



# **iJRASET**

International Journal For Research in  
Applied Science and Engineering Technology



---

# **INTERNATIONAL JOURNAL FOR RESEARCH**

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

---

**Volume:** 13      **Issue:** XI      **Month of publication:** November 2025

**DOI:** <https://doi.org/10.22214/ijraset.2025.75771>

**[www.ijraset.com](http://www.ijraset.com)**

**Call:** ☎ 08813907089

**E-mail ID:** [ijraset@gmail.com](mailto:ijraset@gmail.com)

# Mitigating Bias in AI-Driven Recruitment: A Comparative Study of Machine Learning Algorithms Vs. Human Screeners

Satyam Bajpai<sup>1</sup>, Siddhant Mishra<sup>2</sup>

<sup>1</sup>MBA Scholar, Maharana Pratap Engineering College, Kanpur, Uttar Pradesh, India

<sup>2</sup>Assistant Professor, Maharana Pratap Engineering College, Kanpur, Uttar Pradesh, India

**Abstract:** *The integration of Artificial Intelligence (AI) into human resource management offers a promising avenue for enhancing efficiency, yet concerns regarding algorithmic fairness persist. This study investigates the efficacy of Machine Learning (ML) algorithms compared to human screeners in mitigating bias during the initial candidate screening process. The purpose of this research is to determine whether AI systems reduce demographic biases or merely replicate the historical prejudices embedded in training data.*

*Using a controlled experimental design, the study analyzed a dataset of 5,000 anonymized resumes. A supervised learning model (Random Forest) was pitted against a panel of experienced human recruiters to evaluate candidates based on identical job descriptions. Key findings indicate that while human screeners exhibited significant "affinity bias"—favoring candidates with similar educational backgrounds—the baseline ML model initially perpetuated gender bias found in historical hiring data. However, after applying algorithmic de-biasing techniques, the AI demonstrated a higher degree of consistency and fairness than the human control group. The study concludes that AI is not a silver bullet; however, when rigorously audited, it serves as a crucial objective counterweight to human subjectivity. These results advocate for a "human-in-the-loop" hybrid approach to ensure equitable recruitment practices.*

**Keywords:** *Algorithmic Bias, Machine Learning, Recruitment, Diversity and Inclusion, Human Resource Management.*

## I. INTRODUCTION

### A. Background of the Study

The landscape of Human Resource Management (HRM) has undergone a paradigm shift in the last decade, driven by the rapid digitization of the workforce and the "war for talent." Traditional recruitment—characterized by manual resume reviews and subjective face-to-face interactions—is increasingly being augmented, and in some cases replaced, by Artificial Intelligence (AI).

The rise of AI in HR is primarily a response to volume. With multinational corporations receiving millions of applications annually, the reliance on Machine Learning (ML) algorithms for candidate screening has become a necessity for operational efficiency. These systems, often embedded within Applicant Tracking Systems (ATS), utilize Natural Language Processing (NLP) to parse resumes, match keywords, and predict candidate success probabilities. While initially designed to reduce administrative burdens, these tools are now central gatekeepers to economic opportunity.

### B. Problem Statement

Despite the promise of efficiency, the deployment of AI in recruitment has introduced a critical ethical dilemma: the persistence of algorithmic bias. AI models are typically trained on historical hiring data; if that history contains patterns of discrimination (e.g., favoring male candidates for technical roles), the algorithm often learns and amplifies these prejudices.

The problem is twofold. First, AI systems can inadvertently scale discrimination, rejecting qualified diverse candidates at a speed and volume humans cannot match. Second, while AI is often scrutinized, human screeners are equally flawed, prone to cognitive shortcuts such as the "halo effect," "affinity bias" (preferring candidates like oneself), and fatigue. The core problem this study addresses is not merely that AI is biased, but that we lack a clear understanding of how algorithmic bias specifically compares to and interacts with the inherent biases of human decision-makers.

### C. Research Gap

Current literature is rich with studies on algorithmic fairness in isolation and psychological studies on human bias in isolation. However, there is a distinct comparative research gap.

- Few studies directly compare the output of an ML screening algorithm against a human panel using the exact same dataset.
- There is limited empirical evidence on whether "human-in-the-loop" interventions actually mitigate bias or simply reintroduce human subjectivity into the process.

This study bridges that gap by providing a side-by-side efficacy analysis of machine versus human screening.

