

Tackling Algorithmic Disability Discrimination in the Hiring Process: An Ethical, Legal and Technical Analysis

Maarten Buyl*
Ghent University
Belgium
maarten.buyl@ugent.be

Cristina Frattone*
Roma Tre University
Italy
cristina.frattone@uniroma3.it

Christina Cociancig*
University of Bremen
Germany
chrcoc@uni-bremen.de

Nele Roekens*
Unia
Belgium
nele.roekens@unia.be

ABSTRACT

Tackling algorithmic discrimination against persons with disabilities (PWDs) demands a distinctive approach that is fundamentally different to that applied to other protected characteristics, due to particular ethical, legal, and technical challenges. We address these challenges specifically in the context of artificial intelligence (AI) systems used in hiring processes (or automated hiring systems, AHSs), in which automated assessment procedures are subject to unique ethical and legal considerations and have an undeniable adverse impact on PWDs. In this paper, we discuss concerns and opportunities raised by AI-driven hiring in relation to disability discrimination. Ultimately, we aim to encourage further research into this topic. Hence, we establish some starting points and design a roadmap for ethicists, lawmakers, advocates as well as AI practitioners alike.

CCS CONCEPTS

• **Social and professional topics** → **People with disabilities; Governmental regulations**; • **Computing methodologies** → **Artificial intelligence**.

KEYWORDS

ethics of discrimination, persons with disabilities, reasonable accommodation, Artificial Intelligence Act, equality law, data protection law, automated hiring systems, algorithmic discrimination, social justice

ACM Reference Format:

Maarten Buyl, Christina Cociancig, Cristina Frattone, and Nele Roekens. 2022. Tackling Algorithmic Disability Discrimination in the Hiring Process: An Ethical, Legal and Technical Analysis. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAcT '22)*, June 21–24, 2022, Seoul.

*All authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAcT '22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9352-2/22/06...\$15.00
<https://doi.org/10.1145/3531146.3533169>

Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3531146.3533169>

1 INTRODUCTION

Recent advancements in computer science have made *artificial intelligence* (AI) more attractive than ever for both private and public actors. As a result, AI is highly pervasive nowadays; indeed it is difficult to find any aspect of everyday life that is not affected by it. Inspired by the proposed EU Regulation on AI (AIA) [28], we consider AI as software used to generate predictions, recommendations or decisions, developed with machine learning or other statistical methods. To draw statistical inferences, AI systems are “trained” with the aim of generating predictions that are as accurate as possible. Yet, when predictions affect the lives of individuals, competing interests about the optimization of design and predictions take over [20].

In particular, we study the use of AI as part of automated hiring systems (AHSs), with the help of which employers want to attract and hire the best employees with the lowest costs. Prospective employees, however, ought to be protected from unjust discrimination. AI systems should therefore not only strive to predict which employees are a good fit for a job vacancy, but also strive to avoid *algorithmic discrimination* or *algorithmic bias*. This form of discrimination or bias describes the unjust discriminatory manner in which (partially or fully) automated decisions impact a certain protected group or minority. Discrimination can be caused, e.g., by a disproportionate distribution of prediction errors, a faulty design of the AI architecture, or a historical bias rooted in society and perpetuated by statistical inferences [48, 66]. Consequently, in many ways algorithmic bias is the perpetuation of *human bias*.

Our specific concern is for the algorithmic discrimination experienced by persons with disabilities (PWDs)¹ during hiring processes. In a recent study [89] of three AHSs, algorithmic bias was actively being reduced for some protected traits, yet disability was not mentioned in the examples nor in the validation tests. Even though PWDs find no mention in studies like this, it has been recognized that they face unique forms of discrimination in their everyday life, including when they apply for a job [75]. With the aim of inspiring

¹Considering the importance of language in writing about PWDs, we respect the guidelines laid out in [72], which recommend referring to the person first and the disability second.

further research into this area, we provide an analysis of the ethical, legal and technical challenges and opportunities for promoting social justice for job candidates with a disability.

First in Section 2, we argue that non-discrimination and inclusiveness should be regarded as moral obligations to a distinguishable vulnerable group. We identify and examine the importance of ethics of information, including the decision for or against disclosure of a disability. Concerning automated decision making, we refer to AI Ethics publications. A quantitative study of 84 comprehensive guidelines supports our findings of discriminatory practices towards PWDs. We finally bridge the ethical with the legal perspective, establish commonalities and draw conclusions.

Second in Section 3, we focus on the relevant EU legal framework. Notably, EU equality law [74], which prescribes *reasonable accommodation* for PWDs in hiring procedures, including automated processes. We maintain that failure to provide reasonable accommodation to a PWD by an AHS amounts to *prima facie* discrimination, thus shifting the burden of proof to the defendant. However, while adequate accommodation may require disclosure of disability, the processing of disability data for this purpose needs further legal clarification. Finally, we discuss the challenges and opportunities that stem from the rights and obligations provided by the General Data Protection Regulation (GDPR) [78] and the AIA [28].

Third in Section 4, we present technical challenges and possible avenues of research. As each disability is particular and the existing legal framework is applied on a case-by-case basis, data-driven approaches show little promise to solve discrimination by AHSs. AI would be expected to make problematic assumptions, including about the limitations of each PWD, the necessity of an accommodation, and its degree of reasonableness. Finally, we highlight the opportunities that AI systems present for the assessment of PWDs, such as recommending the extent of accommodations and proving discrimination.

2 ETHICAL ANALYSIS

As a first step we consider the ethical perspective of tackling algorithmic discrimination against PWDs. On the question of morality, it is evident that as a society we need to value non-discrimination and inclusion of minority groups and ensure equal opportunities. However, we aim to raise this value to the level of a moral obligation towards PWDs. In this section, we first lay out arguments for such an amplification in Section 2.1. Moreover, we introduce our findings regarding the ethics of AI specifically in Section 2.2 and finally define the role of information for non-discrimination of PWDs in Section 2.3. We conclude by attempting to develop an interdisciplinary bridge from ethics to law in Section 2.4.

2.1 Non-discrimination and inclusion as moral obligations

While anti discrimination law and societies' values agree on non-discrimination and inclusiveness of minority groups, it has not yet been engraved into our collective moral system as a proactive obligation. As [95] demonstrates and arguments presented in Section 4 support, individual case rulings and verdicts are insufficient to change the living experience of PWDs in the long term, even

though they are compliant with law. It is time to initiate a shift in the morality of society towards an inclusive and non-discriminatory workplace.

We consider PWDs as a distinguishable vulnerable group among other minority groups. The following characteristics shortly outlined in [75] are specific to PWDs and provide a compelling argument for additional measures to protect their right to non-discrimination and inclusion:

- (1) PWDs experience separation and segregation distinct from other minority groups. With the intention of providing accessibility, PWDs are provided with separate entries, buses, bathrooms, classrooms, etc. As a result, PWDs are absent from "normal" classrooms, boardrooms, parliament etc. where diverse representation is needed.
- (2) PWDs mostly do not share their otherness with people in their surroundings, e.g., their families, resulting in an ill-equipped social environment to provide necessary accommodations.
- (3) Disabilities are often connected to pain, fatigue and muscle weakness. To manage bodily discomfort, costly resources like assistive technology are needed.
- (4) PWDs have been perpetually discriminated against by means of language. For PWDs to move out of the sphere of "other", they have made use of this instrument of oppression to reflect an ideological shift of perception from "the handicapped" to viewing the disability as a social construct instead of a defining fault.

As a distinguishable vulnerable group, PWDs not only have the legal right to measures of accessibility, but it should be the society's moral imperative to provide those measures. Measures could include technical solutions as well as reasonable accommodation in the employment context.

2.2 Ethics of AI in relation to disabilities

As we focus on AI and the application thereof, we identified issues of discriminatory practices towards PWDs caused by automated decision making. This form of *digital discrimination* reinforces the existing social inequality PWDs experience on a day-to-day basis [17]. AI describes a set of algorithms that rely heavily on finding patterns in high volumes of data. Patterns, for which those with disabilities once again fall outside of the area of applicability, as we examine from a technical perspective in Section 4.1. This is just one effect of the datafication of humans that AI endorses. Unfortunately, it is easy to find examples of discrimination and questionable practices of number-crunching algorithms.

Even though technology can also assist and mitigate some disabilities, PWDs are affected negatively by AI in multiple areas and significantly more so than people without disabilities [61]. Affected areas are many, as AI is used in more and more versatile ways: humans are using AI to judge a PWD's qualification for an employment, loan, social benefits, etc. As a result, we see PWDs falling through the system more than ever. PWDs themselves use AI in many cases in therapeutic and non-therapeutic contexts. However, even in these contexts intended to make life easier, PWDs also experience discrimination, for example, in forms of voice recognition

technology they cannot communicate with, because it simply does not recognize their voice as *human* [19].

