International Journal of Computational Research in Science and Technology
E-ISSN: 3108-0871, Volume 2 Issue 6
Nov-Dec-2026, PP 01-09

Explainable AI for Fair and Transparent Human Resource Decision- Making

Spoorthi R[1], K.Subhashini[2], Aishwarrya Prakash[3]
[1]Assistant Professor Department Of Commerce
Jnana Vikas Institute Of Management Studies And Commerce, Bidadi, Ramanagar Taluk, Bangalore South District-562109
Spoorthi0110@Gmail.Com
[2]Guest Lecturer Computer Science
Government Arts & Science College, Nannilam.
Subharaniphd@Gmail.Com
[3]Assistant Professor Department Of Commerce And Business Administration / Department Of Mba
United International Business School, Hennur Baglur Road, Bangalore
Aishwarrya13bab@Gmail.Com

Abstract- Human resource (HR) decisions are currently supported by algorithmic processes for hiring, promotions, and retaining employees, the lack of transparency in these models raises concerns about fairness, accountability, and trust of employees towards the process. In this paper, we develop a two-channel explainable artificial intelligence (XAI) model that combines the post hoc explanation methods such as SHAP and LIME with the inherently interpretable classifier for HR decisions. In our methodology, adversarial debiasing is used in the training process to prevent the influence of sensitive attributes on the model performance. For testing the effectiveness of the proposed XAI model, we conduct an experimental evaluation using IBM HR analytics dataset with three HR decision making tasks including promotion prediction, attrition prediction, and candidate selection task. As results, we show that our model achieves reductions in demographic parity disparity without compromising the explanation quality using local accuracy and stability measures. Comparing with baseline models in terms of fairness-utility trade-off consistently yields improvements.

Key Word: Explainable AI, Human Resource Management, Algorithmic Fairness, SHAP, Adversarial Debiasing, Decision Transparency

I. INTRODUCTION

Introduction of Artificial Intelligence into Human Resource Management has revolutionized the way companies recruit new employees, evaluate their performance, and plan for their future. AI algorithms help in sorting resumes, ranking candidates, predicting promotions, and determining whether an employee is at risk of leaving. On the one hand, there is potential for increased efficiency of such processes, while on the other, these systems introduce new challenges related to transparency and interpretability of their conclusions. The more complex these algorithms become, the better their results are, but the harder it becomes to understand what they do [3], [4].

This lack of transparency creates several governance problems. Firstly, the individuals being assessed using these tools have the right to know the reasoning behind decisions that will directly impact their livelihoods. Procedural Justice theory states that people's sense of procedural justice will be based not only on the outcome but also on how transparent and explainable the process is [3]. Secondly, there is an increasing pressure on organizations to protect themselves against legal consequences of discriminatory practices by ensuring that automated HR tools do not discriminate. According to the new EU AI act, which classifies all AI used in hiring process as high-risk, these systems must provide explanations for their recommendations. In the United States, some state regulations require bias assessment for automated employment decision-making tools [3], [5]. Thirdly, without transparent output, it will be impossible for HR managers to provide any oversight [3].

Technical solutions to these problems have been centered on XAI, which involves research into ways in which machine learning models' behavior could be made understandable to humans. Post hoc approaches like SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) have become popular due to their ability to work across different types of models and based on cooperative game theory. SHAP approach also meets desirable criteria of being locally accurate, missingness, and consistent, which makes it appropriate for critical domains where reliable explanation is crucial [1], [2], [4]. Nevertheless, explainability does not necessarily mean fairness since a model can give perfect explanations of its decisions which are unfair and discriminatory anyway [3], [8].

This paper aims to solve this problem through the proposed unified framework, which uses a combination of adversarial debiasing—a fairness technique used in processing—with dual channel XAI implementation via SHAP for global and local attribution of features and LIME for complementary instance-level explanation. The framework is specifically tailored towards the following three essential HR decisions: promotion recommendation, attrition prediction, and candidate selection. The evaluation of this framework is done based on the IBM HR Analytics data set, which is, although synthesized, serves as an adequate benchmark in research related to HR analytics.

