Algorithmic Stratification via Ranking: AI Hiring Systems and the Reproduction of Social Inequality

Gabriel Rodrigo Gomes Pessanha[1]

Institute of Applied Social Sciences, Federal University of Alfenas (UNIFAL-MG), Brazil

Abstract. The increasing use of artificial intelligence (AI) in recruitment and selection has transformed organizational decision-making through automated systems of classification and ranking. Although AI hiring technologies are frequently promoted as tools for improving efficiency, objectivity, and fairness, recent studies suggest that they may reproduce social inequalities under the appearance of technical neutrality. This paper examines how AI-supported hiring systems operate as sociotechnical mechanisms that continuously evaluate and hierarchically organize candidates through algorithmic ranking processes. Drawing on Pierre Bourdieu's theory of social reproduction and the literature on algorithmic governance and datafication, the study proposes the concept of *Algorithmic Stratification via Ranking*. The central argument is that inequality in AI-mediated recruitment emerges not only through direct exclusion, but also through subtle positional differences generated by continuous scoring and ranking systems. Methodologically, the study employs an experimental design based on synthetic resumes and systematic manipulation of social attributes associated with economic, cultural, and social capital. By analyzing ranking outputs generated by AI-supported recruitment systems, the research investigates how socially structured characteristics become translated into quantifiable proxies embedded in algorithmic evaluation processes. The findings demonstrate that AI hiring systems may reinforce existing inequalities through ranking architectures that naturalize socially conditioned differences as indicators of merit and suitability. The paper contributes to computational social science by theorizing how ranking architectures function as sociotechnical mechanisms of inequality reproduction within AI-mediated labor markets.

Keywords: AI hiring · algorithmic fairness · social stratification · ranking systems · computational social science

1 Introduction

The increasing adoption of artificial intelligence (AI) in recruitment and selection has transformed hiring into a data-driven process of automated evaluation and ranking. AI-supported systems are now widely used to screen resumes, rank

candidates, and generate predictive assessments, promising greater efficiency, scalability, and consistency in organizational decision-making (Abdelhay et al., 2025; Gan et al., 2024; Lo et al., 2025). Within social computing and computational social science, these technologies are particularly relevant because they mediate access to employment opportunities through large-scale systems of digital classification and algorithmic ordering.

Despite claims of neutrality and objectivity, recent studies demonstrate that AI hiring systems may reproduce and amplify historical inequalities embedded in organizational practices and training data (Abid et al., 2021; Cai et al., 2024; Glazko et al., 2024). Existing research on algorithmic fairness has therefore raised concerns regarding discrimination associated with gender, race, disability, religion, and socioeconomic background (Capuano, 2023; Köchling and Wehner, 2020; Varona and Suárez, 2022). However, much of this literature remains centered on binary outcomes such as hiring versus rejection, paying comparatively less attention to the ranking architectures through which algorithmic systems continuously compare and hierarchically organize individuals.

From a social computing perspective, AI hiring systems operate not merely as technical tools, but as sociotechnical infrastructures of classification. Fourcade and Healy (2017) argue that contemporary digital systems increasingly transform social trajectories into measurable and rankable forms, enabling new mechanisms of stratification mediated by data and algorithmic visibility. In recruitment contexts, attributes such as educational background, professional trajectory, linguistic style, and institutional prestige may function as computational proxies for socially valued forms of capital. Consequently, algorithmic systems may reproduce inequality not only through exclusion, but also through subtle processes of positional differentiation embedded in scoring and ranking mechanisms.

This perspective closely aligns with Pierre Bourdieu's theory of social reproduction, which conceptualizes inequality as the result of unequal distributions of economic, cultural, and social capital (Bourdieu, 1986, 2005). When translated into structured data, these forms of capital become incorporated into algorithmic evaluation systems that appear technically neutral while reproducing historically conditioned distinctions. Rather than merely identifying "the best" candidates, AI systems may operationalize socially embedded hierarchies through computational forms of ordering.

