

## ARTICLE

# Artificial intelligence and workplace inclusion: a theoretical agenda on neurodiversity, disability, and fairness in recruitment processes

*Inteligência artificial e inclusão no trabalho: uma agenda teórica sobre neurodiversidade, deficiência e justiça nos processos de recrutamento*

*Inteligencia artificial e inclusión laboral: una agenda teórica sobre neurodiversidad, discapacidad y justicia en los procesos de contratación*

**Fábia Aparecida Silva<sup>1\*</sup>**

<sup>1</sup> Pontifícia Universidade Católica de São Paulo (PUC-SP), São Paulo, SP, Brasil. ORCID: <https://orcid.org/0009-0009-5733-9025>.

| SUBMITTED  | ACCEPTED   | PUBLISHED  | DOI        |
|------------|------------|------------|------------|
| 29/03/2024 | 25/10/2024 | 06/12/2024 | in process |

## ABSTRACT

Automated recruitment systems promise efficiency and standardization, but they may reproduce narrow conceptions of competence, communication, and appropriate behavior. This article proposes a theoretical agenda on fairness, disability, and neurodiversity in selection processes mediated by artificial intelligence. It adopts a theoretical-integrative approach that brings together the social model of disability, digital accessibility, organizational justice, and algorithmic governance, in dialogue with international reports on artificial intelligence and work produced by the International Labour Organization and the Organisation for Economic Co-operation and Development. As a result, four conditions for inclusive recruitment are set out: accessibility by design, assessment restricted to the competencies actually required for the position, the provision of reasonable accommodations linked to contestation channels with human mediation, and the participation of persons with disabilities in the governance of these systems. The analysis indicates that statistical fairness metrics are insufficient when the process retains communicational, sensory, or cognitive

barriers. The article also discusses the limits of technical bias-mitigation approaches, the accountability of hiring organizations vis-à-vis third-party software vendors, and the institutional conditions needed to keep the participation of persons with disabilities in algorithmic governance from being reduced to a symbolic gesture. It concludes that the legitimacy of a recruitment system should be judged by its capacity to recognize different legitimate ways of demonstrating professional competence, rather than by the technical sophistication of the predictive models it employs.

**Keywords:** algorithmic recruitment, disability, neurodiversity, inclusion, organizational justice

#### RESUMEN

Los sistemas automatizados de contratación prometen eficiencia y estandarización, pero pueden reproducir concepciones restrictivas de competencia, comunicación y comportamiento. Este artículo propone una agenda teórica sobre justicia, discapacidad y neurodiversidad en los procesos de selección mediados por inteligencia artificial. Se adopta un enfoque teórico-integrador que articula el modelo social de la discapacidad, la accesibilidad digital, la justicia organizacional y la gobernanza algorítmica, en diálogo con informes internacionales sobre inteligencia artificial y trabajo elaborados por la Organización Internacional del Trabajo y la Organización para la Cooperación y el Desarrollo Económicos. Como resultado, se estructuran cuatro condiciones para una contratación inclusiva: accesibilidad desde el diseño, evaluación limitada a las competencias necesarias para el puesto, provisión de ajustes razonables asociados a canales de impugnación con mediación humana, y participación de personas con discapacidad en la gobernanza de los sistemas. El análisis indica que las métricas estadísticas de equidad son insuficientes cuando el proceso mantiene barreras comunicacionales, sensoriales o cognitivas. Se discuten además los límites de los enfoques técnicos de mitigación de sesgos, la responsabilización de las organizaciones contratantes frente a proveedores externos de software y las condiciones institucionales necesarias para que la participación de personas con discapacidad en la gobernanza algorítmica no se reduzca a un gesto simbólico. Se concluye que la legitimidad de un sistema de contratación debe evaluarse por su capacidad para reconocer diferentes formas legítimas de demostrar competencia profesional, y no por la sofisticación técnica de los modelos predictivos que emplea.

**Palabras clave:** contratación algorítmica, discapacidad, neurodiversidad, inclusión, justicia organizacional

---

## INTRODUCTION

Over the past ten years, artificial intelligence systems have come to be embedded in virtually every stage of corporate recruitment: from the initial screening of résumés to the analysis of video interviews, by way of gamified tests, pre-selection chatbots, and candidate-recommendation tools. Human resources technology companies market these solutions with the promise of cutting costs, speeding up processes, and, above all, eliminating the human bias that supposedly contaminates hiring decisions (Raghavan et al., 2020). That promise of neutrality has, however, proven to be more rhetoric than fact. A growing body of research in the sociology of algorithms, data ethics, and organizational studies shows that automated selection systems tend to reproduce, under a new technical guise, longstanding structural inequalities, including those affecting people with disabilities and neurodivergent people (Eubanks, 2018; O'Neil, 2016; Bogen & Rieke, 2018).

The problem has a considerable quantitative dimension. According to estimates from the World Health Organization and the World Bank, around 15% of the world's population lives with some form of disability, and this population faces substantially higher rates of unemployment and underemployment than the population at large (World Health Organization & World Bank, 2011). Reports from the International Labour Organization (ILO) on the future of work and the digitalization of labour relations have drawn attention to the fact that automating people-management processes, unless accompanied by specific safeguards, tends to deepen these disparities rather than correct them (International Labour Organization, 2021). Along the same lines, reports from the Organisation for Economic Co-operation and Development (OECD) on artificial intelligence and the labour market point out that the introduction of predictive systems in human resources requires mechanisms of human oversight, transparency, and accountability that are still far from standard practice in organizations (OECD, 2023).

