a robot. AI-enabled self-checkout systems now use facial recognition to determine whether customers are old enough to purchase age-restricted items like alcohol and cigarettes.⁵² In all of these cases, in order to proceed with a desired action, you need to prove to the machine that you are a certain type of human. We have already seen examples of disabled people being discriminated against based on these systems, from people with amyotrophic lateral sclerosis (ALS) whose speech isn’t recognized by smart assistants⁵³ to disabled clickworkers who are rejected by platforms like Amazon Mechanical Turk due to their inability to pass CAPTCHAs, or due to the speed of their work being below average (speed is often used as a metric to weed out spammers).⁵⁴

Here we can see that AI systems not only reflect normative categories, but also create and enforce these categories, demarcating what it means to be a legible human and whose bodies, actions, and lives fall outside. These systems construct some categories of “ability” and negate others, in the process remapping and calcifying the boundaries of inclusion and marginalization. As the use of these systems expands into an increasing array of sensitive social domains, the risk

⁵⁰ Kafer, *Feminist Queer Crip*, 107.

⁵¹ Nakamura, “My Algorithms Have Determined You’re Not Human: AI-ML, Reverse Turing-Tests, and the Disability Experience,” 1–2.

⁵² “Yoti’s Biometric ID Platform to Provide Proof of Age for Self-Serve Checkout,” *Find Biometrics Global Identity Management*, accessed October 13, 2019, <https://findbiometrics.com/yotis-biometric-id-proof-of-age-self-serve-checkout-510175/>.

⁵³ Shaun K. Kane, Meredith Ringel Morris, Ann Paradiso, Jon Campbell, “‘At times avuncular and cantankerous, with the reflexes of a mongoose’: Understanding Self-Expression through Augmentative and Alternative Communication Devices,” *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (February 2017): 1166–1179, <https://doi.org/10.1145/2998181.2998284>.

⁵⁴ Kathryn Zyskowski, Meredith Ringel Morris, Jeffrey P. Bigham, Mary L. Gray, and Shaun K. Kane, “Accessible Crowdwork? Understanding the Value In and Challenge of Microtask Employment for People with Disabilities,” *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (March 2015); 1682–1693, https://cs.stanford.edu/~merrie/papers/crowdwork_and_disability.pdf.

of not being “seen” as human become greater. As with the examples of autonomous vehicles, if someone doesn’t “look like a pedestrian” when crossing the street, they risk being killed.

Recalling the long history of struggle by people fighting both to be included in and excluded from the *Diagnostic and Statistical Manual of Mental Disorders* (DSM), the standard reference for psychiatric diagnosis, we get a sense of the extremely high stakes—and significant power—that AI systems quietly acquire as they map and remap categories of ability. For example, from 1952 until 1973, being gay was included in the DSM as a psychiatric disorder, and was only removed after sustained protest from gay activists, who “understood quite well the social consequences of being labeled and defined by others.” In working to remove “homosexuality” from the DSM, the Gay Liberation Front and other gay-rights activists were fighting to remove a key justification for their continued mistreatment: “depriving secular society, increasingly dependent on ‘health’ as a moral category, of the ideological justification for many of its discriminatory practices.”⁵⁵ While the legacies of struggle around the DSM are situated within a long history of bias and discrimination, they also tell the story of the rapid pace with which human bias and the culture supporting it can change. There have been radical shifts in the perception of how we treat and understand sickness, health, and people with disabilities over the last century. AI systems, by encoding and fixing perspectives that reflect a specific cultural understanding of ability, threaten to foreclose the opportunity for such significant change by fixing old categories in ways that can’t easily be shifted or rejected.

Work, Disability, and AI

The norms enforced and encoded by AI systems have significant material consequences. Artificial intelligence is rapidly being introduced into the workplace, and is already informing decisions about hiring, management, performance assessment, and beyond.⁵⁶

The example of the AI company HireVue is instructive. The company sells AI video-interviewing systems to large firms, marketing these systems as capable of determining which job candidates will be successful workers, and which won’t, based on a remote video interview. HireVue uses AI to analyze these videos, examining speech patterns, tone of voice, facial movements, and other indicators. Based on these factors, in combination with other assessments, the system makes recommendations about who should be scheduled for a follow-up interview, and who should not

⁵⁵ Ronald Bayer, *Homosexuality and American Psychiatry: The Politics of Diagnosis* (Princeton: Princeton University Press, 1987), 12–13.

⁵⁶ See, for example, Jaden Urbi, “Some Transgender Drivers Are Being Kicked Off Uber’s App,” CNBC, August 8, 2018, <https://www.cnbc.com/2018/08/08/transgender-uber-driver-suspended-tech-oversight-facial-recognition.html>; and Sarah Kurbi, “The New Ways Your Boss Is Spying on You,” *Wall Street Journal*, July 19, 2019, <https://www.wsj.com/articles/the-new-ways-your-boss-is-spying-on-you-11563528604>.

get the job.⁵⁷ In a report examining HireVue and similar tools, authors Jim Fruchterman and Joan Mellea are blunt about the way in which HireVue centers non-disabled people as the “norm,” and the implications for disabled people: “[HireVue’s] method massively discriminates against many people with disabilities that significantly affect facial expression and voice: disabilities such as deafness, blindness, speech disorders, and surviving a stroke.”⁵⁸

Beyond this, it is important to note that a meaningful connection between *any* person’s facial features, tone of voice, and speech patterns, on one hand, and their competence as a worker, on the other, is not backed by scientific evidence—nor is there evidence supporting the use of automation to meaningfully link a person’s facial expression to a particular emotional state.⁵⁹ Even where disabled people are represented in the training data used to instruct such systems, this does not address these underlying questions. Indeed, it raises many more, including concerns about how such data is constructed and how a given dataset’s claims to represent “disability” are validated, in addition to issues of privacy and surveillance.