### D. Research Objectives

To address the identified problem and gap, this study aims to achieve the following objectives:

- 1) To identify the sources and patterns of bias (e.g., gender, name-based, educational) embedded within standard AI recruitment tools.
- 2) To compare the fairness and consistency of AI-driven screening outcomes against those derived from human screening panels.
- 3) To explore and validate technical and procedural strategies (such as algorithmic de-biasing and blinded resumes) for mitigating bias in recruitment.

### E. Research Questions

Guided by the objectives above, this research seeks to answer:

- 1) What types of bias are prevalent in AI-driven recruitment systems?
- 2) How do these biases compare to human decision-making biases?
- 3) What mitigation techniques can reduce bias in AI-based recruitment?

### F. Significance of the Study

The findings of this research hold substantial value for three key groups:

- 1) HR Professionals: It offers evidence-based guidelines on how to responsibly integrate AI tools without compromising diversity and inclusion goals.
- 2) Developers & Data Scientists: It highlights specific pitfalls in training data and model architecture, contributing to the development of "fairness-aware" algorithms.
- 3) Policymakers: As governments (such as the EU with its AI Act) move to regulate high-risk AI applications in employment, this study provides empirical data to inform ethical guidelines and compliance standards.

Ultimately, this study contributes to the broader societal goal of creating equitable access to employment, ensuring that technological advancement does not come at the cost of social justice.

## II. LITERATURE REVIEW

This section synthesizes existing research on the dual challenges of human and algorithmic bias in recruitment, establishing the theoretical and practical foundations for this study.

### A. AI in Recruitment: Overview of Automation

The digitization of Human Resource Management (HRM) has transitioned from simple keyword matching to complex predictive modeling.

- 1) Adoption Rates: As of 2025, approximately 43% of organizations report using AI for HR-related tasks, a significant jump from 26% in 2024.
- 2) Efficiency Gains: Organizations utilizing AI in hiring report an 89.6% increase in hiring efficiency and an 85.3% reduction in time-to-hire.
- 3) Scope: Automation now spans the entire recruitment funnel, from programmatic advertising (placing ads where specific demographics will see them) to resume parsing (extracting skills from unstructured text) and candidate ranking (predicting "cultural fit" or retention likelihood).

### B. Human Bias in Recruitment

Despite the "objectivity" implied by professional standards, human decision-making in recruitment is rife with cognitive shortcuts and implicit associations.

#### 1) Racial & Ethnic Bias:

- In a landmark field experiment by Bertrand and Mullainathan (and replicated in subsequent years), applicants with "White-sounding" names (e.g., Emily, Greg) received 50% more callbacks than applicants with "African-American-sounding" names (e.g., Lakisha, Jamal) despite having identical resumes.
- Critically, improving resume quality (e.g., adding certifications) resulted in a 30% callback increase for White candidates but yielded statistically negligible benefits for African-American candidates.

#### 2) Gender Bias:

- While women constitute roughly 48% of entry-level corporate roles, their representation drops to 29% at the C-suite level, illustrating the "broken rung" phenomenon.
- Male candidates are frequently 1.5 times more likely to advance to initial screening phases in STEM roles compared to equally qualified female candidates.

#### 3) Affinity Bias:

- Approx. 48% of HR managers admit that personal bias (favoring candidates with similar backgrounds, hobbies, or alma maters) influences their hiring choices.

### C. Algorithmic Bias

Algorithmic bias is not an error in the code, but often an accurate reflection of biased history.

#### 1) Causes:

- Training Data: If an algorithm is trained on 10 years of hiring data where men were predominantly hired, it learns that "being male" is a predictor of success.
- Proxy Variables: Even if gender is removed, algorithms find proxies. For example, an algorithm might penalize graduates from "women's colleges" or downgrade applicants who list "softball" as a hobby.

#### 2) Case Studies:

- Amazon (2018): Amazon scrapped an experimental AI recruiting tool after it penalized resumes containing the word "women's" (e.g., "Women's Chess Club Captain"). The model had been trained on resumes submitted over a 10-year period, most of which came from men, leading the system to conclude that male candidates were preferable.
- LinkedIn: Algorithms designed to maximize user engagement were found to show high-paying job ads more frequently to men than women, not due to malicious intent, but because female users were "more expensive" to advertise to (higher market demand), causing the algorithm to optimize for the cheaper demographic (men) to save budget.

### D. Theoretical Frameworks

1) Social Bias Theory: Posits that algorithms are "socio-technical" systems. They do not exist in a vacuum but absorb the pre-existing cultural and institutional expectations of their creators and data sources.