Due to their vulnerability as a social group, which is emphasized by the characteristics in Section 2.1, PWDs “face a greater risk of violation of their fundamental rights and freedoms, which justifies adopting specific approaches based on the principle of equality and non-discrimination” [61]. One would expect at least the AI Ethics community to ensure frameworks are protecting the fundamental rights of PWDs. However, as [61] recognized, PWDs are rarely mentioned in literature as anything other than therapeutic/non-therapeutic users of AI, with hardly any mentions of fairly processing PWDs with AHSs, and completely disregarding accessibility as well as the ability of PWDs to inform AI research or policies themselves. Repeatedly, PWDs are excluded and pushed into the shadows of society.

To support the findings of [61] for academic AI Ethics literature, we conducted a quantitative literary study of AI Ethics guidelines. We sampled AI Ethics publications listed in [53] and were able to successfully source 82 out of 84 listed guidelines. In these documents and websites, we conducted a keyword search for “disab*” to match “disabled”, “disability”, “people with disabilities” and similar phrases. The result provided support and motivation alike: only 22 out of 82 guidelines (26.8%) mentioned “disabled people” or “people with disabilities” as a vulnerable group that should receive special consideration in AI Ethics guidelines. Furthermore, PWDs were not acknowledged as research members and policy makers of AI, despite their unique position and ability to shape future research into (assistive) technology and policy to their demands.

2.3 Ethics of information

Most of the discussion surrounding discrimination of PWDs is related to information and not solely limited to the professional context and while applying AI. This includes the delicate decision for or against disclosure of a disability with the employer and the resulting tension of this decision, as well as the lack of information on both the employer’s and PWD’s sides concerning the rights surrounding reasonable accommodation.

The decision for or against disclosure of a disability is motivated by the expectation of positive as well as negative outcomes of this decision. In a survey conducted in Canada [44] only 51% of participants disclosed a disability at some point during their employment, basing their decision on an established secure and safe feeling within their work environment. Of those who decided against disclosure, most reported that they were able to manage their condition without it affecting their performance. However, 37% did not feel secure enough to disclose their disability, and the same number reported that they were concerned about losing their employment. Numbers from Europe comparing unemployment rates of people with disabilities aged 20-64 (17.1%) to people without disabilities (10.2%) support this concern [65]. Evidently, disclosure requires an appreciative and secure work environment that is built on trust of both employer and employee.

To establish trust between employers and employees, legally and morally informed parties are required. For applicants and employees, this manifests itself as being able to disclose their disability as they see fit and because they feel secure enough to do so, as well as

being enabled to claim accommodation where necessary. For the employer, this implies the moral obligation to establish a secure environment for applicants as well as employees, possibly even proactively offering accommodation, but certainly refraining from obstructing these processes. For both parties, information in the form of disclosure and legal knowledge should be available. Even though in some cases it has been argued that the interests of the employer overrule the right to privacy of an employee [84], we strongly argue against mandatory disclosure as it does not support the value of trust between parties nor a person’s right to self-determination.

2.4 Bridging the ethical with the legal perspective

In bridging the worlds of ethics and law, and considering the algorithmic discrimination of PWDs from both perspectives, we argue that the legal understanding of discrimination is an inherently consequentialist one. Consequentialist morality judges actions as good or bad depending on their consequences only. Intention or character traits and virtues have no power over judgment. In EU law we see a congruent logic, as the intent to discriminate is irrelevant. If the effect of an act or practice is to disadvantage a particular group or person, this will amount to discrimination, irrespective of possible well-intentioned or good faith practices.

From a legal perspective, the consequentialist approach is undisputed, as EU equality law does not aim for any moral verdict of the defendant’s character. For ethical deliberations though, consequentialist literature can have a major impact by re-calibrating our moral compass about discrimination against PWDs. For example, arguing with classic utilitarianist John Rawls, if we were to establish a new society tomorrow, the most inclusive and least biased society would be established under the ignorance of one’s own position within this new society [88]. With everyone in mind, including PWDs, people organically agree on principles that guarantee a maximum of individual liberty without restricting the freedom of others (liberty principle) as well as a distribution of equal opportunity to prosper (difference principle). The latter includes a moral obligation to support those who are less advantaged [88]. In a society built on these principles and obligations, we could apply consequentialist jurisdiction as a natural ethical reinforcement to tackle discrimination.

3 LEGAL ANALYSIS

This section examines the relevant applicable legal EU framework and demonstrates how both obstacles and opportunities for social justice stem from its application to AHSs. First, Section 3.1 discusses a PWD’s right to *reasonable accommodation* under EU equality law, which entails the possibility for accommodating hiring procedures if needed. The contextual assessment of reasonable accommodation is hardly reconcilable with data-driven AHSs. However, the use of AI in recruiting may have some unexpected benefits for PWDs in terms of proving *prima facie* discrimination, as explained in Section 4.3.2. Second, Section 3.2 examines the GDPR [78] and the AIA [28], the latter representing the first attempt by a major regulator to provide a comprehensive framework for AI. We demonstrate that the processing of *disability data* for providing accommodation should be allowed under GDPR. Additionally, we shed light on

some new instruments laid down in the AIA that may help fight disability discrimination, such as processing disability data for *de-biasing*, and the proposed *public database* for high-risk AI systems. Finally, Section 3.3 approaches two recent legislative initiatives in the field of algorithmic discrimination comparatively: the proposed EU Directive on platform work [27] and the New York City bill on automated hiring [29], which aim to provide examples for how algorithmic discrimination should or should not be tackled by the law.

3.1 EU non-discrimination law

The risk of algorithmic discrimination in recruitment is globally recognized as a topic of critical importance. Whilst the risk of discrimination for PWDs is widely discussed², surprisingly little has been published on the right to reasonable accommodation for PWDs in automated hiring. The difficulties in presenting a claim of direct or indirect discrimination under EU equality law have been addressed in literature in recent years [43, 105]. Whereas evidential requirements are difficult to meet in case of algorithmic discrimination, the need of an effective legal framework is even more pressing in this context where existing inequalities tend to be exacerbated.

3.1.1 Legal sources and scope of application. Protection against discrimination in Europe is guaranteed by both EU and Council of Europe law. In this paper we mainly focus on the EU *acquis*, most notably Article 21 of the EU Charter of Fundamental Rights, Articles 2 and 10 of the Treaty on the Functioning of the EU, the Employment Equality Directive (2000/78/EC, hereinafter EED) [74], and jurisprudence of the European Court of Justice (CJEU). Finally, the EU is also a party to the United Nation Convention on the Rights of Persons with Disabilities (CRPD).

Whilst the prohibition of disability discrimination in employment under the EED is well-known, the extension of its scope of application to AHSs warrants further elaboration. The EED also applies to “access to employment [...] including selection criteria and recruitment conditions” (Article 3(a)). This includes advertising job positions and job recommendations [6, 40], as well as the design architecture of AHSs. Furthermore, the CJEU and the national courts have embraced a rather broad interpretation of Article 3 [40], to the extent that any *de facto* employment relationship should fall within the Directive’s scope, including employment administrations, temporary agency, and platform workers [43].

Disability is not defined in any EU legal instruments. However, the concept is of utmost importance as the right to reasonable accommodation only applies to PWDs. Nowadays, the spectrum of disability includes situations that had previously remained on its borders. The CJEU has given an autonomous interpretation of the concept of disability, which is in line with the social approach of the CRPD³, and includes medical or incurable diseases. For that to

²Examples can be found in the debates surrounding the algorithmic profiling system of the Public Employment Service in Austria [15], the French public algorithmic student allocation platform [85], or the private HireVue hiring platform [89].

³Chacón Navas v Euresit Colectividades SA (2006), Case C-13/05; Coleman v Attridge Law (2008), Case C-303/06; Odar v Baxter Deutschland GmbH (2012), Case C-152/11.; Ring v Dansk almennyttigt Boligselskab (2013) DAB, Case C-335/11; Werge v Pro Display A/S, Case C-337/11; Z v A Government Department and the Board of Management of a Community School (2014), Case C-363/12; Glatzel v Cv Freistaat Bayern (2014), Case C-356/12; Kaltoft v Kommunernes Landsforening (2014), Case C-354/13.

be the case, two cumulative conditions must be satisfied. Firstly, the illness must result in a physical, mental, intellectual, or sensory impairment which, in interaction with various barriers, is likely to prevent the person concerned from participating fully, effectively and on an equal footing with other workers in working life. Secondly, the illness must be foreseeable or likely to be long-lasting⁴. This social model is highly context-specific though and thus difficult to embed in automated systems, as explained in Section 4 by [22].