With all such systems, we need to ask who the “user” really is, and whose interests a system’s use will ultimately serve. HireVue licenses their system to large firms like Goldman Sachs and Unilever, who work with HireVue to create a model of a successful worker based on examples of employees who’ve been rewarded and promoted at the firm. In this case, the “user” is very clearly the employer, not the job seeker—the person whose access to opportunity is being shaped by this system. Indeed, the people on the other end of an interview may not realize they are being judged by a third-party AI system at all.⁶⁰ This fits a pattern in which those who bear the costs of AI’s failures and bias are frequently the least powerful actors within a given context. Risk is displaced onto historically marginalized groups, while the benefits of efficiency and profit accrue to the companies producing and licensing such systems.

Troublingly, HireVue also has a patent for an AI system designed to identify disabled people based on their speech, mannerisms, tone, and other physical markers.⁶¹ HireVue’s system proposes to diagnose interview subjects as disabled, without consent, based on methods that are

⁵⁷ HireVue website, *HireVue*, accessed October 23, 2019, <https://www.hirevue.com/>.

⁵⁸ Jim Fruchterman and Joan Mellea, “Expanding Employment Success for People with Disabilities,” Benetech (November 2018), <https://benetech.org/wp-content/uploads/2018/11/Tech-and-Disability-Employment-Report-November-2018.pdf>.

⁵⁹ Feldman Barrett et. al., “Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements.”

⁶⁰ See, for example, recent legislation passed in Illinois requiring disclosure of the use of AI in video interviewing. Illinois General Assembly, “Bill Status of HB2557,” <http://www.ilga.gov/legislation/billstatus.asp?DocNum=2557&GAID=15&GA=101&DocTypeID=HB&LegID=118664&SessionID=108>.

⁶¹ Loren Larsen, Keith Warnick, Lindsey Zuloaga, and Caleb Rottman, “Detecting Disability and Ensuring Fairness in Automated Scoring of Video Interviews,” United States Patent Application Publication, August 20, 2018, <https://patents.google.com/patent/US20190057356A1/en?q=affect&assignee=hirevue&oq=hirevue+affect>. Thank you to AI Now Tech Fellow Genevieve Fried, who unearthed this patent in the course of her research.

not scientifically grounded. It uses a neural network to “detect” similar features across multiple videos of people who identify as disabled, finding patterns of behavior, speech, tone, or other markers that correlate across videos of disabled people. The system is then used to detect disability—or, at least, to identify these same patterns—in subsequent job-candidate videos. While the patent’s authors claim that this system is intended to *reduce* discrimination and bias against disabled people during hiring—by adjusting the automated “job performance score” assigned to people “detected” as disabled—there’s no guarantee it will be used this way, especially if it’s more profitable for employers to use it otherwise.

How such systems can be held accountable to those at risk of harm is an open question. HireVue, like most commercial AI, is a proprietary system. While the company makes claims about its bias mitigation techniques, it does so only in the abstract, and does not open its systems to outside scrutiny. Further, because establishing a pattern of bias or discrimination requires examining the system’s performance across many candidates, those most likely to be harmed by such discrimination (job candidates) lack access to the information they need to bring a suit, while those who do have access (employers) have no incentive to assist. **This makes enforcing antidiscrimination laws like the ADA extremely difficult in the context of workplace and hiring AI.**

These AI systems build on longstanding practices in which “ability” is constructed based on the capacity to engage in waged work. The concept of disability has been defined relative to a person’s aptitude to contribute productively to specific types of industrial workplaces, and in relation to the needs and interests of large employers. It acquired its current connotations in the nineteenth century, following the introduction of workplace compensation, when employers began to use preemployment screenings to eliminate people deemed inefficient, nonproductive, and likely to require extra help and support.⁶²

While the ADA was intended to correct this pattern of discrimination, in practice this has not been the case: comparing employment rates of disabled people before and after ADA, we see that the numbers have actually declined.⁶³ Allison Kafer puts it bluntly: “most disabled people in the United States are un- or underemployed, and . . . almost a third of disabled people live below the poverty line.”⁶⁴ In the context of such pronounced and systemic inequality, the integration of systems trained to reflect and reproduce existing norms, and designed to privilege the interests of those who already possess structural power, is profoundly troubling.

AI and automated systems are also being introduced as replacements for human workers who help provide disabled people access. But as scholar Louise Hickman points out, services like transcription require more than the ability to transcribe spoken language into text. For example, Communication Access Real-time Translation (CART) operators, who provide transcriptions for

⁶²Rose, *No Right to Be Idle: The Invention of Disability, 1840s–1930s*, 223–224.

⁶³ Michelle Maroto and David Pettinicchio, “Twenty-Five Years After The ADA: Situating Disability in America’s System of Stratification,” *Disability Studies Quarterly* 35, no. 3 (September 2015), <http://dsq-sds.org/article/view/4927/4024>.

⁶⁴ Kafer, *Feminist Queer Crip*, 107.

Deaf people, rely on their own “affective capacity to make sense of, and code, speech.” This capacity is developed in relationship with the person whom CART transcription serves, allowing the CART operator to draft a personalized “dictionary” that contains domain knowledge specific to a given person and their needs. The move to replace human workers with AI transcription services is thus “premised on the idea that the means of access is a neutral process of transferring speech and visual data into readable and spoken text, in turn disavowing the centrality of situated and affective labor in the complexity of access work.”⁶⁵ Such “replacements” are premised on the idea that an AI system can replace a human, invoking the “Software Model of Personhood” in which intelligence and ability are viewed as the domain of information processing, and embodied experience, human connection, and care work are excluded.

In the context of disability, AI, and work, it is also important to mention the growing use of AI-enabled worker-management platforms. Amazon warehouses have pioneered the use of such systems, which automatically set workers’ “rate”—the amount of work they have to do that day in order not to be fired. “Amazon has [built] a massive and intricate system of surveillance and control to accelerate the rate of productivity. ‘There is no privacy,’ [an Amazon worker named] Charlie said. . . ‘You should start thinking of it like a prison structure . . . We just presume we’re always being watched.’”⁶⁶ Amazon’s systems are calibrated to extract as much labor from workers as possible, and workers have continually voiced concerns that unrealistic rates and punishing workloads are leading to chronic injury and stress. A former Amazon warehouse Safety Manager tied the high rate of injury to the punishing automated performance targets: “If you had an injury . . . there was no leniency, you were expected to keep that rate.”⁶⁷ A survey by the New York Committee for Occupational Safety & Health found that 80 percent of workers in Amazon’s Staten Island Warehouse felt pressure to work faster and harder, and over 60 percent experience sleep problems as a result of the job.⁶⁸ Amazon presents a model in which the automation of worker performance management, calibrated to push workers to the edge of their abilities in pursuit of profit and efficiency, leads to chronic injury, pain, and psychological stress. In considering the intersections of disability, AI, and work, the ways in which such exploitative AI systems serve to produce harmful and disabling environments for workers must also be accounted for.

