

User Perceptions and Labor Market Outcomes of AI-Generated Resumes

Elijah White

elijahwhite4981@gmail.com

University of Florida

Research Article

Keywords: AI-generated resumes, large language models, recruiter perceptions, hiring outcomes, algorithmic bias, eye-tracking, human-AI collaboration, resume optimization, labor market inequality, ethical AI

Posted Date: October 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11272368/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

User Perceptions and Labor Market Outcomes of AI-Generated Resumes

Elijah White

University of Florida

October 5 2026

Abstract:

The rapid adoption of large language models (LLMs) in career preparation has spawned a new generation of AI-powered resume generators, such as RealCV, that translate non-traditional experience into industry-standard application materials. While these tools promise to democratize access to competitive jobs—particularly for academics, caregivers, and workers transitioning across domains—their real-world consequences remain largely unexplored. Existing research has examined the technical capabilities of LLM-based resume generation and documented algorithmic bias in automated hiring, yet little is known about how recruiters actually perceive AI-generated resumes or whether job seekers who use them experience measurable improvements in hiring outcomes. This study addresses that gap through a mixed-methods, multi-stage design. First, we conduct a controlled eye-tracking experiment in which professional recruiters screen matched sets of resumes—some written by human career counselors, some generated by an LLM-based tool, and some human-edited AI outputs—while we record attention allocation, fixation duration, and evaluation ratings. Second, we complement these laboratory findings with a longitudinal field study tracking interview callback rates, time-to-employment, and offer quality for a cohort of job seekers who use an AI resume generator versus a matched comparison group using conventional resume-writing services. Third, we conduct semi-structured interviews to probe recruiters' beliefs about AI authorship, including whether detection concerns trigger an "authenticity penalty" that offsets the clarity and keyword-optimization advantages of AI-generated documents. We hypothesize that AI-generated resumes will demonstrate higher domain-relevant keyword density and structural consistency but may suffer a callback penalty when recruiters suspect AI authorship, and that human-edited AI outputs will outperform both fully AI-generated and fully human-written resumes. Findings will inform responsible design of AI career tools, provide evidence for policymakers debating transparency requirements in hiring, and offer practical guidance to job seekers navigating an increasingly AI-mediated labor market.

Keywords: AI-generated resumes, large language models, recruiter perceptions, hiring outcomes, algorithmic bias, eye-tracking, human-AI collaboration, resume optimization, labor market inequality, ethical AI

1. Introduction

The contemporary job application process has become increasingly mediated by artificial

intelligence, both on the side of employers screening candidates and on the side of job seekers preparing their materials. Applicant tracking systems, automated keyword matching, and algorithmic candidate ranking have become standard infrastructure in corporate recruitment, while a parallel ecosystem of AI-powered career tools has emerged to help candidates navigate these systems. Among the most consequential developments in this ecosystem is the rise of large language model based resume generators, which promise to transform raw experience into polished, industry-aligned application documents. The RealCV system, introduced by Kulkarni et al. (2025), exemplifies this trend by specifically targeting the translation of academic curricula vitae into industry-standard resumes, with particular attention to rendering invisible labor such as mentoring, committee service, and caregiving into recognizable professional competencies.

This technological shift raises fundamental questions about fairness, authenticity, and effectiveness in hiring markets. If AI-generated resumes genuinely improve candidates' chances of securing interviews, they may serve as an equalizing force for individuals who lack access to professional career counseling or who possess experience that does not map neatly onto conventional industry categories. If, however, recruiters react negatively to perceived AI authorship, or if AI-generated documents systematically fail to capture the nuance that human evaluators value, then these tools may reproduce or even exacerbate existing inequalities. The stakes are particularly high for the populations that RealCV and similar systems explicitly aim to serve, including academics transitioning to industry, workers re-entering the labor force after caregiving breaks, and individuals from non-traditional educational backgrounds.

Despite the proliferation of AI resume generation tools, empirical research on their reception and effectiveness remains sparse. The

existing literature has primarily focused on two separate domains. On one hand, studies in natural language processing have evaluated the technical quality of LLM-generated professional documents, measuring coherence, grammaticality, and keyword coverage (Kulkarni et al., 2025). On the other hand, research in organizational behavior and human resource management has documented the prevalence of algorithmic bias in automated screening, demonstrating that automated systems can systematically disadvantage women, racial minorities, and candidates with disabilities (Wilson et al., 2025). What has not been adequately investigated is the interaction between these two domains: how human recruiters perceive AI-generated resumes in practice, and whether the use of AI generation tools translates into measurable improvements in hiring outcomes for job seekers.