3.1.2 The refusal of reasonable accommodation in automated hiring amounts to discrimination. The prohibition of discrimination requires more than a mere refraining from unequal treatment. It may also require an adjustment of an apparently neutral provision or practice if this creates a particular disadvantage for members of a protected group. In particular, the CJEU [7] has acknowledged a transformative function of equality law [70] by also requiring private actors to make proactive efforts against discrimination, and which matches our interpretation of inclusion as a society’s moral obligation in Section 2.1. This is particularly evident with regard to disability, since refusing to provide reasonable accommodation constitutes a *sui generis* form of discrimination under Article 5 EED.

The obligation of reasonable accommodation entails a duty to take appropriate measures, if needed, to enable a person with a disability to have access to, participate in, or advance in employment, or to undergo training, unless such measures would impose a disproportionate burden on the employer (Articles 6 EED and 2(4) CRPD). The CRPD definition is interpreted broadly by the CJEU: measures might be material (e.g., special office, installation of a stair lift) as well as immaterial (e.g., reduction of hours [5], exemption from certain tasks [4]). Even though Article 5 EED does not explicitly state that violation of the right to reasonable accommodation constitutes a form of discrimination, this is indeed well established. In fact, refusal of reasonable accommodation is considered a *sui generis* category of discrimination by Article 2(3) CRPD - which takes precedence over the EED and must thus inform its interpretation [5] -, the European Court of Human Rights [1], and Member States’ legislation and case law [36]. Accordingly, in the *Ruiz Conejero* case, Advocate General Sharpston reasoned that a dismissal resulting from a failure to respect an accommodation obligation “amount[s] to unlawful discrimination for the purposes of Directive 2000/78”, thus confirming our position [3].

Such obligation also applies to the hiring process, and in a twofold manner. First, the employer is obliged to ensure reasonable accommodation during the application process, e.g., extra time allotted to a person with dyslexia taking a written test. Second, when deciding on a PWD’s application, the employer must assess whether the candidate could perform the job if provided with reasonable accommodation, e.g., by splitting a full-time function into part-time functions for a PWD who cannot work full-time. There are no formal requirements to request reasonable accommodation, which must be provided from the time the PWD requires it. Nonetheless, reasonable accommodation is in principle only offered if the need

⁴Examples of disabilities that have been considered as such by the CJEU or by the Member States include, e.g., auditory impairment, autism, dyslexia, back injuries, HIV, diabetes, severe obesity/obesity, multiple sclerosis, depression, the long-term consequences after cancer treatment.

for accommodation is or could be known, which presupposes disclosure of disability. However, disclosure is challenging from an ethical perspective, as demonstrated in Section 2.3. Furthermore, the question of whether, when and to what extent PWDs have to request reasonable accommodation in the context of AHSs remains unanswered, as pointed out in Section 4.2.

A contextual proportionality test is required to ascertain reasonableness of an accommodation. The UN Committee on the Rights of Persons with Disabilities and the CJEU factor in the cost of the proposed adaptation, the organizational impact, influence on the safety of the working environment and the availability of government subsidies [76]. However, the concept of proportionality is rather vague and highly context-specific due to its dependence on the nature of the company and the job. Also, a wide variety of adjustments is possible. Hence, it is difficult to predict which factors need to be considered and how they would be ranked [89]. As argued in Section 4, it must be assumed that AHSs currently do not provide reasonable accommodation for PWDs. Nevertheless, the mere lack of assessing possible reasonable accommodation for a PWD in a recruitment procedure, can amount to a violation of the right to reasonable accommodation and therefore to discrimination.

In light of the above, a specific application of the duty of reasonable accommodation could be to provide an alternative hiring procedure for PWDs who fear to be treated unequally. In practice, this could be realized by the opportunity to “opt out” of the AI-driven selection procedure, leading to human intervention. However, this raises ethical concerns, since not everyone desires to disclose their disability this way, and “opting out” may paradoxically lead to an unfair advantage as human bias may judge more leniently than its computational counterpart. At the same time, by design AHSs do not take into account this obligation for accommodation, irrespective of a request. The CJEU argued in *Feryn* [9] and *Asociația Accept* [10], and recently restated in *Rete Lenford* [8], that actively deterring job applicants belonging to protected groups from applying is to be considered discrimination, even if no recruitment process is on-going. With this, one could argue that discriminatory job recommendation and advertising breach EU equality law as well if they both *in effect* undermine equal access to the labor market [43]. Consequently, applying this framework to the case at hand, AHSs that do not take into account the obligation of reasonable accommodation should be considered *prima facie* discriminatory, regardless of whether accommodation is requested.

3.1.3 Burden of proof and justification. The assessment of reasonableness of an accommodation is context-sensitive and open to judicial interpretation. The nature of AI applications and more precisely the fact that output is statistically probable yet inevitably uncertain given the inherent opacity of AI, makes it difficult for claimants to realize they are treated unequally, to gather the necessary information, and to seek redress [22]. Such technical challenges are further developed in Section 4.2.1. Compared to traditional forms of discrimination, automated discrimination is more abstract and intangible [105].

A key element in EU anti-discrimination law is the shifting of the burden of proof, introduced to contrast the difficulty of proving discrimination suffered by protected groups and by virtue of the principle of effectiveness. The claimant must demonstrate facts

from which discrimination may be presumed and which the alleged defendant should rebut (Article 10 EED). Traditionally, in discrimination cases, certain elements, e.g., the composition of the comparator group or a particular disadvantage need to be demonstrated in order to meet the threshold to shift the burden of proof. However, these challenging requirements do not apply to reasonable accommodation cases. In these cases, the assumed lack of reasonable accommodation provides sufficient evidence to suggest that discriminatory treatment may have occurred and thus raise a presumption of discrimination. Still, a refusal to make reasonable accommodation can be justified if they result in a disproportionate burden on the employer (Article 5 EED).

Moreover, the CJEU recognized in 1989 in the landmark *Danfoss* case on equal pay that “where an undertaking applies a system [of pay] which is totally lacking in transparency, it is for the employer to prove that his practice [...] is not discriminatory” (par. 16) [2]. Arguably, this principle should be applied to AHSs that lack transparency, so that the principle of equal treatment can be applied effectively to automated hiring processes as well. However, taking into account that a refusal to provide information on reasonable accommodation can raise a presumption of discrimination [6], it might well be that a justification under Article 5 that claims a disproportionate burden will be difficult to provide.

3.2 EU technology law

Whereas EU non-discrimination law applies to any hiring process, EU technology law also shapes the regulatory environment of AI-driven hiring. For the purposes of this paper, we use “EU technology law” to refer to the GDPR [78] and the proposed AIA [28]. The relevance of the GDPR is evident since AHSs process personal data. Moreover, it does not merely safeguard privacy, it rather fulfills a much broader objective, which is “protecting fundamental rights and freedoms” (Article 2), including the right to non-discrimination (mentioned in Recitals 71 and 75). Given its broad definition of AI systems, all AHSs fall within the scope of application of the AIA (Article 3(1) and Annex I). Moreover, AI-driven recruiting is listed as a high-risk AI in Annex III, and thus is subject to a strict regime. A comprehensive study of the relevant provisions would be outside the scope of this paper. The analysis rather aims to assess whether the legal obligations imposed on employers and AI providers effectively pursue social justice objectives [41, 42, 92, 107]. On the other hand, technology may inadvertently facilitate the exercise of legal rights, as argued in Section 4.3.2.

3.2.1 Processing of data related to disability: Problem statement.

Assuming that one of the conditions laid down in Article 6 GDPR for processing personal data is met, AHSs additionally process disability data. As a general rule, Article 9(1) GDPR classifies health data as sensitive data, and forbids processing. Even though its purpose is to avoid unfair discrimination against data subjects, Section 4.2 maintains that “fairness through unawareness” is as simplistic a solution as it is an ineffective one. On the contrary, processing of sensitive data may be necessary for mitigating bias, regardless of how counter-intuitive it might seem [62]. Indeed, inequalities are rooted in society [20, 32, 35, 50, 60, 104], and AI systems merely perpetuate these patterns of injustice [16, 31, 38, 77, 82], “freezing the future and scaling the past” [49]. In particular,

as pointed out in Section 2.2, PWDs experience specific forms of discrimination in both the analog and the digital world, which also justifies framing the inclusion of PWDs as a society's moral obligation.

Yet another problem stems from the algorithmic discrimination resulting from proxies (e.g., searching on the internet for a wheelchair, or participating in disability-related volunteering activities), as thoroughly explained in Section 4.2. Such proxies are not personal data in the meaning of the GDPR, and thus fall outside its scope. Nevertheless, the literature has repeatedly pointed out their adverse impact on protected groups [20, 32, 51, 58, 73, 104]. Consequently, having access to the sensitive information stored in proxies might be useful to recognize and mitigate the effect of them. Finally, an employer needs to be aware of someone's disability to make an informed decision about a reasonable accommodation. Besides it being an ethical challenge riddled with positive and negative consequences, awareness may also be problematic under data protection law, due to the prohibition of the resulting processing of health data which falls under Article 9(1) GDPR.