There is a long history of exploitation of disabled workers, and it continues today. Indeed, section 214(c) of the Fair Labor Standards Act, commonly referred to as 14(c), still allows employers to

⁶⁵ Louise Hickman, “Transcription Work and the Practices of Crip Technoscience,” *Catalyst: Feminism, Theory, Technoscience* Vol 5, no. 1: Special Section on Crip Technoscience (April 2019), <https://doi.org/10.28968/cftt.v5i1.32081>.

⁶⁶ Sam Adler Bell, “Surviving Amazon,” *Logic*, no. 8, August 3, 2019, <https://logicmag.io/bodies/surviving-amazon/>.

⁶⁷ Tonya Riley, “She Injured Herself Working at Amazon. Then The Real Nightmare Began,” *Mother Jones*, March 19, 2019, <https://www.motherjones.com/politics/2019/03/amazon-workers-compensation-amcare-clinic-warehouse>.

⁶⁸ Alexia Fernández Campbell, “The Problem with Amazon’s Speedy Shipping, in One Graphic,” *Vox*, October 18, 2019, <https://www.vox.com/identities/2019/10/18/20920717/amazon-shipping-workers-injuries>.

pay disabled workers below minimum wage under certain conditions, and has provided the pretext for significant exploitation and abuse of disabled workers.⁶⁹ While 14(c) was intended to help encourage employment of disabled people, “the unfortunate truth,” as legal scholar Theresa Golde observes, “is that 14(c) isolates the disabled community with regard to their capability, emphasizing quite literally that different is less.”⁷⁰ In one recent example, a nonprofit licensed to employ disabled people for less than minimum wage was found to be compensating some disabled workers in gift cards instead of money.⁷¹ This pattern of exploitation, and the introduction of harmful and exploitative AI management systems into the workplace, suggest the possibility of new worker coalitions, in which those exploited due to their being diagnosed or identified as “disabled” find common cause with workers whose workplace conditions, including automated management systems, are themselves disabling and exploitative.

Representation, Privacy, and Disability

Given the serious consequences of deploying AI systems that reproduce and amplify systematic exclusion, scholars and advocates have made calls to increase the collection of data representing disabled people.⁷² This is an intuitive response to what can be a serious threat, as the example of autonomous cars makes clear. Disability communities recognize the utility of such data, and have participated in grassroots collection efforts. For example, a rare-illness co-op in France was created to collect data from its members, in the hope of better understanding and addressing their health.⁷³ However, such an open and participatory method of data collection could be risky in contexts where healthcare isn’t guaranteed, because the harms resulting from data disclosures about individuals with rare illnesses could result in higher health-insurance costs, denial of coverage or employment, or other forms of stigma.⁷⁴ Similarly, once collected, it’s difficult to ensure that such data won’t be reused in ways that could cause harm, even if this was not the initial intention.

⁶⁹ US Department of Labor, “The Fair Labor Standards Act of 1938, as Amended,” revised May 2011, accessed November 14, 2019, <https://www.dol.gov/whd/regs/statutes/fairlaborstandact.pdf>.

⁷⁰ Theresa Golde, “Pennies an Hour: Was This Really the Intent behind 14(c) of the Fair Labor Standards Act? A Note Calling for a System Change to an Otherwise Broken System,” *Texas Tech Law Review* 48 no. 2 (January 2016): 459–503.

⁷¹ Alexia Fernández Campbell, “A Loophole in Federal Law Allows Companies to Pay Disabled Workers \$1 an Hour,” *Vox*, May 3, 2018, <https://www.vox.com/2018/5/3/17307098/workers-disabilities-minimum-wage-waiver-rock-river-valley-sel-f-help>.

⁷² Alex Engler, “For Some Employment Algorithms, Disability Discrimination by Default,” *Brookings TechTank*, October 31, 2019, <https://www.brookings.edu/blog/techtank/2019/10/31/for-some-employment-algorithms-disability-discrimination-by-default/>.

⁷³ “Improving the Patients’ Life Path,” *French Foundation for Rare Diseases*, accessed October 13, 2019, <https://fondation-maladiesrares.org/eng/our-%20missions/improve-the-quality-of-life-of-sick-people/>.

⁷⁴ Meredith Ringel Morris, “AI and Accessibility: A Discussion of Ethical Considerations,” *Communications of the ACM* (forthcoming), <https://arxiv.org/ftp/arxiv/papers/1908/1908.08939.pdf>.

Other attempts to gather data about disabled people have also been problematic. Attendees at the workshop raised concerns that efforts to augment datasets to include or represent people with disabilities are done by clickworkers, who construct or label data as being from people who are “disabled” based on what is effectively a hunch.⁷⁵ This echoes the practice of trying to create “racially inclusive” datasets by simulating dark skin, based on images of light-skinned people.⁷⁶ In both cases, in the pursuit of more effective AI systems, simulations and stereotypes stand in for meaningful inclusion. Relying on such simulacra risks constructing categories that effectively exclude many of those they are meant to represent, and likely include many who would not identify as fitting within these categories. Once constructed, it’s also likely that these categories could go on to shape our understanding of disability, and potentially our understanding of ourselves. And as scholar Alison Kafer points out, discussing the politics of classifying disability, “fixed definitions cannot be divorced from the economic effects of such fixing.”⁷⁷

Of course, efforts to collect more, and more accurate, data from disabled people are in tension with efforts to preserve privacy. This is true of all social data collection, but the risks are especially pronounced in the case of disability, given that simply disclosing a disability can pose a significant risk. A workshop participant recounted being barred from purchasing a life insurance policy because their anxiety disorder made them ineligible. This exclusion fits a pattern of insurers discriminating against people with mental-health issues.⁷⁸ Disclosing a disability can put access to healthcare and employment at risk, especially in the US, where access to medical insurance is closely tied to employment, and not universally guaranteed.

