Received: 27.08.2026| **Accepted:** 30.09.2026| **Published:** Oct.2026
Citation: Akeni, W. O., et al. (2026) Explainable Ai Based Decision Support for Automated Hiring and Recruitment Using Shap and Lime. GAIJIST Vol.2.1.2026 pp.49-64 DOI: https://doi.org/10.5281/zenodo.23872278

EXPLAINABLE AI BASED DECISION SUPPORT FOR AUTOMATED HIRING AND RECRUITMENT USING SHAP AND LIME

CHUKWUKA STELLA CHIBUEZE[1], OMOROGBE Osasu Harry (PhD)[2], AMOFORITSE, Fortune Ighotuweyin[3], OBANDE Bonnie Obeka, Ph.D. (CLN)[4] Ozobialu Emmanuel Chukwuma[5]
Igbinedion University Okada, Edo State, Nigeria

Abstract

This study is focused on the design, development, implementation, and evaluation of an Explainable and Fair AI-Based Decision Support Framework for Automated Hiring and Recruitment. The dataset contains 1,500 candidate records and 11 features. Approximately 80% of the dataset were reserved for the purpose of training the model, while the remaining 20% of the dataset were reserved for the purpose of testing the model. In order to ensure a more robust assessment of the performance of the model as well as reduce the tendency of the results being dependent on a single training-test partition, a five-fold cross-validation was implemented during model training. Four models were considered, namely XGBoost, Random Forest, LightGBM, and Logistic Regression. Based on the overall predictive results, XGBoost was selected as the principal prediction model for the framework. The explainability analysis showed that different candidate and recruitment-related characteristics contributed differently to the model's predictions. The global SHAP analysis identified Recruitment Strategy as the feature with the greatest overall influence on the model's predictions. This was followed by Skill Score, Personality Score, Interview Score, and Experience Years. The local SHAP and LIME analyses also showed that the factors influencing one candidate's prediction may differ from those influencing another candidate's prediction. The framework therefore supports a human-in-the-loop approach, where the AI system provides analytical support while the recruiter retains responsibility for the final hiring decision.

Keywords: Automated Hiring and Recruitment; Explainable Artificial Intelligence; SHAP; LIME; XGBoost; Decision Support

*Corresponding author(s). CHUKWUKA, S. C., et al. agree that this article remain permanently open access under the terms of the Creative Commons Attribution License 4.0 International License

1 Introduction

The increasing adoption of Artificial Intelligence (AI) in recruitment and hiring has improved the efficiency of screening, assessing, and selecting job applicants. AI powered recruitment systems enable organizations to process large volumes of applications and support faster, data driven hiring decisions (Black & van Esch, 2020; Köchling & Wehner, 2020). However, despite these benefits, concerns remain about the transparency and fairness of the decisions generated by these systems. Many AI recruitment models use complex algorithms that provide limited explanations for their predictions, making it difficult for recruiters and applicants to understand why some candidates are recommended for selection while others are rejected (Burrell, 2016; Lipton, 2018). This limitation can weaken human oversight and make it difficult for organizations to justify hiring decisions.

Although previous studies have contributed to the development of AI recruitment systems, explainable AI techniques, and algorithmic fairness methods, these areas have often been investigated separately (Adadi & Berrada, 2018; Köchling & Wehner, 2020). Research on predictive modelling commonly emphasises performance measures such as accuracy, precision, and recall, while explainability research focuses on making model predictions understandable. Similarly, fairness research examines how to identify and reduce unequal outcomes among different groups. However, improving one aspect does not necessarily guarantee improvement in the others. A model may produce accurate predictions without providing adequate explanations, while a model that offers explanations may still generate unfair recommendations. Therefore, there is a need to consider predictive performance, explainability, and fairness together when developing AI based recruitment systems.

From a practical perspective, the study provides a web-based application that can assist recruitment team members to review AI-generated hiring recommendations in order to make more informed recruitment decisions. So instead of simply presenting a hiring recommendation, the system provides additional information that helps recruitment teams to understand why a particular recommendation was made and whether there are any concerns about fairness in the outcome. This gives recruiters an opportunity to review the AI-generated recommendation alongside other relevant information before making the final hiring decision. In this way, the system is aimed at supporting the recruitment process rather than replace the judgment and responsibility of the recruitment team.

