Recruitment AI has a Disability Problem: anticipating and mitigating unfair automated hiring decisions

Selin E Nugent* and Susan Scott-Parker

Institute for Ethical AI, Oxford Brookes University
Oxford OX3 0BP, United Kingdom
Scott-Parker International; Business Disability International London, United
Kingdom

Abstract. Artificial Intelligence (AI) technologies have the potential
to dramatically impact the lives and life chances of people with dis-
abilities 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 historically
marginalised groups. Assessments of fairness in Recruitment AI for peo-
ple with disabilities have thus far received little attention or have been
overlooked. This paper examines the impacts to and concerns of disabled
employment seekers using AI systems for recruitment, and discusses rec-
ommendations for the steps employers can take 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.

Keywords: disability, artificial intelligence, recruitment, human resources,
hiring

1   Introduction

Artificial Intelligence (AI) and similar advanced data analytics systems are in-
creasingly sought-after tools for recruitment used to automate time-consuming,
repetitive operational tasks, and expand strategic potential by employers. How-
ever, the complexity of engineering that powers the desirable capabilities of these
systems, also produces downstream difficulties for organisations to validate the
technology for purpose. Uncertain confidence in the decision process and out-
comes of technologies involved in recruitment introduces risk in the hiring prac-
tices and amplifies the probability of unfair treatment toward job seekers. We
posit that without recognition of the limitations of the systems and preventative
governance strategies by employer, AI-powered recruitment technologies have

---

* Corresponding author: snugent@brookes.ac.uk

the potential to dramatically impact the life chances and economic opportunities of people with disabilities seeking employment.

Commercial AI technologies used in recruitment are marketed as highly capable and objective tools for decision making. In contrast, a growing body of research demonstrates a record of inaccurate results as well as inherent disadvantages across historically marginalised people [1], [2], [3]. Assessment of disability fairness in Recruitment AI has thus far received little attention or been overlooked see [4], [5], [6], [7], [8].

These findings accompany a landscape of limited regulation and increasing societal pressure for AI and data analytics systems to be designed with fairness, transparency, and validity. This confluence of pressures 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. Minimising these risks requires employers, human and disability rights campaigners, and academic experts to collaborate to develop approaches to validate systems and the design governance strategies to mitigate risks to the job seekers.

Our aim in this paper takes a consumer protections perspective to understanding and mitigating the risks posed by AI recruitment technology. We review the broad range of technological solutions that support the recruitment process, assess the potential and range of impacts to disabled job seekers. We conclude by discussing the options employers having in preparing for procuring on new systems and in evaluating system currently in use.

2   Disability Marginalisation and Employment

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 does not necessarily capture the complexity and heterogeneity disability, which is a key factor that causes 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 disabilities.

Disability is not completely independent of other features of a person's identity and life experience [9], [10], [11]. The structural forces and social stigmas that impact people with disability are shared and amplified at the intersection

other historically marginalised identities, such those based on as gender, ethnicity, sexuality, and socioeconomic status [12]. 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.

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 able-bodied people [13]. Presently, the sheer scale of the social and economic impacts of the COVID-19 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 [14].

Disability inclusion in the workplace is impacted by number of factors. There is often a qualifications gap between disabled and able-bodied people due to systematic disadvantages in education, training, and previous work experience [15]. Even well-intentioned employers may struggle to recognize how a candidate's previous structural barriers to success impact their present assessment of the candidate's eligibility.

Furthermore, some industries or types of position lack accessibility that can limit employment for people with certain impairments. This maybe due to the employer lacking resources (in the form of advisory programmes) to support the employer in supporting persons with disabilities. Employers may also have negative attitudes and lack confidence or training to support disabled employment seekers [16], [17].

The impact of AI aside, the structural issues affecting people with disabilities in gaining and maintaining employment is a complex and ongoing concern. Therefore layering a complex and opaque system of automated assessment of candidates risk complicating the situation and expanding the risk of harm further.

3   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. AI systems are currently available across a wide range

of recruitment functions, including: Candidate Sourcing, Engagement Candidate Tracking, CV Screening, Pre-Employment Assessments, and Video Interviewing.

The unifying objective for systems operating across these diverse recruitment functions is that they are designed to distill 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 other features, the complexity may be interpreted as abnormality. 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.

4   Exclusion by Design and Discriminatory Use

Recruitment AI risks inadvertently, but adversely impacting 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. When an objective is not specified adequately, the assessment may lead to unintended consequences in the outcome.

Unwanted biases, or biases that treat people negatively, or adversely due to protected characteristics or other features of their identity within an AI system, raise serious risks of discrimination. It is critical to identify and mitigate these potentially harmful biases.