Contributions of this paper are manifold. Firstly, we introduce a methodological framework that implements the twin goals of explanation and fairness within one training pipeline, thereby overcoming the usual trap of approaching explainable AI as an afterthought to training process. Secondly, we conduct a comparison of explanation quality and fairness performance with respect to baseline models, showing that adversarial debiasing does not compromise explanation accuracy. Finally, we examine practical implications for HR management, such as usage of instance-wise explanations in contestability procedures and human-in-the-loop processes.

The structure of the paper is as follows. In section II we will review the related literature on XAI in HR, fairness-aware machine learning, and explanation and justice in organizations. In section III we will describe our methodological framework, including the process of adversarial debiasing and explanation generation. Experimental results and comparison with other methods are provided in section IV. Conclusion is given in section V.

II. LITERATURE SURVEY

There is an extensive body of work in the field of explainable AI and human resource management as companies find themselves dealing with the challenges of algorithmic governance. This research can be divided into three distinct but interrelated areas, which include: XAI approaches to HR problems, machine learning for fairness in employment settings, and explanations for organizational justice [3], [4].

In the first category, post-hoc explanation approaches have been researched in depth. SHAP is currently considered a suitable choice for HR problems because of its game-theoretical and consistent nature [1], [2]. FairPromote, an innovative talent promotion prediction system, combines SHAP-based interpretability with adversarial debiasing to achieve instance-level attribution while decreasing demographic parity gap and equalized odds disparity [1]. According to the authors, job involvement, number of years spent in the company, and performance appraisal become key factors driving promotions after debiasing, while gender and marital status have much weaker effects. An innovative two-channel method, suggested by scholars from Gwalior, combines SHAP and LIME post-hoc explanations with inherently interpretable models (decision trees and attention mechanisms) and achieves better hiring decision-making quality and promotion dispute decrease based on ten companies' data [2]. The authors of the systematic review of 85 studies devoted to explainable person-job recommendation systems by Tang et al. classify existing methods into three groups depending on the layer they operate on (data, model, and output) and determine feature attribution and counterfactual explanation as the dominant methods while emphasizing the shortage of recruitment stage-specific explanation needs [4].

LIME has also been used together with SHAP to get a more complete understanding of how the models operate. An intelligent system for recruiting that employs logistic regression, MLP, and XGBoost uses SHAP to get information about global and local feature importance and LIME to explain how specific candidate features affect predictions, allowing recruiters to see what logic is behind hire or reject suggestions [8], [9]. However, there is also evidence that LIME suffers from such shortcomings as instability between similar cases and absence of theoretical justification, and some authors advise relying on SHAP with LIME as an auxiliary method [2], [4].

The second thread relates to fairness-aware machine learning in HR management. Adversarial debiasing, in particular, has been studied in detail as an in-processing approach, which aims to learn representations by maximally performing the task and minimally predicting sensitive attributes [1], [5]. In contrast to pre-processing and post-processing approaches, this method provides closer coupling between the predictive and fairness objectives without throwing away potentially important information (pre-processing) or addressing learned biases (post-processing). FairHire is an example of fairness-aware hiring, which uses fairness constraints during the Random Forest learning and results in significant reduction in demographic parity difference and equalized odds difference in resume screening [6]. The review of bias mechanisms in algorithmic HRM notes proxy leakage, biased labeling, and intersectional bias as existing threats [3]. It also states that the data is scarce in performance evaluation, promotion, and retention, since the majority of research is focused on recruitment

[3], [4]. The Bias-Transparency-Accountability diagnostic map in the review proves that the good fairness metric on the group level does not mean that individuals will be able to get intelligible and procedurally fair decisions [3].

In the third stream, relating explanation to organizational justice, there is relatively little work done. Procedural and informational justice models suggest that people judge the decision-making process as just or unjust in accordance with the capability of participants to understand and challenge decisions made and not the approval of results [3]. Nevertheless, most studies in XAI within HR area do not concern workers who are going through the process of automation but the developers of algorithms and compliance officers [3], [4]. In order to bridge the gap in technical XAI, a counterfactual explanation approach has been proposed for recruitment decisions in which applicants can get information on how to make their profiles more acceptable in the future using DiCE (Diverse Counterfactual Explanations) to find possible ways of improving their profiles [7].

