Comprehensive Journal of Science
Volume (11), Issue (41), (JUNE, 2026)
ISSN: 3014-6266

مجلة العلوم الشاملة
المجلد (11) العدد 41 يونيو 2026
ردمد: 3014-6266

Predicting Hiring Decisions Using Machine Learning: A Transparent and Fair Approach for Human Resource Management

Farhat M. A. Zargoun [1*], Ramzi Hamid Ganoni[2] and Abobaker Zargoun[3]
[1] Web Technology Department, Faculty of Information Tech., University of Tripoli, Tripoli, Libya
[2] Web Technology Department, Faculty of Information Tech., University of Tripoli, Tripoli, Libya
[3] Control Engineering, Faculty of Electronic Technology, Bani Waleed, Libya
*Corresponding author: f>zargoun@uot.edu.ly

التنبؤ بقرارات التوظيف باستخدام التعلم الآلي: أسلوب يتسم بالشفافية والعدالة لإدارة الموارد البشرية

فرحات امحمد زرقون[1*]، رمزي حميد القانوني [2] ، أبوبكربالحسن .زرقون[3]
[1] قسم تقنيات الانترنت، كلية تقنية المعلومات، جامعة طرابلس، طرابلس، ليبيا
[2] قسم تقنيات الانترنت، كلية تقنية المعلومات، جامعة طرابلس، طرابلس، ليبيا
[3] قسم هندسة التحكم ، كلية التقنية الالكترونية ، بني وليد ، ليبيا

تاريخ الاستلام:2026/05/05 –تاريخ المراجعة:2026/05/28– تاريخ القبول: 2026/06/07– تاريخ للنشر : 2026 /06/21

Abstract: (Times New Roman: size – 10) Abstract must be written in English within 300 words.
This study explores how machine learning can help predict hiring decisions in Human Resource Management (HRM). We tested four different classification algorithms—Logistic Regression, Decision Tree, Random Forest, and XGBoost—using a recruitment dataset from Kaggle. Beyond just predicting who gets hired, our research also focused on making these AI systems understandable (Explainable Artificial Intelligence, or XAI) and fair. Our findings show that advanced machine learning models, particularly ensemble methods, performed better than simpler ones. Crucially, our work on explainability and fairness addresses key limitations found in earlier research, paving the way for more responsible AI in hiring.

Keywords: Human Resource Management, Recruitment, Hiring Decision Prediction, Machine Learning, Explainable Artificial Intelligence, Fairness Analysis, XGBoost, Human Resource Analytics.

الملخص:
تستكشف هذه الدراسة كيف يمكن للتعلم الآلي أن يساعد في التنبؤ بقرارات التوظيف في إدارة الموارد البشرية (HRM). حيث قمنا باختبار أربعة خوارزميات تصنيف مختلفة—وهي: الانحدار اللوجستي (Logistic Regression)، وشجرة القرار (Decision Tree)، والغابة العشوائية (Random Forest)، وXGBoost—باستخدام مجموعة بيانات توظيف من منصة كاجل (Kaggle). وإلى جانب مجرد التنبؤ بمن سيتم توظيفه، ركز بحثنا أيضًا على جعل أنظمة الذكاء الاصطناعي هذه قابلة للفهم (الذكاء الاصطناعي القابل للتفسير ، أو XAI) وعادلة. وتظهر نتائجنا أن نماذج التعلم الآلي المتقدمة، لا سيما أساليب التجميع (Ensemble Methods)، قدمت أداءً أفضل من النماذج الأكثر بساطة. والأهم من ذلك، أن عملنا على قابلية التفسير والعدالة يعالج قيودًا جوهرية كانت موجودة في الأبحاث السابقة، مما يمهد الطريق لاستخدام ذكاء اصطناعي أكثر مسؤولية في مجال التوظيف.
الكلمات المفتاحية:إدارة الموارد البشرية ، التنبؤ بقرارات التوظيف، تعلم الآلة، الذكاء الاصطناعي القابل للتفسير.

1. Introduction

Human Resource Management (HRM) is vital for any organization's success, primarily through its role in recruiting and selecting the right talent. However, as companies face an ever-growing flood of job applications, traditional hiring processes have become incredibly time-consuming and resource-intensive. Manually sifting through resumes and evaluating candidates often demands significant effort from recruiters and can be swayed by personal biases. This is why more and more organizations are turning to data-driven methods to streamline recruitment and make smarter hiring decisions.

The recent leaps in Artificial Intelligence (AI) and Machine Learning (ML) have opened up exciting new possibilities for automating candidate evaluations and forecasting hiring outcomes. By analyzing various candidate traits, qualifications, and assessment results, machine learning models can spot patterns linked to successful hires, thereby helping HR professionals pick the best candidates. Yet, despite these clear advantages, significant challenges remain, particularly concerning the transparency, interpretability, and fairness of AI-assisted recruitment systems [3].