This article starts from a more specific focus: the intersection between artificial intelligence applied to recruitment, on the one hand, and the fields of critical disability studies and neurodiversity, on the other. This is an intersection that remains little explored in the Brazilian literature and still incipient even in the international literature, despite its growing relevance. When a video-interview screening algorithm scores candidates on the basis of sustained eye contact, stable vocal modulation, or conventional facial expressiveness, it is, in effect, encoding a specific pattern of communication and behavior as synonymous with professional competence. Autistic people,

people with ADHD, people with cerebral palsy, people with hearing or speech impairments, and people with other conditions that affect bodily or vocal expression run the risk of being screened out not for lack of technical qualification, but for failing to match a narrow model of "normality" built into the system (Whittaker et al., 2019; Guo et al., 2020).

The question guiding this investigation is the following: how can the principles of the social model of disability, organizational justice, and algorithmic governance be articulated in a theoretically consistent way to ground guidelines for digital recruitment that are genuinely inclusive? The aim of the article is to propose a theoretical agenda that offers answers to this question, structured around four conditions that, it is argued here, need to be present simultaneously for an AI-mediated recruitment system to be considered inclusive: accessibility by design, restricting assessment to the competencies actually required for the position, the provision of reasonable accommodations linked to contestation channels with human mediation, and the participation of persons with disabilities in the governance of these systems.

It is worth noting that the expansion of platform capitalism and the growing digitalization of labour relations make this debate even more urgent. As job applications come to be mediated almost entirely by digital systems rather than by face-to-face contact with a recruiter, design decisions that once seemed merely technical, such as the format of a test or the length of a deadline, come to directly determine who has access to an employment opportunity and who is screened out before any substantive assessment of their professional qualifications even takes place.

The study is justified on three main grounds. First, because the literature on algorithmic bias in recruitment has focused predominantly on categories such as gender and race, leaving disability and neurodivergence in a secondary position, when they are addressed at all (Bogen & Rieke, 2018; Raji et al., 2020). Second, because the most widely used technical metrics of algorithmic fairness, such as demographic parity and equality of opportunity, were not designed to capture accessibility and communication barriers, which produces a significant blind spot in bias audits as currently practiced (Trewin, 2018). Third, because recent regulatory decisions, such as the European Union's Artificial Intelligence Regulation, classify recruitment systems as "high risk" and now require impact assessments, technical documentation, and human oversight, which makes it urgent to have theoretical parameters capable of guiding both regulation and

managerial practice (European Parliament and Council of the European Union, 2024).

The article is organized into five sections beyond this introduction. The second section presents the theoretical framework, divided into six subsections: the trajectory from the medical model to the social model of disability; the concept of neurodiversity and its relationship to the world of work; the distributive, procedural, and interactional dimensions of organizational justice as applied to algorithmic contexts; debates on algorithmic governance and technical opacity; the international landscape on artificial intelligence and work built by bodies such as the ILO and the OECD; and a critical synthesis of the literature on algorithmic bias and disability. The third section describes the methodological path of the integrative review. The fourth section presents the results and discussion, structured around the four conditions of the proposed theoretical agenda, followed by subsections on impact assessment, team training, and the future research agenda. The fifth section presents the closing remarks, including the study's limitations and directions for subsequent empirical research.

## THEORETICAL FRAMEWORK

The theoretical framework of this article draws on five bodies of literature that, although they have developed relatively independently of one another, together offer a robust conceptual basis for analyzing the effects of automation on the labour-market inclusion of people with disabilities and neurodivergent people: critical disability studies, the neurodiversity movement and literature, organizational justice theory, studies on algorithmic governance, and, finally, reports from international bodies on artificial intelligence and the future of work.

### *From the medical model to the social model of disability*

For a long time, disability was understood through a medical or individual model, which locates the problem exclusively in the person's body or mind. From this perspective, disability is seen as a deviation from a standard of functioning considered normal, something to be diagnosed, treated, or rehabilitated (Shakespeare, 2014). The medical model played an important historical role in the development of rehabilitation services and health policies, but it also produced a lasting side effect: it transferred to the individual the entire responsibility for adapting to the environment, exempting social institutions from any obligation to change.

In contrast to this reading, the social model of disability, formulated on the basis of the British disability movements of the 1970s and 1980s and systematized by authors such as Michael Oliver, distinguishes between impairment and disability. Impairment refers to a person's bodily, sensory, intellectual, or neurological condition; disability, in turn, is the result of the interaction between that impairment and a social environment that was not built to accommodate human diversity (Oliver, 2013). A blind person is not "disabled" by blindness itself, but becomes disabled when faced with a website that is not compatible with screen readers, a form with no descriptive text alternative, or a timed recruitment test that assumes fast visual reading. Disability, in this sense, lies not in the body but in the relationship between the body and a hostile or inflexible environment.

This distinction has direct consequences for the analysis of digital recruitment systems. Automated screening platforms, application portals, and virtual testing environments function, in practice, as new built environments, subject to the same kind of criticism historically leveled at staircases without ramps or buildings without elevators. When a job portal is not navigable by keyboard, when a psychometric test imposes a rigid deadline incompatible with certain cognitive conditions, or when a video-interview platform requires a synchronous oral response with no written alternative, the system is erecting barriers equivalent to the architectural barriers of physical space, albeit in digital form (Crip et al., 2022). Exclusion ceases to be read as a personal failing on the candidate's part and comes to be recognized instead as the product of design choices that could have been made differently.

It is worth noting that the social model does not deny the existence of bodily or functional limitations, nor does it suggest that every difficulty is purely social. Authors such as Shakespeare (2014) draw attention to the risks of an overly structural reading that disregards the subjective experience of pain, fatigue, or functional limitation lived by many people with disabilities. What the social model proposes, in a more balanced way, is that the political and institutional response to disability cannot be limited to correcting the individual body; it necessarily needs to include the transformation of the environments, practices, and technologies that produce exclusion.

Applied to algorithmic recruitment, this shift in perspective changes the question that guides the analysis. Instead of asking "why did this candidate fail to complete the test within the set time limit," the relevant question becomes "why was the test designed with a time limit that excludes a



legitimate part of the candidate population." This shift in focus is the theoretical starting point of this article.

