Discrimination in Corporate Hiring Processes from Artificial Intelligence Screeners

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Sciences

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree

Bachelor of Science, School of Engineering

Jayanth Peetla

Spring 2025

On my honor as University of Student, I have neither given not received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Kent Wayland, Department of Engineering and Society

Introduction

In the age of rapid technological advancement, artificial intelligence (AI) has reshaped the way industries operate, often promising greater efficiency, objectivity, and scalability. Among these changes is the widespread adoption of AI in corporate hiring practices—a shift that has fundamentally altered how companies identify, evaluate, and select job candidates. While these tools are marketed as innovative solutions to streamline recruitment, they also introduce serious concerns about fairness, transparency, and accountability.

AI-powered hiring tools, such as automated resume screeners and video interview evaluators, are increasingly relied upon to make high-stakes decisions about people's careers. These systems are trained on historical hiring data, which often reflects past human biases, potentially leading to algorithmic discrimination against marginalized groups. Despite growing evidence of bias—like Amazon's discontinued AI hiring tool that penalized women—many organizations continue to use these systems without fully understanding their implications. This shift raises urgent ethical and societal questions.

This paper investigates the discriminatory nature of AI-powered recruiting tools through the lens of Actor-Network Theory (ANT), focusing on the roles and responsibilities of three key actors: the tech companies that build these tools, the employers who implement them, and the regulatory bodies tasked with oversight. By examining their interactions and influence within this socio-technical network, this research aims to highlight not only how discrimination arises, but what each actor can do to minimize harm and promote more equitable hiring practices in an increasingly automated world.

Background & Context

In the past decade, researchers have made giant strides in the development of artificial intelligence (AI), largely driven by advancements in capabilities such as natural language processing, generative AI, and computer vision. The wide application of this technology has resulted in a rapid increase in adoption across various industries. One such industry that has significantly transformed in recent years due to the integration of AI is recruiting.

Large corporations typically delegate the process of hiring and onboarding new employees to a Human Resources (HR) department. The HR department is responsible for the initial screening of all submitted applications. Workers in this department are typically not technically proficient. As a result, their primary focus is to filter out any employees they think would not be a good fit and send all qualified candidates to higher-ups for further evaluation. For bigger companies, this can be on the order of tens of thousands of applications. Given the volume of applications and repetitive nature of the job, HR departments identified that AI would be a great way to automate their work.

At the forefront of this innovation are companies such as HireVue, Pymetrics, and LinkedIn. These companies have developed AI-driven tools for resume screening, personality assessments, and video interview analysis. These tools are used by large companies as an initial bottleneck to algorithmically eliminate all candidates that would most likely not be good for the company. The machine learning algorithms used in these AI-based filters are trained on historical hiring data, often reflecting the implicit biases of past decisions. For example, Amazon discontinued its internal AI recruiting tool after discovering it discriminated against female candidates - highlighting how these tools can inadvertently perpetuate inequalities. (Simonite 2018)

There are three key actors in this socio-technical situation. First, the tech companies developing these machine learning algorithms. Companies like HireVue have become leaders in this space. Their proprietary algorithms are marketed as offering objective and efficient evaluations. However, the outright lack of transparency as to how exactly these algorithms function has raised ethical concerns.

Second, the companies adopting and implementing these tools in their recruiting processes. Multinational firms use AI to filter resumes, conduct initial assessments, and narrow down candidate pools before any human involvement occurs. This approach is intended to save time and resources, but it often sacrifices nuance and risks reinforcing existing inequalities.

Lastly, regulators play a crucial role in this ecosystem, working to ensure fairness and accountability. Bodies like the Equal Employment Opportunity Commission (EEOC) in the United States have begun scrutinizing AI hiring practices, particularly around their compliance with anti-discrimination laws. Some regulatory authorities have taken action, including the EU's proposed AI Act aiming to classify high-risk AI applications like hiring tools and enforce strict regulations. Despite this, enforcement in many regions remains limited.

This paper will focus on analyzing and answering the following: *How did the various actors of the network contribute to the discriminatory nature of existing AI hiring systems and what actions can they each take to mitigate the negative effects?*

Methods

This paper conducts a literature review to examine the presence and impact of algorithmic discrimination in AI-driven recruitment tools. The research focuses on analyzing documented cases of AI bias, industry reports, and regulatory discussions. In addition, it applies

Actor-Network Theory (ANT) as a framework to understand how different actors—AI developers, employers, and regulators—contribute to and can mitigate discrimination in hiring technologies. ANT, developed by Latour, Callon, and Law, posits that technological systems are shaped by dynamic interactions between human and non-human actors. Rather than viewing AI hiring tools as neutral, this framework highlights how their development, implementation, and oversight are shaped by the actions and interests of multiple stakeholders.