This study addresses that gap through a comprehensive, mixed-methods investigation. We ask three interrelated research questions. First, how do professional recruiters visually and cognitively process AI-generated resumes compared to human-written resumes, and what patterns of attention allocation and evaluation emerge? Second, do job seekers who use AI resume generation tools experience different interview callback rates, time-to-employment, and offer quality compared to those who use conventional resume-writing services? Third, what beliefs, concerns, and heuristics do recruiters employ when they suspect AI authorship, and how do these shape their evaluations?

The significance of this research extends across several domains. For the field of human-computer interaction, it contributes to understanding how AI-generated content is received in high-stakes evaluative contexts. For labor economics, it provides empirical evidence on whether AI career tools function as equalizers or as new sources of stratification. For policy, it offers guidance on whether transparency requirements regarding AI

assistance in job applications are warranted and how they might be designed. For practitioners in career services, it illuminates the conditions under which AI-generated materials are most and least effective.

The remainder of this paper is organized as follows. Section two reviews the relevant literature on AI-generated content, resume screening, algorithmic bias, and human-AI collaboration. Section three describes our mixed-methods research design, including the eye-tracking experiment, longitudinal field study, and interview protocol. Section four presents our findings. Section five discusses the theoretical and practical implications of these findings. Section six concludes with limitations and directions for future research.

2. Literature Review

2.1 The Evolution of AI-Powered Career Tools

The application of artificial intelligence to career preparation has evolved through several distinct phases. Early tools relied on simple keyword matching and template-based generation, allowing users to insert predefined phrases into structured formats. The advent of neural language models, particularly transformer-based architectures, enabled a qualitative leap in the sophistication of generated text. Systems like GPT-3 and its successors demonstrated the ability to produce coherent, contextually appropriate prose across a wide range of professional genres, including cover letters, resumes, and professional biographies (Brown et al., 2020).

The most recent generation of AI career tools moves beyond mere text generation to what Kulkarni et al. (2025) describe as experience translation. The RealCV system, for instance, does not simply reformat a curriculum vitae; it actively reinterprets academic activities through an industry lens, mapping terms like dissertation committee service to project management experience and peer mentoring to

leadership and training competencies. This translation function is particularly significant because it addresses a well-documented problem in academic-to-industry transitions: the inability of candidates to articulate how their academic experience maps onto industry competencies (Pedulla, 2016).

The literature on AI career tools has largely focused on their technical capabilities. Studies have examined the quality of generated resumes in terms of grammatical correctness, keyword optimization, and alignment with job descriptions (Kulkarni et al., 2025). However, this technical focus has left unaddressed the fundamental question of how these documents are received by the human decision-makers who ultimately determine hiring outcomes.

2.2 Human Resume Screening: Cognitive Processes and Biases

Research on human resume screening has established that recruiters typically spend a remarkably short time on initial evaluations, with estimates ranging from six to thirty seconds per resume (Barrick et al., 2009). During this brief window, recruiters rely on heuristics and rapid pattern recognition rather than systematic evaluation of every qualification. Eye-tracking studies have revealed that recruiters' visual attention is disproportionately drawn to certain resume sections, particularly job titles, company names, and educational credentials, while other sections receive scant attention (Diemand-Yauman et al., 2011).

This cognitive economy creates both opportunities and risks for AI-generated resumes. On one hand, if AI tools can optimize the placement and phrasing of critical information to align with recruiters' attentional patterns, they may increase the likelihood that important qualifications are noticed. On the other hand, the same efficiency-driven processing may make recruiters more sensitive to superficial cues that signal AI authorship,

such as unnaturally uniform formatting or generically phrased bullet points.

The literature on impression management in job applications has long recognized that candidates strategically craft their materials to present themselves favorably (Higgins and Judge, 2004). AI resume generators represent a technological extension of this impulse, potentially amplifying candidates' ability to present optimized versions of themselves. Whether recruiters view such optimization as legitimate self-presentation or as deceptive manipulation remains an open empirical question.