Disclosure is not always voluntary. Disturbingly, tech companies are developing AI systems that claim to detect whether people have a disability. An internet search or cursor movements on a screen can now be used to infer disability or diagnosis, creating a spectrum of implicit and explicit disclosure. A recent paper proposed a system for diagnosing Parkinson’s disease from a person’s mouse movements.⁷⁹ The automation of such medical diagnoses and disability “detection” raises fundamental ethical concerns around privacy, consent, and disclosure, even (and perhaps especially) when these systems are marketed as being designed to “reduce bias” against the

⁷⁵ The process by which clickworkers were instructed to label datasets as “disabled” was recounted by workshop participants, based on their knowledge of current industry practice. A similar logic was at work in the Trump administration’s proposal that they would review the public social-media posts of people who receive disability payments and try to verify from their posts if they are really disabled. See Robert Pear, “On Disability and on Facebook? Uncle Sam Wants to Watch What You Post,” *New York Times*, March 10, 2019, <https://www.nytimes.com/2019/03/10/us/politics/social-security-disability-trump-facebook.html>.

⁷⁶ Tom Simonite, “Some Startups Use Fake Data to Train AI,” *Wired*, April 25, 2018, <https://www.wired.com/story/some-startups-use-fake-data-to-train-ai/>.

⁷⁷ Kafer, *Feminist Queer Crip*, 11.

⁷⁸ Louis Schetzer, “Mental Illness and Insurance - Systemic Discrimination That Needs to Change,” *Precedent* 149 (Dec 2018): 45–46, <https://search.informit.com.au/documentSummary;dn=101052307949805;res=IELHSS>.

⁷⁹ Liron Allerhand, Brit Youngmann, Elad Yom-Tov, and David Arkadir, “Detecting Parkinson’s Disease from Interactions with a Search Engine: Is Expert Knowledge Sufficient?,” *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (October 2018): 1539–1542, <https://doi.org/10.1145/3269206.3269250>.

disability community.

Such systems necessarily encode and reproduce a medical model of disability, in which the person being “diagnosed” is pathologized, and context and environment are not accounted for. They also function to construct categories of “disabled” and “non-disabled” based not on people’s lived experience or identity preference, but on assumptions embedded in the data researchers and tech companies have access to. This raises serious questions of validation and accuracy: How are such automated diagnoses confirmed, what are the options to reject or resist such diagnoses, and what are the consequences of being slotted into such a classification, whether accurately or not?

In the case of diagnoses like autism spectrum disorder, the documented racial, cultural, and gender biases in existing nonautomated diagnostic methods are likely to be encoded into automated diagnostic systems in ways that are hard to contest and detect. In one example, a team at Stanford developed a video diagnostic tool that parents could use to receive an autism diagnosis based on video footage of their children. The Stanford team worked with researchers in Bangladesh to apply the tool there. But, as they note, “the people watching the videos [and tagging them for behavior that appeared to indicate autism spectrum disorder] did not speak the same languages as the Bangladeshi children in the videos. Not surprisingly, compared with the team’s prior studies in American children, the algorithms [trained on this tagged data] that predicted which Bangladeshi children had developmental delays were less skewed toward verbal interaction and more toward other behaviors, such as whether the kids made eye contact.”⁸⁰ What such cultural biases and omissions may mean in terms of false positives, false negatives, and the way in which autism is classified across cultures is an open question, but the implications are troubling. They point to a Western-centric view that could exclude those who fall outside of these norms.

Scholars Cynthia Bennett and Os Keyes explore this issue in the context of race and gender, finding that “[d]ependence on diagnostic tools which are based on the experiences of those already diagnosed . . . risk replicating these biases, providing seemingly-objective rigor to determinations that a child presenting inconsistently with (white, assigned male at birth) autistic children cannot be autistic, and should be gatekept out of support systems.”⁸¹

Given that both inaccurate and non-consenting disclosure of a disability can threaten people’s livelihoods and access to medical care (among much else), we need to move beyond a discussion of solving bias by increasing representation to ask larger normative questions: **When is it appropriate or acceptable to make inferences based on data representing (dis)ability? And, who should be tasked with answering this question?** The risk that such information would

⁸⁰ Erin Digitale, “AI-Based Autism Detection Shows Promise across Cultures,” *Scope*, May 8, 2019, <https://scopeblog.stanford.edu/2019/05/08/ai-based-autism-detection-shows-promise-across-cultures/>.

⁸¹ Cynthia L. Bennett and Os Keyes, “What Is the Point of Fairness? Disability, AI and the Complexity of Justice,” *SIGACCESS Newsletter*, no. 125 (October 2019), <http://www.sigaccess.org/newsletter/2019-10/bennet.html>.

be used to exclude and penalize people is high, especially given the relative lack of oversight governing the use of such systems, the incentives that often drive their use (generally profit and efficiency), and the power asymmetries between those in a position to “use” AI and the populations on which AI is “used.”

Designing with, Not For

The credo of the disability rights movement—“Nothing About Us Without Us”—demands that disabled people have ownership of the policies and decisions that affect their lives and communities. In attempting to honor this call in the context of AI, we are faced with many challenges.

