Bias in AI Hiring Tools:
Impacted Groups, Legal Risks, Historical Foundations, and Next Steps

Eesha Bayana

Abstract
This paper investigates the role and influence of artificial intelligence (AI) in applicant tracking systems (ATS) on marginalized groups within the course of the job recruitment process. Although AI-powered ATS may ensure efficiency in recruitment through automated resume screenings and interview analysis, it extends the circle of historic bias, which affects immigrants, persons with disabilities, women, and those with non-Anglo names. These systems tend to screen out qualified candidates for non-standard language, gaps in employment, or characteristics irrelevant to job performance. These practices only further perpetuate economic disparities and psychological harm within already marginalized communities. Notable cases involving such firms as Amazon and Workday demonstrate the legal consequences connected with these discriminatory practices, showcasing the need for organizations to track AI bias to avoid legal liabilities. We call for a review of hiring practices and an inclusive redesign of AI systems that holistically evaluates diverse candidate experiences. It necessitates a paradigm shift in AI ethics, inclusive of active participation of marginalized groups in development, to foster equitable hiring practices through recruitment technologies.

Background

Application Tracking Systems (ATS) are software applications designed to streamline the recruitment process by filtering resumes and analyzing interviews (Fuller et al., 2021). Organizations utilize ATS to efficiently cut down large pools of applicants, many times by scanning for keywords within resumes that match the job description and required qualifications (Akselrod & Venzke, 2023). These keywords generally represent some skills, qualifications, or experience sought by the employer; nonetheless, this may filter out potential qualified candidates who use different words in describing their competencies (Fuller et al, 2021).

Artificial intelligence (AI) embedded into ATS has transformed recruiting. From its early days of use in keyword filtering, the ATS can now be applied in more sophisticated processes such as analyzing video interviews through facial expressions and tone of voice (Fuller et al., 2021). Increasing the speed and efficiency associated with this development suggests that an estimated 70% of organizations and 99% of Fortune 500 companies will adopt this for recruitment through AI-driven ATS (Roth & Delbos, 2024). Such an increasingly sophisticated ATS deliver exactly the outcomes they were engineered for: to minimize the time and costs recruiters spend in finding job candidates. They are not designed to widen the aperture for hiring; their purpose is to maximize the efficiency of the process. Evidence also goes on to illustrate how AI-driven recruitment practices affect the communities that have been historically marginalized. AI systems more often punish candidates with ethnic-sounding names, non-traditional educations, or qualities irrelevant to job performance, such as formatting of resumes or lapses in employment (Pan et al., 2020). Video analysis could further put candidates at a disadvantage based on their accent, facial expressions, or gestures-all of which affects immigrants and non-native English-speaking candidates or those with problems of public speaking. These biases are only derived from AI systems trained on data from biased societies throughout history. The process henceforth favors the archetype of a candidate, usually limited to white males.

Impacted Groups

Immigrants, people with disabilities, women, and those with non-Anglo names are among the many groups that get routinely disadvantaged under the increasingly automated processes of hiring entailed through ATS and AI tools. These systems, devised to make the recruitment process smooth, often fail to contemplate the nuances characteristic of diverse candidates. These would include immigrants and non-native speakers who, because of their accented way of speaking or modes of communication, might be out of place from standardized norms, which AI-driven interview analysis would expect. This can lead to automatic disqualification irrespective of their actual job competencies or potential contributions to the organization. Similarly, people with disabilities may struggle with the non-verbal cues AI systems analyze, such as facial expressions or tone of voice, which have little relevance to job performance but can unfairly skew hiring decisions.

AI systems are often trained on historical data, which inherently includes societal biases that favor white, cisgender males. Because of this, it may filter out resumes with non-Anglo names simply because the AI aligns "success" with certain demographic characteristics (Bertrand et Mullainathan, 2004). ATS systems may be particularly punitive against women who have career interruptions due to caregiving responsibilities by devaluing resumes with employment gaps or non-traditional career paths-even though these gaps in no way affect their ability to perform the job (Goodman, 2018).

These discriminatory practices have deep economic consequences. When qualified applicants are systematically excluded from jobs, the chances of career advancement decrease, so their chances of achieving economic security decrease (Fuller et al., 2021). Contributing to the broadening of wealth disparities between various demographic groups, in this case, sustains poverty and underemployment cycles. In turn, the most sensitive group of citizens is being deprived of the wealth offered by well-paying jobs that come with additional benefits, such as health insurance coverage and retirement plans, which further develop into growing economic inequality.

Psychologically, repeated discrimination in hiring can be extremely injurious. Never-ending rejection for something that one cannot control, such as accent, name, and even gender, may lead to feelings of inadequacy, frustration, and diminished self-appreciation. This, in the long run, may contribute to anxiety, depression, and feelings of alienation (Phillips, 2024). This feeling of being systematically excluded from opportunities gnaws at their confidence and, in turn, can make it harder for them to hold on in a job market that really does seem stacked against them.