2 Related Literature

The use of Artificial Intelligence (AI) in recruitment and hiring has changed the way organizations manage and evaluate job applicants. AI-enabled recruitment tools can be used to screen applications, match candidates with job requirements, rank applicants, and provide recommendations to recruiters (Black & van Esch, 2020; Upadhyay & Khandelwal, 2018). These applications are particularly useful where organizations receive large numbers of applications and need to process candidate information within a limited period. In this sense, AI can serve as a decision-support tool that helps recruiters organize and assess candidate information rather than relying entirely on manual processes.

Another concern is the difficulty of understanding how some AI systems arrive at particular recommendations. Although a recruitment system may indicate that a candidate is suitable or unsuitable for a position, the recruiter may not always be able to determine why the system produced that result. Burrell (2016) identifies opacity as an important challenge in machine-learning systems, while Pasquale (2015) highlights the broader concerns associated with decisions produced by opaque computational systems. In recruitment, this lack of clarity can make it difficult for recruiters to review or justify an AI-generated recommendation and may also reduce the ability to identify potentially unfair patterns in the system's decisions.

Machine learning is also applied to candidate matching and ranking. In this approach, information about a candidate is compared with the requirements of a job, and the system can assign a score or rank to indicate the degree of suitability. Candidate ranking can help recruiters organize large applicant pools and identify candidates for further consideration. Pessach et al. (2020), for example, developed a recruitment analytics framework that used machine learning to predict recruitment success at the level of individual job placements and combined this with an optimization approach to support wider hiring and placement decisions. The study demonstrates how machine learning can be incorporated into a broader decision-support process rather than being used only as a standalone prediction tool.

Another important application is the prediction of hiring outcomes. Machine learning models can be trained using historical candidate and recruitment data to learn relationships between candidate characteristics and previous hiring outcomes. After training, the model can be applied to new candidates and used to estimate whether a candidate is likely to meet a defined hiring outcome. This type of application is closely related to the present study, where machine learning models are developed to predict hiring decisions from candidate and recruitment-related

information. The use of machine learning for personnel selection has also received increasing attention because of its ability to identify patterns in selection data that may be difficult to examine manually (Campion & Campion, 2023).

Explainable Artificial Intelligence (XAI) has become an important area of AI research because of the growing need to understand how AI systems arrive at their decisions. As AI is increasingly used to support decisions that affect individuals and organizations, simply obtaining a prediction from an AI system may not be enough. Users may also want to know why the system produced a particular result and which factors influenced the decision. This concern is particularly important when AI is used in areas such as recruitment, where the outcome of an AI-supported decision may affect a person's employment opportunity (Adadi & Berrada, 2018).

XAI generally refers to methods and approaches that make the behaviour and outputs of AI systems easier for people to understand (Guidotti et al., 2018). It therefore extends the focus of AI beyond prediction alone by considering how the results produced by a model can be explained to its users. Arrieta et al. (2020) explain that explainability is concerned with making the operation and decisions of AI systems more understandable. This can help users examine AI-generated results rather than simply accepting them without further consideration.

Explainability can take different forms. Some models are intrinsically interpretable, meaning that their decision processes can be understood directly from the structure of the model. Other models require post-hoc explanation methods, where additional techniques are applied to explain the outputs of a model after it has generated a prediction (Arrieta et al., 2020). Explanations can also be considered at the global or local level. Global explanations provide information about the factors that generally influence the behaviour of a model, while local explanations focus on why the model produced a particular prediction for an individual case.

The distinction between global and local explanations is particularly relevant to recruitment decision support. A recruiter may want to know which candidate characteristics generally have the greatest influence on the model's hiring predictions, while also wanting to understand why a particular candidate received a specific prediction. XAI techniques can therefore provide information at both levels, helping users to examine the general behaviour of the model as well as individual recommendations.

Another important consideration is that an explanation must be understandable to the person who is expected to use it. A technically accurate explanation may not be useful if it is presented in a form that the

intended user cannot easily interpret. XAI should therefore consider the needs of the users, the decision context, and the type of information required to support meaningful human review (Doshi-Velez & Kim, 2017). In a recruitment environment, explanations should provide useful information to recruiters without requiring them to have advanced knowledge of machine learning. In this study, explainability forms one of the main components of the proposed AI-based decision support framework. The framework uses SHAP and LIME to provide information about the factors influencing the model's hiring predictions. These methods are considered further in the subsequent discussion because they provide the technical basis for the explainability component implemented in the developed system.

