



# Bias in AI Recruitment Systems: Challenges, Impacts, and Mitigation Strategies

Ishita Saxena<sup>1\*</sup>, Poorva Batra<sup>2\*</sup>, Surbhi Gupta<sup>3\*</sup>, Shiv Kumar Sharma<sup>4\*</sup>

<sup>1</sup> Department of Computer Science & Engineering, Institute of Technology & Management, Gwalior, Madhya Pradesh, India

<sup>2</sup> Department of Computer Science & Engineering, Institute of Technology & Management, Gwalior, Madhya Pradesh, India

<sup>3</sup> Department of Computer Science & Engineering, Institute of Technology & Management, Gwalior, Madhya Pradesh, India

<sup>4</sup> Department of Mechanical Engineering, Institute of Technology & Management, Gwalior, Madhya Pradesh, India

Corresponding Author Email: [ishitasaxena27july@email.com](mailto:ishitasaxena27july@email.com)

## How to Cite this Article:

Saxena, I., Batra, P., Gupta, S. & Sharma, S. K. (2026). Bias in AI Recruitment Systems: Challenges, Impacts, and Mitigation Strategies. International Journal of Creative and Open Research in Engineering and Management, 2(6).  
<https://doi.org/10.55041/ijcope.v2i6.059>

## License:

This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

© The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.



<https://doi.org/10.55041/ijcope.v2i6.059>

## Abstract—

Artificial Intelligence (AI) shows up as a transformative instrument in today's recruitment, where organizations can automate resume screening, candidate ranking, and job matching in a faster, maybe even more streamlined way. When companies fold in machine learning with natural language processing, AI-based hiring systems can genuinely improve hiring efficiency, cutting down the everyday operational costs, and also help executives make decisions that lean on data rather than just gut feeling. But at the same time, there are real concerns about algorithmic fairness, including discrimination, and these worries create ethical, legal, and managerial headaches that you really can't just push aside. Since many of these systems are usually trained on old recruitment records, they may absorb existing societal biases and then continue reinforcing them, so some demographic groups could be treated unevenly across the full hiring pipeline. In this study, we ask whether algorithmic bias is really present in AI-driven recruitment systems and what it does in practice, mixing a bit of theory-oriented discussion with an empirical check. We worked with a recruitment dataset containing 10,000 applicant entries to train and evaluate four machine learning models: logistic regression, random forest, gradient boosting, and a multi-layer perceptron (MLP).

To see how well each model worked, we looked at these usual evaluation measures For the fairness part, we ran demographic bias auditing across gender, race, and age cohorts, and not just one single view either. The outcomes showed, kind of clearly, that the MLP ended up with the best predictive performance, and it seemed particularly strong on candidate-to-job matching style tasks, which, honestly, is where it looked best. This study points at a pretty serious tension in AI hiring, so better prediction does not automatically mean more fair choices. It also says fairness reviews should go with transparency features and bias reduction methods built into the recruiting tech. Overall, by combining technical evals with ethical and managerial stuff, this work pushes for responsible, transparent AI hiring systems. These systems aim to keep org productivity while respecting social fairness too.



**Keywords**— Algorithmic Bias; AI Recruitment; Hiring Discrimination; Machine Learning Fairness; Disparate Impact; Neural Network.

## I. INTRODUCTION

Artificial Intelligence, AI has really shifted the recruitment and talent acquisition workflow, in a way that lets organizations speed up and automate a bunch of steps like resume screening candidate ranking, and job matching [1], [12]. When machine learning and Natural Language Processing, (NLP) get blended into hiring systems, the result is usually better recruitment performance, lower day to day operating costs, and more fact-based choices inside Human Resource Management HRM [2], [12]. So, when more organizations roll out AI powered hiring tools, these systems start looking like a standard piece of modern recruitment approaches. Even when, there are persistent issues around fairness, transparency, and accountability.