AI marketing often invokes disabled people, presenting sanguine images of AI as an assistive technology designed to help people navigate their disabilities and the world beyond.⁸² And indeed, there are many examples of technology that has helped and assisted disabled people in one or another context. However, too often, the processes by which such systems are created don’t include disabled people, and the benefits professed in the marketing frequently fail to materialize, manifest as an afterthought, or exist as “assistive pretext” or “disability dongle.” For example, for years developers have (re)invented systems for American Sign Language (ASL) “recognition” that fail to be comprehensive or useful due, in large part, to the exclusion of Deaf people and people who know ASL from system development and design.⁸³

Academia and grant-making incentives contribute to the pattern of technologies and initiatives making claims to help without actually centering disabled people in the process of research and design (reflecting a charity model of disability). For example, the National Science Foundation requires grant seekers to address the “impact” of their proposed work, but does not follow up with grantees to establish whether and how such impact was ever realized. “Helping disabled people” makes for a good impact statement, even when the end goal is to develop technologies designed for wider commercial markets. Such claims are also used by the corporate sector to burnish their image, positioning tech companies as benevolent. Given that tech is a sector where talent is scarce, this is useful: “Businesses trying to recruit from the limited AI talent pool can use social good projects to convince potential hires that their work will contribute in a meaningful way to the

⁸² See, for example, Brad Smith, “Using AI to Empower People with Disabilities,” *The Official Microsoft Blog*, May 7, 2018, accessed October 21, 2019, <https://blogs.microsoft.com/on-the-issues/2018/05/07/using-ai-to-empower-people-with-disabilities/>; Campbell Kwan, “Google Lookout Uses AI to Describe Surroundings for the Visually Impaired,” *ZDNet*, March 13, 2019, <https://www.zdnet.com/article/google-lookout-uses-ai-to-describe-surroundings-for-the-visually-impaired/>; and Jessica Guynn, “Facebook Taps Artificial Intelligence for Users with Disabilities,” *USA Today*, March, 23, 2016, <https://www.usatoday.com/story/tech/news/2016/03/23/facebook-accessibility-people-with-disabilities/82026554/>.

⁸³ Danielle Bragg et al., “Sign Language Recognition, Generation, and Translation: An Interdisciplinary Perspective,” *Proceedings of ASSETS 2019*, August 22, 2019, <https://doi.org/10.1145/3308561.3353774>.

world.”⁸⁴ However, when we examine the substance behind these claims, it becomes clear that the development of technologies for disabled people often remains a fringe concern within these organizations (even though many of the workers dedicated to these efforts are sincere, and do much with limited resources).

For example, the case of certain telephone by-products (e.g., the sound spectrograph) shows that while AT&T used representations of deaf people for promotional and marketing purposes, writing the story of their work with deaf people into the public record as a PR exercise, the company was not in fact invested in serving this population.⁸⁵ Such assistive pretexts can also be seen in the case of experimental technologies that are first tested on disabled people before being deployed to a mass market. And often, people with disabilities are willing to pay higher prices to access such technologies, and thus serve to subsidize research and development, before the technologies are deemed ready for the public at large.

AI-enabled systems, such as machine vision for automatic lip reading, are often introduced to the general public under the guise of offering service to disabled people—in this case, enabling deaf or blind people to converse more easily.⁸⁶ These feel-good stories of “AI for good” overshadow and distract from the significant risks of such technologies while serving to normalize the presence of these systems in our day-to-day lives. In this case, these harms include invasive surveillance that could jeopardize privacy and access to public space. Such surveillance has implications for disabled people who rely on such devices. As scholar Laura Forlano recounts, such devices may supply life-sustaining support around which disabled people structure their own care practices, while simultaneously supplying for-profit companies with incredibly intimate data, making “privacy” something that disabled people aren’t able to choose.⁸⁷ Poet and disability rights scholar and activist Jillian Weise writes about the way in which such surveillance shapes the lives of disabled people who rely on such devices. In an essay that touches on her use of a prosthetic leg instrumented to collect and transmit data, she says, “I’m told by the technicians to maintain an average amount of walking on a daily basis. Don’t go overboard, but don’t be lazy either. Stay in the middle. The insurance company could pull my data and decide I haven’t used my leg enough to justify the next one.”⁸⁸

These systems also raise the question of bystander privacy, and what consent could look like for

⁸⁴ Jackie Snow, “How People with Disabilities Are Using AI to Improve Their Lives,” *Nova Next*, January 30, 2019, <https://www.pbs.org/wgbh/nova/article/people-with-disabilities-use-ai-to-improve-their-lives/>.

⁸⁵ Mara Mills, “Deaf Jam: From Inscription to Reproduction to Information,” *Social Text* 28, no. 1 (March 2010): 35–58, <https://doi.org/10.1215/01642472-2009-059>

⁸⁶ Robin Christopherson, “Lip-Reading with Google’s DeepMind AI: What It Means for Disabled People, Live Subtitling and Espionage!,” *AbilityNet*, March 2, 2017, <https://www.abilitynet.org.uk/lipreading-google-deepmind-future-disabled>.

⁸⁷ Laura Forlano, “Data Rituals in Intimate Infrastructures: Crip Time and the Disabled Cyborg Body as an Epistemic Site of Feminist Science,” *Catalyst: Feminism, Theory, Technoscience* 3, no. 2: Science Out of Feminist Theory Part 2: Remaking Science(s), October 19, 2017, https://catalystjournal.org/index.php/catalyst/article/view/28843/pdf_13.

⁸⁸ Jillian Weise, “Common Cyborg,” *Granta*, September 24, 2018, <https://granta.com/common-cyborg/>.

those who are analyzed and recorded by such systems. So far, there has been little research examining the issue of bystander privacy, although recent work found that “bystanders are more willing to share some kinds of personal information with [visually impaired people] and are willing to share additional information if higher security assurances can be made by improving their control over how their information is shared.”⁸⁹ But the question of how such assurances could be made, and how such consent obtained, remains unanswered.⁹⁰

According to the World Health Organization, only one in ten disabled people has access to so-called assistive technologies, bringing the issues of inequality and cost of access to the foreground.⁹¹ For the ten percent of disabled people who do have access, such technologies are too often difficult to modify, repair, and control. Historically, disability communities have been major contributors to technical efforts like HAM radio systems, and the design and development of wheelchairs. But the ability to participate in shaping core technologies is unavailable when it comes to AI and many other modern computational systems. Companies selling such systems prevent those outside from accessing, understanding, and modifying them, claiming corporate secrecy. Even extremely expensive technologies are closed off to creative forms of modification, leaving the purchaser at the mercy of designers and developers. This is largely because these technologies are produced by for-profit companies, and governed by commercial incentives that ultimately put profit and growth over providing unfettered access. In addition, such networked technologies frequently expose disabled people to surveillance: data collected from devices and sensors is sent back to the corporations that build and patent these technologies, again raising questions of privacy, autonomy, exposure, and risk.