3 Materials and Methods

3.1 Research Design and Dataset

This study adopted the Design Science Research (DSR) methodology proposed by Peffers et al. (2007), which provides a systematic approach for developing and evaluating technological artefacts. The methodology is made up of six interrelated phases: problem identification and motivation, definition of solution objectives, design and development, demonstration, evaluation, and communication. These phases guided the development of the Explainable and Fair AI-Based Decision Support Framework for Automated Hiring and Recruitment. The study commenced with identifying the challenges associated with explainability, fairness, and transparency in AI-assisted recruitment systems through an extensive review of existing literature. Based on the identified problems, the objectives of the proposed framework were established. The framework was subsequently designed, developed, and implemented as a web-based application.

In this study, a publicly accessible recruitment dataset was obtained through the Kaggle repository. The dataset contains 1,500 candidate records and 11 features. The features of the dataset were candidate demographic information, educational qualification, professional affiliation assessment, and recruitment-related information. The dataset was considered appropriate as it contains characteristics fit for the development a machine learning model for the prediction of hiring decisions while at the same time providing support for the assessment of fairness. Table 1 below provides the summary of the dataset characteristics

.

Table 1 Summary of the Dataset Characteristics

Attribute	Description
Dataset Name	Recruitment Data / Hiring Decision Dataset
Source	Kaggle Repository
Total Number of Records	1,500 candidate records
Hired Cases	465
Not Hired Cases	1,035
Number of Predictor Features	10 candidate and recruitment-related features
Hiring Label	1 (Hired)
Non-Hiring Label	0 (Not Hired)
Fairness-Related Attributes	Age and Gender
Data Format	CSV
Learning Type	Supervised Binary Classification
Machine Learning Models	XGBoost, Random Forest, LightGBM, Logistic Regression
Missing Values	None Identified

Source: Adapted from the recruitment dataset

3.2 Data Preprocessing and Model Development

In the data cleaning phase, the dataset was examined with the use of the Python Pandas library for the identification of inconsistencies such as missing values, duplicate records, and inconsistent format among others. The initial data inspection operations revealed that there were no missing values in the dataset. As a result, no missing-value imputation was required. The records were further inspected for duplicate values and other forms of inconsistencies before going over to the next phase of pre-processing.

In the feature encoding phase, all categorical variables such as gender, education level, and recruitment strategy,

were converted into numerical representations using one-hot encoding strategy. The encoding process was facilitated by the one hot encoder function from Python scikit-learn library. Accordingly, each category was represented as a separate binary feature, thus, preserving the categorical information without imposing an artificial numerical relationship among the categories. The encoded variables were then combined with the numerical variables to form the feature matrix that was used during the model development process.

In the feature preparation phase, the dataset was separated into input features (X) and the target variable (y). The essence was to establish the structure required for supervised machine learning. While the candidate demographic information and recruitment-related attributes were specified as the input feature matrix, while HiringDecision was specified as the target variable. The target variable was retained in binary form, with 0 representing a non-hiring decision and 1 representing a hiring decision. Age and Gender were also retained as fairness-related attributes for subsequent assessment of potential disparities in the predictions by the model. The resulting feature matrix and target variable were then prepared for the subsequent model training stage.

In order to obtain an objective evaluation of the machine learning models, thus, reducing the tendency for overfitting, the dataset was divided into training and testing subsets. Approximately 80% of the dataset were reserved for the purpose of training the model, while the remaining 20% of the dataset were reserved for the purpose of testing the model. The training subset of the data was then used to in the development of the machine learning models. Accordingly, the testing data were used in the evaluation of the predictive performance of the model. In order to ensure a more robust assessment of the performance of the model as well as reduce the tendency of the results being dependent on a single training-test partition, a five-fold cross-validation was implemented during model training.