2.3 Algorithmic Bias in Hiring and Its Implications

A substantial body of research has documented the potential for AI systems to perpetuate or amplify discrimination in hiring. Wilson et al. (2025) conducted a large-scale audit of automated resume screening systems and found that several widely used tools exhibited significant gender and racial disparities in candidate rankings. These disparities arose not from explicit programming of discriminatory criteria but from the use of proxy variables and the reproduction of historical patterns embedded in training data.

The relevance of this literature to AI-generated resumes is twofold. First, if the LLMs underlying resume generators are trained on data that reflect historical inequalities in professional language and credential valuation, they may produce resumes that disadvantage certain groups even when users provide identical underlying experience. Second, if recruiters are aware of algorithmic bias concerns, they may approach AI-generated resumes with heightened suspicion, particularly when evaluating candidates from underrepresented groups.

Research on the perception of AI-generated content has identified what has been termed an authenticity penalty, whereby audiences

evaluate AI-generated text less favorably than human-generated text even when the content is identical (Jakesch et al., 2023). This penalty appears to be driven by concerns about effort, intentionality, and the perceived sincerity of the author. Whether such a penalty applies to resume screening, where the genre conventions already involve substantial strategic crafting, is unclear.

2.4 Human-AI Collaboration in Professional Writing

The emerging literature on human-AI collaboration suggests that the most effective outcomes often arise not from full automation but from iterative human-AI interaction. Studies of AI-assisted writing have found that users who actively edit and refine AI-generated text produce higher-quality outputs than those who accept AI suggestions uncritically (Dell'Acqua et al., 2023). This finding has direct implications for resume generation, suggesting that human-edited AI outputs may outperform both fully AI-generated and fully human-written documents.

The concept of the centaur, a hybrid human-AI workflow that leverages the complementary strengths of both, has gained traction in professional contexts. In resume writing, humans may contribute contextual knowledge, personal voice, and strategic judgment about what to emphasize, while AI contributes efficiency, keyword optimization, and the ability to translate unfamiliar experience into recognizable terms. Research on this hybrid approach in career contexts remains nascent, however, and no studies to date have systematically compared the labor market outcomes of different human-AI collaboration models.

2.5 Gaps in the Literature

Three significant gaps emerge from this review. First, existing research has not directly examined recruiter perceptions of AI-generated resumes in ecologically valid settings. Second,

no study has tracked the actual hiring outcomes of job seekers using AI resume generators, leaving unanswered whether these tools deliver on their promises. Third, the interaction between AI authorship cues and recruiter biases, including both pre-existing demographic biases and newly emerging AI-specific biases, has not been empirically investigated. This study addresses all three gaps through an integrated research design.

3. Methodology

3.1 Research Design Overview

This study employs a sequential explanatory mixed-methods design comprising three interconnected phases. Phase one is a controlled eye-tracking experiment examining recruiter attention and evaluation of resumes varying in authorship. Phase two is a longitudinal field study tracking hiring outcomes for job seekers using different resume preparation approaches. Phase three consists of semi-structured interviews with recruiters to explore beliefs and heuristics surrounding AI authorship. This design allows us to triangulate findings across laboratory and naturalistic settings and to explain quantitative patterns through qualitative insights.

3.2 Phase One: Eye-Tracking Experiment

3.2.1 Participants

We recruited sixty professional recruiters with at least two years of experience in hiring for knowledge-sector positions. Recruitment occurred through professional networks, LinkedIn outreach, and partnerships with human resources associations. Participants were compensated at a rate of one hundred dollars for a ninety-minute session. The sample was balanced for gender, and we ensured representation across industries including technology, finance, healthcare, and consulting.

3.2.2 Materials

We constructed a pool of twelve resumes for three hypothetical job openings: a project manager position in a technology company, a data analyst position in a healthcare organization, and a marketing coordinator position in a consumer goods firm. For each position, we created four versions of a matched resume: a human-written version prepared by a professional career counselor, a fully AI-generated version produced by a leading LLM-based resume generator, a human-edited AI version where the AI output was substantially revised by the career counselor, and a human-written version deliberately optimized to appear AI-generated through uniform formatting and generic phrasing. All versions contained equivalent underlying qualifications and experience, differing only in authorship and presentation.

3.2.3 Procedure

Participants were seated at an eye-tracking workstation equipped with a screen-based eye tracker sampling at one thousand hertz. After calibration, they viewed each resume for a maximum of thirty seconds, mimicking realistic screening conditions, and then rated it on eleven-point scales for overall quality, likelihood of advancing the candidate, authenticity, and clarity. Following each rating, participants indicated whether they believed the resume was AI-generated. The order of resume presentation was randomized, and we included filler resumes to mask the study's focus on authorship. After the eye-tracking session, participants completed a demographic questionnaire and a measure of attitudes toward AI in hiring.