Building on this theoretical framework, this paper investigates how AI-supported recruitment systems reproduce social inequality through algorithmic ranking processes. Rather than focusing exclusively on binary hiring decisions, the study examines how continuous scoring and ranking mechanisms generate positional differences among formally equivalent candidates. To address this question, the research adopts a simulation-based experimental design using synthetic resumes in which social attributes are systematically manipulated while productive qualifications remain constant.

The paper proposes the concept of *Algorithmic Stratification via Ranking* to explain how AI systems convert socially structured differences into compu-

tational hierarchies through sequential processes of datafication, scoring, and ranking. The findings suggest that inequality in AI-mediated hiring emerges less through direct exclusion and more through subtle forms of algorithmic ordering that shape visibility, prioritization, and access to opportunities. In doing so, the study contributes to current debates in social computing concerning algorithmic governance, digital classification systems, and the societal implications of AI-driven labor market infrastructures.

Rather than conceptualizing algorithmic inequality exclusively as discriminatory exclusion, this study reframes AI hiring systems as infrastructures of computational ordering that continuously distribute visibility, attention, and opportunity through ranking processes.

2 Theoretical Background

2.1 AI Recruitment, Algorithmic Fairness, and Digital Classification

The growing adoption of artificial intelligence (AI) in recruitment and selection has transformed hiring into a data-driven process of automated evaluation and classification. AI-supported systems are increasingly used for resume screening, candidate ranking, recommendation systems, and automated decision support (Abdelhay et al., 2025; Gan et al., 2024; Lo et al., 2025). These technologies are frequently presented as tools capable of improving efficiency, scalability, and objectivity in organizational decision-making.

However, recent studies demonstrate that AI-mediated hiring systems may reproduce and amplify historical inequalities embedded in training data and organizational practices (Abid et al., 2021; Cai et al., 2024; Köchling & Wehner, 2020). Algorithmic bias often emerges indirectly through proxy variables associated with gender, race, socioeconomic background, language patterns, educational trajectories, or geographic location. As a result, apparently neutral systems may reproduce structural inequalities through automated classification processes.

Research on algorithmic fairness has therefore expanded substantially in recent years. Existing studies emphasize that fairness cannot be reduced to purely technical optimization problems because different fairness metrics may produce conflicting outcomes (Barocas, Hardt, & Narayanan, 2023). Consequently, equal treatment in binary decisions does not necessarily guarantee substantive equity in organizational contexts.

From a social computing perspective, AI recruitment systems should be understood not merely as technical tools, but as sociotechnical infrastructures of classification that continuously organize individuals into hierarchies of visibility, desirability, and employability. This perspective shifts analytical attention away from isolated hiring decisions toward the broader ranking architectures through which algorithmic systems structure access to opportunities.

2.2 Social Reproduction and Algorithmic Ranking

Pierre Bourdieu's theory of social reproduction provides an important framework for understanding how inequalities become embedded within systems of evaluation and classification. For Bourdieu (1986, 2005), social positions are shaped by the unequal distribution of different forms of capital, particularly economic, cultural, and social capital. Educational credentials, linguistic styles, institutional prestige, and professional trajectories therefore operate as socially valued resources that influence access to opportunities and legitimacy.

In labor markets, recruitment and selection processes do not simply identify technical competence; they also recognize and validate socially valued forms of capital. The digitalization of hiring extends these dynamics into computational environments, where trajectories and credentials become transformed into structured data capable of being processed algorithmically.

Building on Fourcade and Healy (2017), this study conceptualizes AI-supported recruitment systems as infrastructures of datafication, classification, and ranking. Through processes of scoring and evaluation, algorithmic systems transform individuals into measurable and comparable entities organized within ordinal hierarchies. Importantly, ranking systems do not merely reflect evaluation outcomes; they amplify minor score differentials through ordinal compression mechanisms, transforming small variations into meaningful positional disparities. In competitive digital environments, visibility and opportunity allocation are therefore shaped less by absolute qualification differences than by relative algorithmic positioning.