It also feeds into larger social problems: the tendency to exclude large parts of society through hiring discrimination keeps marginalized groups out of the workforce or pushes them into lower-paying jobs, threatening diversity of thought and perspective and innovation within organizations. For example, it has been proved that diverse teams perform more successfully and allow for more creative problem-solving (Fuller et al., 2021). Yet, AI-driven hiring practices could block the making of such teams. These discriminatory hiring practices thus deprive an organization of benefits derived from workforces of diverse talent.

In this respect, it is important to note that as applying ATS and AI in hiring is effective, over proportionately, this would render the most vulnerable groups more rigidly biased against and screen out qualified candidates for arbitrary reasons. The downstream effects are tremendous, added-up economic disparities, psychosocially aggregated over time. Such problems will demand no less than a complete reevaluation of how hiring systems are designed and what biases may or may not be embedded into the system.

Legal Risks Caused by AI in Hiring

Artificial Intelligence, in the present day, permeates all aspects of the recruitment process, promising speed and fairness. On the flip side, such systems could continue producing biased results either due to limited data or flawed programming and greatly heighten the possibilities of legal liabilities in hiring and HR matters (Thomas, 2023). This paper examines a few case studies, the legal ramifications of AI bias, and steps an organization can proactively take to minimize these risks.

Case Study #1

Perhaps the most well-known example of this bias is the case of Amazon. The technology giant created an AI tool for hiring, intended to be used to expedite its recruitment process. However, the tool became biased against women. The AI was trained on a decade of resumes that had been submitted to the company, most of them from men (Goodman, 2018). This meant the AI favored male candidates, and Amazon axed the tool even before implementation because of this fact. The described situation aptly illustrates how reliance on historical data perpetuates existing bias into discriminatory hiring practices (Cole, 2018). Since 2018, Amazon has come under sustained fire for discrimination based on gender and race, including a series of lawsuits alleging a hostile work environment and promotion biases. While publicly espousing diversity, equity, and inclusion, Amazon often fails to practice in its internal operations what it preaches.

Ongoing initiatives make hiring practices fair and reduce bias, but the concern about its impact on employees' well-being remains, as discrimination is associated with adverse health consequences. This includes recommendations for improving diversity and Equal Employment Opportunity law compliance training (Lopez et al., 2022). These cases have at least some considerable ramifications. Organizations that use AI in hiring could be held liable for discrimination under several federal laws, such as Title VII of the Civil Rights Act, the Age Discrimination in Employment Act, and the Americans with Disabilities Act.

Case Study #2
The second notable case is against Workday and filed by a job candidate named Derek Mobley. It alleges the company's AI hiring tool auto-rejected applicants based on potentially discriminatory criteria. Mobley also is a Black man over age 40 who describes himself as anxious and depressed. He filed a federal lawsuit against Workday, alleging the AI systems reflected illegal biases, had biased training data, and therefore produced discriminatory outcomes. Most interestingly, however, Mobley's case takes the cake because the court decided that Workday might be held liable as an agent of employers who utilize its tool, thus implicating the firm within the discriminatory outcomes of its software (Phillips, 2024). He said his race was determinable because he had graduated from a historically Black college, his age was determinable based on his year of graduation, and his mental disabilities could be revealed via personality tests. He filed a class-action suit to represent all applicants in the same situation to demonstrate how organizations open themselves up to broad legal liability if they do not address biases in AI recruitment tools. Their inability to reduce the bias found within the algorithms challenges their legal validity, which is burdensome both as a financial penalty and reputational damage (Thomas, 2023). The legal landscape related to AI in hiring is also changing, and companies need to be responsive to ensure they keep pace with emerging laws and regulations.

While AI offers immense benefits for the improvement of the hiring process, it also poses serious legal risks associated with bias and discrimination. Case studies such as Amazon and Workday have shown that there could be some legal consequences if an AI system is not closely monitored or managed. Since these technologies can be quite complex, an organization will have to be proactive in ensuring its recruiting practices are nondiscriminatory and function within the confines of the law. This way, one can benefit from AI while guaranteeing equality and an inclusive workplace culture.

Foundations of ATS Bias
The increasing influence of Machine Learning (ML) in recruitment has reinvented how organizations identify and select candidates, full of both opportunities and challenges simultaneously. Though these technologies promise efficiency and objectivity, they equally raise important concerns about algorithmic injustice. More specifically, ML can solidify historical

inequities at the expense of already vulnerable communities. Where recruitment algorithms become more determinants of who will get hired or not, their implications become increasingly pressing.

