The Institute for Ethical AI

2020 DISABILITY AND AI WHITEPAPER

Recruitment AI has a Disability Problem

Questions Employers Should be Asking to Ensure Fairness in Recruitment

AUTHORS AND CONTRIBUTORS

Chara Bakalis, School of Law, Oxford Brookes University
Nigel Crook, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Paul Jackson, Oxford Brookes Business School, Oxford Brookes University
Kevin Maynard, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Arijit Mitra, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Selin Nugent, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Jintao Long, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Susan Scott-Parker, Business Disability International
James Partridge, Face Equality International
Rebecca Raper, Institute for Ethical Artificial Intelligence, Oxford Brookes University
Alex Shepherd, Institute for Ethical Artificial Intelligence, Oxford Brookes University

SUMMER 2020

This paper is intended to be a guide and living document that will evolve and improve with input from readers and relevant stakeholders. Your feedback is welcome and encouraged. Please share your feedback with us at ethicalAI@brookes.ac.uk .

Cite as:
Nugent, S., Jackson, P., Scott-Parker, S., Partridge, J., Raper, R., Shepherd, A., Bakalis, C., Mitra, A., Long, J., Maynard, K., and Crook, N. (2020). Recruitment AI has a Disability Problem: questions employers should be asking to ensure fairness in recruitment. Institute for Ethical Artificial Intelligence

The Institute for Ethical AI

OXFORD BROOKES UNIVERSITY

business disability international

I Table of Contents

| About the Institute for Ethical Artificial Intelligence

Oxford Brookes University hosts a vibrant and ambitious research environment in the areas of artificial intelligence, computing, and data science. Founding the Institute for Ethical Artificial Intelligence was therefore a natural extension of our vision as a research community to advance knowledge and promote the better understanding of technology and its relationship to business and society in our local and the wider global community. Our mission at the Institute for Ethical Artificial Intelligence is to promote and support the development and deployment of ethical and trustworthy intelligent software solutions for business, organisations, and society.

Our primary focus at the Institute for Ethical Artificial Intelligence is to help organisations working in the professional services to understand and plan for the risks and opportunities that AI and data analysis technologies can bring to their organisation, their stakeholders and society at large. Working with both the users and the providers of AI technology, as well as developing bespoke AI solutions, we research and advise on the ethical impact of AI technology on organisations and individuals.

In order to achieve this, we bring together a diverse group of world-leading experts who together blend knowledge and skills from technology, business, social science and the life sciences. We deliver expertise and independent guidance in areas that include AI and machine learning, disability, psychology, business development, equality and diversity, coaching and mentoring, digital health, and wellbeing.

For more information, please visit our webpage ethical-ai.ac.uk

| Introduction

The purpose of this White Paper is to

→ Detail the impacts to and concerns of disabled employment seekers using AI systems for recruitment, *and*

→ Provide employers with the knowledge and evaluation tools to ensure innovation in recruitment is also fair to all users.

In doing so, we further the point that ***making systems fairer for disabled employment seekers ensures systems are fairer for all.***

…

Artificial Intelligence (AI) and similar advanced data analytics systems are increasingly sought-after tools for recruitment used to automate time-consuming, repetitive operational tasks, and expand strategic potential. However, as engineering of these systems becomes more complex, it is more difficult for organisations to confidently assess whether the technology is functioning in line with their expectations and if employment seekers will be treated fairly.

AI technologies have the potential to dramatically impact the lives and life chances of people with disabilities seeking employment and throughout their career progression. While these systems are marketed as highly capable and objective tools for decision making, a growing body of research demonstrates a record of inaccurate results as well as inherent disadvantages for women and people of colour (Broussard, 2018; Noble, 2018; O'Neil 2017). Assessment of disability fairness in Recruitment AI has thus far received little attention or been overlooked (see Guo et al., 2019; Petrick, 2015; Trewin, 2018; Trewin et al. 2019; Whittaker et al., 2019).