While previous studies have shown how effective machine learning can be in recruitment and candidate screening, many of these approaches tend to focus solely on predictive accuracy. They often provide little insight into how hiring recommendations are actually generated. Furthermore, worries about algorithmic bias and fairness continue to spark important ethical discussions about using AI responsibly in HRM. This study, therefore, set out to answer a critical question: How effectively can machine learning models predict hiring decisions and assist HR professionals in selecting suitable candidates, all while ensuring transparency and fairness [4]

To tackle this, our study aimed to develop and evaluate machine learning models for hiring prediction. We compared Logistic Regression, Decision Tree, Random Forest, and XGBoost algorithms, identified the most influential factors in hiring decisions, enhanced transparency using Explainable Artificial Intelligence (XAI) techniques, and assessed the fairness of recruitment predictions [1][2].

Specifically, we sought answers to several key research questions: Can machine learning accurately predict hiring decisions? Which candidate attributes most strongly influence hiring outcomes? Which machine learning model performs best? How can Explainable AI improve transparency? And finally, can machine learning support fair and responsible recruitment practices?

Our main contribution is a comprehensive recruitment prediction framework that brings together predictive performance, explainability, and fairness assessment into a single decision-support system. By combining machine learning classification, SHAP-based explanations, and fairness evaluation, this framework aims to create more transparent and responsible AI-driven recruitment processes [5].

The rest of this paper is structured as follows: Section 2 reviews existing literature, Section 3 details our research methodology, Section 4 presents the experimental results, Section 5 discusses our findings, Section 6 highlights our contributions, and Section 7 concludes the study while outlining future research directions.

2. Literature Review

The rapid evolution of Artificial Intelligence (AI) and Machine Learning (ML) has profoundly reshaped Human Resource Management, especially in how companies recruit and select staff. Organizations are increasingly leaning on data-driven strategies to boost recruitment efficiency, cut hiring costs, and make more objective decisions. Machine learning algorithms are particularly adept at sifting through vast amounts of candidate information, identifying patterns that lead to successful hires, and ultimately helping recruiters pinpoint qualified individuals.

One of the most common ways machine learning is applied in recruitment is for candidate screening and predicting hiring outcomes. Traditional recruitment often involves manually reviewing resumes and conducting interviews, a process that can be both time-consuming and prone to human bias. Machine learning models offer a way to automate these tasks by analyzing qualifications, experience, and assessment results. Past research has consistently shown that using predictive analytics can make recruitment more efficient and support hiring decisions based on solid evidence.

However, despite the growing adoption of AI-powered recruitment systems, significant hurdles remain, particularly around transparency and accountability. Many machine learning models are often seen as ‘black-box’ systems, meaning HR professionals struggle to understand how hiring recommendations are generated. This lack of clarity makes it difficult to trust and effectively use these systems.

This is where **Explainable Artificial Intelligence (XAI)** comes in. XAI has emerged as a crucial research area focused on making AI systems more transparent and understandable. Techniques like SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) allow users to see how individual features contribute to a model's predictions, fostering more trustworthy decision-making. Recent studies underscore the importance of XAI in building recruiter trust and increasing the acceptance of AI-assisted hiring tools [6].

Another significant challenge in AI-based recruitment is **fairness**. Machine learning systems learn from historical data, which can unfortunately contain existing biases. As a result, recruitment algorithms might unintentionally perpetuate discriminatory patterns or lead to unequal outcomes across different groups of candidates. Recent research highlights the critical need for fairness assessment, bias detection, and ethical oversight in AI-driven recruitment. Experts increasingly recommend integrating fairness evaluations into recruitment frameworks to ensure hiring practices are both responsible and equitable.

2.1 Previous Studies

Several studies have paved the way for our current research:

- **Dadaboyev (2025)** conducted a comprehensive review of AI applications in recruitment, noting that while AI significantly boosts efficiency, concerns about transparency, fairness, and ethics remain major challenges.
- **Zhang et al. (2025)** specifically examined Explainable AI in recruitment, concluding that XAI techniques enhance transparency, traceability, and recruiter trust by providing clear explanations for hiring recommendations. They see XAI as a vital component for future AI-driven recruitment systems.
- Research into **fairness in AI-assisted hiring** has repeatedly shown that algorithms can inherit biases from historical hiring data. This has led to recommendations for fairness auditing, bias monitoring, and the use of fairness-aware machine learning approaches to reduce discriminatory outcomes and improve ethical accountability [10].
- Other studies have demonstrated that **ensemble learning algorithms**, particularly Random Forest and Gradient Boosting methods, often outperform traditional classification models in predicting recruitment outcomes. These models are powerful because they can uncover complex relationships between various recruitment-related factors [11].

2.2 Research Gap

Despite these advancements, there's a noticeable gap in the existing research. While many studies have successfully applied machine learning to recruitment, few have focused on developing hiring prediction models that are *both* interpretable and fair, providing transparent explanations while maintaining high predictive performance. Most research tends to emphasize predictive accuracy, explainability, or fairness in isolation.