3.2.2 Processing of data related to disability: Exceptions. Staying with the GDPR framework, we assess whether it is lawful for employers to process health data regarding job applicants' disability status, in order to provide them with reasonable accommodation, or to correct an unjustly biased rejection of their application. Article 9(2) provides some exceptions that allow processing of sensitive data. In particular, exceptions a), b), f), and g) are the most relevant for our use case.

The exception based on a data subject's "explicit consent" under letter a) should not be favored in this case, since consent can hardly be given freely in the hiring context, due to the imbalanced power relationship between the employer and the applicant [30, 79]. Rather, processing of sensitive data as "necessary for carrying out the obligations [...] in the field of employment" under letter b) is a more suitable option. Arguably, framing reasonable accommodation as a *proactive obligation* for the employer has the implicit consequence of opening the doors to such an exception. Indeed, granting equal opportunities to applicants with a disability is not only an application of the general right to non-discrimination, but also meets the specific duty to provide reasonable accommodation. Of course, this reasoning applies only to discrimination against PWDs, who are the sole beneficiaries of the right to accommodation under EU equality law. As a consequence, employers should be allowed to process an applicant's sensitive data for evaluating a suitable accommodation.

Nonetheless, exception b) is not available in the event of a *fully automated* decision in the meaning of Article 22 GDPR, which restricts the option to exceptions a) (already excluded) and g). The latter allows processing of sensitive data for reasons of "proportionate substantial public interest". Even though EU equality law is meant to protect individuals' rights, it has an undeniable impact on social justice and is thus not only related to but pursues the public interest [64, 93]. Moreover, giving everyone equal opportunities to succeed improves overall productivity, as it gives a chance to highly skilled applicants who might be rejected because of prejudices [102]. However, letter g) requires an explicit authorization by national or Union law. Thus far, the UK seems to be the only

European State (though not EU Member) to have implemented such a law [23, 24, 99] that allows processing of sensitive data under 23 conditions for the reason of public interest, including "equality of opportunity or treatment" (Data Protection Act 2018, Schedule 1, par. 8). Therefore, despite being abstractly a good fit for the presented use case in theory, such an exception is *de facto* unavailable right now throughout Europe, except for the UK.

If we consider fully automated AHSs as unlawful in the first place, the issue might lose its relevance and exception b) would become the most pertinent one. Arguably, none of the three criteria allowing fully automated decision-making under Article 22(2) [81] is met in this case [57, 79, 89]. If such systems were deployed, candidates could thus allege violation of Article 22(1), regardless of the content of the decision. In a different scenario in which applicants exercise their legal rights under the GDPR or EU equality law, processing of sensitive data is permitted under letter f), allowing it for the purposes of "legal claims", which should not be intended as limited to court proceedings.

For what concerns the proposed AIA, Article 10(5) equips providers of high-risk AI systems, such as AHSs, with an additional exception for processing of sensitive data, namely for *debiasing* purposes. This provision may be invoked by the provider of the system, i.e., the same employer or, in most cases, a third party, at a developing stage. Such a broad exception would have both advantages as well as disadvantages [99]. It may nonetheless respond to criticism raised thus far with regard to the ban on collecting equality data for debiasing purposes [13, 39, 100, 108] and hopefully enable AI systems to provide reasonable accommodation as foreseen in Section 4.3.1.

In conclusion, as the law currently stands, processing of data related to applicants' disabilities should be allowed under the exception laid down in Article 9(2)(b) GDPR. A normative intervention, at an EU or a national level, may be needed for the purposes of the exception under letter g) in the event of fully automated AHSs, even though the very same lawfulness of full automation in recruiting is disputable. Finally, Article 10(5) AIA could also provide a further exception, for processing of sensitive data for debiasing purposes during development stages.

3.2.3 Job applicants' rights. Being data subjects for the purposes of the GDPR, job applicants with disabilities have the rights listed in Articles 12-22. In contrast, the AIA does not provide any legal action for individuals affected by an AI system [94], even though a violation of the prescribed standards may be alleged in a national tort law claim [101]. These data protection remedies ultimately safeguard data subjects' fundamental rights (Article 2 GDPR). Therefore, they can be invoked to enforce a PWD's right to non-discrimination. The *right to access processed personal data* (Article 15(1)) might help gather proof for establishing *prima facie* discrimination and thus reverse the burden of proof [48], as already mentioned in Section 4.3.2. On the contrary, non-discrimination law does not equip protected groups with such a right to access such data, hence data protection law proves itself a precious instrument for fighting algorithmic discrimination [6, 48].

Additionally, data subjects are granted the *right to receive meaningful information about the logic involved* (Article 15(1)(h)) as well as the *right to contest the decision*, the *right to express one's point of view*, and the *right to obtain human intervention* (Article 22(3)).

Indeed, these provisions could be exercised by PWDs to contest a decision in violation of their right to non-discrimination, as well as to receive information on the functioning of the AHS for discovering whether it might discriminate and thus the opportunity to demand an alternative evaluation procedure. However, the importance of these rights should not be overestimated. Not only is their meaning highly debated in literature [37, 63, 67, 91, 103], but their application is also restricted to fully automated decisions in the meaning of Article 22 [26, 81]. Nevertheless, fully automated AHSs would be deemed radically unlawful under Article 22(1) for the above-mentioned reasons.

3.2.4 Employers' obligations. Algorithmic disability discrimination constitutes a severe violation of the principles of accuracy and fairness enshrined in Article 5(1) GDPR [48]. Moreover, Art. 35(1) prescribes that data-controllers, i.e. employers, conduct a *data protection impact assessment* (DPIA) whenever a specific type of processing "is likely to result in a high risk to the rights and freedoms of natural persons". Briefly, the DPIA is an assessment of those risks and the measures determined to address them [80]. It is indisputable that AHSs entail a high risk in the meaning of Article 35 and thus require a DPIA. This provision aims to promote further proactive engagement on behalf of the data controller, following a bottom-up approach [34] and ultimately establishing accountability as laid down in Article 5(2), since the specific safeguards are not dictated by law [12]. Rather, it is the controller's duty to implement them. Hence, it can be sustained that both the GDPR and EU disability law aim at promoting a positive effort on behalf of private actors, and ultimately of society [69, 107]. If duly conducted, the DPIA provides an analysis of systemic risks of unfairness and discrimination, hence facilitating the exercise of job candidates' individual rights to contestation, to express their view, and to human intervention [34, 54, 86].

Transitioning to the AIA, we are met with the idea of a public, central database for the registration of "standalone" high-risk AI systems, such as AHSs. Such a database could assist advocacy bodies and national authorities in detecting discriminatory AHSs [101]. Furthermore, being a high-risk AI application, AHSs ought to be "effectively overseen" by a human for "preventing or minimizing the risks to [...] fundamental rights" (Article 14). Arguably, this provision would strengthen the employer's duty to provide reasonable accommodation. In particular, Article 14(4)(d), entrusts the human supervisor with the decision whether to use the high-risk AI system in the first place, or otherwise disregard its output. Therefore, if an AHS posed a risk of algorithmic discrimination against PWDs, the employer would be obliged to provide an alternative evaluation procedure (Articles 5 EED and 14(4)(d) AIA).

3.3 Comparative insights

Lawmakers in Europe and in the US are currently occupied with algorithmic discrimination, specifically when occurring in automated hiring. Across the Atlantic, the debate on algorithmic discrimination in hiring has been fostered by the recent enactment of New York City's bill on automated employment decision tools last December [29]. It reveals a transparency-oriented approach to non-discrimination, since it mandates the disclosure of technology used to evaluate an application. However, it seems to fall short of

protecting PWDs. In fact, it only requires notice after screening of candidates, hence failing to present applicants with disabilities with the opportunity to ask for accommodation. Evidently, this would only be possible if the PWDs discover *in advance* that the AHS might be unjustly biased against them [90]. Therefore, the GDPR seems more far-reaching and fair than the New York bill at first glance.