Presently, a landscape of limited regulation, paired with increasing societal pressure for AI and data analytics systems to be designed with fairness, transparency, and validity, means that organisations face financial, legal, reputational, operational, and ethical risks for implementing them. While there is already much work being done to address the high-level concerns related to artificial intelligence, bias, and fairness, there will inevitably be more challenges ahead that no one company or industry can solve alone. In order to minimise these risks, businesses, human and disability rights campaigners, and academic experts need to collaborate to develop new ways to analyse, validate, and improve these systems and to hold technology developers and suppliers accountable.

Our aim in this paper is to provide a starter toolkit to evaluate organisational and ethical values in relation to the use of recruitment technology, and with regard to vitally important procurement processes. We review the broad technological developments that support recruitment, demonstrate their potential to impact disabled employment seekers in various ways. We then present recommendations for the questions employers should be asking before taking on new technologies and when evaluating currently used systems.

The Institute for Ethical Artificial Intelligence and its partners invite public, third sector and private sector stakeholders to respond to this guidance and to continue discussion toward ensuring fairer recruitment practices for persons with disabilities, and other disadvantaged employment seekers more generally.

I Disability and Employment Discrimination

People with disability have historically and continue to be regularly disadvantaged in seeking and securing employment. Disabled people experience widespread economic and societal exclusion and are *more than twice as likely to be unemployed* as others (Office of National Statistics, 2019). The sheer scale of the social and economic impacts of the COVID pandemic on employment and employability will undoubtedly further disenfranchise people with disabilities. The current climate of instability makes ensuring fair and equal treatment all the more important, given increasing employment among people with disabilities helps raise people out of poverty, improve their life chances, and is a net cultural and economic benefit.

As defined by the United Nations Convention on the Rights of Persons with Disabilities (CRPD), “persons with disabilities include those who have long-term physical, mental, intellectual or sensory impairments which in interaction with various barriers may hinder their full and effective participation in society on an equal basis with others.”

The definition of disability doesn’t necessarily capture the complexity and heterogeneity of people with disabilities, which is a key factor in the complications with AI systems. A disability may be a life-long condition or occur at different life stages or be the result of a major event/change. Disability can have wide-ranging life impacts or be context dependent. Disability may be visible, but most are invisible. Disabilities may include people with hearing, sight and mobility, and dexterity impairments, people with cognitive and intellectual impairments, those with mental health conditions, those with facial disfigurements, those of small stature, and numerous others. Further, individuals may have a combination of multiple factors.

Disability also intersects with other aspects of identity, such as gender, ethnicity, sexuality, an socioeconomic background. Disability is not completely independent of other features of a person's identity and life experience (Collins and Bilge 2020; Parker, 2015; Samuels, 2016). Moreover, the social stigmas attached to disability are intersectional, shared, and amplified with other marginalised identities (Frederick and Shifrer, 2019). In light of the ongoing Black Lives Matter protests against racial violence and injustice, our focus on disability is intended to contribute to a wider discussion of systemic and persistent oppression of marginalized peoples. Recognising and celebrating human diversity is a necessary starting point to design AI systems that fairly and equitably engage with human reality.

…

Disability inclusion in the workplace is impacted by number of factors. There is often a qualifications gap between disabled and non-disabled people due to systematic disadvantages in education, training, and previous work experience (Sayce, 2011). Even well intentioned employers may struggle to recognize how structural barriers to success impact

Some industries or categories of position lack accessibility that can limit employment for people with certain impairments. There are inadequate programmes to support persons with disabilities and those who employment them. Employers may also have negative attitudes/bias and lack confidence or training to support disabled employment seekers (Lindsay et al., 2020; Suter et. al. 2007).