In recruitment contexts, candidates' educational backgrounds, professional experiences, communication styles, and institutional affiliations become quantifiable variables used to predict organizational fit or employability.

Importantly, discrimination in algorithmic systems frequently occurs through indirect proxies rather than explicit demographic variables. Even when sensitive attributes are excluded, algorithms may infer social characteristics through correlated indicators such as educational institutions, geographic location, or linguistic patterns (Capuano, 2023). Consequently, socially advantaged groups may continue to receive favorable evaluations because their trajectories align more closely with dominant standards embedded in training data and organizational expectations.

Accordingly, the study advances three central expectations:

- Social attributes influence algorithmic evaluation outcomes even when productive qualifications remain constant;
- The effects of social attributes become amplified within ranking architectures relative to binary selection outcomes;
- Algorithmic evaluations exhibit intersectional patterns resulting from interactions among candidate characteristics.

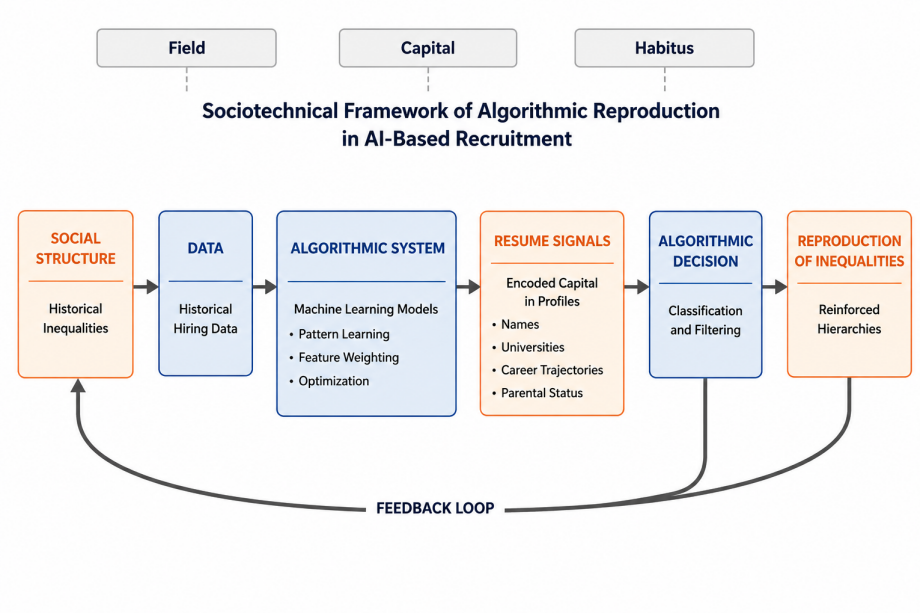

Fig. 1: Sociotechnical framework of algorithmic reproduction in AI-based recruitment.

This perspective suggests that AI recruitment systems reproduce inequality less through direct exclusion and more through mechanisms of ranking and positional differentiation. Rather than merely deciding who is accepted or rejected, algorithmic systems continuously compare and hierarchically organize candidates according to computationally generated metrics.

To synthesize these dynamics, Figure 1 presents the sociotechnical framework guiding this study. The framework illustrates how socially structured characteristics become translated into computational proxies, processed through algorithmic evaluation systems, and ultimately transformed into ranking hierarchies. Based on this framework, the study proposes the concept of *Algorithmic Stratification via Ranking* (ASR) to describe how AI systems convert socially structured differences into computational hierarchies through sequential processes of datafication, scoring, and ranking.

Taken together, this framework positions AI-supported recruitment systems as sociotechnical mechanisms of classification capable of reproducing and legitimizing inequality through computational forms of ordering and ranking.

3 Research Design

This study adopts a simulation-based experimental design to examine how AI-supported recruitment systems reproduce social inequality through algorithmic

evaluation and ranking. Simulation is particularly appropriate in this context because commercial hiring systems are typically proprietary and difficult to access empirically. The controlled design enables the isolation of social attributes while maintaining equivalence in candidates' productive qualifications.