This means there's a clear need for recruitment systems that can simultaneously achieve strong predictive performance, offer understandable explanations for hiring recommendations, and rigorously evaluate fairness across different candidate groups. Our study directly addresses this gap by developing a hiring decision prediction framework that combines machine learning classification, Explainable Artificial Intelligence techniques, and fairness assessment. By integrating these three crucial dimensions—predictive performance, transparency, and fairness—into a single framework, we aim to contribute to the development of more trustworthy and responsible AI-driven recruitment systems.

3. Methodology

Our study used a quantitative predictive analytics approach to investigate how effective machine learning techniques are at predicting hiring decisions. The methodology involved several key stages: preparing the dataset, preprocessing the data, developing the models, evaluating their performance, analyzing their explainability, and assessing fairness. The overarching goal was to create a recruitment prediction framework that balances predictive accuracy with transparency and fairness.

3.1 Dataset Description

For this study, we used the **"Predicting Hiring Decisions in Recruitment Data" dataset**, which we obtained from Kaggle. This dataset was specifically designed for recruitment analytics and tasks related to predicting hiring outcomes. It contains information on 1,500 candidates, with 11 variables covering demographic details, educational background, professional experience, and assessment results.

The dataset includes features such as: Age, Gender, Education Level, Experience Years, Previous Companies, Distance from Company, Interview Score, Skill Score, Personality Score, and Recruitment Strategy. The crucial

target variable, "Hiring Decision," simply tells us whether a candidate was hired or not. An initial look at the data (exploratory data analysis) confirmed its high quality, as it contained no missing values or duplicate records.

3.2 Data Processing

Before feeding the data into our machine learning models, we performed several data preprocessing steps to ensure it was clean and suitable for analysis. First, we thoroughly checked for any missing values, duplicate records, inconsistencies, or other potential quality issues. Our analysis confirmed that the dataset was indeed free of missing values and duplicate observations.

Next, we converted categorical variables (like "Gender" or "Education Level") into a numerical format that machine learning algorithms can understand. After this encoding, the dataset was ready for model training and evaluation. Finally, we split the data into two parts: 80% for training our machine learning models and 20% for testing how well they performed on unseen data.

3.3 Machine Learning Models

We selected and evaluated four supervised machine learning algorithms for this study:

- **Logistic Regression:** We chose Logistic Regression as a baseline model because it's simple, computationally efficient, and easy to interpret. This model estimates the probability of a candidate being hired based on various predictor variables.
- **Decision Tree:** A rule-based classification algorithm, the Decision Tree creates a hierarchical structure of decision nodes. Its ability to generate clear, interpretable rules makes it well-suited for supporting recruitment decisions.
- **Random Forest:** This is an ensemble learning algorithm that combines multiple decision trees. By doing so, it significantly improves predictive performance and helps reduce overfitting, making it a robust and powerful classification tool.
- **XGBoost (Extreme Gradient Boosting):** An advanced ensemble learning technique, XGBoost uses gradient boosting principles to optimize predictive performance. It's known for its excellent performance across many machine learning applications and was included due to its effectiveness in classification tasks.

3.4 Evaluation Metrics

To thoroughly assess the performance of our machine learning models, we used a variety of classification metrics:

- **Accuracy:** This measures the overall proportion of correctly classified observations.
- **Precision:** Precision tells us the proportion of correctly predicted positive instances among all instances predicted as positive.
- **Recall:** Also known as sensitivity, recall measures the model's ability to correctly identify all actual positive cases.
- **F1-Score:** This metric combines precision and recall into a single score, offering a balanced view of model performance.
- **ROC-AUC (Area Under the Receiver Operating Characteristic Curve):** ROC-AUC evaluates how well a model can distinguish between hired and non-hired candidates across different classification thresholds.
- **Confusion Matrix:** Confusion matrices provide a detailed visual breakdown of classification outcomes, showing where correct and incorrect predictions occurred.
- **Statistical Significance Testing:** Beyond traditional metrics, we used **McNemar's statistical test** to determine if the performance differences between the best-performing classifiers were statistically significant. This test is commonly used to compare two classifiers on the same test set and assess the meaningfulness of their performance variations.

3.5 Explainability Analysis

To boost transparency and interpretability, we integrated **Explainable Artificial Intelligence (XAI)** techniques. We performed a **Feature Importance analysis** to pinpoint which variables contributed most significantly to the hiring predictions. Additionally, we used **SHAP (Shapley Additive Explanations)** to provide both a global overview and local, individual explanations of how the model behaved [7].

The **SHAP Summary Plot** helped us understand the overall influence of candidate attributes across the entire dataset, while the **SHAP Waterfall Plot** explained individual predictions, showing how specific features increased or decreased the likelihood of a positive hiring recommendation. These techniques are crucial for improving transparency and helping HR professionals grasp the reasoning behind the model's suggestions [8].