It receives candidate's information in a pre-processed format and uses it in the training of machine learning models to predict candidate's hiring outcome. This module examines the candidate's demographic information based on the patterns learned from the training dataset to generate a binary prediction that indicates whether the candidate is hired or not hired. This component of the framework integrates four predictive models namely XGBoost, Random Forest, LightGBM, and Logistic Regression. These models are evaluated based on their performance and the best-performing model is then adopted for use as the principal prediction model in the framework. Apart from the hiring outcome, this module provides a confidence score which gives the users additional information

about the level of certainty of the model concerning the predictions generated. The output of the prediction and the associated confidence score are then resented to the explainability and fairness module for further assessment.

3.3 Explainability and Decision Support

The module serves as the component that reduces the opacity that is associated with the decision generated by the model such that hiring recommendations are not resented as unexplained outcomes. This module integrates SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME) in the provision of explanations of the predictions of the model. In this module, SHAP is the explanability tool that determines the how individual features contributes to the predictions of the model, providing both global and local interpretations of the behaviour of the model. LIME on the other hand is the explanability tool that generates local explanations by carefully examining the factors that influence the prediction for an individual candidate. This module therefore provides explanation that is in two levels: global and local explanations. The global explanation indicates the features that have the greatest overall influence on the hiring predictions of the model, while local explanations indicates the features that contribute to a particular candidate's prediction. The resulting explanations are presented through the web-based decision support interface, allowing the user of the framework to examine the basis of an AI-generated hiring recommendation.

The decision support module integrates the outputs of the prediction, explainability, and fairness components into a single decision-support environment for the recruiter. Rather than presenting only a binary hiring prediction, the component brings together the predicted hiring outcome, prediction probability, explanatory information, and fairness information associated with the AI-generated recommendation. The information is presented through the web-based decision support dashboard, enabling the recruiter to review the candidate's predicted outcome and examine the factors contributing to the prediction. The dashboard also provides the relevant fairness information, allowing the recruiter to consider potential disparities associated with the model's output. The module follows a human-in-the-loop approach, in which the AI system provides analytical support while the human recruiter has the responsibility of taking the final hiring decision such that the system does not replace human judgement. Rather, it provides the user with additional evidence and information that can be considered alongside professional judgment and other recruitment requirements. Through the integration of prediction, explainability, and fairness information, the

decision support component provides a more transparent and responsible approach to AI-assisted hiring and recruitment.

The framework was implemented as a web-based application in order to provide an interface that allows users to interact with the system. The web interface allows users to enter candidate information into the system, the information is then processed by the AI module, and the hiring prediction output is generated along with explainable information, and fairness indicators. The framework was developed using Python-based computational environment, with Streamlit as the web application framework. Streamlit was used in the development of the graphical user interface as well as connecting the different components of the framework together for proper functionality. The implementation was structured in such as a way that allows all the modules to function within a single application environment. The functional stages of the framework are described below.

4 Results

4 .1 Predictive Performance

The predictive performance of the machine learning models was evaluated using the test dataset. The evaluation focused on Accuracy, Precision, Recall, Specificity, and F1-Score.

Table 2 Comparative Performance of the Machine Learning Models

Model	Accuracy (%)	Precision (%)	Recall (%)	Specificity (%)	F1-Score (%)
XGBoost	93.67	96.25	82.80	98.55	89.02
Random Forest	92.67	93.83	81.72	97.58	87.36
LightGBM	92.33	96.05	78.49	98.55	86.39
Logistic Regression	88.00	85.19	74.19	94.20	79.31

Source: Author's implementation (2026)

Based on the comparative evaluation, XGBoost which emerged as the model with the best overall performance was selected as the principal prediction model for the implemented framework. The selected model was subsequently integrated into the Streamlit application together with the saved preprocessing component. The selection was based on the combined interpretation of Accuracy, Precision, Recall, Specificity, and

F1-Score. This ensured that model selection was not based on a single performance measure. The selected model provided the prediction used by the explainability, fairness, and decision-support components of the framework.

The comparative analysis therefore established XGBoost as the predictive model used in the subsequent implementation of the Explainable and Fair AI-Based Decision Support Framework.