### ***Neurodiversity and work***

The concept of neurodiversity emerged in the late 1990s, in the context of autistic self-advocacy movements, and later expanded to encompass other conditions, such as ADHD, dyslexia, and dyspraxia. The core proposal of the concept is simple to state but radical in its implications: variation in human neurological functioning is part of the natural diversity of the species, not a pathology to be eliminated (Singer, 2017; Armstrong, 2012). Just as genetic diversity confers resilience on a population, cognitive diversity gives society and organizations a broader repertoire of ways to perceive, process information, and solve problems.

This perspective does not deny that neurodivergent people face real difficulties in environments built for the neurotypical majority. What it disputes is the idea that these difficulties stem from an intrinsic deficit rather than from a lack of fit between the person and the environment. In the labour market, this lack of fit becomes particularly acute at the recruitment stage, historically built around conventional social skills: spontaneous verbal fluency, sustained eye contact, the ability to "sell oneself" in a short interview, and quick reading of implicit social cues (Kapp & Aiello, 2022).

The introduction of artificial intelligence into this landscape does not correct this cultural bias; it formalizes it into mathematical rules. Various commercial video-interview analysis platforms use computer vision and audio processing to infer personality traits and "cultural fit" from variables such as the frequency of eye contact with the camera, the symmetry of facial movements, hesitations in speech, and intonation variation (Bogen & Rieke, 2018). For an autistic person who looks away from the camera to concentrate better, or who presents a less varied vocal prosody, these behaviors may be read by the algorithm as signs of disengagement or emotional instability, when in fact they merely reflect a different style of sensory processing and communication.

This point is reinforced by research in organizational psychology documenting specific cognitive strengths often associated with neurodivergent profiles, such as sustained attention to detail, recognition of complex patterns, and the capacity for hyperfocus on tasks of interest (Austin & Pisano, 2017, as cited in Kapp & Aiello, 2022). A screening system that automatically penalizes candidates for failing to match a conventional

communication pattern risks discarding precisely the profiles that could add the most value to certain technical and analytical roles, incurring at once an ethical problem and an efficiency loss for the hiring organization itself.

### ***Organizational justice in algorithmic contexts***

Organizational justice theory, consolidated by Colquitt (2001) on the basis of decades of research in organizational behavior, distinguishes three main dimensions: distributive justice, concerning the perceived fairness of allocated outcomes (who receives the job opportunity, the promotion, the pay raise); procedural justice, concerning the impartiality and transparency of the processes used to arrive at those outcomes; and interactional justice, concerning the quality of the interpersonal treatment received during the application of those processes, including respect, dignity, and the possibility of explanation.

The introduction of automated systems into recruitment affects these three dimensions in distinct ways. In the distributive dimension, the most evident problem is disproportionate exclusion: when an algorithm learns, from historical hiring data, that certain profiles correspond to "professional success," it tends to reproduce the patterns of who was hired in the past, which can mean perpetuating the barriers faced by people with disabilities and neurodivergent people in earlier selection cycles (Raghavan et al., 2020).

In the procedural dimension, the central problem is opacity. Unlike a selection process conducted by a human recruiter, whose criteria can, at least in principle, be questioned and explained, many automated screening systems offer no intelligible explanation for the reasons behind a rejection. This opacity is aggravated when the software is developed by a third-party vendor that treats its inner workings as a trade secret, which prevents both the candidate and the hiring company itself from fully understanding how decisions are made (Pasquale, 2015).

In the interactional dimension, automation tends to reduce the space for dialogue between candidate and organization. A screening chatbot or a one-way video-interview platform, with no possibility of real-time clarification, conveys to the candidate the sense of being evaluated by a system indifferent to their particular situation, which can be especially harmful for candidates who would need to communicate a need for accommodation but find no channel to do so before being rejected (Colquitt & Rodell, 2011).



The sum of these three effects amounts to what might be called a transition from a relational justice, based on human interaction and the possibility of appeal, to a statistical justice, based on population aggregates and performance averages. This transition is not neutral: it shifts the criterion of legitimacy of the selection process away from the quality of the relationship between the parties and toward the supposed objectivity of the numbers alone.

### ***Algorithmic governance and technical opacity***

The third theoretical axis of this article deals with the governance of algorithmic systems, understood as the set of rules, norms, oversight mechanisms, and auditing practices aimed at regulating the life cycle of these systems, from conception to decommissioning (Yeung, 2018). A recurring obstacle to effective governance is the phenomenon described by Pasquale (2015) as the "black box society": machine-learning models, particularly those based on deep neural networks, make decisions from combinations of variables that often even their own developers cannot fully explain.

This opacity has concrete practical consequences for candidates with disabilities. When a person is rejected in an automated screening process, they rarely receive information about which specific variables drove the decision, which precludes any well-founded contestation. Moreover, the technical metrics most commonly used to audit algorithmic bias, such as demographic parity, equality of opportunity, and predictive calibration, were developed predominantly to address categories such as gender and race, and rarely incorporate disability as a relevant protected variable (Trewin, 2018; Guo et al., 2020). This means a system can be certified as "bias-free" according to the usual auditing criteria and still go on systematically excluding people with disabilities, simply because this dimension was never included in the analysis.

The power asymmetry between software developers, hiring companies, and candidates compounds this picture. Human resources technology vendors typically operate under confidentiality agreements that prevent independent external audits, and many hiring companies lack the in-house technical capacity to critically assess the tools they purchase (Kuzior et al., 2020). The result is an accountability vacuum in which, faced with an unfair rejection, it is difficult to identify who should answer for it: the algorithm's developer, the company that deployed it, or both.

