



AI-Powered Employee Attrition Prediction and Workforce Intelligence System

Shreya Gaikwad¹, Omkar Bhosle², Prof. Prajкта Puranik³, Mr. Vaibhav Srivastava⁴

ISBM College of Engineering, Pune, Maharashtra, India

ABSTRACT :

Employee attrition represents a significant financial and operational challenge for modern enterprises. While traditional Human Resources (HR) departments rely on descriptive analytics, recent advancements have shifted focus toward predictive machine learning (ML) models. However, standard ML approaches often function as "black boxes," predicting flight risk without providing actionable or interpretable context, rendering them impractical for HR interventions. This paper proposes an end-to-end Explainable Decision Support System (DSS) that transitions HR from predictive to prescriptive analytics. We evaluated three baseline algorithms (Logistic Regression, Random Forest, and XGBoost) on a custom 12-feature employee sentiment and workload dataset. While all models achieved high baseline accuracy (>95.5%), XGBoost (Accuracy: 95.45%, F1-Score: 95.71%) was selected as the core predictive engine due to its superior capacity to model non-linear human behavioral factors and its native integration with SHapley Additive exPlanations (SHAP). By integrating the XGBoost-SHAP pipeline into a web architecture, the proposed DSS features an automated HR triage queue, local interpretability plots, and a dynamic "What-If" scenario simulator. Furthermore, the system integrates a Generative AI action agent to automatically synthesize targeted retention strategies. This paper represents the Phase I prototype of the system, establishing the core algorithmic architecture.

Keywords: Employee Attrition, Explainable AI, XGBoost, SHAP, Generative AI, Human Resources, Decision Support System.

Introduction

The financial and operational impacts of employee turnover in modern enterprises are substantial. Historically, HR departments have relied on descriptive analytics (e.g., historical turnover dashboards) and reactive strategies. The evolution of predictive machine learning has allowed organizations to forecast flight risk; however, HR managers cannot trust or act on standard probability scores without explainability.

The primary problem this research addresses is the gap between predictive ML output and prescriptive HR action. To bridge this gap, we propose an AI-powered Workforce Intelligence System. This paper details the Phase I development of this system, focusing on building a robust, explainable pipeline utilizing XGBoost, SHAP, and Large Language Models (LLMs) to not only predict who will leave, but explain why, and prescribe how to retain them.

1.1. Research Goal and Objectives

The primary goal of this research is to develop a transparent, end-to-end Decision Support System (DSS) that transitions HR analytics from purely predictive modeling to prescriptive, actionable intelligence. To achieve this, the project defines the following specific objectives:

- To architect a high-accuracy classification pipeline utilizing the 12 most actionable HR metrics to predict individual attrition probability.
- To integrate Explainable AI (SHAP) to isolate the localized, employee-specific drivers of flight risk, overcoming the "black-box" limitations of traditional ML.
- To construct an interactive "What-If" simulator that allows HR management to dynamically test the impact of workplace interventions on risk scores in real-time.
- To orchestrate a Generative AI Action Agent that autonomously synthesizes targeted retention strategies and manager briefing emails based on the simulated diagnostic data.

Literature Review

Recent advancements in human resource analytics have increasingly transitioned from traditional statistical evaluation to robust machine learning frameworks for predicting employee turnover, with foundational studies emphasizing data-driven mitigation and baseline predictive accuracy [3, 9, 12]. To improve classification performance on highly imbalanced HR datasets, researchers have demonstrated that ensemble learning techniques—particularly Extreme Gradient Boosting (XGBoost) and stacked ensembles—significantly outperform linear algorithms in capturing complex, non-linear behavioral interdependencies [1, 7, 11]. However, the inherent "black-box" nature of these advanced models historically limited their practical utility for HR managers, prompting a critical shift towards Explainable AI (XAI). Recent literature extensively validates the integration of SHAP (SHapley Additive

exPlanations) to provide both global feature importance and localized risk scoring, enabling transparent managerial decisions [4, 5, 6, 15]. Despite these diagnostic advancements, current research largely stops at risk identification [20], leaving a significant gap for interactive Decision Support Systems (DSS). This paper addresses that gap by moving beyond diagnostic XAI, combining real-time simulation dashboards with Large Language Model (LLM) orchestration to autonomously synthesize prescriptive retention strategies [21, 24, 27].

Proposed Methodology

1.2. Dataset Engineering and Feature Selection

To ensure reproducibility and eliminate privacy constraints associated with proprietary employee records, this study evaluates the proposed architecture using a standardized, synthetic corporate benchmark based on the IBM HR Analytics framework. To simulate robust enterprise-scale operations, the dataset was expanded to comprise 10,000 employee records.