4.2 Global SHAP Analysis

The global SHAP analysis was used to determine the features that had the greatest overall influence on the hiring predictions produced by the XGBoost model. In this analysis, a higher mean absolute SHAP value indicates that a feature had a greater influence on the model's predictions across the test data. The result showed that Recruitment Strategy 1 had the highest overall influence, with a mean absolute SHAP value of 1.490754. It was followed by Skill Score (0.840794), Personality Score (0.762873), Interview Score (0.761622), and Experience Years (0.647660). These were the features that contributed most strongly to the model's predictions. Figure 1 below shows the graphical representation of the global SHAP

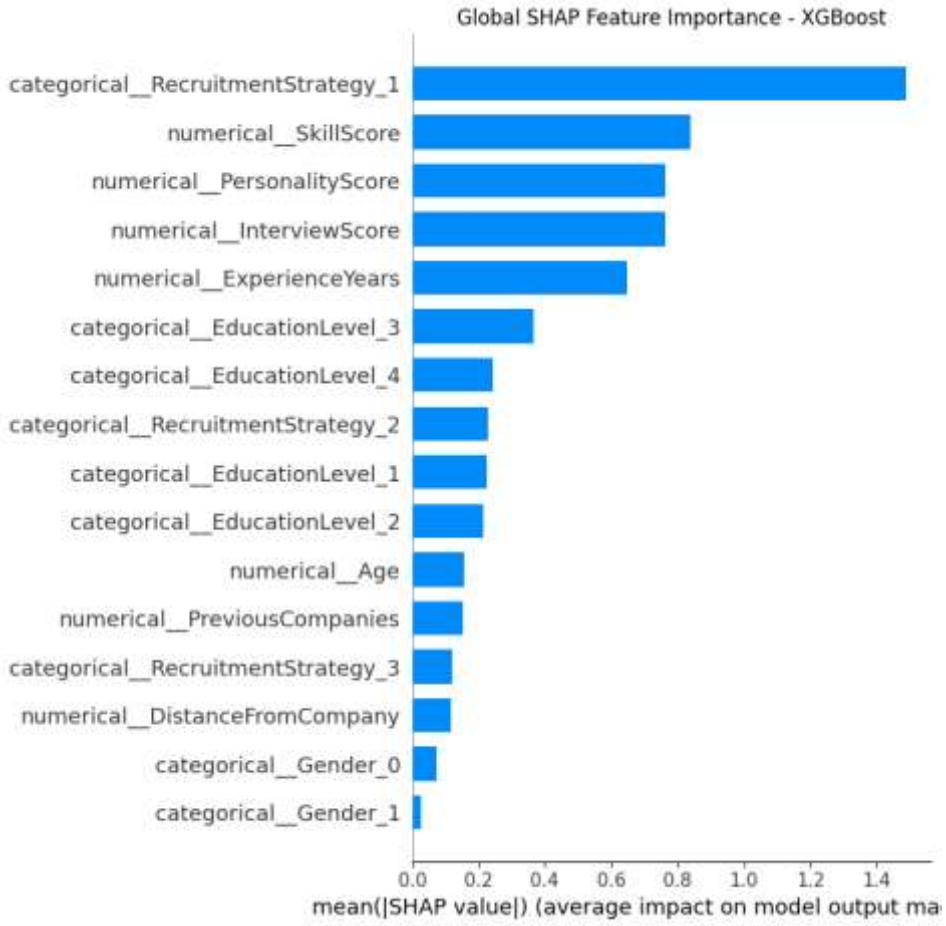

Figure 1 Global SHAP Feature Importance

Source: Author's implementation (2026)

Other features had lower levels of influence.

For example, Age recorded a mean absolute SHAP value of 0.156741, while Gender 0 and Gender 1 recorded values of 0.073634 and 0.025679, respectively. This means that, in terms of overall model influence, the encoded gender variables contributed less than the major recruitment and candidate assessment variables. Figure 1 above show that the model depended more strongly on recruitment strategy and candidate assessment variables than on demographic variables. In particular, skill, personality, interview performance, and experience were among the major contributors to the model's overall behaviour. The relatively smaller values for Age and Gender indicate that these variables had less overall influence on the model than the leading features. However, their lower importance does not by itself prove that the model is fair.