3.6 Fairness Analysis

We conducted a fairness assessment to evaluate whether the machine learning models produced significantly different hiring outcomes across various candidate groups. Our analysis focused on several key metrics:

- **Selection Rates:** The proportion of candidates selected from each group.
- **Disparate Impact Ratio (DIR):** A measure of whether a selection process disproportionately excludes members of a protected group.
- **Statistical Parity Difference (SPD):** The difference in selection rates between different groups.
- **Equal Opportunity Difference (EOD):** The difference in true positive rates between different groups.

This integration of fairness assessment is a distinguishing feature of our study, setting it apart from many previous recruitment prediction studies that often prioritize predictive accuracy while overlooking transparency and ethical accountability [9].

4. Results

This section details the experimental results from implementing our four machine learning models. We evaluated these models using Accuracy, Precision, Recall, F1-Score, ROC-AUC, and Confusion Matrix analysis. Additionally, we performed explainability and fairness assessments to enhance transparency and ensure the responsible use of artificial intelligence.

4.1 Models Performance Comparison

Our initial evaluation, using a train-test split, revealed clear performance differences among the models. Table 1 shows that, the **Random Forest** model achieved the highest classification accuracy (93.33%), while **XGBoost** excelled with the highest precision (96.15%) and ROC-AUC score (0.933). This indicates XGBoost's superior ability to discriminate and its reliability in predicting positive hiring decisions. **Logistic Regression** also delivered competitive results, maintaining a high degree of interpretability, which is valuable when model transparency is crucial. In contrast, the **Decision Tree** model showed the lowest predictive performance.

Table 1. Performance Comparison of Machine Learning Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROCAUC
Logistic Regression	88.00	85.19	74.19	79.31	0.915
Decision Tree	85.67	78.41	74.19	76.24	0.825
Random Forest	93.33	95.06	82.80	88.51	0.931
XGBoost	93.00	96.15	80.65	87.72	0.933

These findings suggest that **ensemble learning methods** (like Random Forest and XGBoost) generally outperform traditional classification approaches in recruitment prediction. Their ability to capture complex relationships within candidate data allows them to generate more accurate hiring recommendations.

To further validate these findings and ensure their robustness, we conducted a **5-fold cross-validation analysis** and **statistical significance testing**. Cross-validation provides a more reliable estimate of a model's generalization performance by evaluating it across multiple data partitions, reducing sensitivity to a single train-test split. We also applied **McNemar's test** to compare the predictive performance of the best-performing ensemble models on the same test set as shown in Table 2 and 3.

Table 2. Cross-Validation Results

Model	Mean CV Accuracy (%)	Standard Deviation (%)
Random Forest	91.33	1.33
XGBoost	93.00	0.94

Table 3. McNemar Test Results

Comparison	RF Better	XGB Better	p-value	Significant
Random Forest vs XGBoost	25	50	0.0052	Yes

The cross-validation results showed that **XGBoost** achieved a higher mean accuracy and less variability across folds compared to Random Forest, indicating superior generalization performance. Furthermore, McNemar's test revealed a statistically significant difference in favor of XGBoost (p = 0.0052). This suggests that Random Forest's initial advantage was likely influenced by the specific data partition rather than a true performance lead. Consequently, **XGBoost was identified as the most reliable model** for predicting hiring decisions in this study.

4.2 ROC Curve Analysis

We generated Receiver Operating Characteristic (ROC) curves to assess each model's classification capability across various decision thresholds. The ROC analysis confirmed that XGBoost and Random Forest demonstrated superior discrimination capabilities compared to Logistic Regression and Decision Tree. Their higher ROC-AUC values signify a stronger ability to differentiate between hired and non-hired candidates.