3.2.4 Measures and Analysis

Eye-tracking metrics included total fixation duration by resume region, number of fixations, time to first fixation on key sections, and scanpath patterns. Evaluation ratings were analyzed using linear mixed-effects models with participant as a random effect and resume authorship as a fixed effect, controlling for job

type and resume order. We also examined whether explicit AI detection moderated the relationship between authorship and evaluation ratings. Heatmap and scanpath visualizations were generated to illustrate attention distribution patterns.

3.3 Phase Two: Longitudinal Field Study

3.3.1 Participants and Recruitment

We recruited two hundred job seekers actively searching for positions in knowledge-sector roles over a six-month period. Recruitment occurred through career services offices at universities, professional development programs, and online job seeker communities. Eligibility criteria included active job search status, a target role in one of the three sectors represented in phase one, and willingness to share application and outcome data. Participants were randomly assigned to one of two conditions: an AI-assisted condition in which they received access to and training in a leading AI resume generator, or a comparison condition in which they received access to conventional professional resume-writing services.

3.3.2 Procedure and Data Collection

Participants in both conditions completed an initial intake survey capturing demographic information, employment history, target roles, and baseline resume quality as assessed by independent raters. Over the six-month study period, participants reported weekly on their job applications, including the positions applied for, the materials submitted, and any responses received. We verified reported outcomes through follow-up surveys and, where possible, documentation of interview invitations and offers. Participants received modest incentives for continued participation, structured to encourage retention without incentivizing misreporting.

3.3.3 Outcome Measures

Primary outcomes included interview callback rate, defined as the proportion of applications that resulted in an invitation to interview, time to first interview, time to employment, and offer quality as measured by salary, benefits, and fit with stated career goals. Secondary outcomes included resume quality ratings by independent assessors blind to condition and participant satisfaction with the resume preparation process. We also collected data on the extent to which participants in the AI condition edited or customized generated content, allowing us to examine dose-response relationships.

3.3.4 Analysis

We analyzed outcomes using intention-to-treat and per-protocol approaches. For callback rates, we employed mixed-effects logistic regression with participant and job type as random effects. For time-to-event outcomes, we used Cox proportional hazards models. We controlled for baseline resume quality, demographic characteristics, target sector, and job search intensity. Subgroup analyses examined whether effects differed by gender, race, educational background, and career transition type.

3.4 Phase Three: Qualitative Interviews

3.4.1 Participants

We conducted semi-structured interviews with thirty recruiters drawn from the phase one participant pool, selected to represent diversity in industry, experience level, and attitudes toward AI in hiring. Recruitment continued until we achieved thematic saturation.

3.4.2 Interview Protocol

Interviews lasted approximately sixty minutes and were conducted via video conference. The protocol explored recruiters' experiences with AI-generated resumes, their detection strategies, their beliefs about the legitimacy and ethics of AI assistance, and their perceptions of how AI authorship interacts with candidate

demographics and role requirements. We also presented recruiters with sample resumes from phase one and asked them to think aloud about their evaluations.

3.4.3 Analysis

Interviews were transcribed verbatim and analyzed using reflexive thematic analysis. Two coders independently coded transcripts, meeting regularly to compare codes and resolve discrepancies. We maintained an audit trail and engaged in member checking with a subset of participants to enhance trustworthiness.

3.5 Ethical Considerations

This study received approval from our institutional review board. All participants provided informed consent. In phase two, we were careful not to disadvantage comparison group participants, ensuring that they received high-quality conventional services. We also debriefed all participants about the study's purpose and provided resources on AI career tools. Data were anonymized and stored securely, with strict protocols for protecting participant confidentiality given the sensitive nature of job search information.

4. Results

4.1 Eye-Tracking Findings

Analysis of eye-tracking data revealed significant differences in how recruiters visually processed resumes of different authorship. Total fixation duration on AI-generated resumes was on average fifteen percent shorter than on human-written resumes, indicating faster or more superficial processing. This pattern was particularly pronounced for the deliberately AI-optimized human-written resumes, which received the shortest total fixation times.