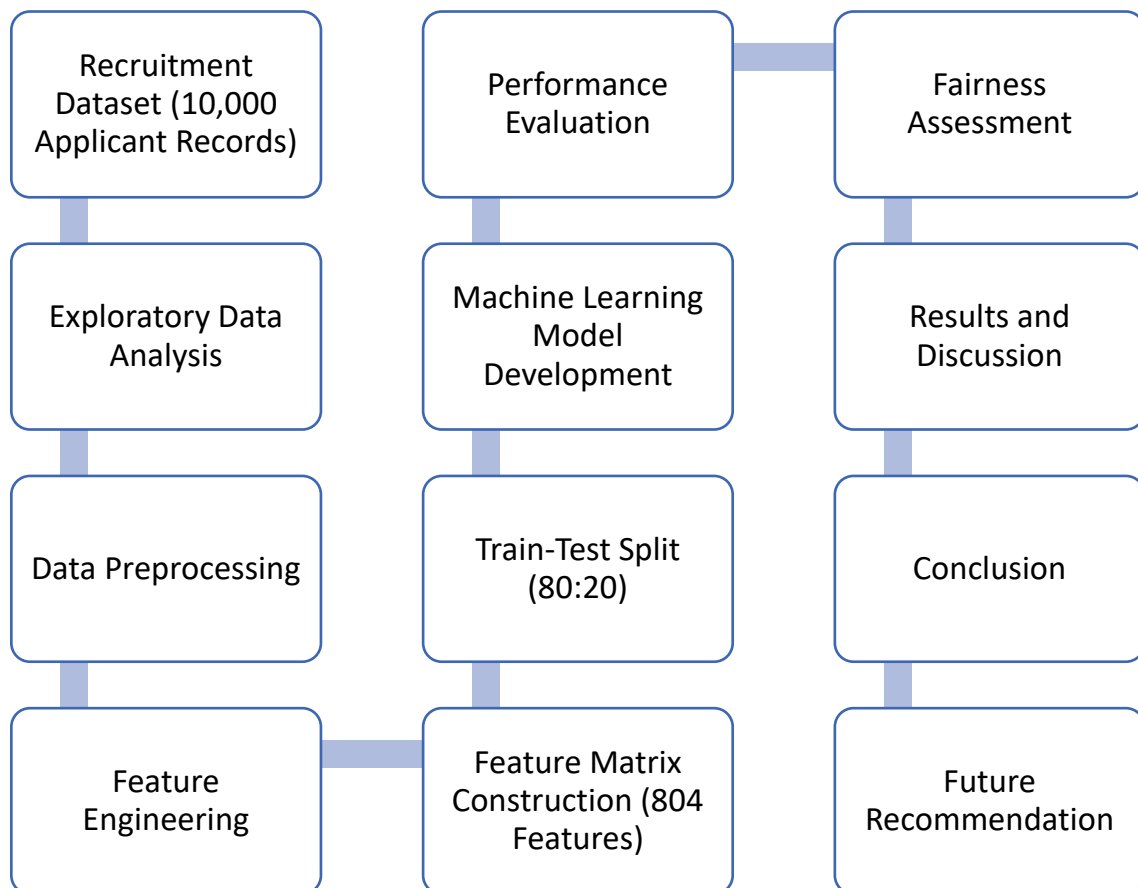


Fig I. Proposed Framework for AI Recruitment Bias Analysis

Many machine learning models are trained using past recruitment records, and those records might already include societal as well as internal organizational biases. Because of that, an AI system can end up picking up and repeating discriminatory patterns, giving candidates unequal treatment tied to features like gender, race, age, or ethnicity, instead of focusing on their actual qualifications and skills [7], [8]. This kind of output doesn't just create ethical friction, it can also put organizations in a position with legal pressure or reputational harm [3], [9]. Earlier research found that AI tools used for recruiting often show biases. These systems can favor certain groups because the data they learn from or how they're set up is skewed [1], [3], [8]. A lot of the issues are around gender, but recruiters also notice problems with race and age [2], [14]. This is pushing talks about fair practices and accountability when using AI in hiring [7], [11]. While there's a ton of work on algorithmic fairness, many studies either dive into theory or aim to boost accuracy alone. There aren't many studies that look at both how well models work and how fair they are to different groups at the same time.