The original framework contains 35 distinct features, many of which are static demographic data (e.g., Age, Marital Status). Relying on unalterable demographics introduces noisy variables that managers cannot control. Therefore, rigorous dimensionality reduction was performed to engineer a highly curated subset of 12 actionable features. Eleven of these features capture psychological and environmental satisfaction on a standardized 1-5 ordinal scale, while the final feature captures quantitative workload capacity (OvertimeHours). Categorical variables were transformed via one-hot encoding, and the dataset was partitioned into an 80% training set (8,000 records) and a 20% testing set (2,000 records) to ensure rigorous evaluation without data leakage.

Table 1 - The 12 Dynamic Features of the Proposed DSS

Feature Name	Data Type	Description
JobSatisfaction	Ordinal (1-5)	Self-reported satisfaction with daily tasks
WorkLifeBalance	Ordinal (1-5)	Perception of personal time boundaries
Workload	Ordinal (1-5)	Manageability of current project volume
StressLevel	Ordinal (1-5)	Current psychological and operational stress
ManagerSupport	Ordinal (1-5)	Effectiveness of direct leadership
RecognitionAppreciation	Ordinal (1-5)	Frequency of positive reinforcement
CareerGrowthSatisfaction	Ordinal (1-5)	Alignment with long-term career goals
LearningDevelopmentOpportunities	Ordinal (1-5)	Availability of upskilling resources
CompensationSatisfaction	Ordinal (1-5)	Satisfaction with current salary and benefits
WorkplaceEnvironment	Ordinal (1-5)	Physical/remote workspace culture
JobRoleSatisfaction	Ordinal (1-5)	Alignment between expected and actual duties
OvertimeHours	Continuous	Average weekly overtime (capped at 30 hours)

1.3. Algorithm Selection and Predictive Modelling

To justify the predictive engine, a comparative analysis of three algorithms—Logistic Regression, Random Forest, and Extreme Gradient Boosting (XGBoost)—was conducted. While Logistic Regression achieved a marginally higher raw accuracy (95.60%), XGBoost (Accuracy: 95.45%, F1-Score: 95.71%) was implemented as the optimal core engine. XGBoost operates by iteratively training decision trees to correct the residual errors of prior trees, making it mathematically superior for capturing the complex, non-linear interdependencies inherent in HR behavioral data (e.g., the compounding effect of high workload and low manager support). The model computes a continuous probability score (P) for employee attrition risk, establishing a strict managerial triage threshold at $P > 0.5$.

1.4. Explainable AI (XAI) Integration

To overcome the "black-box" limitations of standard classification models, the predictive layer is natively coupled with SHapley Additive exPlanations (SHAP). Specifically, a TreeExplainer is applied directly to the XGBoost output to calculate the marginal contribution of each feature for individual employee predictions (local interpretability). This calculation allows the system to instantly rank and isolate the top three root causes of flight risk for any given employee, forming the diagnostic foundation required for the Decision Support System (DSS).

1.5. Explainable AI (XAI) Integration

The final phase translates diagnostic data into prescriptive action through an interactive Streamlit web architecture. The system maps the top three SHAP drivers to dynamic UI sliders, enabling HR managers to adjust parameter values (e.g., reducing OvertimeHours from 30 to 10). These modified arrays

are passed back through the XGBoost engine to simulate hypothetical workplace interventions and recalculate flight risk in real-time. Finally, a programmatic orchestration layer concatenates the baseline risk, isolated SHAP drivers, and simulated targets into a structured prompt. This data is ingested by a Generative AI endpoint, which autonomously synthesizes a professional, context-aware retention action plan for immediate deployment.

1.6. End-to-End System Workflow




Fig. 1 - System architecture and data flow workflow of the proposed AI-powered Decision Support System.

The complete architecture operates through six sequentially integrated layers:

- 1. Data Input Layer (Streamlit UI):** Employees submit their 12-survey metrics via a secure portal.
- 2. Predictive Layer (XGBoost Engine):** The classification model processes the feature array to calculate the Attrition Probability (P).
- 3. Diagnostic Layer (SHAP Explainer):** A TreeExplainer isolates and visualizes the top local risk drivers impacting the prediction.
- 4. Triage & Visualization (HR Dashboard):** The system queues profiles into risk categories (High Risk if $P > 0.5$, Stable if $P \leq 0.5$) based on the diagnostic insights.
- 5. Simulation Layer (What-If DSS):** HR can dynamically adjust parameter controls to recalculate risk probabilities based on hypothetical interventions.
- 6. Prescriptive Layer (GenAI Action Agent):** The system uses LLM orchestration to analyze the simulation targets and generate a targeted manager briefing email.