Additionally, some recent developments in the comparable field of *platform work* can provide further insight on the selected use case. Notably, platform workers are subject to algorithms that automatically determine work conditions, i.a., their earnings and sanctions. Following the Spanish *Ley Rider* law No 9/2021 [11], the EU Commission published its draft EU Directive on platform work last December [27]. Similar to Article 9 GDPR, the proposed Directive prohibits the processing of sensitive data, including the very same exceptions. However, it is interesting to note that the exception based on the worker's explicit consent (Article 9(2)(a) GDPR) is explicitly excluded here, which confirms its unsuitability as a lawful legal basis for data processing in the employment context [30, 79]. Furthermore, the proposal grants platform workers the right to a human review even when the decision is merely supported by an automated tool; this is contrary to Article 22 GDPR that provides such a right only for those subject to fully automated profiling. Such a provision may not only be a source of inspiration for enhancing the protection of PWDs against discrimination, but since AHSs ought not to be fully automated in light of the previous arguments, the right to obtain human intervention should be extended to every and all AHSs for it to be effective.

4 TECHNICAL ANALYSIS

Besides developments in assistive technology [56, 59, 68, 71, 96], AI is generally developed without the needs of PWDs in mind, giving rise to applications discriminating against them [14, 52, 98]. For other protected groups, AI systems can be adjusted to improve the statistical fairness of its output, e.g., by increasing the importance of underrepresented groups [66]. Yet such a data-driven solution is hard to implement for PWDs, due to the heterogeneity of disability, as well as the need for a context-specific assessment of its relevance in any given circumstance [46, 106].

While we acknowledge that technology alone cannot achieve social justice [21], our aim in this section is to assess how AHSs can be made fairer than they currently are for PWDs. First, we propose to reduce the technical obstacles posed by the heterogeneity of disabilities, by using AI to understand disabilities. We refer to this process as *modeling disability* in Section 4.1. Second, in Section 4.2 we assert that AI systems should be able to leverage this understanding in order to provide the necessary accommodation and judge its reasonableness. Finally, even if unsuccessful in these tasks, we remain optimistic about the opportunities that AI can offer as a source of information discussed in 4.3.

4.1 Technical understanding of disability

Most of the success of recent AI systems has been due to their capacity to detect patterns in large amounts of data, with which decisions could then be automated at a large scale. Yet, as evidenced by the case where Amazon's hiring system learned to systematically

undervalue women [33], not all learned patterns are desirable. To prevent discrimination, AI systems are now being modified to avoid patterns involving protected characteristics, again using a data-driven approach [66]. Standard solutions in this area attempt to remove bias against a protected group by introducing a reversed, synthetic bias in the form of positive discrimination for the entire group. However, such a debiasing technique can only be successful under the assumption that discrimination is homogeneous: e.g. that women are discriminated "because they are women".

4.1.1 Heterogeneity. In contrast to the protected groups that are traditionally considered, the discrimination that PWDs experience is mainly heterogeneous [97]. For instance, an applicant that requires assistive technology to complete an online test may score lower on timing metrics in an AHS [98]. Comparably, candidates with autism may avoid eye gaze during an interview, which may be misinterpreted as a lack of confidence in their abilities. Finally, someone who has been recovering from a long-term illness in the past may be judged unfairly due to the gap in their résumé. Even these examples are each just one simplified instance of how the specific 'disability profile' of a person can lead to their undervaluation in an automated assessment. None of these examples, however, was an AHS discriminatory due to the person having a disability *per se*, but rather because the assessment procedure itself was inadequate for the applicant's specific needs.

Instead, AI systems would do well in giving people the opportunity to make the system aware of their disability in an accessible way. Next, the AI should elaborate on such information to determine whether the specific disability would impact the requested skills and abilities. After all, for the purposes of providing reasonable accommodation (for which we refer to Section 4.2), it is not necessary to know the exact disability profile of an applicant. We should only focus on the specific reasons why they require accommodation in order to be assessed justly. For example, if due to a disability an applicant is unable to maintain eye contact during an interview, it is not relevant whether a visual impairment or autism is the determining factor. In both cases, verbal communication may be more informative about their capabilities.

4.1.2 Modeling disability. As a possible technical solution for managing the heterogeneity of disability, we therefore urge for more research into *modeling disability*, i.e. AI systems that can extract a set of features from an individual disability profile, which we refer to as *limitations* and we assume to be transferable to a wide range of cases. The idea is that a different AI (sub)system can rely on such information to propose reasonable accommodation while being unaware of the individual's sensitive disability profile. Examples of such limitations may include "being wheelchair-bound", "unable to communicate verbally" or "unable to work in loud environments".

Inspiration may be found in recent advances in AI-based human resources (HR) tools that aim to model the skills and competencies of users [45]. A well-known example is the use of such a framework by LinkedIn to more efficiently profile job seekers [47, 87]. The main idea of this framework is to construct an ontology of skills, of which some are then assigned to people based on their own reporting or on automatic extraction from their résumés. Unfortunately, little is known about how these frameworks currently operate due to privacy concerns and economic motivations [45]. Still, it may be

possible to better understand the disability of a person by being able to map it, manually or through automation, to an ontology of limitations. Arguably, the technical challenge is in constructing this ontology with the right level of granularity while taking into account that some disabilities only posed limitations in the past, e.g., an illness from which a person has fully recovered.

At the same time, we ought to be aware of the risks in attempting to model disabilities. Though it may help in simplifying the automation of a hiring process, the same simplification may lead to a characterization of limitations that is too vague to provide a true understanding of the disability, or one that is too rigid to provide the nuance that a PWD deserves. Fundamentally, the standardization of a person's abilities poses ethical concerns (already discussed in Section 2.2) that should be carefully balanced against the potential practical advantage that such standardization provides [106]. A system to model disabilities should thus be careful to always put the holistic needs of the individual first.

4.2 Reasonable accommodation

Given the myriad of ways in which data and machine learning models can be biased [66], we should assume that AHSs will not fairly evaluate PWDs unless the process is specifically designed to reduce discrimination. A well-known result from fairness research is that simply removing sensitive data from the algorithm's input (achieving the so-called "fairness through unawareness") is insufficient, since the sensitive information relating to an individual may still be correlated with other, non-protected, features, or even proxies, in a subtle manner that ultimately cannot be ignored [83]. Similar to Section 4.1, we are concerned with the specific heterogeneous risks that PWDs face during these processes. We will explore several types of accommodations and then evaluate their reasonableness.

4.2.1 Providing accommodation. Under the assumption that the AI system has a practical model of a person's disability as discussed in Section 4.1, we distinguish four ways in which a standard AHS may need to be modified in order to tackle the heterogeneity of disabilities and provide reasonable accommodation. First, it may be necessary to adjust the predictions of a model due to a group of persons with a particular disability having been undervalued in the past. Second, some disabilities may render the standard hiring process simply inaccessible, e.g., mutism if verbal communication is part of the assessment. Third, failing to automate accommodation, the only solution may be replacing the AHS with a human. Fourth, even if the candidate is fairly assessed, they may still require accommodation in the future to perform their job.

Adjusting predictions. The adjustment of predictions is the closest to the standard fairness modifications made for other protected groups. It is necessary when the data or model are imperfect [66]. For example, a person that suffered from a long-term illness may have only been employed part-time in the past. Especially if they have now recovered, they should not only be recommended part-time jobs in the future. In a different example, we may see that a data-driven AI system fails to provide any sensible assessment for people who use a wheelchair, simply because the model is unable to learn which jobs are accessible to those individuals and which

are not. The fact that the relevance of limitations is job-specific makes a data-driven approach particularly difficult.

Multimodal input. A hiring process may be setup in such a way that it, in theory, gathers enough information about a person without job-related limitations in order to make a fair assessment. However, if it involves a verbal interview for example, a person needing sign language interpretation can never be fairly judged based on audio data alone. It is thus necessary to design the hiring procedure in such a way that accepts input from multiple sources in different formats, e.g., evaluating either audio or sign language video data. Accommodation is then provided for some PWDs by excusing them from being tested in a format that is inaccessible to them. The problem where people need to receive the same assessment, even though they are ‘seen’ in different types of input format, is referred to as a *multimodal* problem [18]. While perceiving the world in this way is natural for humans, it is not the norm for AI systems and is therefore also not readily accessible for employers to use.

Human intervention. An AI system should have the ability to judge whether it can grant adequate accommodation. Otherwise, a human should replace it. However, this is infeasible for platforms where decisions are automated on a massive scale like LinkedIn [55]. Moreover, it may not improve fairness for PWDs, because there is no guarantee that the human will provide a less discriminatory or biased assessment. Arguably, the main advantage is a legal one: as discussed in Section 3.2, automated decisions are held to a higher standard than those taken by humans.

Future accommodation. We finally mention the problem of providing accommodation for the actual job. A classic example is a person who is wheelchair-bound and thus may require changes to the architecture at the work place to properly perform their job. In order to also automate the end of the hiring process, AI should be able to make recommendations so that the job is made reasonably accessible.