#### 2) Algorithmic Fairness Models:

- Outcome Fairness: Ensuring the statistical likelihood of being hired is equal across groups (e.g., the "80% Rule" used in US employment law).
- Process Fairness: Ensuring the features used to make the decision (e.g., education, experience) are valid and not discriminatory proxies.

3) Human-AI Interaction Theory: Focuses on the "Black Box" problem. Recruiters are less likely to trust or correct an AI's decision if they cannot understand why a candidate was rejected (lack of interpretability), leading to automation bias where humans uncritically accept machine outputs.

### E. Mitigation Techniques

1) Bias Detection & Auditing: Using "adversarial testing," where synthetic resumes are fed into the system to test for disparate impact before deployment.



- 2) Explainable AI (XAI): Techniques like LIME (Local Interpretable Model-agnostic Explanations) and SHAP values help recruiters see exactly which words (e.g., "Python," "Harvard," "Gap Year") caused a resume's score to go up or down, allowing for human intervention if the reasoning is flawed.
- 3) Data Balancing: Techniques such as over-sampling underrepresented groups in the training data to prevent the model from viewing them as "outliers."

### III. RESEARCH METHODOLOGY

This section outlines the systematic approach used to compare the screening efficacy and bias levels of Machine Learning algorithms versus human recruiters.

#### A. Research Design

This study employs a Mixed-Methods Convergent Parallel Design.

- 1) Quantitative Stream: An experimental simulation testing the statistical outputs of ML algorithms against human decisions on a standardized dataset.
- 2) Qualitative Stream: Semi-structured interviews with the human recruiters post-experiment to understand the cognitive rationale behind their screening choices.

#### B. Sample Selection

To ensure a controlled environment, the sample is divided into two categories:

- 1) The Dataset (Candidates):
  - A dataset of N=2,000 anonymized resumes was compiled.
  - The dataset is a modified version of the Kaggle Recruitment Data, stratified to ensure equal distribution of gender (50% Male/Female) and ethnicity proxies, while skill qualifications were randomized.
- 2) The Participants (Human Screeners):
  - N=30 experienced HR professionals (minimum 5 years experience) were recruited from three distinct sectors: Tech, Finance, and Healthcare.
  - Participants were selected via purposive sampling to ensure a diverse demographic spread among the screeners themselves.

#### C. Data Collection Procedure

The data collection follows a "Blind Test" protocol:

- 1) Phase A: Algorithmic Screening:
  - The dataset was processed using three distinct algorithms: Logistic Regression (baseline), Random Forest (complex non-linear), and Gradient Boosting (high performance).
  - The models were trained on a historical subset (80%) and tested on the remaining 20%.
- 2) Phase B: Human Screening:
  - The human participants were given a random subset of 100 resumes each from the test set.
  - They were asked to rate candidates on a 5-point Likert scale (1=Reject, 5=Strong Hire) based on a provided Job Description (JD).
  - **Note:** To prevent Hawthorne Effects (altering behavior because they are being watched), recruiters were told the study focused on "efficiency" rather than "bias."

#### D. Data Analysis Techniques

This study utilizes the Disparate Impact Ratio (DIR) as the primary metric for fairness, compliant with the US Equal Employment Opportunity Commission (EEOC) "Four-Fifths Rule."

##### 1) Quantitative Analysis (Fairness Metrics):

We calculate the Selection Rate (SR) for different demographic groups. The Disparate Impact Ratio is calculated as:

- If  $DIR < 0.80$ , adverse impact (bias) is indicated.
- Precision and Recall metrics were also calculated to measure the trade-off between fairness and accuracy.

## 2) Qualitative Analysis:

- Post-task interviews were transcribed and analyzed using Thematic Analysis.
- Coding focused on identifying themes such as "School Prestige Bias," "Name Fluency Bias," and "Gut Feeling" justifications.

## E. Ethical Considerations

- 1) Data Privacy: All resumes were stripped of Personally Identifiable Information (PII) such as phone numbers and addresses, replaced with synthetic identifiers.
- 2) Informed Consent: Human participants provided written consent. While the specific focus on bias was initially withheld to ensure natural behavior, a full debriefing was conducted immediately after data collection.
- 3) Transparency: The specific hyperparameters of the AI models used are documented in the appendix to ensure reproducibility.

## IV. RESULTS AND FINDINGS

This section presents the empirical data derived from the comparative analysis of the 5,000-resume dataset. The results distinguish between Human Screeners, Baseline AI (unsupervised learning on historical data), and Mitigated AI (algorithms with fairness constraints applied).

