# Learning from the Front: People with Disabilities as Early Adopters of AI

Jeffrey P. Bigham and Patrick Carrington
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15217

Keywords: AI; HCI; Accessibility; Disability

## Abstract

People with disabilities have been among the earliest adopters of Artificial Intelligence (AI) through interactive tools that they use every day. Speculating about how we might adopt AI systems and what problems we might face is difficult. Furthermore, the promise and practice of artificial intelligence are often inconsistent; making claims beyond current capabilities and utility. Instead, we can learn a lot by considering the concrete experiences of people with disabilities who have directly addressed notions of agency while working with AI, being excluded by AI, and trusting AI. As AI is finally approaching the point where many of us may interact with it in our lives, we believe there is a lot to learn from these experiences.

## Introduction

People with disabilities have been among the earliest adopters of AI in interactive tools that they use in their daily lives. Speculating about how we might adopt AI systems and what problems we might face is difficult. Instead, we can consider the concrete experiences of people with disabilities, who have directly addressed notions of agency while working with AI, being excluded by AI, and trusting AI. As AI is finally approaching the point where many of us may interact with it in our lives, we believe there is a lot to learn from these experiences.

## People with Disabilities as Early Adopters

Artificial Intelligence (AI) technology has been motivated by its possible use by people with disabilities since nearly the beginning. Despite the name “intelligence”, much of what AI has been about is perception, and many disabilities affect people’s ability to perceive. Not surprisingly, much of the early work in computer vision focused on applications intended to help people who are blind access a very visual world. While speech recognition technology had found some niche use cases before the recent boom, one of its user groups has been people who cannot type due to, for instance, a motor impairment. Very reasonably, we might assume a grand challenge for the perceptual side of AI should be to provide access to sensory information that someone is unable to access otherwise.

Yet, very commonly the promises of utility to people with disabilities of AI are overstated, not based on real user needs, and often contain profound misunderstandings of user capabilities and needs. For instance, countless applications have promised to tell blind people what the objects are in front of them, even though they can easily discern these identities based on context and touch. Every few months, a new “intelligent cane” is produced offering to imperfectly tell a blind user what they can either already sense much better with the affordances of a cane, or about far away objects that are of little practical use, all the while making the cane unwieldy and more expensive. Sign language recognition projects claim the ability to translate sign language, when in fact they are able to recognize a handful of carefully made individual signs, ignoring the complexities of the language. Mainstream users are now experiencing the exaggerated claims of AI, which although advancing, often fails to live up to the hype.
Gradually, we are seeing the accessibility community put together research that exposes the use cases that matter for people with disabilities. Sometimes this happens via new technological approaches pioneered by us and others that allow Wizard-of-Oz systems to be deployed (Bigham 2010). For instance, our VizWiz crowd-powered visual question answering application collected more than 100,000 visual questions, photos, and answers for blind users, and has recently been turned into a computer vision dataset to appear as a paper at CVPR (Gurari 2018), with the goal of enabling computer vision researchers to work on automatically answering the difficult, real visual questions that blind people have in their everyday lives.

## What happens when you’re expected to trust imperfect AI?

A challenge with putting AI into user interfaces is that AI will make errors, and these need to somehow be conveyed appropriately to users. In AI fields, errors tend to be aggregated along with successes into accuracy numbers. In a study we conducted, we noticed that speech recognition systems often produce nonsensical errors when compared to humans (Lasecki 2012). In one example of transcribing a computer science lecture, both human transcribers and speech recognition made errors. An example of a human error was substituting “someone” for “somebody”, whereas the speech recognition system substituted “in pectoral” for “n-factorial”, yet both errors are treated the same by common metrics in these domains, such as BLEU score and Word Error Rate (WER). People who rely on transcripts do not trust speech recognition and have advocated strongly for human transcription despite its high cost.

Several companies have now experimented with using computer vision to provide labels for social media images. For instance, Facebook now provides automatic labels for all of its images (Wu 2017). A study from Microsoft Research explored how blind users interpreted such labels, especially considering what happens when the descriptions were incorrect (MacLeod 2017). In one example, a picture of Hillary Clinton at a campaign rally was described as, "I am not really confident, but I think it's a man is doing a trick on a skateboard at night.” Despite the stated lack of confidence, some participants came up with justifications for why such a picture would appear on Clinton’s Twitter feed, e.g., maybe she was trying to appear young and hip.

What will all of us do when our interfaces may no longer tell us the truth? How can we design systems that let us take advantage of AI, while not misleading us because of the AI getting it wrong? We are already seeing this sort of problem with AI lulling us into a false sense of security in

self-driving cars[1], and AI helping to convince us that fake news or extreme positions are mainstream perceptions[2].

## Agency and AI

People with disabilities have long fought to have agency over their own lives. It might be no surprise then that people with disabilities have approached how they interact with AI a bit differently than others. For instance, blind people have been interested in the idea of self-driving cars for a long time. Yet, in at least early work in this area, they did not want the cars to actually drive for them; they wanted the cars to be drivable by them[3]. This builds on thinking that differentiates assistive technology from access technology: the former takes the stance on assisting or helping someone do something, whereas the latter takes the stance of providing access to a world that is lacking in some way, given the user's context and abilities. Given the recent news of self-driving cars, it may make sense to revisit this idea of agency, who is in control, and what role the AI serves.