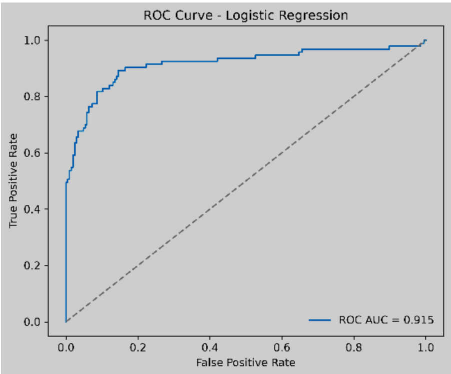

Figure 1. ROC Curve – Logistic Regression

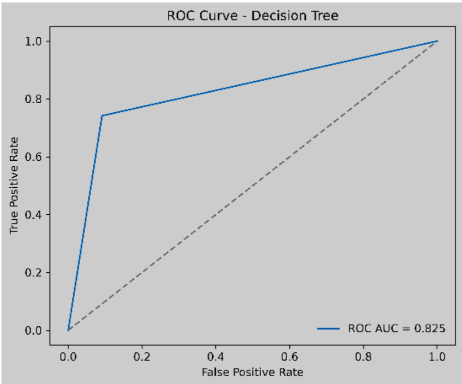

Figure 2. ROC Curve – Decision Tree

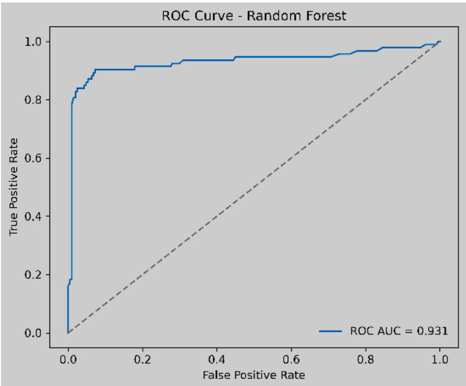

Figure 3. ROC Curve – Random Forest

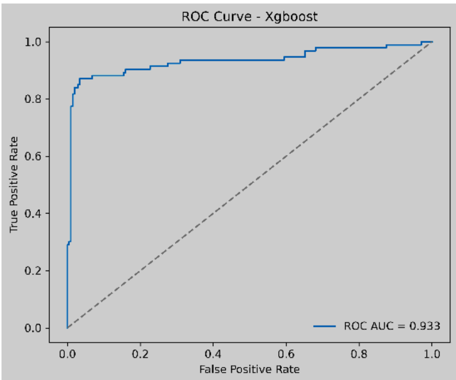

Figure 4. ROC Curve – XGBoost

4.3 Confusion Matrix Analysis

Confusion matrices provided a detailed breakdown of classification outcomes, helping us identify the distribution of correct and incorrect predictions. The results showed that Random Forest and XGBoost correctly classified a larger number of candidates and made fewer errors. This further supports the suitability of ensemble learning models for recruitment prediction applications.

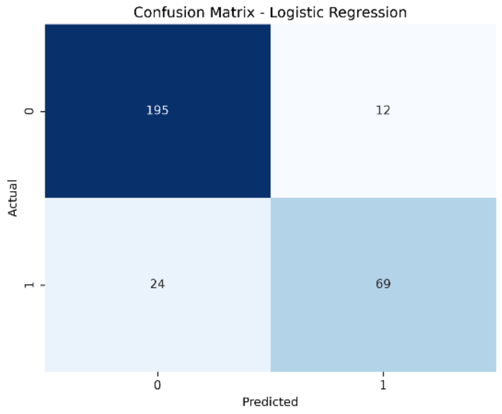

Figure 5. Confusion Matrix – Logistic Regression

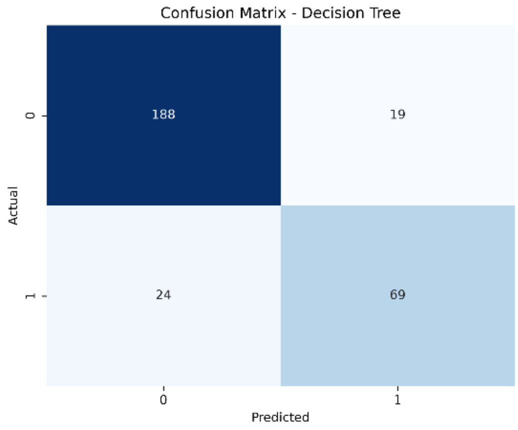

Figure 6. Confusion Matrix – Decision Tree

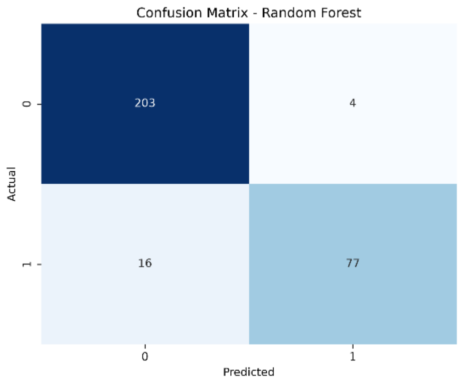

Figure 7. Confusion Matrix – Random Forest

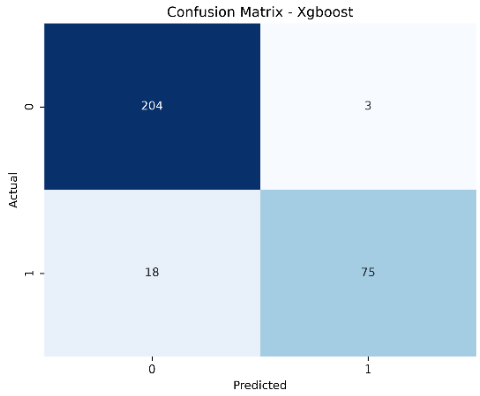

Figure 8. Confusion Matrix – XGBoost

4.4 Feature Importance analysis

To understand which factors most influenced hiring predictions, we conducted a feature importance analysis. The results highlighted **Recruitment Strategy, Skill Score, Personality Score, Experience Years, and Interview Score** as the most significant predictors. This indicates that hiring outcomes are primarily driven by a candidate's competencies, professional qualifications, and recruitment-related factors, rather than demographic characteristics. This aligns with modern recruitment practices that prioritize skills and demonstrated abilities.