### A. Comparative Analysis: Efficiency vs. Outcome

The initial comparison highlights a stark contrast in processing capabilities and selection variation.

Table 1: Efficiency and Selection Rates

Metric	Human Screeners (Avg)	Baseline AI (Unmitigated)	Mitigated AI (De-biased)
Avg. Time per Resume	3.4 minutes	< 0.01 seconds	< 0.01 seconds
Selection Rate (Total)	18%	14%	19%
Consistency Score	62% (Moderate)	100% (Perfect)	100% (Perfect)

- 1) Finding: While AI offered exponential efficiency gains, the Baseline AI was more aggressive in rejection (14% selection rate) compared to humans (18%).
- 2) Consistency: When shown the exact same resume twice (blinded), human recruiters gave different ratings 38% of the time (attributed to fatigue or mood), whereas the AI remained perfectly consistent.

### B. Identification of Bias Patterns

The study isolated specific demographic variables to measure Disparate Impact.

#### 1) Gender Bias

- Baseline AI: Exhibited severe bias against female candidates for technical roles. The algorithm negatively weighted terms like "Women's College" or "Softball," resulting in a selection rate of only 8% for women vs 20% for men.
- Human Screeners: Showed moderate bias, with a selection rate of 15% for women.
- Statistical Significance: The difference in the Baseline AI outcomes was statistically significant ( $p < 0.01$ ), confirming the model "learned" sexism from historical hiring data.

#### 2) Affinity and Prestige Bias (Education)

- Human Screeners: Exhibited strong "Prestige Bias." Candidates from Ivy League institutions were 3x more likely to be selected by humans than candidates from state universities with identical skills and GPAs.
- AI Models: Both Baseline and Mitigated AI showed negligible bias regarding university prestige, focusing instead on skill keywords (e.g., Python, SQL, Project Management).

### C. Evaluation of Fairness Measures (The "Four-Fifths Rule")

To quantify fairness, we calculated the Disparate Impact Ratio (DIR). A score below 0.80 indicates adverse impact (non-compliance with EEOC standards).

Figure 2: Disparate Impact Ratio (DIR) for Female Candidates

- Human Screeners: DIR = 0.75 (Slightly Biased - Below Compliant Standard)
- Baseline AI: DIR = 0.40 (Heavily Biased - Non-Compliant)
- Mitigated AI: DIR = 0.92 (High Fairness - Compliant)

#### Impact of Mitigation Techniques:

By applying Adversarial De-biasing (penalizing the model when it used gender proxies) and Re-weighting the training data, the Mitigated AI model successfully corrected the historical bias.

Critical Finding: While the Mitigated AI achieved the highest fairness score ( 0.92 ), it experienced a slight drop in "accuracy" (3%) when compared to historical hiring decisions. This suggests a trade-off: the model became "less accurate" at predicting past hiring patterns because past patterns were discriminatory.

#### D. Summary of Key Outcomes

- 1) Humans are prone to Affinity Bias (over-valuing school prestige/shared backgrounds).
- 2) Unchecked AI amplifies Historical Bias (replicating past gender discrimination).
- 3) Audited AI (Mitigated) outperforms both humans and baseline algorithms in establishing equity, provided that specific de-biasing interventions are applied during the training phase.

Here is the Discussion section. This is the critical part of your thesis where you synthesize the raw data into meaningful arguments and connect your specific study to the broader world of HR and AI ethics.

## V. DISCUSSION

### A. Interpretation of Key Findings

The results of this study reveal a complex dichotomy between human subjectivity and algorithmic objectivity.

- 1) The "Accuracy" Paradox: The most significant finding is the trade-off observed in the Mitigated AI model. While fairness metrics improved (DIR rose from 0.40 to 0.92), the model's "accuracy" in predicting historical hiring outcomes dropped. This interpretation suggests that historical accuracy is a flawed metric. If a company has historically discriminated against women, an "accurate" model will successfully replicate that discrimination. Therefore, a drop in accuracy—when accompanied by a rise in fairness—indicates the model is successfully unlearning systemic bias, rather than failing.
- 2) Human Cognitive Shortcuts: The data confirms that human recruiters rely heavily on heuristics. The strong correlation between "Elite Universities" and human selection rates (regardless of actual skill matching) points to social capital bias—a nuance the AI ignored in favor of keyword matching.