Recent regulatory frameworks have attempted to fill this vacuum. The European Union's Artificial Intelligence Regulation, approved in 2024, classifies recruitment and selection systems as "high risk" and imposes specific obligations for impact assessment, technical documentation, human oversight, and transparency toward candidates (European Parliament and Council of the European Union, 2024). Although its application is confined to European territory, the regulation has served as a reference point for regulatory debates in other regions, including legislative discussions in Brazil on artificial intelligence and data protection in the workplace.

### ***Artificial intelligence and work in the literature of international bodies***

Beyond the academic literature, technical reports produced by international bodies offer a relevant empirical overview of the effects of automation on the world of work, including its impacts on groups in situations of vulnerability. Over the past decade, the International Labour Organization has published a series of studies on the digitalization of labour relations and the risks associated with the expansion of digital platforms and automated people-management systems (International Labour Organization, 2021). These reports point out that the absence of specific regulation tends to shift onto the individual worker the burden of negotiating with increasingly opaque technical systems, a dynamic that reinforces, rather than reduces, pre-existing inequalities.

The OECD, in turn, has produced comparative analyses of the adoption of artificial intelligence in recruitment processes across different member countries, highlighting the need for meaningful (not merely formal) human oversight and for explainability mechanisms accessible to candidates without a technical background (OECD, 2019, 2023). A relevant finding of these reports is that the pace of adoption of automated human resources tools has, in many sectors, outstripped the pace of development of regulatory frameworks and the technical training of the teams responsible for operating these systems, creating a governance gap that disproportionately affects candidates from historically marginalized groups.

In the specific field of disability, the report "Disability, Bias, and AI," produced by the AI Now Institute, was among the first to systematically map the risks that artificial intelligence systems pose to people with disabilities in the contexts of employment, health, and social assistance, arguing that most debates on algorithmic fairness ignore this population and that automated recruitment tools frequently treat disability as statistical noise to be filtered out rather than as a legitimate form of human diversity

(Whittaker et al., 2019). That report, alongside academic work such as Guo et al. (2020) and Trewin (2018), has helped consolidate a specific research field on algorithmic fairness and disability that is still taking shape but already offers important conceptual parameters for the analysis proposed in this article.

### ***Critical synthesis: the limits of purely technical approaches to bias correction***

The literature reviewed so far allows for a critical synthesis relevant to the objectives of this article. Most of the technical solutions currently proposed for correcting bias in algorithmic recruitment systems focus on statistical adjustments: reweighting variables, balancing training datasets, applying demographic parity metrics across groups. These approaches have value, but they rest on a limited premise, namely that the problem can be solved entirely within the mathematical model itself, without changing the system's architecture, the nature of the variables collected, or the design of the interface through which the candidate interacts with the system (Raghavan et al., 2020).

In the case of disability and neurodivergence, this premise proves particularly fragile. Reweighting variables does not fix an interface incompatible with a screen reader. Balancing training data does not remove the barrier posed by a timed test that ignores cognitive variability. Adjusting predictive-calibration metrics does not confer scientific validity on a pseudoscientific construct, such as inferring emotional stability from facial microexpressions (Bogen & Rieke, 2018). In other words, many of the problems faced by candidates with disabilities in algorithmic recruitment processes are not problems of statistical bias in the strict sense, but problems of accessibility, construct validity, and the absence of contestation channels, issues that call for interventions of a different nature, situated at the level of system design and institutional governance, rather than only at the level of the decision algorithm.

It is this observation that underpins the central proposal of this article: that the inclusion of people with disabilities and neurodivergent people in algorithmic recruitment processes depends on a set of structural conditions, and not merely on isolated technical fixes to the predictive models used.

### ***Barriers to hiring people with disabilities: what the demand-side literature says***

A considerable share of research on employment and disability has historically concentrated on the so-called supply side, that is, on the characteristics, qualifications, and rehabilitation needs of the candidate.

Work published in journals such as the Journal of Occupational Rehabilitation has shifted some of that attention to the demand side, examining directly the attitudes, practices, and organizational barriers that hinder the hiring of people with disabilities, regardless of their qualifications (Chan et al., 2010). These studies identify factors such as lack of knowledge about reasonable accommodations, exaggerated concern over accommodation costs, fear of litigation, and stereotyped beliefs about productivity as central determinants of low hiring rates, more so than any real limitation in candidates' performance.

This shift in focus, from the individual to the organization, is consistent with the social model of disability discussed above and reinforces an important point for the analysis of automated recruitment systems: if the barriers to hiring were already, before automation, predominantly organizational and attitudinal, it is reasonable to assume that the introduction of screening algorithms, far from neutralizing this picture, tends to formalize these same barriers into decision rules that are less visible and harder to contest. A human recruiter with implicit biases can, at least in theory, be trained, supervised, or confronted with evidence of their own bias. An algorithm that has learned equivalent patterns from historical hiring data operates consistently and at scale, without there being, in most cases, any comparable instance of oversight.

The literature on vocational rehabilitation also draws attention to the importance of interventions aimed at the work environment, and not only at the individual, including adapting workstations, flexibilizing schedules, and revising performance-evaluation processes. These findings reinforce the central thesis of this article: the labour-market inclusion of people with disabilities depends primarily on changes to the design of processes and organizational environments, and only secondarily on interventions centered on the individual adaptation of the worker or candidate to the prevailing standard.

### ***Intersectionality between disability, neurodivergence, and other social markers***

It must be acknowledged that people with disabilities and neurodivergent people do not constitute homogeneous groups, and that their experiences of exclusion in algorithmic recruitment frequently intersect with other social markers, such as gender, race, age, and class. A Black autistic woman, for example, may face simultaneous barriers related to racial bias, gender bias, and bias associated with neurodivergence, and these barriers do not simply

add up but combine in specific ways that intersectionality literature describes more precisely than analyses that treat each category in isolation.