4 3 Local SHAP Analysis

While the global SHAP analysis explains the general behaviour of the model, it does not explain why a particular candidate received a particular prediction. Therefore, a local SHAP analysis was carried out for an individual candidate. For the selected candidate (candidate index 0), the model predicted Not Hired. The local SHAP results showed that Recruitment Strategy 1 had the strongest negative contribution, with a SHAP value of −1.543190. This was followed by Personality Score, which had a negative contribution of −1.056366. However, not all factors contributed negatively. Interview Score had a positive SHAP value of 0.783032, while Skill Score contributed 0.709015 and Experience Years contributed 0.372534 towards the prediction. This means that some characteristics of the candidate supported a positive hiring outcome, while other characteristics pushed the prediction in the opposite direction. The local analysis therefore provides a clearer explanation of how the model arrived at the particular prediction for this candidate. Instead of simply showing "Not Hired", the system can show the recruiter which factors supported or opposed that prediction.

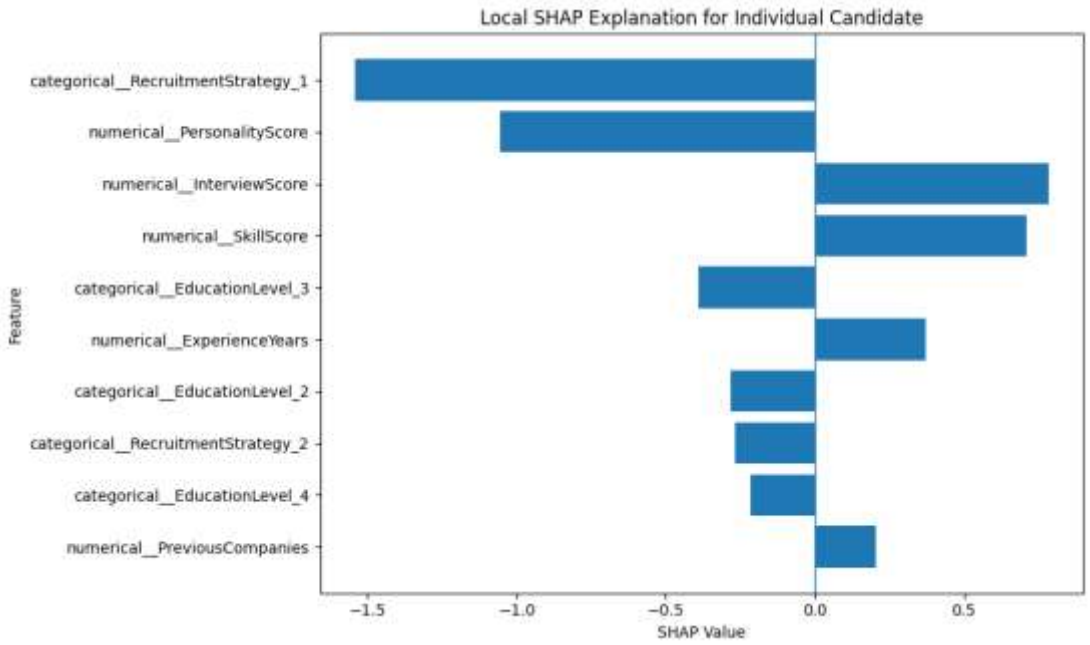

Figure 2 Local SHAP Explanation for the Selected Candidate

Source: Author's implementation (2026)

The local SHAP explanation as indicated in figure 2 above visually shows the direction and strength of the individual feature contributions such that features with negative contributions pushes the prediction towards Not Hired, while positive contributions pushed it towards Hired.

4 4 LIME Analysis

LIME was also applied to the selected candidate to provide a second local explanation of the model's prediction. The LIME analysis classified the candidate as Not Hired, with a predicted probability of 91.61%, compared with 8.39% for Hired. The result showed that Recruitment Strategy 1 ≤ 0.00 had the strongest negative contribution, with a LIME weight of −0.481352. On the other hand, Interview Score > 0.87 contributed positively with a weight of 0.224181. Other conditions, including Education Level 4, Education Level 3, and Personality Score, also contributed to the final prediction.