Historical hiring practices are inextricably linked to the very datasets that create these algorithms. Most recruitment models are built from data illustrating past hiring decisions, which are themselves often polluted with biases on race, gender, and socioeconomic status. The ideal candidate profile spit out from AI systems is frequently one of systemic oppression, favoring those attributes that have been tied to privilege throughout history (Isaac et al., 2022). This is both a situation that enforces and perpetuates existing inequality and moves toward the marginalization of any who do not fit into this narrow mold.

In practice, algorithmic fairness tends to operate within the technical fixes of improving model accuracy, diversifying datasets used, and the use of bias mitigation techniques. While these are vital initial steps, they fall short when trying to solve more deeply entrenched societal and structural problems that algorithms perpetuate. Such framing of complex social issues like "acceptable" behavior and "normal" body types in neutral technical terms camouflages the inherently political nature of such tasks. Consider predictive policing algorithms; trumpeted as objective methods for preventing crime, they tend to reproduce and amplify existing racial biases in the data. Unless we also critically question the assumptions and relations of power behind the algorithms, we risk using sticking-plaster approaches to treat fundamentally deep-seated problems. Technical fixes cannot redress the historical inequities and social arrangements underpinning algorithmic decision-making (Birhane, 2021).

Most of the current predictive tools deployed in areas related to data science and AI will use an algorithm that finds data patterns and deploys the same for predictions. However, these tools operate with little or no actual understanding of the context behind the data (Zajko, 2022). They simply use internal mathematical models to process inputs and spew out outputs, based on correlations that arise rather than querying the underlying reasons behind them. Thus, the systems may create biases and structural inequalities found in the data, adding to the historical injustices without questions and tending their roots.

As predictive models built from biased datasets, such as historic discrimination in hiring practices or criminal justice systems, it is very real that they run a real risk of replicating those biases in their predictions (Birhane, 2021). These tools are all focused on efficiency and accuracy, building models that optimally predict an outcome without considering broader implications of the patterns they detect. The "computer brain" follows logical and data-based patterns without capacity for critical examination of ethical or human consequences of the output.

Further, these problems are exacerbated by the lack of diversity in the groups working on artificial intelligence development. If teams are homogenous, they are less apt to take into account the nuances that surround various experiences and needs of different candidates. In this respect, lack of diverse insights may lead to algorithms failing to consider the value brought about by non-traditional qualifications and life experiences, which entrenches further the present biases (Zajko, 2022).

What really is required to meet such challenges is a paradigm shift in our thought process regarding ethics and justice as far as AI and data science are concerned. Birhane (2021) argues that AI ethics needs to be conceived of relationally, not as a technical fix that can encompass the complexity of personhood, of data, and of justice. A relational approach underscores the active inclusion of those who are most impacted by algorithmic systems. Marginalized users should be at the center of the development and implementation process of AI technologies.

It, therefore, appeals to scholars and practitioners to engage in deep reflection regarding the implications of ML in recruitment. These burning issues cannot be resolved solely with technical remedies but must be deeply embedded within the core of AI and data practices. With a deeper understanding of how technology and our social values are intertwined, it's possible to work toward hiring practices that are more equitable and reflect the true diversity and intricacy of our society.

The approach towards constructing predictive tools should instead be shifted from pattern detection to understanding the deeper context behind those patterns. Instead of merely building models that predict, based on the correlations present in data, reasons why certain patterns exist should be dug out. This involves going beyond the numbers to interrogate the social, historical, and structural factors that shape the data into critical questions of how inequalities and biases might have worked their way into the outcomes we see.

Deepening our understanding of the context in which inequality is happening helps us to design systems that predict inequality but also work to actually remedy some of the root causes. (Birhane, 2021). This would mean that we do not inadvertently continue structural injustices via our tools. Consider building models such as flag biased hiring or biased policing. It's better to ensure the system questions and corrects the biases to make the outcomes more balanced.

Empathy, ethical reasoning, and critical thinking add the human touch and must be ingrained during the development of predictive tools. We should go beyond the "computer brain" and turn on the "human brain," with efficiency and logic as a focus, and seriously ponder making our systems socially responsible, rather than just effective. It is in the development of this tool that the human touch can contribute, keeping design considerations toward fair and just outcomes, creating a future in which technology predicts not only outcomes but contributes toward a more equitable society.

Next Steps

Work can be an inclusive space by being actively and consciously aware of educating oneself about neurodiversity and AI employment bias. It would be appropriate if companies took steps in balancing hiring processes to avoid discrimination. This would be in the form of the following process.