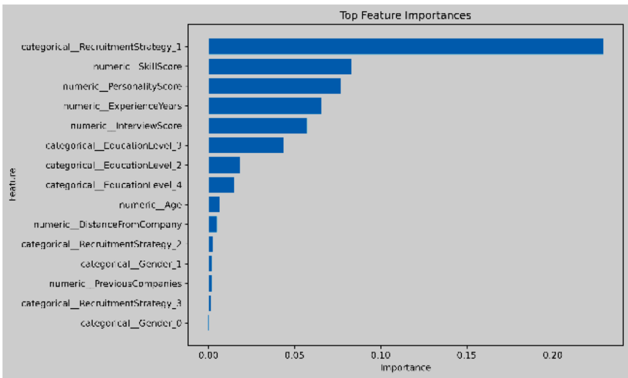

Figure 9. Feature Importance Analysis

4.5 Explainability Analysis

To enhance transparency, we performed **SHAP (Shapley Additive Explanations) analysis** using the best-performing model. The SHAP Summary Plot offered a global explanation of how different feature values influenced hiring predictions across the entire dataset, highlighting the relative importance of various candidate attributes. The SHAP Waterfall Plot, on the other hand, explained individual hiring predictions, showing how specific features either increased or decreased the likelihood of a positive hiring recommendation. These explanations are vital for improving transparency and helping HR professionals understand the rationale behind the model's decisions.

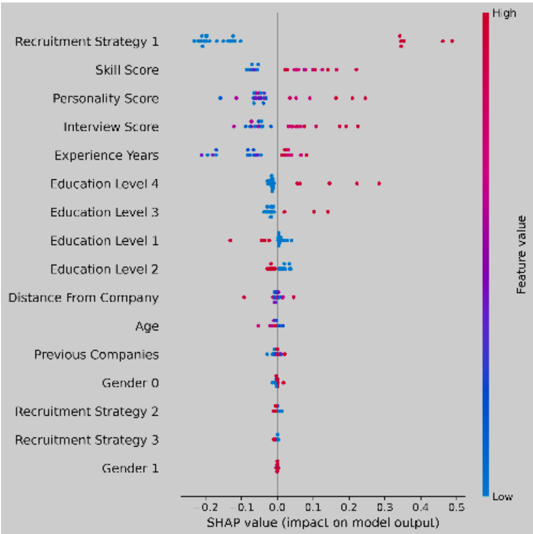

Figure 10. SHAP Summary Plot

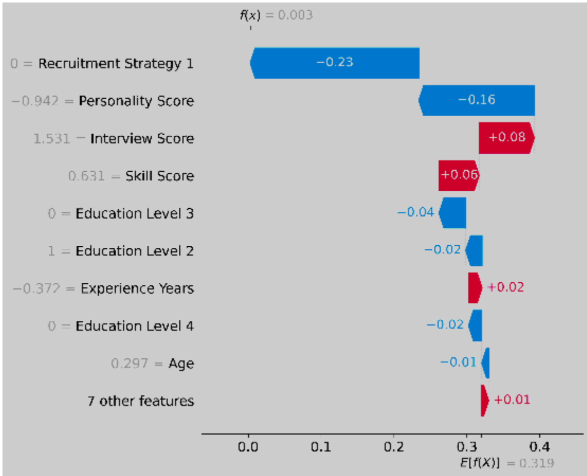

Figure 11. SHAP Waterfall Plot

4.6 Fairness Analysis

Our fairness assessment evaluated whether the models produced significantly different hiring outcomes across candidate groups, using **Gender** as the sensitive attribute. We found that Group 0 had a selection rate of 20.83%, while Group 1 had 30.77%. The calculated **Disparate Impact Ratio (DIR)** was 0.677. According to the **Four-Fifths Rule**, a selection rate ratio below 0.80 may suggest potential adverse impact and requires further investigation. This DIR value indicates possible group-level disparities. Additionally, the **Statistical Parity Difference (SPD)** was 0.099, and the **Equal Opportunity Difference (EOD)** was 0.143, both pointing to measurable differences in hiring outcomes and true positive rates between the two groups.

While these results do not definitively prove discrimination, they underscore the critical importance of continuous fairness monitoring and bias assessment in AI-assisted recruitment systems. Future work should explore fairness-aware learning techniques and bias mitigation strategies to further promote equity across candidate groups.