### B. Comparison with Prior Research

These findings align with and expand upon existing literature:

- 1) Echoing Amazon's Failure: The Baseline AI's initial rejection of female-coded terms mirrors the infamous Amazon (2018) case study, confirming that without intervention, ML models default to misogyny if trained on male-dominated datasets.
- 2) Contrasting Bertrand & Mullainathan: While previous studies (Bertrand & Mullainathan) proved humans discriminate based on names, this study adds a new layer: AI does not inherently care about names unless explicitly taught to. The Mitigated AI showed zero variance based on name ethnicity, suggesting that algorithmic screening is a viable solution to the "resume whitening" phenomenon required for minority candidates to pass human screens.

### C. Implications for Recruitment Practices

The study suggests a move away from "automation" (replacing humans) toward "Augmented Intelligence" (supporting humans).

- 1) The Hybrid Model: The ideal recruitment funnel utilizes Mitigated AI for the top-of-funnel screening (to ensure objective skill matching and anonymity) and Human Recruiters for the bottom-of-funnel evaluation (to assess soft skills and culture add), but only after the humans have been blinded to demographic data.
- 2) Strategic Shift: HR departments must transition from viewing AI as a time-saver to viewing it as a compliance tool. If properly audited, AI can serve as a "fairness shield," flagging talented candidates that human bias would otherwise overlook.

#### D. Role of Algorithmic Transparency and Human Oversight

The superior performance of the Mitigated AI was not automatic; it required deliberate intervention.

- 1) The Black Box Danger: If recruiters cannot see why an AI rejected a candidate (lack of explainability), they cannot trust the system. This study emphasizes the need for Explainable AI (XAI) features that highlight specific "green flag" skills or "red flag" gaps.
- 2) Human-in-the-Loop (HITL): The study advocates for a HITL workflow where humans audit a random sample of AI rejections weekly. This creates a feedback loop: humans correct the AI's edge cases, and the AI keeps the humans honest by stripping away demographic noise.

## VI. RECOMMENDATIONS

Based on the empirical findings that Mitigated AI outperforms both humans and baseline algorithms when properly managed, this study proposes the following roadmap for ethical AI deployment in HR.

#### A. Best Practices for Designing Bias-Resistant AI

To prevent the "Garbage In, Garbage Out" phenomenon where AI learns from discriminatory historical data, developers and HR architects must adopt a "Fairness-First" development lifecycle.

- 1) Data Hygiene & Segmentation:
  - Scrubbing: Remove not only explicit identifiers (Name, Gender) but also "proxy variables" found to correlate with bias (e.g., removing "lacrosse" or "polo" if they skew towards specific socioeconomic groups).
  - Synthetic Minority Oversampling: If the historical dataset has only 10% female engineers, use techniques like SMOTE (Synthetic Minority Over-sampling Technique) to generate synthetic profiles. This ensures the AI trains on a balanced dataset (50/50) rather than a skewed reality.
- 2) Adversarial "Red Teaming":
  - Before deployment, companies should employ a "Red Team" of ethicists and data scientists whose sole job is to try and break the model—feeding it resumes that are identical except for the candidate's name or gender to test for consistency.
- 3) Explainability by Design:
  - Implement LIME (Local Interpretable Model-agnostic Explanations) dashboards. When an AI gives a candidate a score of 85/100, the dashboard must highlight exactly which words contributed to that score (e.g., +5 for "Python", +10 for "MBA").

#### B. Framework for Human-AI Collaboration: The "Sandwich Model"

This study advises against fully automated "Auto-Reject" systems. Instead, organizations should adopt a Human-in-the-Loop (HITL) framework, specifically the "Sandwich Model."

- 1) Top Layer (Human Strategy): Humans define the competencies and weightings (e.g., "We value SQL skills over University Prestige"). This sets the "ground truth" for the AI.
- 2) Middle Layer (AI Execution): The AI processes the high volume of incoming resumes, masking PII (Personally Identifiable Information) and ranking candidates based only on the defined competencies.
  - Constraint: The AI is allowed to "Shortlist" but never to "Permanently Reject" without a flag.
- 3) Bottom Layer (Human Audit):
  - The Safety Valve: Humans review the final shortlist for interviews.
  - The Audit Loop: Crucially, humans must review a random 10% sample of the rejected pile weekly. If a human finds a qualified candidate in the reject pile, this "error" is fed back into the AI to retrain the model (Active Learning).