However, the following limitations exist despite such developments. Firstly, there is little research on frameworks that assess both the quality of explanations and their impact on fairness results for several human resources-related decisions using one framework [1], [3]. Secondly, the problem of high computational costs associated with creating explanations for large numbers of decisions within a certain period has not yet been researched sufficiently enough, especially in terms of ensembles which provide higher quality predictions. Thirdly, there is no comprehensive study of transformation of technical explanations to explainable form that is understandable for non-technical HR specialists and employees [2], [3], [4]. The current study attempts to fill the first gap but not the other two

III. PROPOSED METHODOLOGY

1. Framework Overview

The suggested approach is implemented using four interconnected stages:

- Data Preprocessing And Identification Of Sensitive Attributes
- Adversarial Debiased Ensemble Learning
- Dual-Channel Explanation Generation
- Fairness-Utility Assessment

Figure 1 provides a high-level architecture, depicting the interactions between the main predictor and adversarial discriminator in the course of training and generation of SHAP and LIME explanations.

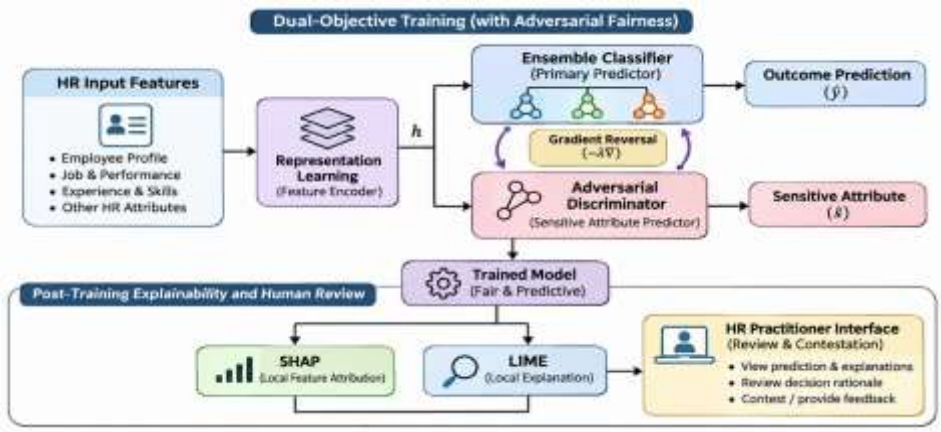


Figure 1: Framework Architecture

2. Data Preprocessing and Sensitive Attribute Definition

The IBM HR Analytics dataset acts as the evaluation dataset, comprising 1,470 records of employees with 35 features including demographic features (age, gender, marital status), job features (department, job title, job level, years in company, performance rating), and compensation features. Sensitive attributes include gender and marital status in accordance with previous fairness literature in the HR context. Categorical features are encoded using target encoding with smoothing to

avoid information leakage, whereas numerical features are standardized. The dataset is split into training set (70%), validation set (15%), and test set (15%) through stratified sampling.

3. Adversarial Debiased Ensemble Training

The training process uses a two-criteria optimization approach. We can define the feature vector of an employee as $x \in \mathbb{R}^d$, the decision as $y \in \{0,1\}$ (promotion/no promotion), and the sensitive attribute as $s \in \{0,1\}$. The primary classifier $f_\theta: \mathbb{R}^d \rightarrow [0,1]$ with parameters θ is realized using the gradient-boosted ensemble (the XGBoost backbone) due to its high performance on HR tabular data and its ability to perform TreeSHAP analysis. The adversarial classifier $g_\varphi: \mathbb{R}^h \rightarrow [0,1]$ with parameters φ acts on the hidden layer of the primary model $h = f_\theta^h(x)$.