Within this framework, AI developers function as key actors who design and train machine learning models based on historical hiring data. Employers act as intermediaries who adopt these tools, relying on them to automate decision-making and streamline the hiring process. Regulators, such as the Equal Employment Opportunity Commission (EEOC), represent another critical actor group aiming to impose fairness standards and legal accountability. However, as ANT suggests, these actors do not operate independently; their interactions shape the ultimate impact of AI hiring tools on the job market.

To analyze the discriminatory effects of AI in hiring, this paper reviews peer-reviewed academic research, government reports, and industry case studies from sources such as IEEE Xplore, ACM Digital Library, Google Scholar, and legal publications. The research examines historical cases of AI hiring bias, such as Amazon's discontinued AI hiring tool and HireVue's video interview assessments, alongside legal decisions like Mobley v. Workday, which question whether AI vendors should be held accountable for discrimination.

The analysis is conducted by identifying patterns of bias, regulatory responses, and proposed solutions across different sources. By mapping the relationships between AI developers, employers, and regulators, this paper assesses how each actor influences and

reinforces discrimination in hiring, as well as what steps they can take to mitigate these issues. The application of ANT enables a socio-technical perspective, revealing that algorithmic bias is not an isolated technological failure but a product of interconnected decisions and systemic structures.

Owners of AI Hiring Tools

Tech companies play a foundational role in shaping the functionality and ethics of AI-powered hiring systems. They are responsible for designing the algorithms that analyze resumes, assess video interviews, and predict candidate fit—all based on the data they are trained on. A central critique of these systems is that they inherit biases from the historical data fed into them. Harvis-Nazzario (2022) argues that "there is no way to avoid bias when using algorithmic decision making tools; systemic bias will still be present." Because hiring algorithms are trained on past decisions—many of which reflect unequal access to opportunities—they often replicate and amplify these same inequalities.

A well-known example is Amazon's discontinued AI recruiting tool, which was trained on ten years of resumes submitted to the company, most of which came from male applicants due to the tech industry's gender imbalance. As a result, the system penalized resumes that included the word "women" or referenced all-female institutions (Dastin, 2018). This case highlights the real risk of training algorithms on past data that has likely been subject to decades of subjective discrimination by humans. Decades of gender discrimination done by humans found its way to algorithmic decision-making.

Another example involves HireVue, whose video-interview analysis platform uses facial recognition and voice analysis to assess candidates. Harvis-Nazzario (2022) notes that the

Electronic Privacy Information Center (EPIC) filed a complaint to the Federal Trade Commission (FTC), alleging that the tool was biased against people with mental health conditions and racial minorities. These critiques emphasize that AI companies are not neutral creators of technology - they are active participants shaping the hiring experience.

Nicholas Tilmes (2022) adds that machine learning systems often treat disability as a simple binary variable, ignoring the wide spectrum of individual differences. This reduction flattens complex human traits into easily digestible metrics, ultimately harming candidates who do not fit dominant normative molds. As such, tech companies must take greater responsibility not only for their data practices but also for how they define and encode "worthiness" into their systems.

As shown, tech companies have unintentionally managed to create discriminatory algorithms by using historically biased data and deploying poorly audited algorithms. To reverse these effects, these companies must be transparent and embed accountability throughout the development lifecycle. This can be in the form of diversifying and de‑biasing training data, implementing rigorous third‑party audits, publishing transparent model documentation, and ensuring meaningful human oversight at every decision checkpoint. By adopting these practices, owners of AI tools can shift from perpetuating inequality to fostering a more honest hiring process.

Users of AI Hiring Tools

Employers are the ones deploying these tools, often with the goal of increasing efficiency and cutting down on the time and cost required to process thousands of applications. However, in doing so, they often delegate critical decisions to systems they do not fully understand. Many HR departments lack technical expertise, making them ill-equipped to question the outcomes of AI

hiring tools. Consequently, algorithmic discrimination may go unnoticed or unchallenged within an organization's hiring pipeline.

Because AI systems can be opaque, it's not always clear why certain candidates are rejected, which further limits employers' ability to intervene or offer recourse. This is problematic not just from an ethical standpoint, but also a legal one. Employers can be held liable if the tools they use have discriminatory effects, even if the bias is unintentional. Todd Horn (2024) warns that employers are increasingly being named as co-defendants alongside AI vendors in hiring discrimination lawsuits, signaling a shift in how legal accountability is assigned.