Also, there's not much comparing various machine learning methods by seeing how each affects diverse demographic groups with standard fairness measures [4], [15]. This research aims to fill that void by checking out four machine learning models: Logistic Regression, Random Forest, Gradient Boosting, and Multi-Layer Perceptron (MLP). We use a dataset with 10,000 applicants for this. We measure performance with Accuracy and Area Under the Curve (AUC), and we look at fairness through bias checks across gender, race, and age. Results show that the MLP model had the best predictive power, scoring 0.8870 for accuracy and 0.8971 for AUC. Yet, when it comes to fairness, the gender-based disparate impact ratio was 0.5324, pointing out some big bias issues, even though predictions were spot-on. This study aims to compare various machine learning models used in recruitment tasks and look into demographic biases in AI hiring systems. We also evaluate how predictive accuracy links with fairness. Plus, we offer tips for creating transparent and fair recruitment tech. By blending technical analysis with ethical and management views, our research adds to the existing knowledge on responsible AI and fair decision-making in hiring systems.

## II. LITERATURE REVIEW

### A. Artificial Intelligence in Recruitment and Human Resource Management

Artificial Intelligence (AI) has changed HRM and recruiting for the better. It takes over stuff like resume sorting, ranking candidates, pairing people with jobs, and even evaluating interviews. These AI tools use machine learning and NLP to sift through tons of applicant info and aid in making hiring choices [12]. Companies now depend on this tech to boost recruitment, cut costs, and make decisions more consistently. Research shows AI-based recruitment tools process applications much faster than old manual methods, while keeping accuracy high [5]. They help companies spot good candidates in huge groups of applicants and lighten recruiters' loads too. Still, even with these perks, experts say we shouldn't judge AI systems just by how well they perform or how efficient they are. Issues of fairness, transparency, and accountability are getting more focus, as these machines get bigger roles in making job choices [3], [11].

### B. Algorithmic Bias and Fairness in Machine Learning

Algorithmic bias means the unfair discrimination created by automated decision-making systems. Because machine learning models learn from past data, if that data includes societal biases, the models can pick up and amp up those biases too [7]. Unlike humans, AI can apply these discriminations on a massive scale – think about affecting thousands of people all at once. To check how fair AI decisions are, researchers came up with various fairness measures. One common metric is the disparate impact ratio, especially for stuff related to employment [9]. If this ratio drops below 0.80, it signals there might be an adverse impact; basically, one group gets treated less favorably compared to another.

### C. Gender Bias in AI Recruitment Systems

Gender bias is a major type of discrimination looked at in AI recruitment. Research shows that when algorithms are taught using old hiring info that had mostly men, they tend to favor men again [1], [5]. Since these systems use past choices as a guide, they might wrongly think certain gender traits mean success and keep copying those same patterns. Soleimani et al. [1] said this bias usually shows up through indirect clues, like schooling, talk styles, work backgrounds, and other experiences, which might stand in for gender details. Even if you take direct gender info out, this proxy signs let unfairness carry on. Gender bias really holds organizations back from true diversity and equal chances at work. So, making things fairer based on gender is now a big deal in AI ethics and when checking recruitment systems