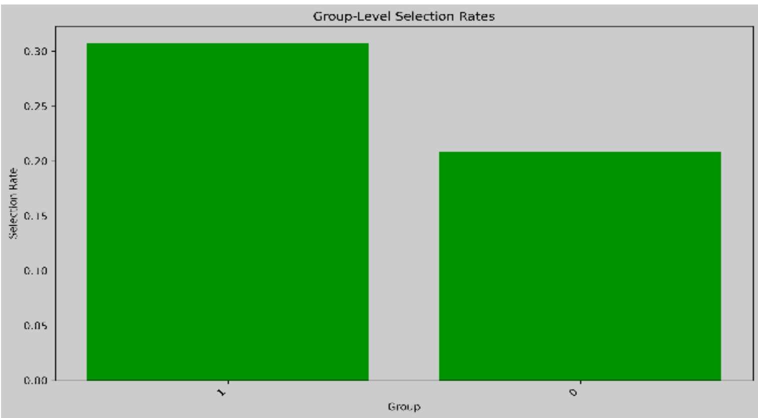

Figure 12. Group-Level Selection Rates by Gender

5. Discussion

Our study's findings clearly demonstrate that machine learning can effectively support recruitment decision-making and predict hiring outcomes. Consistent with previous research, **ensemble learning algorithms** (specifically Random Forest and XGBoost) significantly outperformed traditional classification methods like Logistic Regression and Decision Tree across most evaluation metrics. This indicates that ensemble methods are better equipped to capture the intricate relationships present within recruitment data. Further statistical analysis confirmed the observed differences between these ensemble models were indeed statistically significant.

The **feature importance analysis** revealed that factors such as **Recruitment Strategy, Skill Score, Personality Score, Experience Years, and Interview Score** were the most influential in predicting hiring decisions. This suggests that hiring outcomes are primarily driven by a candidate's qualifications, competencies, and performance indicators, rather than demographic characteristics. This observation aligns well with modern recruitment practices that prioritize skills and demonstrated abilities during candidate evaluation.

A significant contribution of this research is the integration of **Explainable Artificial Intelligence (XAI)** techniques. While many prior recruitment studies focused mainly on predictive accuracy, our work shows that hiring predictions can also be made transparent and justifiable. The **SHAP analysis** provided detailed explanations of how the models arrived at their recommendations, allowing recruiters to understand the underlying reasoning. Such transparency is crucial for building trust in AI-assisted recruitment systems.

Furthermore, our **fairness analysis** directly addressed one of the most pressing concerns surrounding AI in Human Resource Management. The assessment identified measurable differences in hiring outcomes across the evaluated groups. Using Gender as the sensitive attribute, the calculated **Disparate Impact Ratio (DIR)** of 0.677 fell below the commonly accepted **Four-Fifths Rule** threshold of 0.80. This result suggests a potential adverse impact and indicates that group-level disparities might be present. Additionally, the **Statistical Parity Difference (0.099)** and **Equal Opportunity Difference (0.143)** further highlighted differences in selection outcomes and true positive rates between the groups. While these findings do not definitively prove discrimination, they emphasize the critical need for continuous fairness monitoring, bias assessment, and responsible AI governance in recruitment systems.

Statistical analysis further solidified the superiority of **XGBoost** over Random Forest. Cross-validation results showed greater stability and better generalization performance for XGBoost, and McNemar's test indicated a statistically significant difference ($p = 0.0052$) in its favor. These findings reinforce our conclusion that XGBoost is the most reliable model for hiring decision prediction within the framework we proposed.

6. Research Contribution

This study makes several important contributions to the growing fields of Human Resource Analytics and AI-assisted recruitment:

1. **Demonstrated Effectiveness of Machine Learning:** We show how effective machine learning techniques are for predicting hiring decisions using candidate data relevant to recruitment. By comparing Logistic Regression, Decision Tree, Random Forest, and XGBoost, our research provides empirical evidence on which machine learning approaches are best suited for recruitment applications.
2. **Enhanced Transparency with Explainable AI:** We integrated Explainable Artificial Intelligence techniques to significantly improve transparency and interpretability. Our use of SHAP analysis allows recruiters to understand the reasoning behind hiring recommendations, thereby increasing trust in AI-assisted decision-making processes.
3. **Integrated Fairness Assessment:** This research incorporates fairness assessment directly into the recruitment prediction process. By evaluating hiring outcomes across different candidate groups, our study addresses critical ethical concerns related to algorithmic bias and promotes responsible AI deployment.
4. **Unified Decision-Support Framework:** Our proposed framework uniquely combines predictive performance, explainability, and fairness within a single recruitment decision-support system. This integrated approach directly fills an identified research gap, contributing to the development of more transparent, accountable, and trustworthy AI-driven recruitment systems.

7. Conclusion and Future Work

7.1 Conclusion

This study explored how machine learning techniques can be applied to predict hiring decisions in Human Resource Management. Our goal was to build a recruitment prediction framework that combines strong predictive performance with explainability and fairness, ultimately supporting transparent and responsible hiring decisions.

We implemented and evaluated four machine learning algorithms: Logistic Regression, Decision Tree, Random Forest, and XGBoost, using the "Predicting Hiring Decisions in Recruitment Data" dataset. Our experiments clearly showed that **ensemble learning methods** delivered superior predictive performance compared to traditional classification approaches.

While Random Forest initially achieved the highest accuracy in the first train-test split, further cross-validation and statistical significance testing revealed that **XGBoost** demonstrated superior overall performance and greater stability. This led us to identify XGBoost as the most effective model for predicting hiring decisions within our proposed framework.