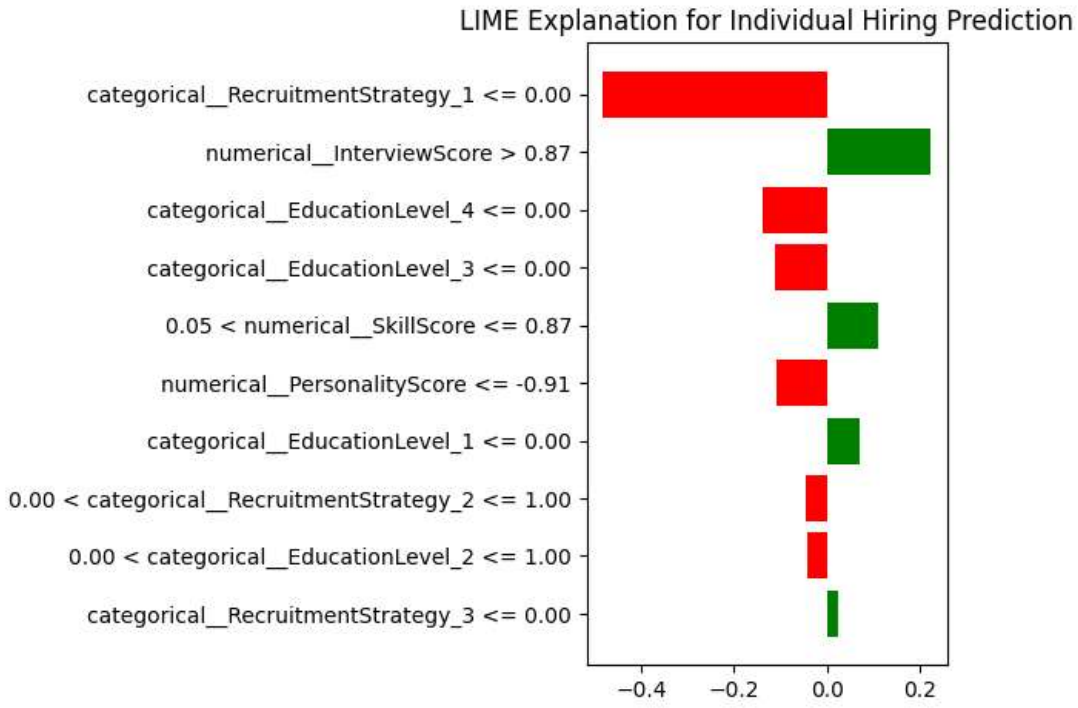

Figure 3 LIME Explanation for the Selected Candidate

Source: Author's implementation (2026)

The LIME result as indicated in figure 3 above provides a candidate-specific view of the prediction. The negative weights indicate conditions that contributed towards the Not Hired outcome, while the positive weights indicate conditions that supported the opposite direction. The probability of 91.61% Not Hired shows that the model had a strong preference for the negative class for this particular candidate.

5 Discussion

The combined SHAP and LIME results show that the developed framework was able to explain the model's predictions at two levels. The global SHAP analysis showed the features that generally influenced the model most across the test data. The local SHAP analysis then showed how specific features contributed to the prediction of an individual candidate. LIME provided an additional local explanation of the same candidate's result. For example, although Skill Score and Interview Score were among the important features globally, their effects were not necessarily the same for every candidate. In the selected candidate's case, both contributed positively to the prediction, but Recruitment Strategy 1 and Personality Score contributed more strongly in the negative direction. This demonstrates why examining individual predictions is important in an AI-assisted recruitment system. Overall, the explainability results support the objective of making the AI-based hiring process more transparent. The system does not only provide a hiring prediction; it also provides information that allows the recruiter to examine the factors behind that

prediction. This supports the human-in-the-loop approach adopted in the study, where the AI system provides analytical support while the recruiter retains responsibility for reviewing the recommendation and making the final hiring decision.

Explainability can also contribute to the assessment of fairness. By providing information about the factors influencing model predictions, explanation methods can help users identify patterns that may require further investigation (Mehrabi et al., 2021). However, an explanation does not by itself guarantee that a model is fair. Rather, it provides additional information that can help users examine the model's behaviour and identify possible concerns. This distinction is important because explainability and fairness address related but different aspects of responsible AI.