Even people with disabilities who have the training to develop their own systems are still faced with significant barriers to adoption and use. While international standards were adopted in the mid-1990s, meant to ensure that websites were accessible to disabled people, a 2019 study of the top one million websites by WebAIM found that 97.8 percent of website home pages had detectable failures to comply with the Web Content Accessibility Guidelines.⁹² And a recent investigation into Facebook found that it was chronically accessible for blind users.⁹³ In the AI space, the conferences and environments in which AI is developed are often not designed to include the perspectives and presence of disabled people. This is something that some in the

⁸⁹ Tousif Ahmed, Apu Kapadia, Venkatesh Potluri, and Manohar Swaminathan, “Up to a Limit? Privacy Concerns of Bystanders and Their Willingness to Share Additional Information with Visually Impaired Users of Assistive Technologies,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, no. 3 (September 2018): article no. 89, <https://dl.acm.org/citation.cfm?doid=3279953.3264899>.

⁹⁰ Leah Findlater, Steven Goodman, Yuhang Zhao, Shiri Azenkot, and Margot Hanley, “Fairness Issues in AI Systems that Augment Sensory Abilities,” SIG Access Newsletter, 125 (October 2019), <http://www.sigaccess.org/newsletter/2019-10/findlater.html>.

⁹¹ World Health Organization, “Assistive Technologies Fact Sheet,” May 18, 2018, <https://www.who.int/en/news-room/fact-sheets/detail/assistive-technology>.

⁹² “The WebAIM Million: An Accessibility Analysis of the Top 1,000,000 Home Pages,” WebAIM, accessed October 21, 2019, <https://webaim.org/projects/million/>.

⁹³ April Glaser, “When Things Go Wrong for Blind Users on Facebook, They Go Really Wrong,” *Slate*, November 20, 2019, <https://slate.com/technology/2019/11/facebook-blind-users-no-accessibility.html>

disability community have begun to contest. One such group, {Dis}Ability in AI, describes itself as “supporting and advocating for disabled people” in AI, including providing support at major AI conferences NeurIPS, ICML, and ICLR.⁹⁴ Similarly, programs like Access Computing are working to provide training and computer science skills to disabled people, to address the significant barrier to entry faced by disabled people wishing to participate in creating computational technologies.⁹⁵

APIs and toolchains that are essential for modern software development are often themselves inaccessible, not interoperable with screen readers and failing other basic accessibility requirements. Here, again, disabled people continue to be regarded as consumers and passive recipients, and not as the creators of technology. However, even if tech companies were to make APIs and toolchains more accessible, this wouldn’t address the core issue: such developer technologies are still controlled by the companies who provide them, not by those who make use of them to develop their own projects. Even developer tools and interfaces are consumer-facing products produced by tech corporations, which ultimately control how, whether, and for whom such products function.

This is true for AI systems across the board. Developing and deploying AI at scale requires significant resources, including massive computational infrastructure, large-scale data (and the capacity to continually collect and label such data), and rare and highly trained engineering talent. Access to these resources requires significant capital and market reach, which in practice means that only about five companies in the US right now are able to “create” such AI, and that these systems are ultimately calibrated to serve the interests of the corporations that design and deploy them, whether directly or not. In many cases, AI models are licensed by these corporations to third parties in an “AI as a service” business arrangement. For example, Amazon licenses facial recognition as a service to police across the US, and Amazon, Google, Microsoft, and Facebook all sell such AI services, in most cases without disclosing to whom, and for what purpose, such technologies are being sold.

Given the significant concentration of power in the AI sector, and the significant resources required to build and maintain AI systems at scale, how we might envision “developing with, not for” in the context of AI systems remains a troubling provocation. These systems are currently being used throughout core social institutions, making significant determinations, from who gets a job⁹⁶ to how medical care is allocated⁹⁷ to which school a child can attend.⁹⁸ The people whose

⁹⁴ {Dis}Ability in AI, accessed November 1, 2019, https://elesa.github.io/ability_in_AI/.

⁹⁵ AccessComputing: The Alliance for Access to Computing Careers, accessed November 12, 2019, <https://www.washington.edu/accesscomputing/>.

⁹⁶ Miranda Bogen and Aaron Rieke, “Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias,” December 2018, <https://www.upturn.org/static/reports/2018/hiring-algorithms/files/Upturn%20--%20Help%20Wanted%20-%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf>.

⁹⁷ Jonah Comstock, “UK’s NHS Will Test Babylon’s Triage Chatbot to Replace Non-emergency Hotline,” *MobiHealthNews*, January 4, 2017,

lives are affected, whether they are disabled or not, have little to no control over how the system perceives them, and little ability to push back on a bad or discriminatory determination. Further, even if AI systems were opened for scrutiny, in many cases examining the inner workings of a given algorithm won't provide answers to questions of bias, fairness, and efficacy. In this sense, such systems are essentially "unhackable."

Challenges and Interventions

Centering disability in discussions of AI bias and justice can help us refocus and refine our approach to remediating AI's harms, moving away from a narrow focus on technology and its potential to "solve" and "assist," and toward approaches that account for social context, history, and the power structures within which AI is produced and deployed.

By attending to the ways in which ability has been defined (and redefined) through technologies of measurement and classification, we can begin to develop a deeper understanding of "AI bias" that examines the normative assumptions at the core of AI systems, and attends to whom these assumptions include, and whom they exclude. Such an approach illuminates the way aberrance is often used as a justification for disparate treatment, and connects AI's logics with histories of exclusion (and the pseudoscience constructed to support exclusion). How are the determinations about "worth" and "ability" produced by such systems justified, and what are the consequences of such determinations?

Such questions are difficult, and satisfactory answers will not be produced simply through innovative engineering, or tweaks to data and algorithms. These issues expose the need for approaches to AI bias that go well beyond the technical, and that focus on issues of justice, recognizing the issues raised by AI's rapid proliferation as social and political. How can those harmed by these systems contest the harmful determinations made about them? How do those cast to the margins of normative classification systems contest their marginalization, and what coalitions might such marginalization forge? How do we move from "inclusion" to "agency and control" given the increasingly proprietary nature of the technologies being created, and the centralization inherent in the current form of AI?