4.2.2 Estimating reasonableness of accommodations. For every type of accommodation to the hiring process listed in Section 4.2.1, we can question whether it is reasonable to make such adjustments. There is no established legal standard as to what is reasonable, since every accommodation should ideally be judged on a case-by-case basis. To judge the reasonableness of accommodations at the workplace, AI systems could be used to predict its economic cost for the employer. If this cost is accurately estimated as unreasonable, then the result could be used by the employer as motivation to not grant that accommodation. Building and training a prediction algorithm is again challenging due to the heterogeneity of disability and of possible tasks and assignments. Additionally, the reasonableness of such accommodations also changes over time as technology improves. It should also be questioned which adjustments to the AHS itself are reasonable. The adjustment of predictions for every kind of limitation or the possibility of learning from multimodal input are open research problems. Even if possible with the state-of-the-art, such accommodation carries with it a significant added complexity of the AI system and additional data requirements.

4.3 AI as a source of information

So far, our technical analysis has mainly consisted of listing the challenges and risks in providing reasonable accommodation in AHSs. However, some of the accommodations we deem ethically and legally necessary may not yet be possible to automate with present technology, due to the heterogeneity of disability that is so incongruent with the current paradigm of machine learning that is consistently data-driven.

Though we hope for further research into ways to model disabilities and provide reasonable accommodation, we also point out two ways in which AI systems can already improve the experience of PWDs in the job market. First, we can leverage data-driven approaches to recommend reasonable accommodations to job seekers, employees and employers. Second, we can benefit from the (partial) automation of hiring processes to protect the rights of PWDs.

4.3.1 Recommending accommodation. Most hiring tools today are not fully automated (and should not be, as argued in Section 3.2.2). Instead, AI-based tools for hiring are developed to replace specific intermediate steps in the end-to-end hiring process [25]. In a semi-automated hiring process where humans still make the most important decisions, AI tools need not provide reasonable accommodation directly.

Rather, if a candidate chooses to disclose their disability, an AI-based tool could signal the need for accommodation to a human decision maker. If the disability is understood well enough by the AI system (e.g., using a technique discussed in Section 4.1), then the AI can only report the relevant limitations while keeping the sensitive disability information hidden from the human decision maker. If they agree that accommodation is needed, AI systems could then be used to recommend accommodations. Of the examples listed in Section 4.2, the simplest accommodation to suggest is that of *future* accommodation, as it does not require any modification to the assessment procedure. As discussed in Section 2.3, prospective employees also benefit from such recommendations, as information can help establish trust between them and their employers. AI can provide an invaluable benefit to both parties by making the information on reasonable accommodations more easily accessible.

4.3.2 Protecting rights. If a part of the hiring process is automated through the use of a sufficiently transparent AI, then the use of it may inadvertently facilitate proving discrimination. While gathering proof is arguably one of the most burdensome tasks for PWDs under EU non-discrimination law, in the future it may be gathered from the *ex post* analysis of the automated decision-making process. That way it can, e.g., be established whether individuals with a certain limitation are consistently and unjustly undervalued. As data subjects under EU technology law (see Section 3.2), PWDs are also facilitated in exercising their right to access and right to contest. As such, the AI itself may be the surprise witness of an alleged discrimination.

5 CONCLUSION

Fundamentally, PWDs form a distinguishable vulnerable group that experience separation and discrimination in a manner unlike other protected minorities. As job seekers, PWDs are uniquely disadvantaged in that the disclosure of their disability is difficult to control

and may often not be in their best interest, causing a lack of trust in prospective employers. Conversely, these employers ought to feel morally obligated to prevent algorithmic discrimination if they employ AHSs, even though their good faith does not excuse breaches of EU equality law by virtue of its consequentialist approach.

The attention of lawmakers directed towards algorithmic discrimination in general, and to AHSs specifically, further demonstrates the urgency of the matter. Our analysis of current EU non-discrimination law concludes that the right of PWDs to reasonable accommodation also applies to cases in which AHSs are used. EU technology law provides instruments to ensure reasonable accommodation in automated hiring, paving the way to utilizing disability data for the purposes of improving social justice, and imposing further obligations on employers that may help enforce the applicants' right to non-discrimination. Though there might be no determined legal solution to algorithmic discrimination, the proposed AIA (and the proposed Directive on platform work) reveal(s) the EU Commission's willingness to address algorithmic bias.

In practice, AI systems struggle to supply reasonable accommodation due to a heterogeneity that is unique to PWDs as a protected group, which makes a data-driven solution infeasible. To some extent, this heterogeneity can be tackled by employing AI to better understand disabilities. Yet even with this understanding, the technical adjustments needed to provide accommodation are varied and challenging, presenting obvious risks for PWDs and therefore a demand for further research and testing in the long term. In the short term, AI can offer opportunities for PWDs by acting as a source of information during the hiring process, either by assisting in providing reasonable accommodation in hiring decisions that are not automated, or by presenting proof of the lack of accommodation in hiring decisions that are automated.

Our analysis has shed some much-needed light on the many forms in which PWDs experience algorithmic discrimination by AHSs. The use of data-driven tools for hiring purposes deserves careful consideration, given their impact on an individual's life. Decent work and employment are vital for the well-being and dignity of all.

ACKNOWLEDGMENTS

We thank Mark Dempsey for his fruitful comments. This research was supported by the Flemish Government (AI Research Program), the BOF of Ghent University (PhD scholarship BOF20/DOC/144) and Roma Tre University.

REFERENCES

- [1] 51500/08, *Çam v Turkey*, European Court of Human Rights, Judgement of 23 February 2016. [Online]. Available: https://www.legislationline.org/download/id/7067/file/ECHR_Cam_v_Turkey_2016_en.pdf.
- [2] C-109/88, *Handels- og Kontorfunktionærernes Forbund I Danmark v Danfoss* represented by Dansk Arbejdsgiverforening, Judgement of 17 October 1989. [Online]. Available: <https://curia.europa.eu/juris/liste.jsf?language=de&num=C-109/88&td=ALL>.
- [3] C-270/16, *Carlos Enrique Ruiz Conejero v Ferroservicios Auxiliares SA and Ministerio Fiscal*, Judgement of 18 January 2018. [Online]. Available: <https://curia.europa.eu/juris/document/document.jsf?text=&docid=198527>.
- [4] C-312/11, *European Commission v Republic of Italy*, Judgement of 4 July 2013. [Online]. Available: <https://curia.europa.eu/juris/document/document.jsf?text=&docid=139105>.
- [5] C-335/11 and C-337/11, *Jette Ring*, §29-30, Judgement of 11 April 2013. [Online]. Available: <https://curia.europa.eu/juris/document/document.jsf?docid=136161&doclang=EN>.
- [6] C-415/10, *Galina Meister v Speech Design Carrier Systems GmbH*, Judgement of 19 April 2012. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1490004730485&uri=CELEX:62010CJ0415>.
- [7] C-43/75, *Gabrielle Defrenne v Société anonyme belge de navigation aérienne Sabena*, Judgement of 8 April 1976. [Online]. Available: <https://curia.europa.eu/juris/liste.jsf?language=de&jur=C,T,F&num=C-43/75>.
- [8] C-507/18, *Associazione Avvocatura per i diritti LGBTI v Rete Lenford*, Corte suprema di cassazione, Judgement of 23 April 2020. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=ecli:ELCI%3AEU%3AC%3A2020%3A289>.
- [9] C-54/07, *Centrum voor gelijkheid van kansen en voor racismebestrijding v Feryn NV*, Arbeidshof te Brussel, Judgement of 10 July 2008. [Online]. Available: <https://eur-lex.europa.eu/legal-content/DE/TXT/?uri=CELEX%3A62007CJ0054>.
- [10] C-81/12, *Asociația Accept v Consiliul Național pentru Combaterea Discriminării, Curte de Apel București*, Judgement of 25 April 2013. [Online]. Available: <https://eur-lex.europa.eu/legal-content/DE/TXT/?uri=CELEX%3A62012CJ0081>.
- [11] *Real Decreto-ley 9/2021, Modificación de la Ley del Estatuto de los Trabajadores*, Judgement of 11 May 2021. [Online]. Available: <https://www.boe.es/eli/es/rdl/2021/05/11/9>.
- [12] Alberto Alemanno. 2012. *Regulating the European Risk Society*. In *Better Business Regulation in a Risk Society*. Springer, 37–56.
- [13] Katayoun Alidadi. 2017. Gauging process towards equality? Challenges and best practices of equality data collection in the EU. *Eur Equality Law Rev* 2017, 2 (2017), 15–27.
- [14] Robin Allen and Dee Masters. 2020. Artificial Intelligence: the right to protection from discrimination caused by algorithms, machine learning and automated decision-making. In *ERA Forum*, Vol. 20. Springer, 585–598.
- [15] Doris Allhutter, Florian Cech, Fabian Fischer, Gabriel Grill, and Astrid Mager. 2020. Algorithmic profiling of job seekers in Austria: how austerity politics are made effective. *Frontiers in big Data* 3 (2020), 5.
- [16] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias: There's software used across the country to predict future criminals. *And it's biased against blacks*. *ProPublica* 23 (2016), 77–91.
- [17] Xavier Ferrer Aran, Tom van Nuenen, Jose Such, Mark Coté, and Natalia Criado Pacheco. 2020. Bias and Discrimination in AI: a cross-disciplinary perspective. *IEEE Technology and Society Magazine* 40, 2 (2020).
- [18] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* 41, 2 (2018), 423–443.
- [19] Francisco Jose Bariffi. 2021. Artificial Intelligence, Human Rights and Disability. *Pensar-Revista de Ciências Jurídicas* 26, 2 (2021).
- [20] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [21] Cynthia L Bennett and Os Keyes. 2020. What is the point of fairness? Disability, AI and the complexity of justice. *ACM SIGACCESS Accessibility and Computing* 125 (2020), 1–1.
- [22] Reuben Binns. 2020. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 514–524.
- [23] Reuben Binns and Reuben Kirkham. 2021. How Could Equality and Data Protection Law Shape AI Fairness for People with Disabilities? *ACM Transactions on Accessible Computing (TACCESS)* 14, 3 (2021), 1–32.
- [24] European Data Protection Board. 2020. Guidelines 05/2020 on consent under Regulation 2016/679. [Online]. https://edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202005_consent_en.pdf
- [25] Miranda Bogen and Aaron Rieke. 2018. Help wanted: An exploration of hiring algorithms, equity, and bias. <https://apo.org.au/sites/default/files/resource-files/2018-12/apo-nid210071.pdf>. (2018).
- [26] Lee A Bygrave. 2019. *Minding the Machine v2. 0: The EU General Data Protection Regulation and Automated Decision-Making*. In *Algorithmic Regulation*. Oxford University Press, 248–262.
- [27] European Commission. 2021. Commission proposals to improve the working conditions of people working through digital labour platforms. https://ec.europa.eu/commission/presscorner/detail/en/ip_21_6605
- [28] European Commission. 2021. Proposal for a Regulation laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. *COM(2021) 206 final* (2021). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>
- [29] New York City Council. 2021 [Online]. A Local Law to amend the administrative code of the city of New York, in relation to automated employment decision tools, 2021/144. <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9>
- [30] European Data Protection Supervisor European Union Agency for Fundamental Rights Council of Europe, European Court of Human Rights. 2018. *Handbook on European data protection law*. Publications Office of the European Union, Luxembourg.
- [31] Kate Crawford. 2021. *The Atlas of AI*. Yale University Press.