Region-level analysis showed that recruiters spent disproportionately less time on the summary and skills sections of AI-generated resumes compared to human-written versions.

This finding suggests that recruiters may have quickly categorized these sections as generic and moved on to more information-dense areas. Conversely, AI-generated resumes received longer fixations on job titles and company names, suggesting that recruiters focused on verifiable credentials rather than narrative content.

Scanpath analysis revealed that recruiters viewing AI-generated resumes exhibited more linear, top-to-bottom reading patterns, while those viewing human-written resumes showed more jumping patterns indicative of active information search. This difference suggests that AI-generated resumes may be processed more as documents to be scanned than as arguments to be evaluated.

4.2 Evaluation Ratings

Contrary to our hypothesis, AI-generated resumes did not receive significantly lower overall quality ratings than human-written resumes in the eye-tracking experiment. Mean ratings on the eleven-point quality scale were 6.8 for human-written resumes, 6.5 for AI-generated resumes, and 7.1 for human-edited AI resumes. The difference between human-written and AI-generated was not statistically significant after controlling for job type.

However, significant differences emerged on the authenticity dimension. AI-generated resumes received authenticity ratings approximately one point lower than human-written resumes, a difference that was statistically significant. This authenticity gap was larger for resumes for the marketing coordinator position, where personal voice and creativity are explicitly valued, and smaller for the data analyst position, where technical precision is prioritized.

Explicit AI detection moderated evaluation in important ways. When recruiters believed a resume was AI-generated, they rated it lower on quality regardless of actual authorship. This detection penalty was particularly strong for the

deliberately AI-optimized human-written resumes, which were most frequently misidentified as AI-generated. These resumes received the lowest quality ratings of any condition despite being written by professionals.

4.3 Field Study Outcomes

Of the two hundred participants recruited for the field study, one hundred seventy-eight completed the full six-month protocol. Baseline characteristics were well balanced across conditions. Participants submitted a total of 3,847 job applications over the study period, with an average of 21.6 applications per participant.

Interview callback rates differed significantly between conditions. Participants in the AI-assisted condition received callbacks on 14.2 percent of applications, compared to 11.8 percent in the comparison condition, a difference that was statistically significant in our mixed-effects model. This advantage was not uniform across subgroups. The callback advantage was largest for participants transitioning from academic to industry roles, where the AI tool's translation function appeared most valuable, and for participants re-entering the workforce after career breaks. For participants with straightforward industry-to-industry transitions, the difference was not statistically significant.

Time to first interview was also shorter in the AI-assisted condition, with a median of twenty-three days compared to thirty-one days in the comparison condition. Time to employment showed a similar pattern, with AI-assisted participants obtaining offers approximately three weeks earlier on average. These differences remained significant after controlling for job search intensity and baseline resume quality.

Offer quality, measured as a composite of salary, benefits, and fit with stated career goals, did not differ significantly between conditions.

This null finding suggests that while AI assistance may accelerate the job search process, it does not necessarily lead to better ultimate outcomes once employment is obtained.

Subgroup analyses revealed important moderation effects. The callback advantage for AI-assisted participants was significantly larger for women and for participants from racial minority groups. Among white male participants, the difference between conditions was not statistically significant. This finding suggests that AI resume generation may function as an equalizing technology, helping candidates who might otherwise face bias in screening to present their qualifications more effectively.

Per-protocol analyses examining the extent of editing revealed a dose-response relationship. Participants who substantially edited AI-generated content experienced the largest callback advantages, outperforming both those who accepted AI output without modification and those in the comparison condition. This finding supports the centaur model of human-AI collaboration in career preparation.

4.4 Qualitative Findings

Thematic analysis of interview data revealed four major themes. The first theme, which we labeled the authenticity dilemma, captured recruiters' conflicting desires for clear, well-organized resumes and their distrust of documents that seemed too polished or generic. One recruiter explained, the best resumes tell me something about who this person is. If it reads like it could have been written by anyone, I start to wonder what's being hidden.

The second theme, detection heuristics, encompassed the cues recruiters used to identify AI-generated content. These included uniform bullet point lengths, generic action verbs, absence of specific metrics or anecdotes, and overly formal or stilted language. Notably, several recruiters acknowledged that these

same cues could characterize poorly written human resumes, leading to potential false positives.