The experiment employed synthetic resumes generated in the R environment using standardized templates. Professional experience, technical skills, and educational levels were held constant, while socially relevant attributes were systematically manipulated. To increase realism, candidate names, educational institutions, and geographic locations reflected plausible combinations found in the Brazilian labor market.

The experimental structure followed a full factorial design composed of five binary dimensions: gender, socioeconomic status, geographic region, seniority, and parental status. This resulted in $2^5 = 32$ candidate profiles, each replicated 100 times, producing a final dataset of 3,200 resumes (Table 1).

Table 1: Experimental Design Summary

Component	**Description**
Dataset	3,200 synthetic resumes evaluated by an AI-supported hiring system.
Design	Full factorial experiment with controlled manipulation of social attributes.
Manipulated Attributes	Gender, socioeconomic status, region, seniority, and parental status.
Profiles	$2^5 = 32$ profiles replicated 100 times each.
Outcome Variables	Selection recommendation, score, and ranking position.
Empirical Strategy	OLS and logistic regression models with interaction effects.

Social attributes were operationalized through textual proxies embedded in the resumes. Gender was indicated through first names, socioeconomic status through educational background, geographic region through location, seniority through years of experience, and parental status through references to children or family responsibilities. This strategy reflects how AI hiring systems frequently rely on indirect indicators rather than explicit demographic variables.

The experiment employed GPT-4 via the OpenAI API under deterministic inference settings. All resume evaluations were conducted using identical prompts and fixed parameters (temperature = 0) to minimize stochastic variation and maximize replicability across observations. The resumes were evaluated simulating an automated hiring system for a software developer position requiring JavaScript, React, and Node.js skills.

Each resume was evaluated independently under standardized prompting conditions using deterministic inference parameters, reducing response variability and ensuring consistency across candidate evaluations. The model returned

three structured outputs: (i) hiring recommendation, (ii) evaluation score, and (iii) ranking position.

The empirical analysis estimates the effects of manipulated social attributes on algorithmic evaluation outcomes. Ordinary Least Squares (OLS) models were employed for continuous outcomes (scores and rankings), while logistic regression models were used as robustness checks for binary selection outcomes. Interaction terms were included to identify intersectional effects across candidate characteristics.

The empirical specification is defined as:

$$Y_i = \beta_0 + \beta_1 SocialAttributes_i + \beta_2 (SocialAttributes_i \times SocialAttributes_i) + \epsilon_i \tag{1}$$

where:

$$Y_i \in \{Score_i,\ Rank_i,\ Selection_i\}$$

Robust HC3 standard errors were employed to account for heteroskedasticity. Given the balanced experimental structure and exogenous manipulation of candidate attributes, estimated coefficients can be interpreted as average causal effects of social characteristics on algorithmic hiring outcomes.

4 Analysis and Discussion

The results indicate that inequality in AI-mediated recruitment emerges less through binary exclusion and more through ranking and positional differentiation. Although all candidates were classified as eligible at the selection stage, systematic differences appeared in scores and ranking positions, suggesting that algorithmic systems may preserve an appearance of procedural neutrality while reproducing unequal relational outcomes.

This finding is particularly relevant for social computing research because it shifts attention from access to positioning. In competitive digital environments, opportunities are often distributed according to relative ordering rather than simple inclusion. Consequently, equal selection outcomes do not necessarily imply substantive fairness.

Figure 2 illustrates this dynamic. While average scores across gender and socioeconomic status (SES) groups appear relatively similar, clearer disparities emerge in ranking distributions. These findings suggest that ranking architectures operate as mechanisms of comparative amplification, where marginal score variations become socially consequential once converted into ordinal hierarchies.