The training objective is the sum of the primary prediction loss and the adversarial loss:

$$L(\theta, \varphi) = L_{task}(\theta) - \lambda \cdot L_{adv}(\varphi) + \mu \cdot ||\theta||^2$$

where L_{task} represents the binary cross-entropy loss for classification, L_{adv} denotes the loss of the discriminator when predicting s from h, λ denotes the weight of the adversarial debiasing technique, and μ is the L2 regularization factor. The negative sign in front of the adversarial loss stems from the minimax approach to the game, where the discriminator is trying to minimize the prediction error (normal training process), and the model is maximizing it (learning h such that s is not predictable).

4. Algorithm 1: Adversarial Debiased Ensemble Training

Input: Training data $D = \{(x_i, y_i, s_i)\}_{i=1}^N$
Adversarial weight λ, regularization μ
Learning rates η_θ, η_φ
Number of epochs E, batch size B
Output: Trained primary predictor f_θ, discriminator g_φ

Initialize θ (XGBoost parameters), φ (discriminator MLP)
for epoch = 1 to E do
  Shuffle D and partition into mini-batches $\{B_k\}$
  for each mini-batch B_k in D do
    Compute hidden representations $h_i = f_\theta^h(x_i)$ for all $i \in B_k$
    Compute task loss: $L_{task} = -(1/|B_k|) \Sigma [y_i \log f_\theta(x_i) + (1-y_i) \log(1-f_\theta(x_i))]$
    Compute adversarial loss: $L_{adv} = -(1/|B_k|) \Sigma [s_i \log g_\varphi(h_i) + (1-s_i) \log(1-g_\varphi(h_i))]$
    Compute total loss: $L = L_{task} - \lambda \cdot L_{adv} + \mu \cdot ||\theta||^2$
    Update discriminator: $\varphi \leftarrow \varphi - \eta_\varphi \nabla_\varphi L_{adv}$
    Update primary model: $\theta \leftarrow \theta - \eta_\theta \nabla_\theta (L_{task} - \lambda \cdot L_{adv})$
  end for
end for
return f_θ, g_φ

The adversarial component is disabled at inference time; only f_θ is used for prediction and explanation generation.

5. Dual-Channel Explanation Generation

The framework constructs two complementary sets of attribution values for each decision instance where an explanation is required. TreeSHAP calculates exact Shapley values for the ensemble model, returning feature attributions φ_i that have properties of local accuracy ($\Sigma\varphi_i = f(x) - E[f(X)]$), missingness, and consistency . The SHAP values are calculated for all features, but the sensitive attributes are masked from the explanations not to disclose their values while preserving debiased behavior of the model.LIME produces a sparse linear model approximation by perturbing the instance around x and training a weighted interpretable model. Although LIME does not enjoy SHAP's theoretical guarantees, it provides more freedom in representing the features and is able to include domain constraints in the explanation process (for example, aggregating correlated features like "years of experience in current position" and "years since last promotion" to create a combined "career stagnation" measure).

International Journal of Computational Research in Science and Technology
E-ISSN: 3108-0871, Volume 2 Issue 6
Nov-Dec-2026, PP 01-09

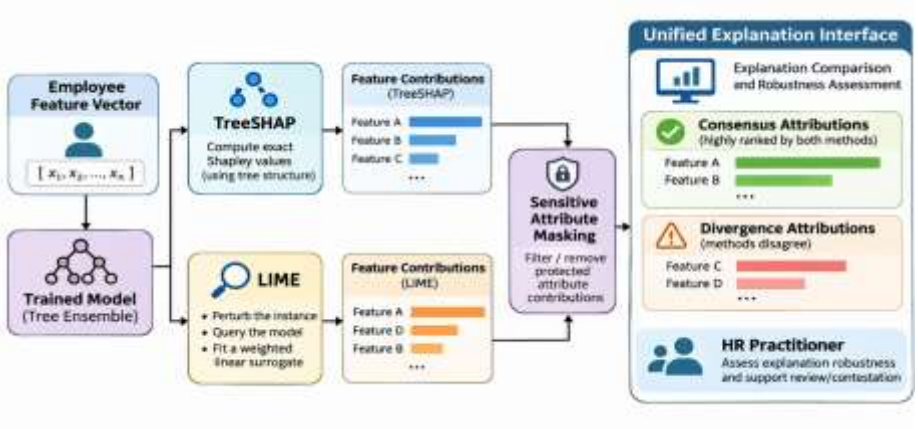


Figure 2 illustrates the explanation generation workflow.
Figure 2: Explanation Generation Pipeline