System Implementation

1.7. Development Environment and Software Stack

The end-to-end architecture was developed using Python, selected for its extensive ecosystem in both machine learning and rapid web application prototyping. Data manipulation and preprocessing of the 10,000-record dataset were executed using pandas and numpy. The predictive pipeline utilized scikit-learn for data partitioning and the xgboost library for model training. The diagnostic layer was implemented using the shap framework. The interactive frontend and backend orchestration were entirely built and deployed using streamlit.

1.8. Model Training and Persistence

During the offline training phase, the XGBoost classifier was trained on the 8,000-record training split, specifically optimized to minimize false negatives in flight-risk identification. Once evaluated against the 2,000-record testing split, the final model and the fitted SHAP TreeExplainer object were serialized and saved as static artifacts (e.g., via joblib or pickle). This architectural decision ensures that the live application infers predictions from pre-compiled models rather than retraining on the fly, guaranteeing the sub-second latency required for the interactive simulator.

1.9. Application Orchestration and State Management

The live Decision Support System (DSS) relies heavily on Streamlit's session state management. When an employee profile is loaded into the HR Dashboard, the system initializes a session state containing their specific 12-metric array. As a manager interacts with the What-If Simulator sliders,

Streamlit captures these real-time adjustments (e.g., modifying the OvertimeHours integer). The application instantly passes the modified array back through the serialized XGBoost and SHAP artifacts, recalculating the Attrition Probability (\$P\$) and rendering updated visualizations without requiring page reloads.

1.10. Generative AI Endpoint Integration

The Prescriptive Layer was implemented by integrating an external Large Language Model (LLM) API. A programmatic prompt template was engineered within the Python backend to automatically concatenate the system's outputs. When the user triggers the Action Agent, the system extracts the baseline risk score, the top three localized SHAP drivers, and the final simulated target values from the Streamlit session state. These variables are injected into the prompt and passed to the LLM, which is constrained by system instructions to output a strictly formatted, professional retention strategy email.

Experimental Results and Discussion

1.11. Predictive Performance

The XGBoost model successfully minimized false negatives—a critical business metric in HR, as failing to identify an at-risk employee is far more costly than falsely flagging a stable one. The ROC-AUC (Receiver Operating Characteristic - Area Under Curve) analysis confirmed the model's high capability in distinguishing between the two classes (Attrition vs. Retention), proving its reliability as the foundational layer of the DSS.




Fig. 2 - Receiver Operating Characteristic (ROC) curve comparing the predictive performance of baseline algorithms.




Fig. 3 - Confusion matrix evaluating the final XGBoost predictive classification model.

1.12. Explainability and Prescriptive Analytics

The application of SHAP successfully un-black-boxed the XGBoost predictions. Global SHAP summary plots confirmed that dynamic variables—specifically Overtime Hours, Stress Level, and Job Satisfaction—were the primary drivers of attrition risk across the dataset. By feeding these isolated, localized SHAP values into the Streamlit UI, the system successfully demonstrated real-time prescriptive capabilities. During simulation tests, adjusting a single high-impact variable (e.g., reducing OvertimeHours from 30 to 10) correctly propagated through the XGBoost engine, dynamically lowering the predicted flight risk and instantly triggering the GenAI Action Agent to output a modified retention email.




Fig. 4 - SHAP summary plot illustrating global feature importance and their impact on the model's attrition risk output.

1.13. System Implementation and User Interface




Fig. 5 - Streamlit data input layer displaying the bi-monthly feedback survey utilized for real-time ingestion of the 12 dynamic evaluation metrics.



Fig. 6 - Explainable AI (XAI) diagnostic dashboard illustrating a local SHAP waterfall plot to isolate the top localized drivers of an individual's attrition risk.




Fig. 7 - Decision Support System (DSS) What-If Simulator enabling dynamic parameter adjustments to test targeted workplace interventions and project simulated risk outcomes.




Fig. 8 - Generative AI Action Agent user interface synthesizing a customized retention strategy and manager briefing email based on the simulated intervention data.

Conclusion and Future Scope

1.14. Conclusion

This research successfully developed an end-to-end Explainable Employee Attrition Prediction and Workforce Intelligence System. By transitioning HR analytics from diagnostic observations to prescriptive interventions, the proposed architecture bridges the gap between predictive ML models and actionable HR management. The integration of XGBoost, SHAP, and a Generative AI Action Agent provides a novel, transparent framework that enables managers to simulate retention strategies and automate intervention planning.

1.15. Future Scope (Phase II)

This paper represents the Phase I prototype developed during the current academic semester. The project will be expanded in the subsequent semester (Phase II) to address current limitations and scale the architecture. Future enhancements will include:

- Migrating the static Streamlit session state to a persistent, cloud-hosted relational database (e.g., PostgreSQL or SQLite).
- Implementing a financial ROI (Return on Investment) calculator to quantify the cost-savings of the AI-generated retention strategies versus the cost of employee replacement.
- Conducting live user acceptance testing (UAT) with HR professionals to refine the GenAI prompt engineering and validate system usability.