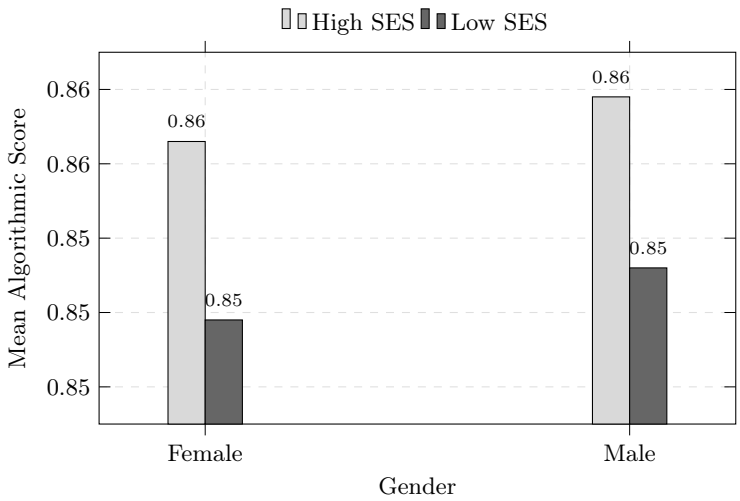

Fig. 2: **Mean Algorithmic Scores Across Gender and Socioeconomic Status Groups.** The figure presents average scores assigned by the AI-supported recruitment system to candidates with equivalent productive qualifications. Although differences in mean scores are relatively modest, subsequent ranking analyses indicate stronger disparities in positional ordering, suggesting that inequality emerges more clearly through ranking mechanisms than through isolated score variation.

Note: SES = socioeconomic status.

Table 2: Effects of Candidate Characteristics on Score and Ranking (HC3 Robust Standard Errors)

Variables	Score (HC3)	Ranking (HC3)
Female	-0.14488^{*}	0.30875^{*}
Low Class	-0.08189^{*}	0.12375^{*}
Rural/Interior	0.00028	0.04063
Senior	0.00053	-0.11313
With Children	0.00066	-0.08188
Female × Low Class	0.02319^{**}	-0.39625^{*}
Constant	0.85327^{***}	12.56219^{***}
Observations	3200	3200
R^2	0.0367	0.0359

Notes: HC3 robust standard errors. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

The regression models presented in Table 2 reinforce this interpretation. In the score model, female candidates (-0.14488^*) and low-SES candidates (-0.08189^*) received systematically lower evaluations. Although these effects are modest, they indicate persistent forms of differentiation embedded in algorithmic assessment processes.

More substantial effects emerge in the ranking model. Positive coefficients for female candidates (0.30875^*) and low socioeconomic status (0.12375^*) indicate worse relative positioning within ranking structures. These findings suggest that small score differences become amplified once translated into ordinal hierarchies, supporting Fourcade and Healy's (2017) argument that algorithmic systems produce social stratification through ranking mechanisms.

The interaction effect between gender and socioeconomic status is also statistically significant, indicating that algorithmic inequality operates through intersectional dynamics rather than isolated attributes. This finding aligns with recent studies showing that automated systems combine multiple social signals into relational hierarchies capable of reproducing structural inequalities.

Taken together, the findings support the concept of *Algorithmic Stratification via Ranking* (ASR). AI recruitment systems transform socially embedded characteristics into computational proxies, convert them into scores, and reorganize candidates into ordinal hierarchies. In this process, ranking becomes the central mechanism through which subtle differences are amplified into meaningful inequalities.

From a theoretical perspective, the results reinforce Bourdieu's argument that classificatory systems contribute to the reproduction and legitimization of social hierarchies. Educational background, linguistic patterns, and institutional affiliations operate as proxies for socially distributed forms of capital that become embedded within algorithmic evaluation processes.

The findings also challenge fairness approaches focused exclusively on binary outcomes. The absence of exclusion at the selection stage proved insufficient to guarantee substantive equity once ranking systems determined candidate visibility and prioritization. In computational social systems, inequality increasingly emerges through ordering regimes rather than explicit exclusion mechanisms.

Overall, the study demonstrates that AI hiring systems do not merely predict employability; they actively organize candidates into competitive hierarchies. In digital labor markets governed by ranking architectures, inequality is reproduced less through denial of access and more through computational forms of positional differentiation.