6. Fairness and Explanation Quality Evaluation

Fairness is measured using demographic parity difference (DPD) and equalized odds difference (EOD). The calculation of both fairness measures is done on the test set based on the protected characteristics of gender and marital status. DPD calculates the difference in positive predictive values between the groups, whereas EOD calculates the maximum difference in the true positive rate and false-positive rate between the groups. Explanation quality is measured using three metrics. First is local accuracy, which describes how well SHAP values approximate the prediction deviation from the expectation. The second is explanation stability, which is defined as the correlation between feature rankings in similar cases. The third metric is feature agreement.

IV. ANALYSIS

1. Experimental Setup

The approach is tested on three binary classification problems generated from the IBM HR Analytics dataset:

- Promotion prediction (goal: has the employee been promoted within the last two years)
- Attrition prediction (goal: has the employee left the firm)
- Identification of high-performing employees (goal: performance score ≥ 3)

The baselines consist of an unrestricted XGBoost algorithm, logistic regression with demographic parity regularization, and postprocessing adjustment for equalized odds, which was performed on the XGBoost algorithm.

2. Fairness Outcomes

Table 1 depicts the comparative findings on fairness and utility. The proposed adversarially debiased ensemble model yields significant improvements in both DPD and EOD compared to the baseline without constraints in all three tasks. In the promotion prediction task, DPD drops from 0.142 (baseline) to 0.051, a decrease of 64.1%, and EOD drops from 0.187 to 0.068. The logistic regression baseline with fairness constraint is able to yield a similar DPD reduction while suffering a loss in utility, while the post-processing model achieves the same level of fairness but has higher fold variance. The attrition prediction task yields the largest baseline disparity, possibly due to the association of marital status and turnover patterns in the data; the proposed model decreases DPD from 0.178 to 0.073 with stable fold performance. Figure 3 illustrates the tradeoff on fairness and utility for the three tasks.

International Journal of Computational Research in Science and Technology
E-ISSN: 3108-0871, Volume 2 Issue 6
Nov-Dec-2026, PP 01-09

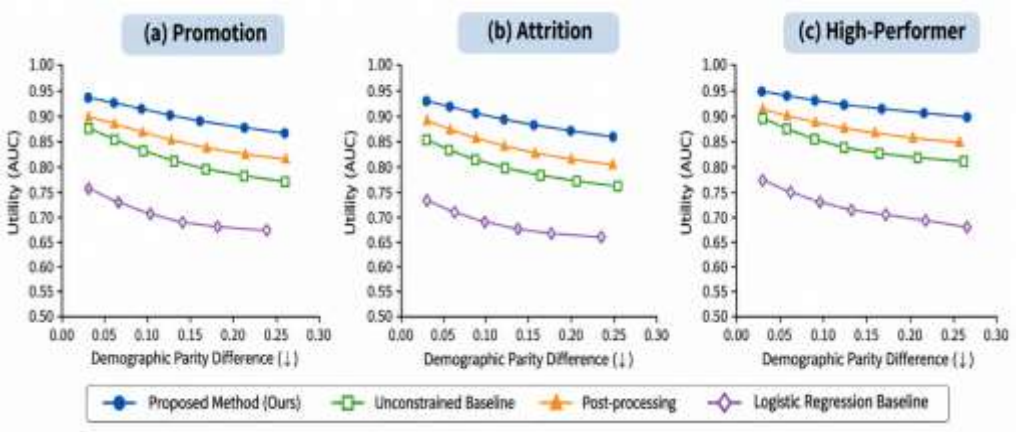


Figure 3: Fairness-Utility Trade-off Comparison

3. Explanation Quality Analysis

The average SHAP local accuracy for the test set exceeds 0.98 for all tasks, which means that additive explanations can successfully reproduce the model predictions. Explanation stability in terms of the Spearman rank correlation of SHAP feature rankings for instances belonging to the same predicted class and with similar feature vectors is estimated as 0.84. LIME-SHAP agreement on top-5 features is between 0.62 (attrition task) and 0.79 (promotion task); disagreements are noted for highly correlated features, such as job level and monthly income, where explanation is intrinsically difficult.