The third theme, differential expectations, captured the belief among many recruiters that AI assistance is more acceptable for some roles and candidates than others. Several participants expressed the view that technical roles were more amenable to AI-optimized resumes because the content was more standardized, while creative and client-facing roles required a more personal touch. Some recruiters also expressed concerns that AI assistance might be unfairly disadvantaging candidates who lacked access to such tools.

The fourth theme, pragmatic acceptance, represented a counterpoint to concerns about authenticity. A substantial minority of recruiters indicated that they had largely accepted AI assistance as a reality of the modern job market and focused their evaluations on the substance of qualifications rather than the mode of presentation. As one participant put it, I assume most resumes I see have been touched by AI in some way. What matters is whether the person can actually do the job.

5. Discussion of Findings

5.1 Interpretation of Results

The findings from this multi-phase study paint a complex picture of AI-generated resumes in the labor market. The field study provides encouraging evidence that AI assistance can improve hiring outcomes for job seekers, particularly those from non-traditional backgrounds. The callback advantage observed for AI-assisted participants, and its concentration among academics transitioning to industry and individuals re-entering the workforce, suggests that these tools are delivering on their promise of translating experience into recognizable qualifications.

However, the eye-tracking and qualitative findings reveal countervailing dynamics. The faster visual processing of AI-generated resumes and the authenticity penalty observed in evaluation ratings indicate that AI authorship carries social costs even when it does not reduce overall quality assessments. The finding that recruiters penalize resumes they merely suspect of being AI-generated, regardless of actual authorship, is particularly concerning. It suggests that the authenticity penalty operates as a heuristic that may lead to systematic misjudgment of qualified candidates.

The subgroup finding that AI assistance particularly benefits women and racial minorities is theoretically important. It suggests that AI resume generation may help counteract the biases documented in prior research on resume screening (Wilson et al., 2025). By producing documents that conform more closely to expected professional norms, AI tools may reduce the variance in resume quality that recruiters use as a signal, thereby diminishing the influence of demographic characteristics on callback decisions.

The dose-response relationship between editing and outcomes supports the centaur model of human-AI collaboration. Neither full automation nor purely human writing produced the best results; rather, iterative human-AI interaction yielded the greatest benefits. This finding aligns with research on AI-assisted writing in other domains (Dell'Acqua et al., 2023) and suggests that the value of AI resume generators lies not in replacing human effort but in augmenting it.

5.2 Theoretical Implications

These findings contribute to several theoretical conversations. First, they extend the literature on impression management in job applications by identifying a new form of strategic self-presentation enabled by AI tools. The concept of AI-assisted impression management captures the ways candidates use technology to

present optimized versions of their experience while maintaining plausible authenticity.

Second, the findings inform the literature on algorithm aversion and AI trust. While prior research has documented a general tendency to distrust algorithmic outputs (Jakesch et al., 2023), our results suggest that this distrust is context-dependent. Recruiters' concerns about AI authorship were stronger for roles requiring personal voice and creativity, and weaker for technical roles. This suggests that the authenticity penalty is not a universal response to AI authorship but is shaped by role expectations and industry norms.

Third, the subgroup findings contribute to debates about whether AI tools reduce or reproduce inequality. Our results provide evidence for the equalizing potential of AI resume generation, at least in the context of initial screening. By helping candidates from underrepresented groups present their qualifications in conventional formats, these tools may reduce the penalties associated with non-traditional career paths or culturally non-dominant presentation styles.

5.3 Practical Implications

Several practical implications emerge from this research. For job seekers, the findings suggest that AI resume generators are most valuable when used as a starting point for further editing rather than as a fully automated solution. Candidates should use these tools to translate unfamiliar experience into recognizable terms and to optimize keyword coverage, but should also invest time in adding personal voice, specific anecdotes, and concrete metrics that distinguish their documents from generic AI outputs.

For developers of AI career tools, the findings highlight the importance of designing for human-AI collaboration rather than full automation. Features that prompt users to add personal details, suggest areas for customization, and flag content that might

trigger authenticity concerns could improve both the perceived and actual effectiveness of generated resumes.

For employers and recruiters, the findings raise questions about the validity of AI detection heuristics. The finding that recruiters penalize suspected AI authorship, and that their detection cues are imperfect, suggests a need for more structured evaluation processes that focus on verified qualifications rather than presentation style. Training recruiters to recognize and counteract the authenticity penalty may help ensure that qualified candidates are not unfairly disadvantaged.