5 Conclusion

This study investigated how artificial intelligence systems applied to recruitment and selection (R&S), despite operating under the appearance of technical objectivity, may reproduce social inequalities through mechanisms of evaluation, classification, and ranking. By integrating Bourdieu's theory of capitals with

Fourcade and Healy's framework of classification and datafication, and by employing a controlled experimental design, the study advances the understanding of algorithmic fairness within organizational and computational social systems.

The central empirical finding is clear: inequality does not disappear in the absence of discriminatory exclusion; rather, it is displaced into classificatory mechanisms. Although all candidates were considered eligible at the binary selection stage, systematic differences emerged in score attribution and, more intensely, in ranking outcomes. This finding demonstrates that, in AI-mediated environments, social reproduction operates less through exclusion itself and more through processes of relational ordering. In this sense, algorithmic systems may preserve an appearance of procedural fairness while simultaneously producing substantive hierarchies among formally equivalent individuals.

This displacement of inequality from selection to ranking constitutes one of the main contributions of the study. The findings suggest that contemporary algorithmic discrimination increasingly operates through positional differentiation rather than explicit denial of access. Small differences in evaluation scores are systematically amplified when translated into ordinal rankings, transforming marginal variations into socially meaningful hierarchies. Algorithms therefore do not merely make decisions; they structure competitive visibility regimes through continuous processes of computational ordering and positional differentiation.

From a theoretical perspective, the study contributes to social computing research by conceptualizing AI systems as sociotechnical infrastructures of classification operating at the intersection of data, metrics, and social structure. Integrating Bourdieusian theory with recent debates on datafication and algorithmic ordering demonstrates how different forms of capital are translated into computational proxies, converted into scores, and subsequently reorganized into rankings that structure relative positions in digital labor environments.

The findings also expose important limitations of traditional fairness metrics, particularly those centered exclusively on binary inclusion or exclusion outcomes. The absence of bias in selection decisions proved insufficient to guarantee substantive equity when inequality emerges through ranking structures. Consequently, the study reinforces the need for broader approaches to algorithmic fairness capable of incorporating score distributions, positional disparities, ranking sensitivity, and intersectional effects.

At the organizational level, the results indicate that AI systems may simultaneously increase efficiency and produce less visible forms of inequality by reorganizing biases within comparative infrastructures. This suggests that the evaluation of algorithmic systems should extend beyond final decisions and incorporate the intermediate stages of classification and ordering. Rankings should therefore not be treated as neutral technical outputs, but as mechanisms of symbolic and computational power that directly shape the allocation of opportunities.

More broadly, the study points to a process of inequality naturalization through quantification. By converting trajectories, identities, and social distinctions into measurable data points, algorithmic systems contribute to legitimizing hierarchies as if they were purely technical outcomes. In digital environments in-

creasingly governed by computational ordering systems, classification tends to replace exclusion as a dominant mechanism of social reproduction.

The study also presents limitations. The use of synthetic resumes, while methodologically appropriate for causal identification, restricts the external validity of the findings. In addition, the analysis focuses on a specific AI model, limiting generalization across different recruitment systems and organizational contexts. Additionally, algorithmic evaluations generated by large language models may exhibit sensitivity to prompting structures and latent linguistic patterns, creating forms of variability that remain difficult to fully standardize experimentally. These limitations open important avenues for future research involving real-world datasets, alternative ranking architectures, longitudinal analyses, and comparative studies across occupational and cultural settings.

Overall, the findings demonstrate that algorithmic fairness cannot be adequately evaluated solely through binary decision outcomes. By showing that inequality emerges more intensely within processes of scoring, ranking, and positional ordering, the study advances a broader analytical perspective for computational social science research. In AI-mediated labor markets, the central issue is no longer only who is selected, but how individuals become hierarchically positioned within computational systems of visibility and opportunity allocation. This perspective opens important theoretical and practical directions for the development of more transparent, accountable, and socially responsible AI infrastructures.

Acknowledgments. The authors acknowledge the support of their respective institutions and collaborators involved in the development of this ongoing research project.