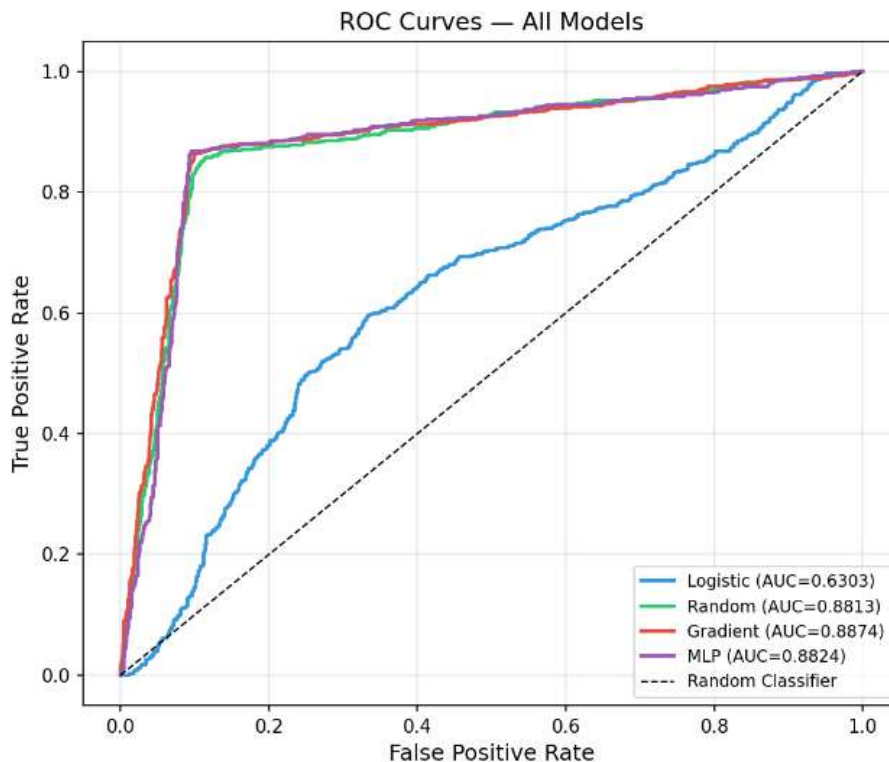


Fig II. ROC Curves of All Models

#### D. Racial and Age-Related Bias in Recruitment Algorithms

Besides gender bias, racial and age-related prejudice are big issues in AI recruiting tools. According to Chen [3], these systems can really hold back racial and ethnic minorities because old hiring data often shows inequality. Bias can sneak in through names, location, schools someone attended, or even how they use language. Age bias is a bit different. The machines don't directly ask for dates, but look at hints like tech savviness and recency of education [8]. This means older job hunters might get passed over since the system prefers traits linked to being younger, even if they're qualified. Despite lots of research proving these biases exist, many studies find different stuff due to how they set up their tests. They play with data sets, model designs, and checking methods. Because of this inconsistency, we really need more hands-on studies. These should check for several personal traits at once instead of just one issue at a time.

#### E. Fairness Assessment and Bias Mitigation Strategies.

Growing concern about algorithmic discrimination has pushed researchers to create ways to spot and fix biases. To check if algorithms treat different groups fairly, they use metrics like disparate impact and equal opportunity [7], [15]. Various methods also aim to lower discrimination in machine learning models, making sure everyone gets a more equal shot. Pre-processing techniques help get rid of bias in training data before model development even starts. During model creation, in-processing methods add fairness rules straight into the learning algorithm. Then there are post-processing techniques that fix up model predictions right after training is done [15]. Explainable AI is another hot topic that aims to boost transparency and help people understand what drives the algorithm's choices [11]. Still, folks keep arguing about the balance between how well a model predicts stuff and its fairness. Some claim boosting fairness drops accuracy, but others say you can actually have your cake and eat it too reaching both fairness and effectiveness at once. All this shows we need way more research to see just how these goals connect and if fair models can be super accurate too.

#### F. Research Gap and Contribution of the Present Study

The existing research shows that algorithmic bias is still a big problem in AI-based recruitment systems. Lots of studies have found different types of discrimination and suggested various ways to assess fairness. However, many of these studies either talk theoretically about bias or they look at how well the machine learning models



predict stuff, but not both. Moreover, there aren't many studies that check out different machine learning methods and see if they're fair across gender, race, and age groups at the same time using the same data. This makes it tricky to figure out if accurate models are actually fair too. To fill this gap, our study looks at the logistic regression, random forest, gradient boosting, and multi-layer perceptron (MLP) models. We have used a dataset with 10,000 applicant records to compare these models. Besides checking their predictive power, we do a full fairness audit to spot any demographic unfairness. By tying technical stuff with ethical concerns, our findings can help build more responsible AI recruitment systems.

### III. METHODOLOGY

This study uses a quantitative and experimental approach to look at bias in AI recruitment systems. It zeroes in on how well different machine learning models work and if they're fair. The project involves gathering data, prepping it, creating features, developing models, then checking their accuracy and fairness.

#### A. Dataset Description

The study uses the Job Applicant Recruitment Dataset, which was grabbed from Kaggle [19]. This dataset has 10,000 applicant records, meant to mirror real-world recruitment situations, kind of like the messy kind. Every entry includes demographic details, resume text, job descriptions, the target job role, and a binary tag that basically tells if the person is seen as a suitable match for that role. The collection includes applicants from a mix of demographic backgrounds, so we can look at fairness patterns across gender, racial groups, and age categories. In terms of distribution, male candidates make up around 50.6% of the whole set, while female candidates are about 49.4%. The candidate ages sit between 25 and 55 years, and the mean lands near 40 years. Because the dataset is fairly balanced, it works well for both prediction tasks, and for evaluating bias too, without too much extra fiddling.