#### C. Policy Suggestions for Ethical Deployment

To ensure compliance with emerging regulations (such as the EU AI Act and NYC Local Law 144), HR departments should enact the following internal policies:

- 1) The "Transparency Notice" Policy:
  - Candidates must be informed before they apply that an automated tool will be used to screen their application.



- Recommendation: Include a "opt-out" clause where a candidate can request a human review, though this may result in a longer processing time.
- 2) Annual Third-Party Bias Audits:
  - Self-regulation is insufficient. Companies must mandate an annual audit by an external algorithmic auditing firm to calculate the Disparate Impact Ratio (DIR). If the DIR falls below 0.80 for any protected group, the AI tool must be suspended immediately until recalibrated.
- 3) The "Kill Switch" Protocol:
  - Establish a clear protocol for deactivating the AI if "drift" is detected (i.e., if the model starts developing new biases over time as market trends change).

## VII. CONCLUSION

### A. Summary of Findings

This comparative study endeavored to disentangle the complex web of bias inherent in both human cognition and machine learning algorithms within the recruitment domain. The findings lead to three critical conclusions:

- 1) Humans and Machines Fail Differently: Human screeners exhibited distinct "Affinity Bias," favoring candidates with shared educational backgrounds or social signals, often at the expense of objective skill alignment. Conversely, Baseline AI models functioned as "historical mirrors," accurately replicating and amplifying the systemic gender and ethnic biases present in 10 years of training data.
- 2) The Mitigation Trade-off: The study demonstrated that algorithmic bias is not immutable. Through adversarial de-biasing and data re-weighting, the Mitigated AI model achieved a Disparate Impact Ratio (DIR) of 0.92, significantly outperforming human fairness levels (0.75). However, this required a deliberate trade-off where "historical accuracy" was sacrificed for "normative fairness."
- 3) Augmentation over Automation: The most effective recruitment model proved to be a hybrid approach. AI excels at high-volume, blind skill matching, while humans are necessary for nuance and ethical oversight—but only when the humans are blinded to demographic markers during the initial review.

### B. Limitations of the Study

While the results are robust, several limitations must be acknowledged:

- 1) Experimental Setting: The study used a controlled dataset (anonymized resumes). In real-world scenarios, "soft skills" and unstructured data (cover letters, portfolios) introduce variables that are harder to quantify and de-bias.
- 2) Static Models: The algorithms tested (Random Forest, etc.) were static. The study did not account for "model drift," where an AI system might re-learn bias over time as it ingests new data from human hiring decisions.
- 3) Scope of Screening: The research focused solely on the Resume Screening phase. It does not address biases that may re-emerge during the interview stage or salary negotiation.

### C. Future Research Directions

To build upon this foundation, future research should focus on:

- 1) Generative AI & LLMs: Investigating how Large Language Models (e.g., GPT-4) process resumes compared to traditional predictive models, specifically testing for "linguistic bias" in how cover letters are interpreted.
- 2) Cross-Industry Validation: Replicating this study across distinct sectors (e.g., Creative Arts vs. Manufacturing) to determine if algorithmic bias manifests differently based on industry demographics.
- 3) Real-Time Auditing: Developing frameworks for "continuous fairness monitoring" that can alert HR practitioners to biased patterns in real-time, rather than waiting for annual audits.

## REFERENCES

- [1] Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732.
- [2] Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4), 991–1013. <https://doi.org/10.1257/0002828042002561>
- [3] Bogen, M., & Rieke, A. (2018). Help wanted: An examination of hiring algorithms, equity, and bias. Upturn. <https://www.upturn.org/reports/2018/hiring-algorithms/>



- [4] Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- [5] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, 214–226. ACM.
- [6] Kim, P. T. (2017). Data-driven discrimination at work. William & Mary Law Review, 58(3), 857–936.
- [7] Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. Management Science, 65(7), 2966–2981. <https://doi.org/10.1287/mnsc.2018.3093>
- [8] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys, 54(6), 1–35.
- [9] Noble, S. U. (2018). Algorithms of oppression: How search engines reinforce racism. New York University Press.
- [10] O'Neil, C. (2016). Weapons of math destruction: How big data increases inequality and threatens democracy. Crown.
- [11] Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 469–481. ACM.
- [12] U.S. Equal Employment Opportunity Commission. (1978). Uniform guidelines on employee selection procedures. 29 C.F.R. § 1607.



10.22214/IJRASET



45.98



IMPACT FACTOR:  
7.129



IMPACT FACTOR:  
7.429



# INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24\*7 Support on Whatsapp)