Unwanted biases relevant to marginalised people, including people with disability, are primarily introduced by historical hiring decisions. Since people with disabilities are already twice as unlikely to be unemployed, they are consequently less likely to be represented in data on past successful employees. These biases may be introduced into systems through two 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 1) defining the objective or problem the developer determines and 2) selecting the parameters that define the operation of the system at an optimal level [18].

A scenario in which this may occur is if automated CV screener is programmed to predict the best qualified candidate based on an exclusionary parameter, such as having attended a top-tier university. The prestige of an institution may be one factor in a successful employee, but that parameter also disadvantages historically marginalised people, including people with disabilities, who already face systemic barriers to be equally represented in prestigious institutions.

The training data is the preliminary data from which an system learns how to apply the model and produce results in application [18]. The model operates

as well as the supplied training data. Bias may be introduced in multiple points along the process before learning begins, from decisions made at data collection to data cleaning to selection for the purpose of training.

Building on the previous scenario, problems may arise with an automated CV screener that was trained on data that did not include the profiles of successful, yet historically marginalised employees. 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 or "unusual" 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 [19]. 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-offer 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 them 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 offered employment 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–or at least marketed as–working perfectly.

The use of a rigid, standardised recruitment processes that cannot be adequately adjusted to enable candidates with disabilities to compete fairly are inherently discriminatory [20]. 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. 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 or fearful of interacting with AI and, at worst, discriminating against such candidates. Expecting disabled employment seekers to go through standardised processes is analogous to asking a wheelchair user to take the stairs to the interview room.

5   AI on the Market: the risks of discrimination

Recruitment AI encompasses a wide array of technologies functioning at different points in the recruitment process. We outline 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.

5.1   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 referred to consider future job openings.

We consider these systems together because they share similar impact risks 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 [4]. People with difficulties related to dexterity or visual impairment are disproportionately affected. 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.

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 as a reflection of previous denial of educational and employment opportunities.

5.2   CV/ Resume Screeners

CV screening is a major feature of the recruitment innovation powered by AI, addressing the need for processing high application volumes. Automated screeners detect characteristics in the CV content, such as key phrases and words 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 systems also supplement data in CVs 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.

5.3 Conversational Agents

Recruitment conversational agents, or chatbots, are designed to mimic human conversational abilities during the recruitment process. These technologies use 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 [21].

However, if not designed and implemented considering disability, 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 cannot 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.

5.4 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.

There is general recognition that these assessments are often not reliable as one-size-fits-all approaches [22]. The generalisability of psychometric tests for people with disabilities, as well as numerous other backgrounds, is unreliable [23]. 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 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 [24].

5.5   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 [25].

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 [2]. The implementation of these technologies for

recruitment risks legitimising the flawed methodological premise in a way that perpetuates historic disadvantages and exclusion for marginalised peoples.

6   Intervention Recommendations

Designing and implementing Recruitment AI systems that treat persons with disabilities and by that extent, all job 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 be prepared with the tools and language needed 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 employers can take to procure and implement technological systems that adapt to the values and expectations of their organisation and societal stakeholders toward disabled job seekers. This process begins by introducing disability into governance plans across departments. Involving multiple stakeholders in asking disability-conscious questions of technology suppliers increases the likelihood of identifying and mitigating risks before the system begins making critical decisions about job seekers.

6.1   Executive Level Stakeholders

Starting at the executive level, addressing disability inclusion in recruitment may be assessed as part of wider strategic and corporate vision considerations. Suppliers can be assessed for suitability based on if

i. the recruiting system supports organizational goals to increase diversity and representation

ii. the recruiting system follows any existing policies with regard to the ethical and responsible development and implementation of artificial intelligence?

iii. the supplier actively engages in learning how to adapt to match organisational values and needs as a business and stakeholders?

iv. which employees should be involved in the governance process which determines how to investigate, procure, apply and monitor recruitment systems so that they do not adversely impact disadvantaged job seekers?

6.2   Human Resources Stakeholders

Recruitment AI supports tasks that would otherwise be undertaken by HR or recruitment professionals. These individuals are well-suited to evaluate if the automated process mirrors protocols that would be expected of human practitioners. Considerations for disability-inclusive suitability of the system include:

i. the benefits and risks of the system for disabled and other disadvantaged employment seekers

ii. that the supplier has a shared understanding of inclusivity and fairness—with specific reference to eliminating the root causes of disability related discrimination—designed into the system.

iii. if implementing the system requires alternative evaluation routes to enable people with different impairments to be recruited on the basis of individual capability and potential.

iv. if recruitment system enables candidates to readily request adjustments, in a non-stigmatising manner, at every stage of the process.