Our explainability analysis pinpointed **Recruitment Strategy, Skill Score, Personality Score, Experience Years, and Interview Score** as the most influential factors in hiring decisions. Furthermore, SHAP analysis provided detailed insights into how the model works, allowing for transparent interpretation of hiring recommendations and boosting trust in AI-assisted recruitment systems.

We also integrated a fairness assessment into our framework to evaluate potential disparities across candidate groups. The results indicated measurable differences in hiring outcomes between the evaluated groups. Specifically, the calculated **Disparate Impact Ratio (DIR)** of 0.677 fell below the **Four-Fifths Rule** threshold of 0.80, suggesting a potential adverse impact. This highlights the critical importance of fairness monitoring in AI-assisted recruitment systems. While these findings don't conclusively prove discrimination, they demonstrate the immense value of incorporating fairness assessment into recruitment analytics to promote ethical and responsible AI adoption.

Overall, this study confirms that machine learning can effectively support recruitment decision-making while maintaining transparency and accountability. More importantly, our proposed framework successfully addresses the identified research gap by integrating predictive performance, explainability, and fairness into a unified recruitment decision-support system.

7.2 Limitations

It's important to acknowledge some limitations of our study. We used a single, publicly available recruitment dataset from Kaggle. While suitable for evaluating machine learning techniques, this dataset might not fully represent the complexities of real-world recruitment processes across different industries, organizations, and geographical regions. Consequently, the generalizability of our findings may be limited.

Additionally, the dataset does not originate from an actual organizational recruitment system, meaning it might contain synthetic or simulated patterns that differ from those observed in genuine hiring environments. As a result, the relationships identified by our machine learning models might not fully capture the intricate nature of real-world recruitment decision-making.

Furthermore, our fairness analysis was constrained by the available candidate attributes in the dataset, particularly using Gender as the sensitive attribute. While we evaluated fairness metrics like Disparate Impact Ratio, Statistical Parity Difference, and Equal Opportunity Difference, real-world recruitment bias is often influenced by more complex organizational, social, and historical factors that may not be present in publicly available datasets. Therefore, our fairness findings should be interpreted cautiously and ideally validated with real organizational recruitment data in future studies.

7.3 Future Work

Despite the promising results of our framework, several exciting opportunities exist for future research:

- Future studies could explore integrating **Large Language Models (LLMs)** and **Natural Language Processing (NLP)** techniques for automated resume screening and more sophisticated candidate-job matching. Such approaches could enable a more comprehensive analysis of unstructured recruitment data, including resumes, cover letters, and job descriptions.
- Additional research could investigate **fairness-aware machine learning algorithms** and **bias mitigation strategies** to further enhance ethical recruitment practices. Implementing continuous fairness

monitoring and bias auditing mechanisms could also help organizations maintain equitable hiring processes over time.

- Furthermore, future work should evaluate our proposed framework using diverse, real-world organizational recruitment data to confirm its applicability and effectiveness across various contexts.

REFERENCES

[1] S. M. Lundberg and S. I. Lee, (2017) “A Unified Approach to Interpreting Model Predictions,” in Advances in Neural Information Processing Systems (NeurIPS), vol. 30.

[2] F. Pedregosa et al., (2011) “Scikit-learn: Machine Learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830.

[3] L. Breiman, (2001) “Random Forests,” Machine Learning, vol. 45, no. 1, pp. 5–32.

[4] T. Chen and C. Guestrin, (2016) “XGBoost: A Scalable Tree Boosting System,” in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), pp.785–794.

[5] G. Zhang, (2025) “Explainable Artificial Intelligence in the Talent Recruitment Process,” Cogent Business & Management,.

[6] S. M. U. Dadaboyev, (2025) “Role of Artificial Intelligence in Employee Recruitment,” Discover Artificial Intelligence.

[7] S. Fabeyo, (2025) “Explainable AI in Employment Decision-Making: A Systematic Review of Transparency Methods in Hiring Algorithms,” Issues in Information Systems,.

[8] C. Malin et al. (2025), “Rejected by an AI? Comparing Job Applicants’ Fairness Perceptions of Artificial Intelligence and Humans in Personnel Selection,”.

[9] D. F. Mujtaba and N. R. Mahapatra (2024), “Fairness in AI-Driven Recruitment: Challenges, Metrics, Methods, and Future Directions”.

[10] Alnnale, T. (2026). Predictive Governance in Digital Enterprises: An LSTM-Enhanced Deep Learning Framework for Economic Optimization of IT Incident Management Using Enriched Process Logs. *Al-Farooq Journal of Sciences*, *2*(3), 86-113.

[11] A. Fabris et al., (2023) “Fairness and Bias in Algorithmic Hiring: A Multidisciplinary Survey,”.

[12] R.H. Elghanuni, M. B. Swidan2, A. A. Diaf, A. A. Elhoni, (2026) “Evaluating the Impact of PCA-Based Dimensionality Reduction on Bitcoin Transaction Forecasting: A Comparative Study of XGBoost, LSTM, and GNN”, LOUJAS, Volume 2-Issue 1 -2026-Pages 119-126

\