Most currently available algorithmic bias-auditing metrics were not built to capture these combinations. When an audit separately assesses demographic parity by gender, by race, and by disability, it may fail to detect a pattern of exclusion that only emerges at the intersection of these categories. This is a further argument in favor of the condition of active participation of people with disabilities in the governance of these systems, discussed in the results section: representatives with direct lived experience of multiple forms of exclusion are better placed to identify intersectional patterns that escape statistical audits segmented by a single category at a time.

## METHODOLOGICAL PATH

This study is characterized as theoretical research of an integrative nature (Torraco, 2016). The integrative review is a method recognized in the applied social sciences for allowing the critical synthesis of literature drawn from distinct disciplinary fields, which is particularly well suited to the object of this article, situated at the intersection of human resources management, artificial intelligence ethics, critical disability studies, the sociology of algorithms, and organizational psychology.

The development of the study followed four stages. In the first stage, the scope of the investigation was delimited and gaps were mapped in the recent literature on the impacts of artificial intelligence tools on recruitment, with particular attention to the under-representation of debates on disability and neurodiversity in this field. In the second stage, a matrix of analytical axes was constructed, cross-referencing the social model of disability (Oliver, 2013), the classic dimensions of organizational justice (Colquitt, 2001), the principles of algorithmic governance (Yeung, 2018; Floridi et al., 2018), and the landscape produced by international bodies such as the ILO and the OECD. In the third stage, critical analysis and conceptual synthesis were carried out, examining how market-driven premises of efficiency and predictive standardization collide with rights to accessibility, reasonable accommodation, and workplace dignity. In the fourth and final stage, the propositional theoretical agenda that structures the results section of this article was formulated.

The literature search considered publications indexed in databases such as Web of Science, Scopus, SciELO, and Google Scholar, prioritizing peer-

reviewed articles published over the last decade, complemented by seminal books in the sociology of technology and disability studies, as well as technical reports from international bodies and applied research institutes such as the AI Now Institute. The decision to include grey literature, that is, technical reports not published in peer-reviewed journals, is justified by the recent nature of the subject: a substantial part of the output on artificial intelligence and disability in work contexts still circulates primarily in institutional reports and public-policy documents, rather than solely in traditional academic articles.

One point of methodological attention concerns the construct validity of instruments sold in the human resources technology market. Work psychology establishes that selection tests must demonstrate an empirically proven relationship with performance on job tasks, what the literature calls criterion-related or task validity. Many multimodal video-analysis platforms, however, are sold with no independent validation subjected to peer review, relying instead on theories of facial expression that have already been superseded by contemporary psychology (Bogen & Rieke, 2018). This validation gap is not merely a technical problem: it represents a labour-rights issue, insofar as candidates can be excluded on the basis of inferences with no consistent scientific backing.

Finally, it is important to acknowledge the limits of the method adopted. An integrative review does not permit statistical generalizations, nor does it replace empirical field studies. Its value lies in offering a conceptual synthesis capable of guiding future investigations and grounding practical guidelines, but the propositions formulated in this article should be understood as theoretical hypotheses to be tested, not as definitive conclusions about the workings of any specific systems.

## **RESULTS AND DISCUSSION: A THEORETICAL AGENDA FOR INCLUSIVE RECRUITMENT**

Bringing together the social model of disability, the literature on neurodiversity, organizational justice theory, and studies on algorithmic governance makes it possible to structure four conditions that, it is argued here, need to be present jointly for an AI-mediated recruitment system to be considered inclusive. The central thesis guiding this section is that algorithmic fairness cannot be resolved through isolated statistical adjustments to biased models; it requires structural changes in the design, development, and governance of selection processes.



---

***Condition 1: accessibility by design***

The first condition is that every digital interface involved in the selection process, including application portals, psychometric tests, video-interview platforms, and candidate-tracking systems, be designed with accessibility built in from the initial design phase, rather than retrofitted afterward (Crip et al., 2022). Historically, the human resources software market has prioritized users without disabilities, resulting in complex gamified tests, interfaces with no alternative text, portals incompatible with screen readers, and timed deadlines that ignore candidates' cognitive variability.

Accessibility by design means native compatibility with assistive technologies, adequate color contrast, full keyboard navigation, and the provision of multiple communication channels (Groom et al., 2021). It also means not imposing a single mode of expression: for candidates with hearing or speech impairments, for example, the requirement of a synchronous audio response can be replaced by real-time transcription or by asynchronous submission of written answers. This condition turns accessibility from a mere item of legal compliance into a core parameter of software quality, on a par with information security or technical performance.

***Condition 2: construct validity and restricting assessment to essential competencies***

The second condition demands rigor in validating the constructs assessed by these systems. Video-interview tools frequently infer personality traits, emotional stability, or "cultural fit" from microexpressions, vocal modulation, and speech rhythm, drawing on psychological theories whose scientific robustness is contested (Bogen & Rieke, 2018). These metrics disproportionately penalize people with cerebral palsy, people on the autism spectrum, people who stutter, and people with hearing or speech impairments, whose nonverbal communication diverges from the statistical patterns of the neurotypical majority.

The agenda proposed here establishes that algorithmic assessment should be restricted to competencies directly related to the performance of the job's duties, with empirically demonstrated validity, excluding the use of automated predictors of internal emotional states or generic personality traits. If a competency cannot be directly linked to the exercise of the role, its automated measurement constitutes an indirect form of discrimination, even if unintentional.

---

***Condition 3: proactive provision of accommodations and contestation channels with human mediation***

The third condition addresses the procedural and interactional dimensions of organizational justice, and establishes the need for reasonable accommodations to be offered proactively, rather than only upon a request that is difficult for the candidate to formulate. This includes extended deadlines, alternative assessment formats, and the support of interpreters or human mediators when needed.