6.3　Procurement Stakeholders

Procurement is the first entry point for recruitment systems into an organisation. Procurement professionals are in a position assess disability inclusivity of a recruitment system as part of the risk management process. Considerations may include:

i. if this supplier can demonstrate that their product is 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?

6.4　Information Technology Stakeholders

Finally IT professionals will also need to be prepared to consider the role of disability in the technical functionality of recruitment systems. They will be in a position to evaluate:

i. if they have been provided with the appropriate explainability and interpretability resources to assess outputs and impacts on employment seekers' disabilities.

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

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

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

disfigurements. We appreciate the contributions and insights provided by Julien Burnett, Nigel Crook, Paul Jackson, and Rebecca Raper.

References

1. Broussard, M.: Artificial Unintelligence: How Computers Misunderstand the World. MIT Press, Cambridge (2018)
2. O'Neil, C.: Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Penguin Random House, New York (2017)
3. Noble, S.: Algorithms of Oppression: How Search Engines Reinforce Racism. New York University Press, New York (2018)
4. Guo, A., Kamar, E., Vaughan, J.W., Wallach, H., Morris, M.R.: Toward Fairness in AI for People with Disabilities: A Research Roadmap. arXiv preprint arXiv:1907.02227 (2019)
5. Petrick, E.R.: Making Computers Accessible: Disability Rights and Digital Technology. Baltimore: Johns Hopkins University Press. (2015)
6. Trewin, S.: AI fairness for people with disabilities: Point of view. arXiv preprint arXiv:1811.10670. (2018)
7. Trewin, S., Basson, S., Muller, M., Branham, S., Treviranus, J., Gruen, D., Hebert, D., Lyckowski, N. and Manser, E.: Considerations for AI fairness for people with disabilities. AI Matters, 5(3), 40–63. (2019)
8. Whittaker, M., Alper, M., Bennett, C.L., Hendren, S., Kaziunas, L., Mills, M., West, M.: Disability, Bias, and AI. AI Now Institute, November (2019)
9. Collins, P.H., Bilge, S.: Intersectionality. John Wiley and Sons (2020)
10. Parker, A.M.: Intersecting Histories of Gender, Race, and Disability. Journal of Women's History 27, 1 (2015)
11. Samuels, E.: Fantasies of Identification: Disability, Gender, Race. New York University Press, New York. (2016)
12. Frederick, A., Shifrer, D.: Race and disability: From analogy to intersectionality. Sociology of Race and Ethnicity, 5(2), 200–214 (2019)
13. Office for National Statistics, Disability And Employment, UK: 2019, https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/disability/bulletins/disabilityand
14. Schur, L.: The difference a job makes: The effects of employment among people with disabilities. Journal of Economic Issues, 36(2), 339–347 (2002)
15. Sayce, L.: Getting in, staying in and getting on: Disability employment support fit for the future (Vol. 8081). The Stationery Office (2011)
16. Lindsay, S., Leck, J., Shen, W., Cagliostro, E., Stinson, J.: A framework for developing employer's disability confidence. Equality, Diversity and Inclusion: An International Journal (2019)
17. Suter, R., Scott-Parker, S. and Zadek, S.: Realising potential: disability confidence builds better business. Employers' Forum on Disability (2007)
18. Russell, S., Norvig, P.: Artificial intelligence: a modern approach (2nd. ed.). Pearson Education (2003)
19. Davis, L.: Disabilities in the workplace: recruitment, accommodation, and retention. Aaohn Journal, 53(7), 306–312 (2005)
20. Hamraie, A.: Building Access: Universal Design and the Politics of Disability. Minneapolis: University of Minnesota Press. (2017)
21. Følstad, A., Brandtzæg, P.B., Feltwell, T., Law, E.L., Tscheligi, M., Luger, E.A.: SIG: chatbots for social good. In Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems, 1–4. (2018)

22. Edenborough, R.: Assessment methods in recruitment, selection performance: a manager's guide to psychometric testing, interviews and assessment centres. Kogan Page Publishers. (2005)
23. Smith, K., Abrams, S.S.: Gamification and accessibility. The International Journal of Information and Learning Technology (2019)
24. Cook, D.A., Beckman, T.J.:Current concepts in validity and reliability for psychometric instruments: theory and application. The American journal of medicine, 119(2), 166–167 (2006)
25. Leslie, D.: Understanding bias in facial recognition technologies. arXiv preprint arXiv:2010.07023 (2020).