Related examples occur in privacy and technology use. For instance, we observed that many VizWiz users would take pictures that one might consider private, with the tradeoff of getting something done. We know that many people with disabilities trade off privacy for independence numerous times in their lives, for everything from help getting to the correct bus, to giving a pin number to a grocery clerk because the checkout kiosk is inaccessible (Ahmed 2016). For the most part, people with disabilities know when they are making such tradeoffs, and are much more comfortable with making the choice than they would be with trusting an always observing assistant. Yet, we wonder to what extent they are truly making a choice.

We also find a myriad of tradeoffs with regard to assistive robotics and autonomous wheelchairs for people with motor impairments. Advances in AI fuel the capability of robotics to support physical tasks, such as those involved in mobility or interacting with the environment. While AI has a profound potential to enable robots to automate tasks, there is still the issue of the extent to which the user remains in control of the process. Autonomous wheelchairs leverage AI but involve important tradeoffs between a user's mobility, agency, and sense of independence.

Similarly, as systems that use AI become pervasive, to what extent will we have a choice, and to what extent will we know that we are making a choice? For instance, it seems plausible that many Facebook users do not realize the extent to which their interaction on the platform is being mediated by AI. For many services, it is unclear whether a choice is truly being made.

## What happens when AI limitations exclude you?

AI also has the potential to exclude, as people with disabilities are finding. Many smart devices are now being sold with voice as the primary (and often only) interaction modality. This includes devices like Amazon Echo and Google Home, but also one-off products, such as light bulbs,

---

[1] https://www.bloomberg.com/news/articles/2018-03-19/uber-crash-is-nightmare-the-driverless-world-feared-but-expected

[2] https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html

[3] http://www.blinddriverchallenge.org/

toasters, headphones, electrical outlets, in-car systems, etc. Thus these devices exclude people who cannot speak or hear. Recently, we have created an application that speaks on behalf of Deaf and hard of hearing people, so they can do things like turn on the lights in a place outfitted with these "smart" devices (Bigham 2017). What will happen if such exclusionary devices become pervasive?

Far from a question only for people with disabilities, we have already seen AI systems developed that exclude people. For instance, speech recognition systems work poorly for people who do not speak the target language or speak it with an accent. They work worse for women[4]. We have famously seen an HP camera fail to recognize black faces[5], and recently seen the iPhone fail to tell some Chinese faces apart[6]. AI systems are being actively developed with the goal of differentiating people, for instance attempting to tell whether someone is gay by their photos (Wang 2017).

AI increases our abilities to exclude, both on purpose and because of bias in the systems. While it can be difficult to know what it might exclude based on next, the experiences of the past are likely informative of how it will happen. Furthermore, labels have recently resurfaced as a point of contention in society. For AI, a technology class that continues to excel at labeling, this is in direct conflict with the impetus put forth to be more conscientious about the labels that people use in society. Despite the best of intentions, the complex meanings that labels hold, and their various interpretations, may be used for either benevolent or maleficent means.

## Summary

AI will soon be in most of our interactive technology. This offers many opportunities for benefit, but, as the experiences of people with disabilities have shown, many opportunities for harm. We would be smart to build from the lessons learned from those who have led the way in adoption of intelligent interactive systems -- people with disabilities.

---

[4] https://www.dailydot.com/debug/google-voice-recognition-gender-bias/

[5] https://gizmodo.com/5431190/hp-face-tracking-webcams-dont-recognize-black-people

[6] https://gizmodo.com/how-apple-says-it-prevented-face-id-from-being-racist-1819557448

## References:


Ahmed, T., Shaffer, P., Connelly, K., Crandall, D., and Kapadia, A. Addressing Physical Safety, Security, and Privacy for People with Visual Impairments. SOUPS 2016.

Bigham, J.P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R.C., Miller, R., Tatarowicz, A., White, B., White, S., and Yeh, Tom. VizWiz: Nearly Real-time Answers to Visual Questions Bigham. UIST 2010.

Bigham, J.P., Kusahlnagar, R., Huang, T.H., Flores, J.P., and Savage, S. On How Deaf People Might Use Speech to Control Devices. ASSETS 2017.

Gurari, D., Li, Q., Stangl, A.J., Guo, A., Lin, C., Grauman, K., Luo, J., and Bigham, J.P. VizWiz Grand Challenge: Answering Visual Questions from Blind People. CVPR 2018.

Lasecki, W.S., Miller, C., Sadilek, A., Abumoussa, A., Borrello, D., Kushalnagar, R., Bigham, J.P. Real-time Captioning by Groups of Non-experts. UIST 2012.

MacLeod, H., Bennett, C., Morris, M.R., and Cutrell, E. Understanding Blind People's Experiences with Computer-Generated Captions of Social Media Images. Proceedings of CHI 2017.

Wu, S., Wieland, J., Farivar, O., Schiller, J. Automatic Alt-text: Computer-generated Image Descriptions for Blind Users on a Social Network Service. CSCW 2017.

Wang, Y., and Kosinski, M. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. PsyArXiv, 2017-10-16.