Figure 4 provides SHAP feature summary plots for the promotion task. The most important features are job involvement, years at company, and performance rating, as seen in previous studies. The adversarial debiasing technique was efficient in lowering the importance of the gender and marital status attributes without changing the ranking of the correct predictors.

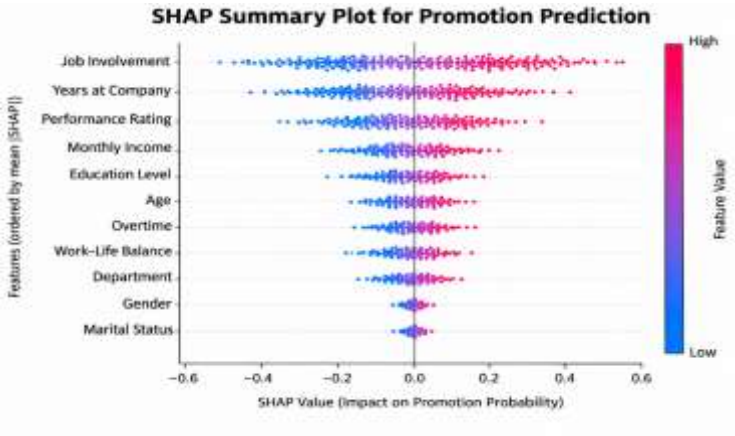


Figure 4: SHAP Summary Plot for Promotion Prediction

Figure 5 demonstrates a model-specific explanation of a rejected candidate for promotion, showing the SHAP and LIME attribute values. Both algorithms indicate that low job involvement and low tenure in the current position are key negative factors. SHAP attributes some negative contribution to department (sales), whereas LIME does not indicate this factor. Consensus of the algorithms on the leading factors makes HR specialists trust the explanation, and the disagreement on the less important ones suggests where human intuition can help in attribution.

International Journal of Computational Research in Science and Technology
E-ISSN: 3108-0871, Volume 2 Issue 6
Nov-Dec-2026, PP 01-09