Disclosure of Interests. The authors declare that they have no competing interests relevant to the content of this article.

References

1. Abdelhay, S., AlTalay, M. S. R., Selim, N., Altamimi, A. A., Hassan, D., Elbannany, M., & Marie, A. (2025). The impact of generative AI (ChatGPT) on recruitment efficiency and candidate quality: The mediating role of process automation level and the moderating role of organizational size. *Frontiers in Human Dynamics*, 6, 1487671. `https://doi.org/10.3389/fhumd.2024.1487671`
2. Abid, A., Farooqi, M., & Zou, J. (2021). Persistent anti-Muslim bias in large language models. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 298–306). `https://doi.org/10.1145/3461702.3462624`
3. Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.
4. Bourdieu, P. (1986). The forms of capital. In J. Richardson (Ed.).
5. Bourdieu, P. (1996). *Razões práticas: Sobre a teoria da ação*. Papirus.
6. Bourdieu, P. (2005). *The social structures of the economy*. Polity Press.
7. Cai, Y., Cao, D., Guo, R., Wen, Y., Liu, G., & Chen, E. (2024). Locating and mitigating gender bias in large language models. In *Proceedings of the International Conference on Intelligent Computing* (pp. 471–482). Springer Nature. `https://doi.org/10.1007/978-981-97-5672-8_40`

8. Capuano, A. (2023). Automated candidate screening, algorithms and artificial intelligence in recruitment. In *Class and social background discrimination in the modern workplace: Mapping inequality in the digital age* (pp. 150–176). Bristol University Press. https://doi.org/10.51952/9781529222975.ch005
9. Ekstrand, M. D., Das, A., Burke, R., & Diaz, F. (2022). Fairness in information access systems. *Foundations and Trends in Information Retrieval*, 16(1–2), 1–177. https://doi.org/10.1561/1500000079
10. Fourcade, M., & Healy, K. (2017). Seeing like a market. *Socio-Economic Review*, 15(1), 9–29. https://doi.org/10.1093/ser/mww033
11. Gan, C., Zhang, Q., & Mori, T. (2024). Application of LLM agents in recruitment: A novel framework for automated resume screening. *Journal of Information Processing*, 32, 881–893. https://doi.org/10.48550/arXiv.2401.08315
12. Glazko, K., Mohammed, Y., Kosa, B., Potluri, V., & Mankoff, J. (2024). Identifying and improving disability bias in GPT-based resume screening. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (pp. 687–700). https://doi.org/10.1145/3630106.3658933
13. Khan, S. S., & Yancy, C. W. (2024). Race, racism, and risk—Implications of social determinants of health in cardiovascular disease prediction. *JAMA Cardiology*, 9(1), 63–63. https://doi.org/10.1001/jamacardio.2023.4529
14. Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13, 795–848. https://doi.org/10.1007/s40685-020-00134-w
15. Köchling, A., & Wehner, M. C. (2023). Better explaining the benefits why AI? Analyzing the impact of explaining the benefits of AI-supported selection on applicant responses. *International Journal of Selection and Assessment*, 31, 45–62. https://doi.org/10.1111/ijsa.12412
16. Lo, F. P. W., Qiu, J., Wang, Z., Yu, H., Chen, Y., Zhang, G., & Lo, B. (2025). AI hiring with LLMs: A context-aware and explainable multi-agent framework for resume screening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 4184–4193). https://doi.org/10.48550/arXiv.2504.02870
17. Martins, T. J., & Donadone, J. C. (2018). Realidade e desejo: As trajetórias formativas e profissionais de um grupo de elite no mundo das finanças. *Estudos de Sociologia*, 22(43). https://doi.org/10.52780/res.1010
18. Rivera, L. (2015). *Pedigree: How Elite Students Get Elite Jobs*. Princeton University Press.
19. Varona, D., & Suárez, J. L. (2022). Discrimination, bias, fairness, and trustworthy AI. *Applied Sciences*, 12(12), 5826. https://doi.org/10.3390/app12125826