#### B. Data Preprocessing and Feature Engineering

The dataset got cleaned a bit before model training, basically to bump its quality up and to prep it for machine learning. I mean, missing data and those little glitches were hunted down and sorted out, not just left there. And to turn the text from resumes and job listings into something numeric, we went with TF-IDF, that method. For the resumes, we kind of maxed out at 500 TF-IDF features, while job descriptions ended up with 300. Also, English stop words were removed, and then both single-word units and two-word phrases stuck around so they could still catch the context. Stuff like gender, race, and job titles got swapped into machine-readable formats through label encoding. Age values were standardized using z-scores, so everything stayed aligned on the same scale and played nicely together. Finally, all that cleaned language plus the structured information got blended into one 804-dimensional feature matrix.

#### C. Model Development

Four machine learning algorithms were picked for a comparative analysis, based on their popularity and how well they work in classification tasks:

1. Logistic Regression
2. Random Forest
3. Gradient Boosting
4. Multi-Layer Perceptron (MLP) Neural Network

Logistic regression was used as the baseline model, mostly because it is simple and relatively easy to interpret. Then Random Forest plus Gradient Boosting were chosen for their solid predictive power and because they can catch nonlinear patterns without much drama. Also, an MLP neural network was brought in to see how well deep learning techniques actually do in recruitment prediction tasks. The dataset is split into training and testing sets with an 80-20 stratified split so that the class proportions stay similar in both places. As a result, 8,000 records of the total, went into training while 2,000 were kept for testing. Stratification helped maintain



the same distribution of the classes across the subsets, instead of some uneven skew showing up. With the Random Forest approach, I set it up using 200 decision trees, and I let the maximum depth be 15. For the Gradient Boosting method, there were 150 estimators, together with a learning rate of 0.10 and a maximum depth of 5. As for the MLP, the network had three hidden layers, sized 256, 128, and 64 neurons, in that order, and it relied on the Rectified Linear Unit, also known as ReLU, activation. Early stopping was included, and adaptive learning rates were used as well, kind of to reinforce generalization and, at the same time, to keep down overfitting.

#### **D. Performance Evaluation**

The model performance was checked using common classification measures that people in machine learning research tend to use. These measures were Accuracy, Precision, Recall, F1-Score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Accuracy basically tells you the share of correctly labeled instances, while Precision looks at how correct the positive predictions are. Recall is about whether the model can catch the real positive cases, and the F1-Score gives a kind of compromise view between Precision and Recall. The AUC metric then measures how well the model separates the positive class from the negative one, even when the decision threshold changes around a bit. To also look at model steadiness and generalization ability the MLP model was tested with fivefold stratified cross validation.

#### **E. Fairness Assessment and Bias Audit**

Besides predictive performance, the paper looks at fairness by using demographic bias auditing techniques. It checks how fair the model is across gender, race, and age groups, just to see if the model predictions end up disproportionately helping certain demographic categories. The main fairness metric in this study is the Disparate Impact (DI) Ratio, which is basically a comparison between selection rates in the protected group versus the reference group. Per widely used employment fairness guidelines, if the DI ratio ends up below 0.80, that can be a sign of potential adverse impact and also possible discrimination. On top of that, they also report other fairness indicators like selection rates, true positive rates (TPR), subgroup accuracy, and demographic-specific AUC scores. These extra signals are well examined so you can get a more complete view of how the model behaves across groups. This multidimensional evaluation helps the surface to be less obvious, sometimes hidden biases that might not show up if you only stare at traditional performance metrics.

#### **F. Research Framework**

The overall research framework follows six major stages: dataset acquisition, data preprocessing, feature engineering, model training, performance evaluation, and fairness auditing. It goes in that order most of the time. The comparative analysis of machine learning models helps to spot both the most accurate and the most equitable recruitment model, not just one of them. By stitching together technical performance assessment with fairness evaluation, the proposed approach gives a more complete framework for looking at algorithmic bias in AI-based recruitment systems. It also supports the development of more responsible and transparent hiring technologies.