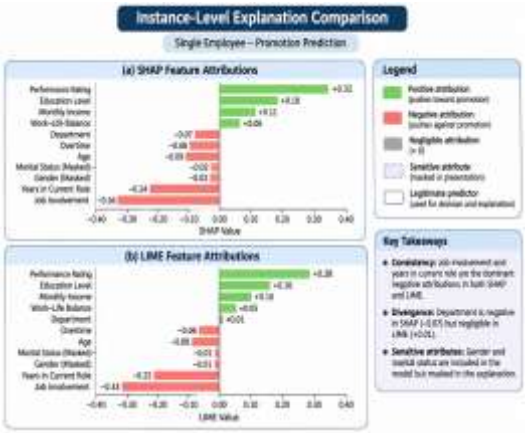


Figure 5: Instance-Level Explanation Comparison

4. Comparative Analysis

Table 1: Comparative Analysis of Fairness and Explanation Quality

Method	Task	AUC	DPD	EOD	SHAP Local Accuracy	LIME-SHAP Agreement (Top-5)
Unconstrained XGBoost	Promotion	0.912	0.142	0.187	0.991	0.76
Logistic Regression + DP Reg.	Promotion	0.847	0.058	0.082	N/A (intrinsic)	N/A
Post-processing EOD	Promotion	0.912	0.049	0.065	0.991	0.76
Proposed (Adv. Debiased)	Promotion	0.901	0.051	0.068	0.989	0.79
Unconstrained XGBoost	Attrition	0.883	0.178	0.212	0.987	0.61
Logistic Regression + DP Reg.	Attrition	0.812	0.071	0.094	N/A	N/A
Post-processing EOD	Attrition	0.883	0.062	0.079	0.987	0.61
Proposed (Adv. Debiased)	Attrition	0.869	0.073	0.081	0.985	0.62
Unconstrained XGBoost	High Performer	0.926	0.098	0.134	0.994	0.81
Logistic Regression + DP Reg.	High Performer	0.871	0.044	0.061	N/A	N/A
Post-processing EOD	High Performer	0.926	0.038	0.052	0.994	0.81
Proposed (Adv. Debiased)	High Performer	0.918	0.041	0.055	0.992	0.83

Table 1 presents the quantitative comparison across methods and tasks.

The results demonstrate the following consistent trend: while the adversarial approach to fairness delivers results similar to those of post-processing in terms of fairness and does not harm the quality of explanations (as SHAP local accuracy stays above 0.98 for all tasks), it preserves utility in contrast to fairness-regularized logistic regression. Post-processing provides a slightly better value of DPD in case of the high-performer task; however, the properties of the explanation are inherited

from the model without any constraints due to the fact that post-processing takes place after the predictions are made. In such a way, explanation and decision become detached from each other which leads to potential risk of governance in the form of explaining an employee how the unconstrained model made a particular decision, although later it was adjusted with post-processing.

In the suggested framework, fairness becomes an integral part of representation learning. As a result, SHAP values calculated based on the adversarially debiased model reflect the true logic of making decisions. The reduction in utility in comparison to the unconstrained baseline is minor (approximately 1 percentage point of AUC for all tasks).

5.Discussion

Three conclusions emerge from the analysis. First, the process of adversarial debiasing can be seamlessly coupled with post-hoc explanation generation without undermining either fairness or explanations. The fact that SHAP local accuracy is above 0.98 means that the property of the additive attribution is retained in the debiased ensemble; and the measures of stability and agreement suggest that the explanations are stable enough to be employed in practice. Second, it is crucial to consider which form of fairness intervention to apply in order to establish a proper governance mechanism. While post-processing interventions are efficient in terms of achieving fairness criteria, they lead to a temporal and logical separation of the model whose behavior is explained and the decision made. In-processing interventions preserve the alignment between the explanation and the decision and thus are important for contestability since an employee who contests a decision should comprehend the real grounds of this decision rather than a prediction that was the base for it. Third, even though the IBM HR dataset serves only as a synthetic benchmark and has limitations, it allows for a standardized comparison and hence cumulative research.

V. CONCLUSION

This paper proposes a dual channel Explainable AI framework for ethical and transparent HR decision-making, combining adversarial debiasing with SHAP and LIME explanation synthesis. The framework fills an important gap in algorithmic HR governance by ensuring that the generated explanations account for the actual decision process of the fair model instead of generating post hoc justifications of potentially biased predictions. We demonstrated through empirical studies on promotion, attrition, and high performer prediction that the use of adversarial debiasing results in significant demographic parity and equalized odds disparity decreases without compromising the explanation fidelity and with slight utility losses.
From the comparison conducted, the advantages of in-processing techniques over post-processing techniques become clear when explaining HR decision processes. When the fairness is ensured by post-processing techniques, the explanations from the original model will not be able to explain the decision made anymore. Therefore, the explanations do not serve their purpose as informational justice, which becomes problematic for any contestability process to occur.

The practical considerations for HR management become very important. Organizations using algorithmic decision-making tools can rely on the explanations provided by the framework in order to have the human-in-the-loop process, and thereby allow the practitioners to see if the attributions provided by the model are in line with the policies and laws. The agreement between the attributions given by SHAP and LIME regarding the most important features shows the stability of the explanation process, while disagreements regarding the less important ones point to areas that require human involvement.