Some of Global Disability Facts

There are more than 1.3 billion people with disabilities worldwide and the number is growing with an aging population and advances in medical science. *WHO*

…

80% of this 1.3 billion people live in the developing world. *WHO*

…

15-18% of any country's population will have a disability and/or chronic health condition *WHO*

…

Circa 80% of people with disabilities have impairments that are not immediately visible. Check

…

1 in 5 women will have a disability UN

…

1 in 3 people aged 50-64 will have a disability regardless of their ethnicity Check

…

People who live to the age of 70 are likely to have at least 10 years of lived experience of disability Check

…

In any large organisation, from 10-12% of the workforce are like to have a disability and/or chronic health condition UK Labour Force Survey

…

At least 1 in 3 consumers will either be disabled or will have someone who has a disability in their immediate circle European Commission

I Recruitment AI

As organisations increase in scale and receive larger volumes of job applicants, they are under pressure to balance often competing interests in recruiting and retaining the talented candidates, optimising workflow efficiency and productivity, and managing costs. This means that employers are increasingly turning to automated tools to support the employee’s journey from recruitment to retirement.

Artificial Intelligence (AI) has featured prominently in these developments. AI is a subfield of computer science, focused on training computers to perform traditionally human tasks. For additional reference, a glossary of relevant AI terms is provided at the end of this document.

AI systems are currently available across a wide range of recruitment functions, including:

→ Candidate Sourcing / Engagement
→ Candidate Tracking
→ CV/ Resume Screening
→ Pre-Employment Assessments
→ AI Interviewing

We will discuss each of these categories of technology in relation to their potential to impact people with disabilities in greater detail below.

The unifying objective for systems operating across these diverse recruitment functions is that they are designed to distil the vast array of information about applicants down to a few select predictable features for the purpose of making quantifiable and easily comparable decisions. However, when systems need to cope with the reality of human diversity, whether it pertains to disability, ethnicity, gender, and/or other features, they often interpret complexity as an abnormality, or outlier. In this case predictability may come at the expense of the life chances of disabled people who are already faced with systematic disadvantages and unfair discrimination in securing employment.

| Exclusion by Design and Discriminatory Use

Recruitment AI may inadvertently adversely impact employment seekers with disabilities via two major routes: biased systems and discriminatory processes.

Biased Systems

The design of an AI system involves first specifying an objective and then specifying how the system achieves and optimizes achieving that objective. Humans are often not skilled at specifying objectives. If an objective is not specified appropriately, the outcome may have unintended consequences.

Unwanted biases, or biases that treat some people negatively, or adversely due to protected characteristics or other features of their identity, raise serious risks of discrimination. It is critical to identify and mitigate these potentially harmful biases. And to prevent and mitigate bias, it is necessary to understand how humans introduce biases into an AI system.

Developing this knowledge begins with defining what biases exist within a system *and* where they exist, or have a potential to exist. Disability-related biases in AI systems are heavily influenced by historical hiring decisions. Since people with disabilities are twice as unlikely to be unemployed, they are simply less likely to be represented in data on past successful employees. These biases may be introduced into systems through two primary mediums: the *algorithmic model* and the *training data*.

...

The algorithmic model is the mathematical process by which an AI system performs a certain function. Designing this model involves defining the objective or problem the developer wishes to address and selecting the parameters that define the system's operation at what they determine is an optimal level (Russell and Norvig, 2003).

How can this go wrong? For instance, an automated CV screener is programmed to predict the best qualified candidate based on the ("optimal") parameter of having attended a top-tier university. Someone who has worked hard to achieve success, right? The prestige of an institution may be one factor in a successful employee, but that parameter also disadvantages people with disabilities, different socioeconomic backgrounds, and/or underrepresented ethnicities, who already face systemic barriers to be equally represented in prestigious institutions.

The training data is an initial set of data used to help a program learn how to apply the model and produce sophisticated results in application (Russell and Norvig, 2003). The model operates as well as the training data that goes in. The sampling strategy used to collect the training data and the representativeness of the data is a conscious decision by the developer.

Building on the previous example, what if the automated CV screener was trained on data that **did not** include the data profiles of successful employees who have a non-English name, went to state school, participate in disability-related volunteering activities, had a break in employment due to family or illness, or have an address in an economically disadvantaged area? These are simple, seemingly innocuous features that will be represented in a CV. Interacting with information in a CV that the programme has not previously encountered means that the system may be more likely to reject a candidate. This is because these novel features do not fit the prescribed collection of features that is modelled to represent the 'ideal' employee. These novel features may be innocuous, but they may also be indirectly related to the experience of being disadvantaged on the job market.