Such questions return us to the disability rights movement, and the history of the fight for self-determination and agency in the face of social and political systems that served to justify mistreatment and marginalization. We are reminded that AI systems allocate the power to classify the world to a small handful of companies and privileged individuals. Ed Roberts, a

<https://www.mobihealthnews.com/content/uks-nhs-will-test-babylons-triage-chatbot-replace-non-emergency-hotline>.

⁹⁸ DJ Pangburn, "Schools Are Using Software to Help Pick Who Gets In. What Could Go Wrong?," *Fast Company*, May 17, 2019,

<https://www.fastcompany.com/90342596/schools-are-quietly-turning-to-ai-to-help-pick-who-gets-in-what-could-go-wrong>.

disability rights movement leader, recalls the stakes of such a profound asymmetry between those with the power to label the world, and those being labeled: “If we have learned one thing from the civil rights movement in the US, it’s that when others speak for you, you lose.”⁹⁹

To truly tackle the core justice issues that the rapid proliferation of AI systems raises, the community currently examining AI bias cannot simply adopt the ideas and language of disability studies and activism. Disabled people, along with other affected communities, must be at the center of any approach, defining the terms of engagement, the priorities of the debate, and retelling the story of AI from the perspective of those who fall outside of its version of “normal.”

Further Reading

Disability Studies: Core Texts and Concepts

Shakespeare, Tom. “The social model of disability.” In *The Disability Studies Reader*, edited by Lennard J. Davis, 197–204. 2nd ed. New York and London: Routledge, 2006.

DasGupta, Sayantani. “Medicalization.” In *Keywords for Disability Studies*, edited by Rachel Adams, Benjamin Reiss, and David Serlin, 120–121. New York: New York University Press, 2015.

Siebers, Tobin. “Disability as Masquerade.” *Literature and Medicine* 23, no. 1 (Spring 2004): 1–22.

Zola, Irving Kenneth. “Medicine as an Institution of Social Control.” *The Sociological Review* 20, no. 4 (November 1972): 487–504.

McRuer, Robert. “Compulsory Able-Bodiedness and Queer/Disabled Existence.” In *The Disability Studies Reader*, edited by Lennard J. Davis, 88–99. 2nd ed. New York and London: Routledge, 2006.

Linton, Simi. *Claiming Disability: Knowledge and Identity*. New York and London: New York University Press, 1998.

Linton, Simi. “Disability Studies/Not Disability Studies.” *Disability & Society* 13, no. 4 (1998): 525–539.

Conrad, Peter. “Medicalization and Social Control.” *Annual Review of Sociology* 18, no. 1 (1992): 209–232.

⁹⁹ James I. Charlton, *Nothing About Us Without Us: Disability Oppression and Empowerment* (Oakland: University of California Press, 1998), 3.

Siebers, Tobin. "Disability in Theory: From Social Constructionism to the New Realism of the Body." *American Literary History* 13, no. 4 (2001): 737–754.

Garland-Thomson, Rosemarie. "The Case for Conserving Disability." *Journal of Bioethical Inquiry* 9, no. 3 (2012): 339–355.

Snyder, Sharon L. and David T. Mitchell. *Cultural Locations of Disability*. Chicago: University of Chicago Press, 2010.

Ginsburg, Faye and Rayna Rapp. "Disability Worlds." *Annual Review of Anthropology* 42 (2013): 53–68.

Disability and Technology

Alper, Meryl. *Giving Voice: Mobile Communication, Disability, and Inequality*. Cambridge: MIT Press, 2017.

Baynton, Douglas. "Disability and the Justification of Inequality in American History." In *The New Disability History: American Perspectives*, edited by Paul K. Longmore and Lauri Umansky, 33–57. New York: New York University Press, 2001.

Cryle, Peter and Elizabeth Stephens. *Normality: A Critical Genealogy*. Chicago and London: University of Chicago Press, 2017.

Davis, Lennard J. *Enforcing Normalcy: Disability, Deafness, and the Body*. London and New York: Verso, 1995.

Dolmage, Jay Timothy. *Disabled Upon Arrival: Eugenics, Immigration, and the Construction of Race and Disability*. Columbus: Ohio State University Press, 2018.

Eglash, Ron. *Appropriating Technology: Vernacular Science and Social Power*. Minneapolis: University of Minnesota Press, 2004.

Forlano, Laura. "Hacking the Feminist Disabled Body." *Journal of Peer Production* no. 8 (2016). <http://peerproduction.net/issues/issue-8-feminism-and-unhacking/peer-reviewed-papers/hacking-the-feminist-disabled-body/>.

Gerut, Elianna, Sarah Levin, Daniel Rabinovitz, Gabe Rosen, and Ben Schwartz. "It's Time for a National Museum of Disability." *New York Times*, September 5, 2018.

Hamraie, Aimi. *Building Access: Universal Design and the Politics of Disability*. Minneapolis: University of Minnesota Press, 2017.

Hamraie, Aimi and Kelly Fritsch. "Crip Technoscience Manifesto." *Catalyst: Feminism, Theory, Technoscience* 5, no. 1 (2019).

Hirschmann, Nancy and Beth Linker, eds. *Civil Disabilities: Citizenship, Membership, and Belonging*. Philadelphia: University of Pennsylvania Press, 2015.

Jackson, Liz. "We Are the Original Lifehackers." *New York Times*, May 30, 2018. <https://www.nytimes.com/2018/05/30/opinion/disability-design-lifehacks.html>

Kafer, Alison. *Feminist, Queer, Crip*. Bloomington: Indiana University Press, 2013.

Koester, Nicole. "Choosing the Right Words." National Center on Disability and Journalism, 2019. <https://ncdj.org/wp-content/uploads/2019/01/Choosing-the-Right-Words.pdf>.

Mills, Mara. "Deaf Jam: From Inscription to Reproduction to Information." *Social Text* 28 no. 1 (2010): 35–58. <https://doi.org/10.1215/01642472-2009-059>.

Mills, Mara. "Hearing and the History of Electronics Miniaturization." *IEEE Annals of the History of Computing* 11 (2011): 24–44.

Mills, Mara and Rebecca Sanchez. "Design with Disability," *Public Books*, October 2, 2018.