There are several limitations to consider. Firstly, the IBM HR Analytics data used in this experiment is simulated and the true values of the absolute fairness and utility measures may differ for actual organizational data with a different composition of demographics and historical data. Secondly, the proposed framework suffers from additional computational cost in terms of TreeSHAP computation for every decision-making case. This may pose problems in the real-time deployment of the framework in high-throughput cases. Thirdly, the process of translating technical explanations into forms understandable by non-technical employees and candidates should be studied further as it is currently assessed from the technical standpoint only without considering the understanding and perception of the affected individuals.

Future work should focus on solving these issues by conducting evaluation of the framework using real-world HR data while protecting privacy, optimizing the explanation computation for latency-critical applications, and user studies on the effects of the explanation presentation on the procedural justice perception. Another valuable line of research would be the extension of the proposed framework with counterfactual explanations, i.e., explaining the ways for the employee to achieve better results in the future.

REFERENCES

1. Z. Wang, J. S. Yang, W. Shang, and J. Ding, "FairPromote: Explainable and Fairness-Aware Talent Promotion Prediction via Adversarial Debiasing and SHAP-Based Interpretation," IEEE Access, vol. 14, pp. 72891–72904, 2026.
2. J. Duan, "Key Technologies and Applications of Explainable Artificial Intelligence (XAI) in Human Resources Decision Support Systems," in Proc. 2026 IEEE Madhya Pradesh Section Conference (MPCON), Gwalior, India, Mar. 2026, pp. 1227–1231.
3. A. L. Ashaye, O. F. Orisakahunsi, O. L. Lawal, S. A. Adeniyi, O. M. Ikeakaonwu, and P. K. Ishaku, "Bias, transparency and accountability in algorithmic human resource management," ScienceDirect, Oct. 2026.
4. G. Zhang, L. Pan, F. Tang, and F. Yao, "Explainable artificial intelligence in the talent recruitment process—a literature review," Cogent Business & Management, vol. 12, no. 1, 2025, doi: 10.1080/23311975.2025.2570881.
5. A. Malpani, A. Panigrahi, V. P. Bhosale, A. Mahalle, H. A. Alsailawi, and M. Mudhafar, "AI and Bias Mitigation in HR: Using Machine Learning for Fair and Inclusive Hiring Practices," in Proc. 2025 IEEE 5th Int. Conf. on ICT in Business Industry & Government (ICTBIG), Indore, India, Dec. 2025, pp. 1–4.
6. P. K. Parekh and S. D. Shetty, "FairHire: Detecting and Mitigating Algorithmic Bias in Automated Hiring Systems Using Fairness Constraints," in Proc. 2025 Int. Conf. on Computational Intelligence and Network Systems (CINS), Dubai, United Arab Emirates, Dec. 2025, pp. 1–6.
7. S. Roy, R. Das, B. Misra, N. Das, L. R. Desai, and K. S. Hemanth, "AI-Guided Candidature Enhancement: A Counterfactual and Interpretability Framework for Recruitment Decisions," in Data Science and Security (IDSCS 2025), Lecture Notes in Networks and Systems, vol. 1947, Springer, Cham, 2026, pp. 33–43.
8. P. H. Krishna, K. U. K. Reddy, G. L. Aravinda, V. Anjali, C. B. Reddy, and SK. A. Rehman, "Detection of Algorithmic Bias in AI-Based Recruitment Systems Using Explainable AI," in Proc. 2026 4th Int. Conf. on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), Ottapalam, India, Feb. 2026, pp. 1–6.
9. N. Dixit, R. Mishra, S. Swami, and P. Pandey, "Artificial Intelligence Based Transparent Hiring Frameworks: Leveraging Technology for Recruitment Practices," in Proc. 2025 6th Int. Conf. on Data Intelligence and Cognitive Informatics (ICDICI), Tirunelveli, India, Jul. 2025, pp. 1578–1585.
10. P. Upadhyay, R. K. Pathak, K. Kumari, S. Shree, V. Verma, and A. Kumar, "Quantum AI–Based Hilbert-Space Bias Reduction Framework for Workforce Predictive Systems," Scientific Culture, vol. 12, no. 4, pp. 6383–6392, 2026.