Improper Implementation and Use

Even as systems become more technically sound with regard to acknowledging and mitigating bias in design, risks for applicants with disability may be generated and/or amplified by improper use and implementation of the technology.

Most recruiters recognize that no single assessment method is suitable and fair for all applicants. However, the marketed reliability and the ease of automated adaptations of recruitment processes has resulted in many cases where AI tools are being used in isolation of other measures of suitability and human decision makers in the application package. In some organisations, a single product may be the sole gate of entry into employment.

Moreover, AI assessment fails to factor in the likelihood that the employer would make the adjustment post job hire that would *determine* if a particular disabled candidate was 'right' for the job. For example, a qualified, visually impaired, cybersecurity expert will only be the best candidate *if* the employer enables her to use specialized software.

Acknowledging and monitoring uncertainty in AI systems is critical to making fair and adequate decisions as sensitive and life changing as whether a person is employed or not. The life chances of job seekers precariously intersect with the computational complexities related to disability, the inherent challenges of bias, and the uncertainty around automated decision-making. No system should be expected to work perfectly.

...

The use of rigid, standardised recruitment processes that cannot be adequately adjusted to enable candidates with disabilities to compete fairly are inherently discriminatory (Hamraie, 2017). Candidates may have the option to request accommodations to these systems – although some developers expect this is the role of the employer to deliver such adjustments. However, unless candidates are given explicit assurances that they may request and be provided with equally-evaluated, alternative routes, the employer risks, at best, making disabled users uncomfortable/fearful of interacting with AI and, at worst, discriminating against such candidates. Expecting disabled employment seekers to go through standardised processes is akin to asking a wheelchair user to take the stairs to the interview room.

| Tech on the Market: the dangers of discrimination

Recruitment AI encompasses a wide array of technologies functioning at different points in the recruitment process. This section outlines the broad categories currently in use, detailing the impact potential for people with disabilities. This list is by no means exhaustive, but highlights major technologies used in the candidate sourcing and selection phases of recruitment.

ATS and CRM Systems

Applicant Tracking Systems (ATS) are platforms where recruiters can conduct each step in the hiring process from posting position openings to collecting applications to screening candidates to evaluation and selection. Candidate Relationship Management (CRM) systems maintain a connection between recruiters and employment seekers so that desirable candidates may be easily referred to future job openings.

We consider these systems together because share similar potential impacts on people with disabilities. They are likely to utilise automated outlier detection tools, such as CAPTCHAs, that when insufficiently trained can flag people with disabilities as not human, or a spammer (Guo et al. 2019). The difference between human and non-human may come down to a few seconds delay in response, a minor slip in highlighting the correct answer, or misinterpreting an obscured set of letters. People with difficulties related to dexterity or visual impairment are disproportionately affected.

Further, the skills and qualification gap for disabled people due to systemic inequalities likely disadvantages these candidates when evaluated against the standard person specification as well as historic hiring decisions. These systems are not designed with the flexibility that would take into account that some candidates appear less qualified only due to systemic denial of education and employment opportunities.

CV/ Resume Screeners

CV screening is a major driver of the recruitment innovation powered by AI systems, addressing the need for processing high application volumes. Automated screeners detect characteristics in the CV content, such as key phrases, proper nouns to evaluate employability against criteria for the position. These criteria are determined by either the job description or by evaluating the features of previously successful candidates. They may go further to interpret characteristics of the applicant, such as personality, sentiment, and demographics. Some also supplement data in CVs to with information about the candidate from public data sources, social media, and information about their previous employers.

Once again, the skills and qualification gap for disabled people due to systemic inequalities is likely to disadvantage these candidates when evaluated against a standard job description as well as historic hiring decisions. These systems are not designed with flexibility that considers some appear less qualified due to systemic lack and denial of education and employment opportunities.