Explainability also has implications for human oversight and trust. A recruiter may be more willing to consider an AI recommendation when the system provides understandable information about the basis of that recommendation. At the same time, the explanation gives the recruiter an opportunity to question or reject the recommendation rather than accepting it automatically. This is important because the purpose of an AI-based decision-support system is not necessarily to remove human judgement, but to provide additional information that can assist the person responsible for the decision (Adadi & Berrada, 2018). It is important to recognise, however, that explainability alone cannot make an AI decision ethical. A system may provide a clear explanation of a prediction and still produce an unfair outcome. Ethical AI decision-making therefore requires explainability to be considered alongside fairness assessment, appropriate data practices, accountability, and human oversight. In the context of recruitment and hiring, these elements are particularly important because AI-supported decisions can influence people's access to employment opportunities.

6 Limitations

First, the study relied on a single publicly available recruitment dataset. As a result, the findings may not fully represent recruitment practices and candidate characteristics across different organizations, industries, or geographical locations.

Third, the developed framework was implemented as a prototype web-based decision-support application. Its performance in an actual organizational recruitment environment may differ from the results obtained from the experimental dataset.

Future research may also examine how recruiters interact with AI explanations and fairness information during actual recruitment activities. This could provide useful evidence on whether explainable and

fairness-aware decision-support systems improve recruiter understanding, confidence, and decision quality.

7 Conclusion

Overall, the study concludes that combining prediction, explainability, fairness assessment, bias control, and human oversight provides a more responsible approach to AI-assisted recruitment than relying on prediction alone. The developed framework provides recruiters with additional information for reviewing AI-generated hiring recommendations while keeping the final decision under human control. In this way, the framework is positioned as a decision-support tool rather than an autonomous hiring decision-maker.

Explainability should be incorporated into AI-based recruitment systems. Techniques such as SHAP and LIME can help recruiters understand the factors influencing AI-generated recommendations and provide greater transparency during the recruitment process.

References

Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). IEEE Access, 6, 52138–52160.

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., … Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115.

Black, J. S., & van Esch, P. (2020). AI-enabled recruiting: What is it and how should a manager use it? Business Horizons, 63(2), 215–226.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. Big Data & Society, 3(1), 1–12.

Campion, M. A., & Campion, E. D. (2023). Machine learning applications to personnel selection: Current illustrations, lessons learned, and future research. Personnel Psychology, 76(4), 993–1009.

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. ACM Computing Surveys, 51(5), 1–42.

Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. Business Research, 13, 795–848.

Lipton, Z. C. (2018). The mythos of model interpretability. Queue, 16(3), 31–57.

Omorogbe, O. H., Eduje, A. I., Anyanwu, L., Obeka, O. B., Ukaoha, K. C., Izevbuwa, O. G., Ighotuweyin, A. F., & Arenvbaguehita, O. D. (2025). The impact of artificial intelligence (AI) on fraud detection in banks in Nigeria. Oil and Gas Journal, 2(9), 10.

https://www.researchgate.net/publication/395988619_The_Impact_of_Artificial_Intelligence_AI_On_Fraud_Detection_in_Banks_in_Nigeria

Omorogbe, O. H., Amoforitse, F. I., & Ozobialu, E. C. (2026). Explainable reinforcement learning for adaptive autonomous agents: A framework for interpretable decision-making in dynamic environments. Pacadpub (IJIS).

https://www.pacadpub.com/index.php/ijis/article/view/49

Omorogbe, O. H., Amoforitse, F. I., Ozobialu, E. C., Eduje, A. I., & Obande, B. (2026). Resolving the scalability–explainability–adaptability trilemma in hybrid optimization systems for disaster management. International Journal of Data Science and Analytics (IJDSA), 2(1).

https://www.pacadpub.com/index.php/ijdsa/article/download/51/36

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.

Pasquale, F. (2015). The Black Box Society: The Secret Algorithms That Control Money and Information. Harvard University Press.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, 24(3), 45–77.

Pessach, D., Singer, G., Avrahami, D., Chalutz Ben-Gal, H., Shmueli, E., & Ben-Gal, I. (2020). Employees recruitment: A prescriptive analytics approach via machine learning and mathematical programming. Decision Support Systems, 134, 113290.

Upadhyay, A. K., & Khandelwal, K. (2018). Applying artificial intelligence: Implications for recruitment. Strategic HR Review, 17(5), 255–258.