- [32] Natalia Criado and Jose M Such. 2019. Digital discrimination. In *Algorithmic Regulation*. OUP.
- [33] Jeffrey Dastin. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters* (Oct. 2018). <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- [34] Giovanni De Gregorio and Pietro Dunn. 2021. Profiling under Risk-based Regulation: Bringing together the GDPR and the DSA. (2021). https://assets.ctfassets.net/iapmw8ie3ije/5EuxLPaUlsGt7R6PgeuFK/c9269e55e10bb2a7a0b392624c08f4d0/De_Gregorio_Dunn_My_Data_is_Mine_1_.pdf Unpublished paper.
- [35] Virginia Dignum. 2021. The Myth of Complete AI-Fairness. In *International Conference on Artificial Intelligence in Medicine*. Springer, 3–8.
- [36] L. Dixon, N. Clancy, B. M. Miller, S. Hoegberg, M. M. Lewis, and B. Bender et al. 2021 [Online]. Reasonable Accommodation for Persons with Disabilities: Exploring Challenges Concerning its Practical Implementation. An Equinet Publication. <https://equineteurope.org/wp-content/uploads/2021/03/Reasonable-Accommodation-Disability-Discussion-Paper.pdf>
- [37] Lilian Edwards and Michael Veale. 2018. Enslaving the algorithm: From a “Right to an Explanation” to a “Right to Better Decisions”? *IEEE Security & Privacy* 16, 3 (2018), 46–54.
- [38] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.
- [39] European Union Agency for Fundamental Rights (FRA). 2013 [Online]. FRA Opinion on the situation of equality in the European Union 10 years on from initial implementation of the equality directives. <https://fra.europa.eu/en/publication/2013/fra-opinion-situation-equality-european-union-10-years-initial-implementation>
- [40] European Union Agency for Fundamental Rights (FRA) and European Court of Human Rights. 2018. *Handbook on European non-discrimination law*. Publications Office of the European Union, Luxembourg.
- [41] Iason Gabriel. 2020. Artificial intelligence, values, and alignment. *Minds and machines* 30, 3 (2020), 411–437.
- [42] Iason Gabriel. 2021. Towards a Theory of Justice for Artificial Intelligence. *arXiv preprint* (2021).
- [43] Janneke Gerards and Raphaela Xenidis. 2021. Algorithmic Discrimination in Europe: Challenges and Opportunities for Gender Equality and Non-Discrimination Law. (2021).
- [44] Monique AM Gignac, Arif Jetha, Kathleen A Martin Ginis, and Selahadin Ibrahim. 2021. Does it matter what your reasons are when deciding to disclose (or not disclose) a disability at work? The association of workers’ approach and avoidance goals with perceived positive and negative workplace outcomes. *Journal of Occupational Rehabilitation* (2021), 1–14.
- [45] Akshay Gugnani and Hemant Misra. 2020. Implicit Skills Extraction Using Document Embedding and Its Use in Job Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 13286–13293.
- [46] Anhong Guo, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach, and Meredith Ringel Morris. 2020. Toward fairness in AI for people with disabilities: a research roadmap. *ACM SIGACCESS Accessibility and Computing* 125 (2020).
- [47] Viet Ha-Thuc, Ganesh Venkataraman, Mario Rodriguez, Shakti Sinha, Senthil Sundaram, and Lin Guo. 2015. Personalized expertise search at LinkedIn. In *IEEE International Conference on Big Data (Big Data)*. IEEE, 1238–1247.
- [48] Philipp Hacker. 2018. Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under EU law. *Common Market Law Review* 55, 4 (2018).
- [49] Mireille Hildebrandt. 2020. Code-driven Law: Freezing the Future and Scaling the Past. *Is Law Computable?: Critical Perspectives on Law and Artificial Intelligence* (2020), 67.
- [50] Mireille Hildebrandt. 2021. Discrimination, Data-driven AI Systems and Practical Reason. *European Data Protection Law Review* 7 (2021), 358–366.
- [51] Mireille Hildebrandt. 2021. The Issue of Proxies, And why EU law matters for recommender systems. (2021). <https://osf.io/preprints/socarxiv/45x67/> (submitted for publication).
- [52] Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Unintended machine learning biases as social barriers for persons with disabilities. *ACM SIGACCESS Accessibility and Computing* 125 (2020), 1–1.
- [53] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [54] Margot E Kaminski and Gianclaudio Malgieri. 2020. Multi-layered explanations from algorithmic impact assessments in the GDPR. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 68–79.
- [55] Krishnaram Kenthapadi, Benjamin Le, and Ganesh Venkataraman. 2017. Personalized job recommendation system at linkedin: Practical challenges and lessons learned. In *Proceedings of the eleventh ACM conference on recommender systems*. 346–347.
- [56] Łukasz Kidziński, Carmichael Ong, Sharada Prasanna Mohanty, Jennifer Hicks, Sean Carroll, Bo Zhou, Hongsheng Zeng, Fan Wang, Rongzhong Lian, Hao Tian, et al. 2020. Artificial intelligence for prosthetics: Challenge solutions. In *The NeurIPS’18 Competition*. Springer, 69–128.
- [57] Pauline T Kim and Matthew T Bodie. 2021. Artificial Intelligence and the Challenges of Workplace Discrimination and Privacy. *Journal of Labor and Employment Law* 35, 2 (2021), 289–315.
- [58] Joshua A Kroll, Joanna Huey, Solon Barocas, Edward W Felten, Joel R Reidenberg, David G Robinson, and Harlan Yu. 2017. Accountable Algorithms. *University of Pennsylvania Law Review* 165 (2017), 633–705.
- [59] Marco Leo, G Medioni, M Trivedi, Takeo Kanade, and Giovanni Maria Farinella. 2017. Computer vision for assistive technologies. *Computer Vision and Image Understanding* 154 (2017), 1–15.
- [60] Jonas Lerman. 2013. Big data and its exclusions. *Stan. L. Rev. Online* 66 (2013), 55.
- [61] Aspen Lillywhite and Gregor Wolbring. 2020. Coverage of artificial intelligence and machine learning within academic literature, Canadian newspapers, and twitter tweets: The case of disabled people. *Societies* 10, 1 (2020), 23.
- [62] Timo Makkonen. 2016. *European Handbook on Equality Data*. Publications Office of the European Union, Luxembourg. https://ec.europa.eu/info/sites/default/files/european_handbook_on_equality.pdf
- [63] Gianclaudio Malgieri and Giovanni Comandè. 2017. Why a right to legibility of automated decision-making exists in the general data protection regulation. *International Data Privacy Law* (2017).
- [64] Alessandro Mantelero. 2017. From group privacy to collective privacy: towards a new dimension of privacy and data protection in the big data era. In *Group Privacy*. Springer, 139–158.
- [65] Marie Lecerf. 2020. Employment and disability in the European Union. Briefing PE 651.932. [https://www.europarl.europa.eu/RegData/etudes/BRIE/2020/651932/EPRS_BRI\(2020\)651932_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2020/651932/EPRS_BRI(2020)651932_EN.pdf). Accessed: 2021-01-7.
- [66] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [67] Isak Mendoza and Lee A Bygrave. 2017. The right not to be subject to automated decisions based on profiling. In *EU Internet Law*. Springer, 77–98.
- [68] Nafisa Mostofa, Kelly Fullin, Sharare Zehetabian, Safa Bacanlı, Ladislav Bölöni, and Damla Turgut. 2020. IoT-Enabled Smart Mobility Devices for Aging and Rehabilitation. In *IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
- [69] Elise Muir. 2013. The transformative function of EU equality law. *European Review of Private Law* 21, 5/6 (2013), 1231–1253.
- [70] Elise Muir. 2018. *EU Equality Law: The First Fundamental Rights Policy of the EU*. Oxford University Press.
- [71] Davide Mulfari. 2018. A TensorFlow-based assistive technology system for users with visual impairments. In *Proceedings of the 15th International Web for All Conference*. 1–2.
- [72] ADA National Network. 2018. Guidelines for Writing About People With Disabilities. <https://adata.org/factsheet/ADANN-writing>. Accessed: 2021-01-7.
- [73] Selin Nugent, Paul Jackson, Susan Scott-Parker, James Partridge, Rebecca Raper, Chara Bakalis, Alex Shepherd, Arijit Mitra, Jintao Long, Kevin Maynard, et al. 2020. Recruitment AI has a Disability Problem: questions employers should be asking to ensure fairness in recruitment. (2020).
- [74] Council of the European Union. 2000. Council Directive 2000/78/EC of 27 November 2000 establishing a general framework for equal treatment in employment and occupation. *Official Journal of the European Union* L 303 (2000). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32000L0078>
- [75] Rhoda Olkin. 2002. Could you hold the door for me? Including disability in diversity. *Cultural Diversity and Ethnic Minority Psychology* 8, 2 (2002), 130.
- [76] Committee on the Rights of Persons with Disabilities. 2018. General comment no. 6 on equality and non-discrimination. In *Proceedings of the 19th session, Geneva*. <https://digitalibrary.un.org/record/1626976>
- [77] Cathy O’neil. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, New York.
- [78] European Parliament and Council. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* L 119 (2016), 1–88. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [79] Data Protection Working Party. 2001. Opinion 8/2001 on the processing of personal data in the employment context. 5062/01/EN/Final, WP 48. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2001/wp48_en.pdf
- [80] Data Protection Working Party. 2017. Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is “likely to result in a high risk” for the purposes of Regulation 2016/679. WP 248 rev.01. https://ec.europa.eu/newsroom/just/document.cfm?doc_id=47711
- [81] Data Protection Working Party. 2018. Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. WP251rev.01. <https://ec.europa.eu/newsroom/article29/redirection/document/49826>