AI screener systems that have not been trained on CV data from users with diverse cognitive and intellectual abilities may have additional challenges with linguistic flexibility. For screeners that analyse personality and emotion from texts, further problems may arise. For example, people with neuro- and cognitive diversity may express emotion in writing in a style previously not encountered by the AI system, resulting in incorrect classifications about their emotional state or personality. . And many pre-lingually Deaf individuals speak the official spoken language of their country as a second language.

Conversational Agents

Recruitment conversational agents, or chatbots, are designed to mimic human conversational abilities during the recruitment process. These technologies use an approach termed natural language processing (NLP) to analyse questions and comments and to respond effectively. Conversational agents are desirable additions to the recruitment process as a means of increasing communication with employment seekers in order to answer frequently asked questions, collect information on candidates, ask screening questions, and schedule interviews or meetings with a human recruiter.

Conversational agent systems have the potential to be helpful in some circumstances where they are designed with accessibility in mind. Agents that augment text with visual illustration (i.e. highlight key words, spelling and grammar check, text suggestion), speech functionality, and dictation tools can enhance accessibility and usability for a wide range of users.

However, if not thoughtfully designed and implemented, conversational agents may also not respond appropriately, or in a hateful manner, and unfairly screen out candidates. Depending on the nature of the agent's function this can at best lead to poor user experience and at worst discriminatory candidate screening.

Conversational agents are often not trained on language data gathered from people with cognitive, intellectual, physical and linguistic diversity or those from neuro-diversity groups. Undertrained agents may be unable to correctly interpret spellings or phrases they haven't previously encountered, such as messages from people who have physical difficultly typing or have dyslexia, autism, dysphagia, dyspraxia, ADHD, among numerous others. Moreover, agents that do not support communications methods beyond writing, such as text-to-speech and dictation, limit or exclude many individuals from participating in communication and being competitive in the recruiting process.

Pre-Employment Assessments

A range of candidate aptitude assessments, such as cognitive ability, technical skills, personality, and decision making, are a commonly used to quantitatively measure and compare job applicants for a particular role. Broadly, these tests are aimed at gauging a candidate's ability to think quickly, solve problems, and interpret data.

Many recruiters recognise that these assessments are often not reliable as one-size-fits-all approaches. The generalisability of psychometric tests for people with disabilities—as well as many populations who are not from WEIRD (western, educated, industrialized, rich, and democratic)— backgrounds is unreliable (Cook and Beckman, 2006). There is a degree of uncertainty about whether any assessed candidates, never mind those with disabilities, are indeed able to successfully learn and perform the duties of the role or not. Furthermore, many psychometric tests are in themselves inaccessible to a wide range of disabled candidates. These assessments must be balanced by other measures in the recruitment process.

Gamified assessments raise additional concerns related to dexterity, vision impairment, and response time. Games often involve tasks that are assessed based on speed of reaction to prompts and precision of responses, which may affect people with motor limitations, who need extra time or assistance to complete dexterity tasks. People with visual impairment may require magnification and colour adjustment and additional time. Furthermore, people with cognitive diversity may require language adjustment and additional time to read prompts.

AI Interviewing

AI powered interviewing includes facial analysis tools and speaking conversational agents—aka robot recruiters (refer above to limitations of Chatbots). These tools evaluate employability from the language, tone, and facial expressions of candidates when they are asked an identical set of question in a standardised process. Candidates are assessed based on a variety of facial, linguistic, and non-verbal measures. 'Ideal' measures often are those that most closely align with the same measures from historically successful candidates for any given role.

As with previous examples, systems that are not trained on a diverse range of potentially successful candidates, face challenges in fairly assessing people with facial features, expressions, voice tone, and non-verbal communication that it has not previously encountered.

For instance, facial analysis software may inaccurately assess and potentially exclude people with facial disfigurement or paralysis as well as conditions such as Down syndrome, achondroplasia, cleft lip/palate, or other conditions that result in facial differences. Further, people with blindness may not face the camera or make eye contact in a manner acceptable to the system's parameters. Moreover, issues may exacerbated by differences in eye anatomy and dark glasses. People who need captions due to hearing loss, or who lip read may struggle to hear or interpret the questions.