This was further emphasized in the case of *Mobley v. Workday*, where the court determined that AI vendors may not simply be "neutral facilitators," but can be seen as joint decision-makers in hiring processes (See, 2024). This decision underscores the need for employers to rigorously vet the AI systems they implement and treat them not as plug-and-play solutions, but as high-stakes tools that can affect real lives.

Adam Hamel (2023) recommends that companies treat AI as a supplement to, rather than a replacement for, human judgment. He also encourages employers to maintain transparency with job seekers about the use of AI in evaluations, arguing that this helps build trust and reduces the risk of backlash. As algorithmic hiring becomes more common, companies must take an active role in understanding and mitigating the unintended consequences these systems may cause.

Regulators of AI Hiring Tools

As algorithmic hiring systems have become more prevalent, regulators have increasingly stepped in to address concerns about fairness and discrimination. One of the earliest foundational

protections in this space is Title VII of the Civil Rights Act of 1964, which prohibits discrimination in employment based on race, color, sex, or national origin. While this law was originally created to combat explicit discrimination, its principles are now being extended to algorithmic systems. As Harvis-Nazzario (2022) notes, the law also prohibits facially neutral practices that result in disparate impact, which is exactly the kind of bias often embedded in algorithmic tools.

The Equal Employment Opportunity Commission (EEOC) has taken a more proactive role in recent years, publishing technical assistance documents to guide employers on the responsible use of AI in employment decision-making. In 2023, the EEOC launched an initiative to ensure that the use of algorithmic tools complies with existing anti-discrimination laws (EEOC, 2023). The agency has also emphasized that employers cannot shield themselves from liability by blaming third-party vendors, a position reaffirmed by recent legal decisions.

On an international level, regulatory frameworks like the European Union's proposed AI Act take this even further by categorizing hiring tools as "high-risk" applications of AI. If passed, the Act would impose strict requirements on transparency, documentation, and risk assessment for such systems (European Commission, 2021). These emerging policies reflect a growing awareness that without meaningful oversight, AI can quietly entrench systemic inequality under the guise of efficiency.

Despite these efforts, enforcement remains inconsistent. Many regions still lack the resources or legal clarity needed to hold companies accountable for algorithmic bias. As a result, continued pressure from watchdog organizations, public interest groups, and legal scholars is essential to ensure that regulation keeps pace with technological development. Without persistent effort from

these parties, the discriminatory outcomes of AI hiring tools will remain a deeply entrenched issue.

Conclusion

The development and deployment of AI-powered hiring tools have revealed how easily biased outcomes can become embedded in systems presented as objective. As demonstrated by multiple cases and legal challenges, responsibility cannot be pinned solely on the algorithms themselves but must also include the developers who build them, the companies that deploy them, and the regulators who fail to provide sufficient oversight. Each actor plays a role in shaping the technology and its consequences. Owners of AI-hiring tools should make efforts to make their systems more transparent and equitable. Users of these systems should be wary of their discriminatory nature and use it in a responsible or equitable way. Policies must be strengthened to ensure regular auditing, accessible appeals processes, and inclusive data practices. If all of the actors in this network work together, there is hope for an AI-based recruiting system that is fair and equal.

References

Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. Retrieved from https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G

Hamel, A. (2023, May 19). Will artificial intelligence revolutionize the hiring process?. New Hampshire Business Review, 1–3.

Harvis-Nazzario, L. (2022). It's not the algorithms, it's the people: Preventing bias in automated hiring tools starts with humans. Rutgers Computer and Technology Law Journal, 49(1), 138–175.

Horn, T. J., Fucci, S. A. (2024). Artificial intelligence hiring tools: ACLU warns employers and AI companies it will not stand for deceptive "bias-free" promises. Employee Relations Law Journal, 50(3), 37–40.

See, R. V. (2024). Mobley v. Workday: Court holds artificial intelligence service providers could be directly liable for employment discrimination under "agent" theory. Employee Relations Law Journal, 50(3), 41–43.

Simonite, T. (2018, October 11). Amazon's AI recruiting tool was biased against women, and no one caught it. Wired. Retrieved from https://www.wired.com/story/amazons-ai-recruiting-tool-biased-against-women/

Tilmes, N. (2022). Disability, fairness, and algorithmic bias in AI recruitment. Ethics and Information Technology, 24(2), 1–15.

https://www.proquest.com/docview/2652409779?accountid=14678&sourcetype=Scholarly%20Journals