In addition, every automated recruitment system needs to provide contestation and review channels with effective human mediation, the so-called human-in-the-loop. The absence of any mechanism for appeal against an automated rejection constitutes a violation of the principles of organizational justice and due process (Pasquale, 2015; Colquitt, 2001). It is important to stress that this human mediation cannot be reduced to a formality in which a reviewer merely rubber-stamps the algorithm's decision without actually examining it; it needs to be a genuine review mechanism, with the authority to reverse exclusions when technological or communicational barriers in the process are shown to exist.

***Condition 4: active participation in algorithmic governance***

The fourth condition returns to the principle "Nothing about us without us," historically central to the international disability rights movement. The governance of artificial intelligence systems applied to human resources, including the design of training datasets and the conduct of periodic audits, should not be confined to homogeneous technical teams (Raji et al., 2020). The paid participation of people with disabilities, neurodivergent people, disability studies specialists, and civil society representatives is needed at every stage of the system's life cycle, from defining data requirements through to compliance audits.

Symbolic participation, in which representatives are consulted only at the end of the process or with no real power to influence decisions, does not satisfy this principle. The literature on participatory technology governance suggests that early consultation reduces the costs of later corrections and increases the legitimacy of decisions made, in addition to making it possible to identify problems that developers or managers without direct experience of the barriers in question would not otherwise notice (Floridi et al., 2018).

The synthesis of these four conditions supports the central argument of this article: statistical metrics of algorithmic fairness, such as demographic

parity calculated from historical data, are insufficient when the process retains communicational, sensory, or cognitive barriers in its structure. Algorithmic fairness needs to be understood more broadly, incorporating accessibility, construct validity, the possibility of contestation, and participation in governance as inseparable dimensions of any assessment of fairness in recruitment systems.

### ***Practical implications for people management and public policy***

The proposed theoretical agenda has direct implications both for managerial practice and for public-policy formulation. At the corporate level, human resources departments cannot fully offload the ethical responsibility for personnel selection onto software vendors under the guise of technological neutrality. Organizations that adopt automated screening systems need to establish internal committees responsible for discriminatory-impact testing and accessibility audits before deploying any tool, and to maintain that oversight throughout the entire period the system is in use, not only at the stage of contracting the software.

At the level of public policy, labour and data-protection regulatory bodies need to advance in defining legal obligations for algorithmic explainability and restrictions on the use of automated psychometric tests lacking scientific validation. The European Union's Artificial Intelligence Regulation offers a relevant benchmark in this regard, by classifying recruitment systems as high risk and requiring impact assessment and meaningful human oversight (European Parliament and Council of the European Union, 2024). In Brazil, the legislative debate on artificial intelligence and data protection at work is still ongoing, and the European experience, together with the ILO and OECD reports, can help inform the construction of regulatory frameworks adapted to the national context.

### ***Assessment criteria and impact documentation***

The adoption of automated recruitment systems should be preceded by a documented impact assessment capable of identifying the tool's purpose, the groups potentially affected, the data used, the competencies actually being measured, and the less invasive alternatives available. The purpose of this assessment is not only to verify the system's technical performance, but to understand whether the process creates barriers or produces disproportionate exclusion for particular groups of candidates.

This documentation needs to accompany the system's entire life cycle. Version changes, new training datasets, and changes in criteria can

significantly alter outcomes and therefore need to be recorded and periodically revisited. The hiring organization remains responsible even when it uses technology supplied by third parties; software licensing contracts need to guarantee access to sufficient information to allow for auditing, error correction, and, if necessary, discontinuation of the system.

Impact indicators should combine quantitative data with experiences reported by candidates themselves. Approval rates broken down by group are informative, but on their own they do not reveal difficulties navigating the platform, discomfort during the process, the inability to request an accommodation, or the absence of an explanation for a rejection. Interviews with candidates, accessibility tests conducted by real users, and systematic analysis of complaints complement the purely statistical assessment and offer a fuller picture of the experience lived through the selection process.

### ***Training recruitment teams and shared governance***

Inclusion also depends on training recruiters, managers, and the technical teams responsible for operating and overseeing these systems. Human resources professionals need to understand the foundations of the social model of disability, recognize different legitimate forms of communication, and be able to distinguish essential competencies from behavioral patterns that are merely conventional. Without such training, the human review intended as a safeguard tends to reproduce the very biases already encoded in the automated system, hollowing out its corrective purpose.

Technology governance committees in human resources should include people with disabilities and neurodivergent people on a paid basis and with real power to influence decisions, not merely as consultative guests with no capacity to veto or redirect course. Symbolic participation does not satisfy the co-creation principle discussed under condition 4; it requires that participants have adequate access to the system's technical information, sufficient time for analysis, and a formal means of recording disagreements that are then actually taken into account in the decision-making process.

This shared governance, beyond its ethical justification, strengthens the technical quality of the selection process itself. People who experience everyday barriers identify problems that developers or managers without that direct experience would be unlikely to notice. Early consultation, before a system is definitively implemented, tends to reduce the costs of later correction and to increase the legitimacy of decisions made throughout the selection process.

---

### ***Theoretical propositions and research agenda***

The theoretical agenda proposed in this article makes it possible to formulate propositions that can guide future empirical studies. The first proposition holds that systems designed with accessibility from the outset show a lower drop-out rate among candidates with disabilities over the course of the selection process, compared to systems adapted after the fact. The second proposition indicates that assessments restricted to competencies directly related to the position tend to reduce outcome differences associated with patterns of expression and social interaction, rather than real differences in technical qualification.

The third proposition relates the provision of reasonable accommodations to the perception of procedural justice: processes that offer clear accommodation options, accessible communication channels, and effective human review tend to be perceived as more legitimate by candidates, including those who are not ultimately selected. The fourth proposition suggests that the participation of people with disabilities in the governance of recruitment systems improves the identification of specific barriers and the quality of the corrective measures adopted by organizations.