- [82] Frank Pasquale. 2015. *The black box society*. Harvard University Press.
- [83] Dino Pedreshi, Salvatore Ruggieri, and Franco Turini. 2008. Discrimination-aware data mining. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 560–568.
- [84] Anders J Persson and Sven Ove Hansson. 2003. Privacy at work—ethical criteria. *Journal of Business Ethics* 42, 1 (2003), 59–70.
- [85] Robin Allen QC and Dee Masters. 2020 [Online]. Regulating for an Equal AI: A new Role for Equality Bodies. Meeting the new Challenges to Equality and Non-Discrimination from Increased Digitisation and the Use of Artificial Intelligence. An Equinet Publication. https://equineteurope.org/wp-content/uploads/2020/06/ai_report_digital.pdf
- [86] Claudia Quelle. 2018. Enhancing compliance under the general data protection regulation: The risky upshot of the accountability-and risk-based approach. *European Journal of Risk Regulation* 9, 3 (2018), 502–526.
- [87] Rohan Ramanath, Hakan Inan, Gungor Polatkan, Bo Hu, Qi Guo, Cagri Ozcaglar, Xianren Wu, Krishnaram Kenthapadi, and Sahin Cem Geyik. 2018. Towards deep and representation learning for talent search at LinkedIn. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 2253–2261.
- [88] John Rawls. 1971. *A theory of justice*. Harvard University Press.
- [89] Javier Sánchez-Monedero, Lina Dencik, and Lilian Edwards. 2020. What does it mean to 'solve' the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 458–468.
- [90] Matt Scherer and Ridhi Shetty. 2021. NYC draft bill on AI in hiring needs higher and clearer hurdles. <https://cdt.org/insights/nyc-draft-bill-on-ai-in-hiring-needs-higher-and-clearer-hurdles/>
- [91] Andrew Selbst and Julia Powles. 2018. Meaningful Information and the Right to Explanation. In *Conference on Fairness, Accountability and Transparency*. PMLR, 48–48.
- [92] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
- [93] Nathalie A Smuha. 2021. Beyond a human rights-based approach to AI governance: Promise, pitfalls, plea. *Philosophy & Technology* 34, 1 (2021), 91–104.
- [94] Nathalie A Smuha, Emma Ahmed-Rengers, Adam Harkens, Wenlong Li, James MacLaren, Riccardo Piselli, and Karen Yeung. 2021. How the EU Can Achieve Legally Trustworthy AI: A Response to the European Commission's Proposal for an Artificial Intelligence Act. *SSRN* (2021).
- [95] Kathie Snow. 2004. The Moral Imperative of Inclusion. <https://nebula.wsimg.com/f091343f852ba6a73471bcf8d76e6f88?AccessKeyId=9D6F6082FE5EE52C3DC6>. Accessed: 2021-01-7.
- [96] Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 2020. "Person, Shoes, Tree. Is the Person Naked?" What People with Vision Impairments Want in Image Descriptions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [97] Shari Trewin. 2018. AI fairness for people with disabilities: Point of view. *arXiv preprint* (2018).
- [98] Shari Trewin, Sara Basson, Michael Muller, Stacy Branham, Jutta Treviranus, Daniel Gruen, Daniel Hebert, Natalia Lyckowski, and Erich Manser. 2019. Considerations for AI fairness for people with disabilities. *AI Matters* 5, 3 (2019), 40–63.
- [99] Marvin van Bekkum and Frederik Zuiderveen Borgesius. 2021. Using special category data to prevent discrimination: does the GDPR need a new exception? Digital Legal Talks by Tilburg University. (2021). https://www.sectorplands.nl/wordpress/wp-content/uploads/2021/11/DLT-2021_paper_35.pdf Unpublished paper.
- [100] Jozefien Van Caeneghem. 2019. Ethnic Data Collection: Key Elements, Rules and Principles. In *Legal Aspects of Ethnic Data Collection and Positive Action*. Springer, 155–257.
- [101] Michael Veale and Frederik Zuiderveen Borgesius. 2021. Demystifying the Draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International* 22, 4 (2021), 97–112.
- [102] Sandra Wachter. 2021. How Fair AI Can Make Us Richer. *European Data Protection Law Review* 7, 3 (2021), 367–372. <https://doi.org/10.21552/edpl/2021/3/5>
- [103] Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. 2017. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law* 7, 2 (2017), 76–99.
- [104] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2020. Bias preservation in machine learning: the legality of fairness metrics under EU non-discrimination law. *W. Va. L. Rev.* 123 (2020), 735.
- [105] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2021. Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. *Computer Law & Security Review* 41 (2021).
- [106] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kazunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, bias, and AI. *AI Now Institute* (2019).
- [107] Josephine Wolff. 2021. How Is Technology Changing the World, and How Should the World Change Technology? *Global Perspectives* 2, 1 (2021).
- [108] Indrè Žliobaitė and Bart Custers. 2016. Using sensitive personal data may be necessary for avoiding discrimination in data-driven decision models. *Artificial Intelligence and Law* 24, 2 (2016), 183–201.