Awareness and Education: Firstly, at an organizational level, make sure that steps are taken to enlighten not only themselves but also their teams on how to overcome certain challenges and learn best practices to accommodate neurodiverse employees. They will be allowed to access and use playbooks and guides on neurodiversity at work, for example, cases of companies that have successfully implemented neurodiversity programs (Phillips, 2024). These would provide insights into workplace experiences and challenges while informing strategies to create an enabling environment. It is also important that organizations conduct anti-bias training regularly, which will raise awareness over unconscious bias and help employees recognize how certain biases affect neurodiverse people and other underrepresented groups.

Community Commitment: The elimination of AI employment bias requires all organizations to be aware of the potential risks of AI-powered recruiting tools. For instance, it is suggested that companies should conduct complete due diligence well in advance of deploying the solutions and fully understand what kinds of data are being used to train their AI systems and how that data might affect different demographics. In this case, businesses should hold AI vendors accountable with specified language in contracts that demands nondiscriminatory and transparent AI practices, especially around fairness in neurodiverse candidates and other underrepresented groups (Livingston, 2020). Audits will also be crucial to weed out biases from these AI systems regularly to make sure they are not in violation of anti-discrimination laws and to instill confidence in the recruitment process through continuous monitoring and improvement.

Regular Auditing: The organization should regularly audit its artificial intelligence systems to eliminate bias. Auditing offers a guarantee of compliance from anti-discrimination laws, while continuous auditing will help retain confidence in the recruitment process (Phillips, 2024).

The inclusion of AI must be balanced with sound governance and proactive steps toward minimizing bias. AI systems should be designed to consider demographic variation, and any biases that arise should be identified and addressed as soon as possible. This would serve to advance the cause of diversity, equity, and inclusion of candidates during the hiring process and will be in tune with both legal requirements and company values (Phillips, 2024). In turn, this will also allow companies to go the extra mile in creating a neurodiverse-friendly environment that dampens AI bias in employment and promotes equity in the workplace for all employees.

Conclusion

The paper goes into great detail regarding the biases built into AI-powered ATS and how such systems disproportionately bear down on the marginalized groups when applicants apply for a vacancy. While these systems make life so much easier with the automation of resume screening, they often reinforce existing inequities in favor of traditional qualifications and criteria. This is to say that those from the most diverse backgrounds-a case in point being immigrants, people with disabilities, women, and all those with non-Anglo names-stand to face staggering disadvantages. Further, the study discusses the economic repercussions of such bias and the psychological consequences on the candidates who suffered because of this, pressing for reform. By taking a closer look at legal cases that involve companies like Amazon, the paper spells out the risk of discriminatory practices in hiring technologies. The whole thing basically requires a rethink in the design of ATS, encouraging them to be more inclusive of minorities' voices for a non-discriminatory recruiting culture that will drive on-field workplace diversity and inclusion.

Works Cited

Derek Mobley v. Workday Inc (2024) | Findlaw, caselaw.findlaw.com/court/us-dis-crt-n-d-cal/116378658.html. Accessed 4 Jan. 2025.

Embracing Neurodiversity at Work, www.aei.org/wp-content/uploads/2024/04/Embracing-Neurodiversity-at-Work-Unleashing-Americas-Largest-Untapped-Talent-Pool.pdf. Accessed 4 Jan. 2025.

Evandro Gigante, Edward Young. “Job Applicant’s Algorithmic Bias Discrimination Lawsuit Survives Motion to Dismiss.” *Law and the Workplace*, 26 July 2024, www.lawandtheworkplace.com/2024/07/job-applicants-algorithmic-bias-discrimination-lawsuit-survives-motion-to-dismiss/.

Hidden Workers: Untapped Talent, www.hbs.edu/managing-the-future-of-work/Documents/research/hiddenworkers09032021.pdf. Accessed 4 Jan. 2025.

Navigating the AI Employment Bias Maze: Legal Compliance Guidelines and Strategies, www.americanbar.org/groups/business_law/resources/business-law-today/2024-april/navigating-ai-employment-bias-maze/. Accessed 4 Jan. 2025.

Olga Akselrod, Cody Venzke. “How Artificial Intelligence Might Prevent You from Getting Hired: ACLU.” *American Civil Liberties Union*, 15 Nov. 2023, www.aclu.org/news/racial-justice/how-artificial-intelligence-might-prevent-you-from-getting-hired.

Sarah Roth and Becca Delbos, Opinion Contributors. “Ai Is Causing Massive Hiring Discrimination Based on Disability.” *The Hill*, The Hill, 7 Apr. 2024, thehill.com/opinion/technology/4576649-ai-is-causing-massive-hiring-discrimination-based-on-disability/.

Workday Must Face Novel Bias Lawsuit over AI Screening Software | Reuters, www.reuters.com/legal/litigation/workday-must-face-novel-bias-lawsuit-over-ai-screening-software-2024-07-15/. Accessed 4 Jan. 2025.