Facial analysis tools that go further to interpret emotion and personality from facial expressions pose alarmingly high risks. Beyond issues of accuracy and algorithmic bias, the fundamental scientific concepts behind personality assessments derived from facial feature measurements, is not supported -and is rooted in pseudoscientific race studies (Noble, 2018). The implementation of these technologies for recruitment risks legitimising the flawed methodological premise with ill-informed buyers in a way that can only perpetuate historic disadvantages and exclusion for marginalised peoples.

| Intervention Recommendations

Designing and implementing Recruitment AI systems that treat persons with disabilities and by that extent, ***all*** employment seekers fairly requires the engagement of all stakeholders—technology suppliers, purchasers, and users alike. Our aim is to facilitate purchasers in joining the discussion and to collaborate with us to prepare the tools and language needed to initiate the conversation that asks: How do we assess if any given Recruitment AI system is‘safe’ for employment seekers with disabilities and others disadvantaged in any labour market?.

There are a number of actions a forward-thinking organisation can take to support those technology suppliers who share the values and expectations of the organisation and its clients toward applicants with disability. This process begins by asking the right questions of technology developers and suppliers.

Vision, Strategy & Corporate Governance Stakeholders

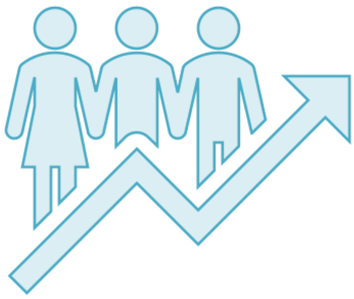

i. Does this technology align with our organizational strategy to increase diversity and representation?

ii. Does use of this technology reflect our organisation’s strategic policies with regard to the ethical and responsible development and implementation of artificial intelligence?

iii. Is this supplier actively engaged in learning more about how to adapt to match our values and needs as a business and those of our stakeholders?

iv. Who in this organisation should be involved in the governance process which determines how we investigate, procure, apply and monitor HR tech systems so that at the very least they do not adversely impact disadvantaged job seekers?

Human Resources and Operations Stakeholders

i. What are the benefits and risks of this technology for disabled and other disadvantaged employment seekers?

ii. Was a shared understanding of inclusivity and fairness—with specific reference to eliminating the root causes of disability related discrimination—designed into this technology?

iii. Will implementing this technology require alternative evaluation routes to enable people with different impairments to be recruited on the basis of individual capability and potential?

iv. Does the AI recruitment tool enable candidates to readily request adjustments, in a non-stigmatising manner, at every stage of the process

Procurement Stakeholders

i. Has this supplier proved their products are safe for disabled and other disadvantaged employment seekers before you purchase?

ii. How has the supplier actively involved people with disabilities to test and validate its products?

iii. Was a shared understanding of inclusivity and fairness—with specific reference to eliminating the root causes of disability related discrimination—designed into this technology?

iv. Do contractually defined performance standards require the supplier to track the experience of job seekers with disabilities – particularly those who have requested disability related adjustments?

v. Can they evidence that they have actively consulted and involved persons with disabilities as expert advisors and potential users in their product development life cycle?

Information Technology Stakeholders

i. Will our organisation be provided with the appropriate explainability and interpretability resources to assess outputs and impacts on employment seekers' disabilities?

ii. Does the relevant, quality data exist to support this technology in performing effectively for persons with disabilities?

iii. What are the appropriate oversight mechanisms to evaluate the performance of the system and can the system withstand scrutiny by disabled employment seekers?

iv. Can the supplier demonstrate how the processes will adapt so as to ensure equal opportunities for disabled employment seekers?

| References

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

Collins, P.H. and Bilge, S., 2020. *Intersectionality*. John Wiley & Sons.

Cook, D.A. and Beckman, T.J., 2006. Current concepts in validity and reliability for psychometric instruments: theory and application. *The American journal of medicine*, *119*(2), pp.166-e7.