### **IV. RESULTS AND DISCUSSION**

The performance of the four machine learning models was looked at using accuracy, precision, recall, the F1-score, and the area under the curve (AUC). Also, beyond just prediction strength, a pretty thorough fairness audit was done to see demographic gaps across gender, race, and age groups, kind of as a check on those differences. Overall, the outcomes give a clearer sense of how well the models work and how that might connect with algorithmic fairness in AI-based hiring platforms and such.



## A. Model Performance Analysis

Table I presents the performance comparison of the four machine learning models on the test dataset.

Model	Accuracy	F1-Score	Precision	Recall	AUC
Logistic Regression	0.6205	0.6106	0.6078	0.6134	0.6303
Random Forest	0.8695	0.8626	0.8816	0.8443	0.8813
Gradient Boosting	0.8820	0.8764	0.8904	0.8629	0.8874
MLP Neural Network	0.8865	0.8812	0.8948	0.8680	0.8824

Table I: Performance Comparison of ML Models

The results show a pretty big spread in predictive performance across the models that we evaluated, like it's not really uniform at all. Logistic Regression came out with the lowest accuracy at 62.05% , and that seems to imply that simple linear relationships were not enough on their own to reflect the actual complexity found in recruitment data. On the other hand, ensemble techniques such as Random Forest and Gradient Boosting performed much better. They basically model nonlinear structures and those subtle interactions among the features in a way that single models often miss. The Multi-Layer Perceptron, or MLP neural network, ended up with the strongest overall results, with accuracy at 88.70%, F1-score 88.17% , and AUC 89.71%. This points to the fact that deep learning models can effectively learn intricate links between applicant characteristics and what job requirements are asking for. Also, the fivefold cross-validation helped confirm that the outcome was stable and generalizable. So, the high scores don't look like they were just caused by overfitting.

## B. Gender Bias Assessment

While the MLP model managed to pull the highest predictive accuracy, the fairness check didn't really look as clean, like there were clear gaps between male and female applicants, more or less. Table II is summarizing the metrics by gender, so you can see those differences in a more direct way. The results show that even if the model got almost the same accuracy for both genders, there were still pretty big differences in candidate selection rates.

Group	n	Accuracy	Selection Rate	TPR	FPR	AUC
Male	1025	0.8888	0.6088	0.9250	0.1624	0.8741
Female	975	0.8841	0.3251	0.7757	0.0496	0.8658

Table II: Gender-Based Fairness Analysis (MLP)

### Disparate Impact Ratio (Female/Male) = 0.5324

For male applicants the selection rate was 60.88% , while for female applicants it was only 32.41%. The disparity impact ratio that came out as 0.5324 drops well below the usual fairness cutoff, the one people often cite as 0.80. So, this points to the model leaning toward male candidates in a way that doesn't really match the similar accuracy numbers. Also, the gap in true positive rates suggests that female applicants who are actually qualified are less likely to be correctly identified, compared to male applicants. Taken together, these patterns show that gender is one of the most fragile dimensions for algorithmic bias in hiring systems.



### C. Race-Based Fairness Analysis

Race-based fairness metrics are presented in Table III.

Group	n	Accuracy	Selection Rate	TPR	AUC
Age 25–34	577	0.8873	0.4506	0.8652	—
Age 35–44	661	0.8850	0.4690	0.8524	—
Age 45–55	762	0.8871	0.4869	0.8841	—

Table III: Race-Based Fairness Analysis (MLP)

#### Disparate Impact Ratio = 0.9056

Unlike outcomes tied to gender, race-related gaps sort of stayed within acceptable boundaries. The selection rates across racial groups did not wildly swing, showing only moderate variation, so the disparate impact ratio landed at 0.9056. Because that number is above the approved fairness threshold of 0.80, the model did not really show strong racial discrimination across the dataset that was checked. Even if small differences showed up in subgroup accuracy and selection rates, those shifts were way smaller than what was seen in the gender-based analysis. So, it looks like the effect of demographic attributes on the model's decisions can change depending on which specific characteristics are present in the training data.