Future research could compare different commercial platforms, economic sectors, and job types, combining accessibility tests conducted by real users with in-depth qualitative interviews. Experimental studies could assess how different interfaces, deadlines, and response formats affect the performance of candidates with diverse profiles under controlled conditions, making it possible to isolate the specific effect of each barrier identified in this article. Longitudinal studies, finally, could follow cohorts of candidates across several hiring cycles, assessing whether changes to the design of recruitment systems produce, over time, measurable changes in the hiring and retention rates of people with disabilities and neurodivergent people.

### ***The role of technology vendors and the limits of self-regulation***

Much of the discussion on inclusion in algorithmic recruitment tends to focus on the obligations of hiring companies, leaving in the background the responsibility of the software vendors that develop and market these tools. This omission is problematic, because many of the design decisions most decisive for the accessibility and construct validity of what is being assessed are made during the product-development stage, before any hiring company even comes into contact with the system.

The experience of other regulated sectors suggests that corporate self-regulation, isolated from any external oversight, tends to produce inadequate results when a vendor's commercial incentives are not aligned with the interests of the candidates evaluated by its product. A video-screening software vendor has a commercial incentive to maximize the speed and volume of candidates processed, but not necessarily to invest in independent scientific validation of its predictive constructs or in extensive accessibility testing with real users of assistive technologies. For this reason, the theoretical agenda proposed in this article argues that liability for harm arising from indirect algorithmic discrimination should fall jointly on both the hiring company and the technology vendor, which tends to create more balanced incentives for accessibility and scientific validation to be treated as market requirements rather than optional differentiators.

Independent certifications of accessibility and psychometric validity, issued by bodies external to the vendor itself, could play a role similar to that of quality seals already well established in other industries. Without this kind of external verification, the claim that a system is "bias-free" or "internally audited" carries little probative value, given that the very company that developed the system is also the one responsible for evaluating it.

### ***Tensions between operational efficiency and inclusion: a false dilemma***

A recurring argument in corporate discourse on recruitment automation holds that additional requirements around accessibility, scientific validation, and human review would undermine the efficiency that justified adopting these technologies in the first place. This argument deserves careful examination, because it rests on an opposition that does not always hold up empirically.

First, many of the barriers discussed in this article do not stem from automation itself, but from avoidable design choices, such as excessively rigid deadlines or the absence of screen-reader compatibility, which could be corrected without any loss of processing speed. Second, excluding qualified candidates because of communication or accessibility barriers is itself a form of inefficiency, insofar as it shrinks the pool of talent actually evaluated by the organization and can generate downstream costs related to turnover, labour litigation, and reputational damage. Third, human review of contested cases, when well designed, need not involve the entirety of processed candidates, but only a relatively small fraction of questioned decisions, which limits its impact on the overall scale of the process.



The point, then, is not to deny that real costs are associated with implementing the four conditions proposed in this article. It is to recognize that these costs need to be weighed not against a hypothetical scenario of pure efficiency with no safeguards whatsoever, but against a realistic scenario that includes the often-invisible costs of the systematic exclusion of qualified candidates and the legal and reputational risks associated with discriminatory practices identified after the fact.

## CLOSING REMARKS

The incorporation of artificial intelligence into recruitment and selection processes represents a significant shift in how organizations assess candidates, with the potential either to expand or to restrict employment opportunities for people with disabilities and neurodivergent people. Automated tools do indeed offer gains in scale and standardization, but they also carry the real risk of institutionalizing narrow forms of competence, communication, and behavior under the guise of technical neutrality.

This article has sought to contribute to this discussion by proposing a theoretical agenda that brings together the social model of disability, the literature on neurodiversity, organizational justice theory, and studies on algorithmic governance, in dialogue with reports from international bodies on artificial intelligence and work. The four conditions set out in this article, accessibility by design, restricting assessment to essential competencies, offering reasonable accommodations linked to contestation channels with human mediation, and the participation of people with disabilities in system governance, together offer a set of parameters capable of guiding both the design of new human resources technologies and the formulation of specific regulatory policies.

The central argument of this article can be summarized directly: the legitimacy of a recruitment system should not be measured by the technical sophistication of its predictive models, but by its capacity to recognize different legitimate ways of demonstrating professional competence. Purely statistical fairness metrics, calculated from historical data, are not sufficient when the process retains structural barriers of accessibility, communication, or cognition. Algorithmic fairness, in this sense, needs to be understood as a broader problem of institutional design, not merely as a technical problem to be resolved within the algorithm itself.

It is also important to recognize that the experience of exclusion described in this article is not confined to the moment of initial screening. Many of the

same tools discussed here are increasingly being extended to later stages of the employment relationship, such as periodic performance reviews, productivity-monitoring systems, and recommendation algorithms for internal promotion. If the barriers identified in this theoretical agenda are not corrected at the recruitment stage, it is reasonable to assume that they will carry over into these subsequent stages, affecting not only initial access to employment but also the retention and progression of people with disabilities and neurodivergent people throughout their entire professional trajectory. This continuity reinforces the importance of treating algorithmic inclusion as a structural and ongoing problem, rather than as an issue confined to a single decision point in the life cycle of the employment relationship.

This study has limitations that should be clearly acknowledged. It is a theoretical and essayistic investigation, which did not involve primary data collection or empirical analysis of specific commercial algorithms. The propositions formulated in the previous section represent hypotheses to be tested, not definitive conclusions about the workings of particular systems. In addition, the literature reviewed is predominantly produced in Global North contexts, which limits the direct generalization of its conclusions to different labour-market realities, including the Brazilian context, where regulation on artificial intelligence and data protection at work is still being consolidated.

It is recommended that future research pursue empirical investigations, including case studies of organizations that already use automated screening systems, technical analyses of specific commercial platforms, and experimental studies capable of isolating the effect of different accessibility barriers on candidate performance. Studies that directly incorporate the perspective of people with disabilities and neurodivergent people, through participatory methodologies, would be especially valuable for testing and refining the propositions presented in this article, and for verifying to what extent adopting the four proposed conditions actually transforms the rates of attraction, hiring, retention, and talent development among these groups in today's labour market.