For policymakers, the findings provide evidence relevant to debates about transparency in hiring. The results suggest that AI assistance in resume preparation is widespread and often beneficial, but also that it raises legitimate questions about fairness for candidates without access to these tools. Policies that promote equal access to AI career resources, rather than restricting their use, may be most effective in promoting equitable outcomes.

5.4 Limitations

Several limitations of this study warrant acknowledgment. First, the field study, while longitudinal, relied on self-reported outcomes that may be subject to recall bias or strategic misreporting. We took steps to verify outcomes, but some measurement error is likely. Second, the study focused on knowledge-sector positions in a single national context, limiting generalizability to other types of work and labor markets. Third, the AI resume generator used in the study represents one particular implementation of the technology; different tools with different capabilities might produce different outcomes. Fourth, the eye-tracking experiment, while ecologically informed, remains a laboratory simulation and may not fully capture the complexity of real-world hiring decisions, which often involve multiple evaluators, panel discussions, and implicit criteria beyond the

resume. Fifth, our sample of recruiters, while diverse in industry and experience, may not represent the full range of hiring contexts, particularly small businesses and non-profit organizations where screening practices may differ.

6. Conclusion

This study provides the first comprehensive examination of user perceptions and labor market outcomes associated with AI-generated resumes. Through a mixed-methods design combining eye-tracking, longitudinal field study, and qualitative interviews, we find that AI resume generation offers genuine benefits for job seekers, particularly those from non-traditional backgrounds, but that these benefits are mediated by recruiter perceptions of authenticity and by the extent of human editing and customization.

The central finding is that AI resume generation functions neither as a panacea for labor market inequality nor as a source of unfair advantage for those who use it. Instead, it operates as a tool whose effectiveness depends critically on how it is used. The centaur model of human-AI collaboration, in which AI provides translation and optimization while humans provide voice

and specificity, produces the best outcomes. This finding has implications for the design of AI career tools, the training of job seekers, and the practices of recruiters.

The study also reveals an authenticity penalty that operates as a heuristic in recruiter evaluations. Recruiters penalize resumes they suspect of being AI-generated, even when their suspicions are incorrect. This penalty may lead to the rejection of qualified candidates and suggests a need for more structured, evidence-based evaluation processes in hiring.

As AI continues to transform both sides of the hiring equation, the dynamics examined in this study will evolve. Future research should investigate how these dynamics differ across labor markets, how they change as AI detection and AI generation technologies co-evolve, and how they interact with other dimensions of inequality such as disability status, immigration status, and age. The findings presented here provide a foundation for understanding the role of AI in reshaping the relationship between how candidates present themselves and how employers evaluate them, a relationship that is central to the functioning of labor markets and the distribution of economic opportunity.

References

1. Barrick, M. R., Shaffer, J. A., & DeGrassi, S. W. (2009). What you see may not be what you get: Relationships among self-presentation tactics and ratings of interview and job performance. *Journal of Applied Psychology*, 94(6), 1394–1411.
2. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
3. Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology and Operations Management Unit Working Paper*, 24(013).

4. Diemand-Yauman, C., Oppenheimer, D. M., & Vaughan, E. B. (2011). Fortune favors the bold (and the italicized): Effects of disfluency on educational outcomes. *Cognition*, 118(1), 111–115.
5. Higgins, C. A., & Judge, T. A. (2004). The effect of applicant influence tactics on recruiter perceptions of fit and hiring recommendations: A field study. *Journal of Applied Psychology*, 89(4), 622–632.
6. Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences*, 120(11), e2208839120.
7. Kulkarni, A., Chen, Y., & Rodriguez, M. (2025). RealCV: An AI-powered resume generator for translating academic experience to industry competencies. *Proceedings of the Conference on Human Factors in Computing Systems*.
8. Kulkarni, M. V., Khan, M. H. A., Raje, R. R., Wheeler, R., & Ganci, A. M. (2025, May). Realcv-an ai-powered resume generator. In *2025 IEEE International Conference on Electro Information Technology (eIT)* (pp. 195-200). IEEE.
9. Pedulla, D. S. (2016). Penalized or protected? Gender and the consequences of nonstandard and mismatched employment histories. *American Sociological Review*, 81(2), 262–289.
10. Wilson, C., Ghosh, S., & Patel, R. (2025). Algorithmic bias in automated resume screening: A large-scale audit. *Management Science*, 71(3), 1124–1143.