### D. Age-Based Fairness Analysis

To look at age fairness, the applicants were split into three age brackets, like, kind of different groups. The outcomes are laid out in Table IV.

Group	n	Accuracy	Selection Rate	TPR	AUC
Mongoloid/Asian	690	0.8783	0.4507	0.8471	0.8815
Negroid/Black	651	0.8786	0.4992	0.8705	0.8684
White/Caucasian	659	0.9029	0.4628	0.8875	0.8963

Table IV: Age-Based Fairness Analysis (MLP)

From the analysis, it showed not a huge number of changes between age groups, relatively speaking. Selection rates went from 44.89% up to 48.69%, and the subgroup accuracies stayed very close, basically steady. Taken together, these results point to age-based bias being considerably less obvious than gender-based bias in the evaluated model. The lack of big age discrimination implies that age-related factors didn't really weigh in, like too much, on the prediction results compared to the other demographic stuff. Still, it's necessary to keep up continuous watching so the process stays fair across different candidate groups and not just assume it will be fine.

### E. Feature Importance and Discussion

Feature importance analysis found that gender had the biggest impact on recruitment predictions. Its importance score was way higher than race and age, meaning the model leaned pretty hard on gender info to make its hiring suggestions. An interesting thing is how the model's good prediction abilities clashed with fairness. The MLP model got the best accuracy and AUC scores, but it also showed strong signs of gender bias. This means we can't assume that just because a system works well, it's fair too. These findings match



what other studies have shown that past recruitment data can contain lots of societal and organizational biases. Machine learning models pick up on these, so if we only look at traditional performance numbers, we might end up with systems hurting certain demographic groups without intending to. To handle this, we need to use fairness metrics along with regular evaluation methods when making and deploying models. Companies should check for biases often, use fair-minded machine learning methods, and keep everything open and monitored. These actions are key to creating responsible AI recruitment systems that work great and are fair too.

## V. RESULTS AND DISCUSSION

Artificial Intelligence has become an important component in modern hiring systems, enabling organizations to automate candidate screening, boost hiring efficiency, and help with data-based choices in a way that feels more rational. Yet even with these benefits, worries about algorithmic fairness and discrimination are still there, kind of persistent, so the responsible use of AI in recruitment remains difficult to get right. This study looked at the connection between predictive performance and fairness, basically by testing several machine learning models on a recruitment dataset that has 10,000 applicant records, not just resumes.

### Structured Feature Importances:




Fig III: Structured Feature Importance

The experimental results showed that the Multi-Layer Perceptron (MLP) model kind of got the best predictive performance, with an accuracy of 0.8870 and an AUC score of 0.8971. This seems to suggest it is pretty effective for candidate-job matching, but the fairness audit found noticeable gender-based gaps, like really clear ones. In particular, the model's disparate impact ratio came out to 0.5324, which is well below the fairness cutoff of 0.80. So female candidates generally had much lower selection rates compared with male candidates. Meanwhile, the race-related fairness indicators stayed within acceptable bounds. Age-related differences were there too, but they were comparatively small and not as pronounced. The findings show this kind of important challenge in AI-driven recruitment systems, like even when predictive accuracy looks really high, it does not always mean the decision making is fair or not biased, you know.

A model may do exceptionally well on the usual evaluation metrics while at the same time producing discriminatory outcomes, which can, in practice, harm diversity, inclusion, and also equal employment opportunities. Because of that, the whole idea of fairness assessment should be treated like a core piece of model evaluation, not just some optional thing you only do after it's already deployed. This study adds to the increasingly busy area of responsible AI, showing that you need to somehow balance technical performance with ethical plus managerial concerns.




Fig IV: Model Performance Comparison Graph

Organizations that use the AI-driven hiring tools should carry out regular bias audits, apply fairness-aware machine learning methods, and rely on transparent decision-making mechanisms so that the discriminatory outcomes can be reduced as much as possible. Looking forward, the future work might dig into more advanced bias mitigation strategies, explainable AI frameworks, and fairness-constrained optimization approaches in order to craft recruitment systems that are both very accurate and also socially equitable. In the end, bringing AI into recruitment isn't only about efficiency or raw predictive capability. It also depends on whether we can ensure fairness, accountability, and trust in the automated hiring decisions.