Frederick, A. and Shifrer, D., 2019. Race and disability: From analogy to intersectionality. *Sociology of Race and Ethnicity*, *5*(2), pp.200-214.

Guo, A., Kamar, E., Vaughan, J.W., Wallach, H. and Morris, M.R., 2019. Toward Fairness in AI for People with Disabilities: A Research Roadmap. *arXiv preprint arXiv:1907.02227*.

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

Lindsay, S., Leck, J., Shen, W., Cagliostro, E. and Stinson, J., 2019. A framework for developing employer's disability confidence. *Equality, Diversity and Inclusion: An International Journal*.

Parker, Alison M. 2015 "Intersecting Histories of Gender, Race, and Disability." Journal of Women's History 27, 1.

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

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

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

Office for National Statistics, 2019. *Disability And Employment, UK: 2019*.

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

Russell, S., & Norvig, P. (2003). Artificial intelligence: a modern approach (2nd. ed.). Pearson Education.

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

Sayce, L., 2011. *Getting in, staying in and getting on: Disability employment support fit for the future* (Vol. 8081). The Stationery Office.

Suter, R., Scott-Parker, S. and Zadek, S., 2007. Realising potential: disability confidence builds better business.

Trewin, S., 2018. AI fairness for people with disabilities: Point of view. *arXiv preprint arXiv:1811.10670*.

Trewin, S., Basson, S., Muller, M., Branham, S., Treviranus, J., Gruen, D., Hebert, D., Lyckowski, N. and Manser, E., 2019. Considerations for AI fairness for people with disabilities. *AI Matters*, *5*(3), pp.40-63.

Whittaker, M., Alper, M., Bennett, C.L., Hendren, S., Kaziunas, L., Mills, M. and West, M., 2019. Disability, Bias, and AI. *AI Now Institute, November*.

| Recommended Resources

Crawford, Kate, Roel Dobbe, Theodora Dryer, Genevieve Fried, Ben Green, Elizabeth Kaziunas, Amba Kak, Varoon Mathur, Erin McElroy, Andrea Nill Sánchez, Deborah Raji, Joy Lisi Rankin, Rashida Richardson, Jason Schultz, Sarah Myers West, and Meredith Whittaker. AI Now 2019 Report. New York: AI Now Institute, 2019

- https://ainowinstitute.org/AI_Now_2019_Report.html.

European Commission, 2020. White paper on artificial intelligence–a European approach to excellence and trust.

- https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf

Leslie, D., 2019. Understanding artificial intelligence ethics and safety. *arXiv preprint arXiv:1906.05684.*

- https://arxiv.org/pdf/1906.05684.pdf

Office of Artificial Intelligence Guidelines, 2020. Guidelines for AI procurement.

- https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/890699/Guidelines_for_AI_procurement__Print_version_.pdf

Whittaker, M., Alper, M., Bennett, C.L., Hendren, S., Kaziunas, L., Mills, M. and West, M., 2019. Disability, Bias, and AI. *AI Now Institute, November*.

- https://wecount-cms.inclusivedesign.ca/wp-content/uploads/2020/06/Disability-bias-AI.pdf

World Economic Forum, 2019. White Paper: Guidelines for AI Procurement.

- http://www3.weforum.org/docs/WEF_Guidelines_for_AI_Procurement.pdf

| Glossary

Algorithm

A formula or set of rules that determines the process by which the machine goes about finding answers to a question or solutions to a problem.

Artificial Intelligence (AI)

A field of computer science focused on the study of computationally supported intelligent decisions and problem solving.

Augmented Intelligence

Complementing and supporting, rather than replacing, human tasks and intelligence.

Autonomous AI

An AI system that doesn’t require input from a human operator to function and complete tasks.

Data mining

The process of identifying patterns within large sets of data with the intention of deriving useful information about the data.

Deep learning

An approach in machine learning that models and examines complex structures and relationships among data by employing algorithms.

Machine learning

A field of AI focusing employing algorithms that learn automatically from experience for analytical modelling.

Natural language processing (NLP)

A field of AI that reads and interprets human languages in order to derive meaning from them.