Ott, Katherine. "Prosthetics." In *Keywords for Disability Studies*, edited by Rachel Adams, Benjamin Reiss, and David Serlin, 140–43. New York: New York University Press, 2015.

Linker, Beth and Whitney Laemmli. "Half a Man: The Symbolism and Science of Paraplegic Impotence in World War II America." *Osiris* 30, no.1 (2015): 228–249.

Lombardo, Paul A. *Three Generations, No Imbeciles: Eugenics, The Supreme Court, and Buck v. Bell*. Baltimore: Johns Hopkins University Press, 2008.

Parker, Alison M. "Intersecting Histories of Gender, Race, and Disability." *Journal of Women's History* 27, no. 1 (Spring 2015): 178–186.

Puaca, Laura Micheletti. "The Largest Occupational Group of All the Disabled: Homemakers with Disabilities and Vocational Rehabilitation in Postwar America." In *Disabling Domesticity*, edited by Michael Rembis, 73–102. New York: Palgrave Macmillan, 2017.

Petrack, Elizabeth R. *Making Computers Accessible: Disability Rights and Digital Technology*. Baltimore: Johns Hopkins University Press, 2015.

Pinch, Trevor and Nelly Oudshoorn. *How Users Matter: The Co-Construction of Users and Technology*. Cambridge: MIT Press, 2005.

Rodas, Julia Miele. *Autistic Disturbances: Theorizing Autism Poetics from the DSM to Robinson Crusoe*. Ann Arbor: University of Michigan Press, 2018.

Rose, Sarah F. *No Right to Be Idle: The Invention of Disability, 1840s–1930s*. Chapel Hill: University of North Carolina Press, 2017.

Samuels, Ellen. *Fantasies of Identification: Disability, Gender, Race*. New York and London: New York University Press, 2016.

Schalk, Sami. *Bodyminds Reimagined: (Dis)ability, Race, and Gender in Black Women's Speculative Fiction*. Durham, NC and London: Duke University Press, 2018.

Schalk, Sami. "Reevaluating the Supercrip." *Journal of Literary & Cultural Disability Studies* 10, no. 1 (2016): 71–86.

Stern, Alexandra Minna. *Eugenic Nation: Faults and Frontiers of Better Breeding in Modern America*. Oakland: University of California Press, 2005.

Trent, James W. *Inventing the Feeble Mind: A History of Intellectual Disability in the United States*. Oxfordshire and New York: Oxford University Press, 2017.

Wendell, Susan. "Unhealthy Disabled: Treating Chronic Illnesses as Disabilities." *Hypatia* 16, no. 4 (2001): 17–33.

Williamson, Bess. *Accessible America: A History of Disability and Design*. Vol. 2. New York: New York University Press, 2019.

Yergeau, Melanie. *Authoring Autism: On Rhetoric and neurological queerness*. Durham: Duke University Press, 2018.

Disability Language Style Guide, National Center on Disability and Journalism (revised 2018): <https://ncdj.org/style-guide/>.

AI Bias and the Politics of Large-Scale Technical Systems

Benjamin, Ruha. *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity Press, 2019.

Bowker, Geoffrey C. and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. Cambridge Massachusetts and London: MIT Press, 2000.

Broussard, Meredith. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge: MIT Press, 2018.

Desrosières, Alain. *The Politics of Large Numbers: A History of Statistical Reasoning*. Cambridge Massachusetts and London: Harvard University Press, 1998.

Eubanks, Virginia. *How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press, 2018.

Franklin, Ursula M. *The Real World of Technology: Revised Edition*. Toronto: House of Anansi Press, 1999.

Gray, Mary L. and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston and New York: Houghton Mifflin Harcourt, 2019.

Hacking, Ian. *The Taming of Chance*. Cambridge: Cambridge University Press, 1990.

Hicks, Mar. "Hacking the Cis-tem," in *IEEE Annals of the History of Computing* 41 (2019) 20–33, <https://ieeexplore.ieee.org/document/8634814>.

Irani, Lilly. *Chasing Innovation: Making Entrepreneurial Citizens in Modern India*. New Jersey and Oxfordshire: Princeton University Press, 2019.

Noble, Safiya. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Penguin Random House, 2017.

Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press, 2015.

Pinch, Trevor J. "Opening black boxes: Science, technology and society." *Social Studies of Science* 22, no. 3 (1992): 487–510.

Rankin, Joy Lisi. *A People's History of Computing in the United States*. Cambridge Massachusetts and London: Harvard University Press, 2018.

Reynolds, Kim and Seeta Peña Gangadharan. "Data Brokers and Opting Out." *Our Data Bodies*, 2019. <https://www.odbproject.org/2019/02/05/data-brokers-and-opting-out/>.

Suchman, Lucy A. *Human-Machine Reconfigurations: Plans and Situated Actions 2nd Edition*.
Cambridge: Cambridge University Press, 2007

Workshop Participants

Meryl Alper, Northeastern University
Shiri Azenkot, Cornell University
Mary Ballard, Microsoft Inc.
Cynthia Bennett, University of Washington
Jeff Bigham, Carnegie Mellon University
Wendy Chisholm, Microsoft Inc.
Kate Crawford, AI Now Institute
Roel Dobbe, AI Now Institute
R. Luke DuBois, New York University
Leah Findlater, University of Washington
Jim Fruchterman, Benetech
Faye Ginsburg, New York University
Gerard Goggin, University of Sydney
Sara Hendren, Olin College
Matt Huenerfauth, Rochester Institute of Technology
Amy Hurst, New York University
Liz Kaziunas, AI Now Institute
Walter Lasecki, University of Michigan
Megan Lawrence, Microsoft Inc.
Varoon Mathur, AI Now Institute
Joanna McGrenere, University of British Columbia
Mara Mills, New York University
Sarah Myers West, AI Now Institute
Karen Nakamura, University of California, Berkeley
Meredith Ringel Morris, Microsoft Research
Emily Rogers, New York University
Marcel Salas, New York University
Jutta Treviranus, OCAD University
Shari Trewin, IBM Inc.
Meredith Whittaker, AI Now Institute