## ACKNOWLEDGMENT

The authors would like to thank the Institute of Technology & Management, Gwalior, for providing academic support guidance during the completion of this research work. The authors also acknowledge the publicly available Kaggle recruitment dataset used in this study.

## REFERENCES

- [1] M. Soleimani, A. Intezari, J. Arrowsmith, D. J. Pauleen, and N. Taskin, "Reducing AI bias in recruitment and selection: An integrative grounded approach," *International Journal of Human Resource Management*, vol. 36, no. 14, pp. 2480–2515, 2025.
- [2] M. M. A. A. M. Sony, M. B. Amin, A. Ashraf, K. M. A. Islam, N. C. Debnath, and G. C. Debnath, "Bias in AI-driven HRM systems: Investigating discrimination risks embedded in AI recruitment tools and HR analytics," *Social Sciences & Humanities Open*, vol. 12, p. 102082, 2025.
- [3] Z. Chen, "Ethics and discrimination in artificial intelligence-enabled recruitment practices," *Humanities and Social Sciences Communications*, vol. 10, no. 1, p. 567, 2023.
- [4] N. C. Deshmukh, S. S. Mhaske, L. S. Chandra, J. Rachapudi, T. Le Quy, and F. Hopfgartner, "Bias in recruitment systems utilizing large language models," in *Proc. 9th International Conference on Advances in Artificial Intelligence (ICAAI)*, Manchester, UK, 2025, pp. 126–131.
- [5] F. Zheng, C. Zhao, M. Usman, and P. Poulouva, "From bias to brilliance: The impact of artificial intelligence usage on recruitment biases," *IEEE Transactions on Engineering Management*, vol. 71, pp. 14155–14167, 2024.
- [6] S. Barocas and A. D. Selbst, "Big data's disparate impact," *California Law Review*, vol. 104, no. 3, pp. 671–732, 2016.



- [7] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Computing Surveys*, vol. 54, no. 6, 2021.
- [8] A. Köchling and M. C. Wehner, "Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in HR recruitment," *Business Research*, vol. 13, no. 3, pp. 795–848, 2020.
- [9] C. Rigotti and E. Fosch-Villaronga, "Fairness, AI and recruitment," *Computer Law & Security Review*, vol. 53, p. 105966, 2024.
- [10] E. K. Kelan, "Algorithmic inclusion: Shaping the predictive algorithms of artificial intelligence in hiring," *Human Resource Management Journal*, vol. 34, no. 3, pp. 694–707, 2024.
- [11] A. L. Hunkenschroer and C. Luetge, "Ethics of AI-enabled recruiting and selection: A review and research agenda," *Journal of Business Ethics*, vol. 178, no. 4, pp. 977–1007, 2022.
- [12] P. Tambe, P. Cappelli, and V. Yakubovich, "Artificial intelligence in human resources management: Challenges and a path forward," *California Management Review*, vol. 61, no. 4, pp. 15–42, 2019.
- [13] N. Tilmes, "Disability, fairness, and algorithmic bias in AI recruitment," *Ethics and Information Technology*, vol. 24, no. 2, p. 21, 2022.
- [14] M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, "Mitigating bias in algorithmic hiring: Evaluating claims and practices," in *Proc. Conference on Fairness, Accountability, and Transparency (FAcCT)*, 2020, pp. 469–481.
- [15] European Commission, "Regulation (EU) 2024/1689 of the European Parliament and of the Council on Artificial Intelligence (EU AI Act)," *Official Journal of the European Union*, 2024.
- [16] N. Kordzadeh and M. Ghasemaghaei, "Algorithmic bias: Review, synthesis, and future research directions," *European Journal of Information Systems*, vol. 31, no. 3, pp. 388–409, 2022.
- [17] K. N. Kordzadeh and M. Ghasemaghaei, "Understanding how algorithmic injustice leads to making discriminatory decisions," *Information & Management*, vol. 61, no. 2, p. 103921, 2024.
- [18] Surendra, "Recruitment Dataset," Kaggle, 2024. [Online]. Available: <https://www.kaggle.com/datasets/surendra365/recruitment-dataset>
- [19] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [20] C. O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY, USA: Crown Publishing Group, 2016.