Finally, it should be noted that the discussion proposed in this article should not be read as a rejection of the use of artificial intelligence in recruitment processes. The argument put forward here is more specific: automation technologies can, in principle, expand employment opportunities for people with disabilities and neurodivergent people, for example by enabling assessment formats that are more flexible, asynchronous, and adaptable to

different cognitive and sensory profiles. That positive potential, however, depends directly on the design and governance choices examined throughout this article. Without accessibility by design, without demonstrated construct validity, without effective contestation channels, and without the participation of people with disabilities in system governance, automation tends to reproduce and amplify existing barriers. With these conditions in place, it can indeed become an instrument for expanding opportunity. The difference between the two scenarios lies not in the technology itself, but in the institutional decisions that shape its development and its practical application.

#### REFERENCES

- Armstrong, T. (2012). Neurodiversity in the classroom: Building strength-based success in the regular classroom. ASCD.
- Bogen, M., & Rieke, A. (2018). Help wanted: An examination of hiring algorithms, equity, and bias. Upturn.
- Chan, F., Strauser, D., Gerver, R., & Lee, E. J. (2010). Introduction to demand-side factors related to employment of people with disabilities. *Journal of Occupational Rehabilitation*, 20(4), 407-411. <https://doi.org/10.1007/s10926-010-9243-7>
- Colquitt, J. A. (2001). On the dimensionality of organizational justice: A construct validation of a measure. *Journal of Applied Psychology*, 86(3), 386-400. <https://doi.org/10.1037/0021-9010.86.3.386>
- Colquitt, J. A., & Rodell, J. B. (2011). Justice, trust, and trustworthiness: A longitudinal analysis integrating three theoretical perspectives. *Academy of Management Journal*, 54(6), 1183-1206. <https://doi.org/10.5465/amj.2007.0572>
- Crip, T., Ellis, K., & Goggin, G. (2022). Digital accessibility and disability in the platform economy. *Media International Australia*, 183(1), 5-19.
- Eubanks, V. (2018). Automating inequality: How high-tech tools profile, police, and punish the poor. St. Martin's Press.
- European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People: An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707. <https://doi.org/10.1007/s11023-018-9482-5>

- Groom, C., Peters, S., & Owens, K. (2021). Inclusive design in human resources technology: Principles and practice. *Journal of Business and Technical Communication*, 35(1), 88-107.
- Guo, A., Kamar, E., Vaughan, J. W., Wallach, H., & Morris, M. R. (2020). Toward fairness in AI for people with disabilities: A research roadmap. *ACM SIGACCESS Accessibility and Computing*, (125), Article 9. <https://doi.org/10.1145/3386296.3386298>
- International Labour Organization. (2021). *World employment and social outlook 2021: The role of digital labour platforms in transforming the world of work*. International Labour Office.
- Kapp, S. K., & Aiello, R. E. (2022). Neurodiversity, autism, and the workplace: A critical review of employment barriers. *Journal of Occupational Rehabilitation*, 32(1), 1-14.
- Kuzior, A., Kettler, K., & Rąb, Ł. (2020). Digitalization of work and human resources processes: Challenges for responsible management. *Administrative Sciences*, 10(3), 53. <https://doi.org/10.3390/admsci10030053>
- OECD. (2019). *Artificial intelligence in society*. OECD Publishing. <https://doi.org/10.1787/eedfee77-en>
- OECD. (2023). *OECD employment outlook 2023: Artificial intelligence and the labour market*. OECD Publishing. <https://doi.org/10.1787/08785bba-en>
- Oliver, M. (2013). The social model of disability: Thirty years on. *Disability & Society*, 28(7), 1024-1026. <https://doi.org/10.1080/09687599.2013.818773>
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469-481. <https://doi.org/10.1145/3351095.3372828>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33-44. <https://doi.org/10.1145/3351095.3372842>
- Shakespeare, T. (2014). *Disability rights and wrongs revisited*. Routledge.
- Singer, J. (2017). *NeuroDiversity: The birth of an idea*. CreateSpace.
- Torraco, R. J. (2016). Writing integrative literature reviews: Using the past and present to explore the future. *Human Resource Development Review*, 15(4), 404-428. <https://doi.org/10.1177/1534484316671606>

---

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

Whittaker, M., Alper, M., Bennett, C. L., Hendren, S., Kaziunas, L., Mills, M., Morris, M. R., Rankin, J., Rogers, E., Salas, M., & West, S. M. (2019). Disability, bias, and AI. AI Now Institute.

World Health Organization & World Bank. (2011). World report on disability. World Health Organization.

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505-523. <https://doi.org/10.1111/regg.12158>

#### **AUTHOR CONTRIBUTION**

Fábia Aparecida Silva: conceptualization, methodology, investigation, formal analysis, data curation (not applicable, as the study did not produce primary data), writing of the original manuscript, review and editing, supervision, and funding acquisition (not applicable, as the study received no external funding).

#### **BIOGRAPHICAL NOTE**

Fábia Aparecida Silva is affiliated with the Pontifícia Universidade Católica de São Paulo (PUC-SP) and conducts research on labour-market inclusion, diversity, people management, disability, neurodiversity, and the organizational impacts of artificial intelligence.

#### **ETHICAL APPROVAL**

Not applicable. The study is theoretical-integrative in nature and did not involve human participants, identifiable personal data, or animal experimentation.

#### **DATA AVAILABILITY**

Not applicable. The article did not produce a primary dataset.

#### **EDITORIAL DECLARATIONS**

Funding: the study received no external funding.

Conflict of interest: the author declares no conflict of interest.

Use of artificial intelligence: not declared by the author.

© 2026 The author. Published by Noemica under the Creative Commons Attribution 4.0 International License (CC BY 4.0).