REFERENCES

- [1] K. Konar, S. Das, S. Das, and S. Misra, "Employee Attrition Prediction Using Bayesian Optimized Stacked Ensemble Learning and Explainable AI," SN Computer Science, vol. 6, 2025.
- [2] "HR Analytics: Predicting Employee Turnover using Machine Learning Models," IEEE ICDISS, 2025.
- [3] N. E. Vijayan, "Mitigating Attrition: Data-Driven Approach Using Machine Learning and Data Engineering," arXiv, 2025.
- [4] M. Al-Ali et al., "Integrating Machine Learning and Explainable AI for Employee Attrition Prediction in HR Analytics," Scientific Reports, 2026.
- [5] H. Jia, "Explainable AI for Employee Turnover Prediction: A SHAP-Based Intelligent Analytics Approach," Scientific Reports, 2026.

- [6] M. S. Pavithran and S. M. Vadivel, "Explainable Attrition Risk Scoring for Managerial Retention Decisions in Human Resource Analytics," *Frontiers in Big Data*, 2026.
- [7] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016.
- [8] S. M. Lundberg and S. I. Lee, "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774, 2017.
- [9] M. Fallucchi, M. Coladangelo, R. Romeo, and E. W. De Luca, "Predicting Employee Attrition Using Machine Learning Techniques," *Computers*, vol. 9, no. 4, p. 86, 2020.
- [10] S. Jain, A. Nayyar, and F. S. Al-Mubarak, "Predictive Analytics in Human Resource Management: A Review of Machine Learning Approaches," *IEEE Access*, vol. 10, pp. 45672–45685, 2022.
- [11] D. K. Sharma and S. Gupta, "A Decision Support System for Workforce Planning using Extreme Gradient Boosting," *Decision Support Systems*, vol. 165, 2023.
- [12] Y. Wang, "Machine Learning for Human Resource Management: A Comprehensive Review," *IEEE Transactions on Computational Social Systems*, vol. 10, pp. 1245–1258, 2023.
- [13] P. R. Kumar and V. Ravi, "Employee Attrition Prediction using Ensemble Learning and Deep Learning," *Neural Computing and Applications*, vol. 35, pp. 8921–8935, 2023.
- [14] J. H. Choi, Y. Kim, and S. Park, "Beyond Prediction: Actionable HR Analytics using Causal Machine Learning," *Journal of Applied Psychology*, vol. 109, no. 2, pp. 210–225, 2024.
- [15] R. P. Srivastava and S. K. Singh, "Explainable AI in Employee Retention: Bridging the Gap Between Analytics and HR Strategy," *Expert Systems with Applications*, vol. 238, 2024.
- [16] A. Q. O'Brien and H. Kim, "Generative AI in Human Resources: Opportunities and Ethical Considerations," *Journal of Business Research*, vol. 172, 2024.
- [17] H. N. Nguyen and T. Tran, "Integration of Explainable AI and Streamlit for Real-Time Predictive Modeling," *SoftwareX*, vol. 25, 2024.
- [18] E. T. C. Goh, "Visualizing Machine Learning Models with SHAP: A Guide for Non-Technical Stakeholders," *Information Visualization*, vol. 23, no. 1, pp. 45–59, 2024.
- [19] C. F. Chien et al., "A Data-Driven Framework for Employee Turnover Prediction and Proactive Retention," *Computers & Industrial Engineering*, vol. 188, 2024.
- [20] A. Patel and R. Mehta, "From Diagnostic to Prescriptive Analytics: The Next Frontier in Human Resource Management," *Harvard Data Science Review*, vol. 7, no. 1, 2025.
- [21] K. E. Smith and T. Jones, "Large Language Models as Action Agents in Enterprise Systems," *ACM Transactions on Intelligent Systems and Technology*, vol. 16, no. 3, 2025.
- [22] S. K. A. Yildiz and A. B. Yildiz, "The Impact of Generative Artificial Intelligence on Human Resource Decision Support Systems," *Technological Forecasting and Social Change*, vol. 210, 2025.
- [23] L. V. Pratt and J. E. Adams, "Interactive Dashboards for HR Analytics: Evaluating the Impact of Real-Time Data on Managerial Decision Making," *Computers in Human Behavior*, vol. 162, 2025.
- [24] M. Zhao, C. Wang, and J. Liu, "Automated Retention Strategy Generation via LLM-Driven Decision Support," *International Journal of Human Resource Studies*, vol. 15, no. 2, 2026.
- [25] A. Sharma, R. K. Singh, and M. Gupta, "Integrating Large Language Models with Tabular HR Data for Automated Prescriptive Analytics," *IEEE Transactions on Artificial Intelligence*, vol. 6, pp. 312–325, 2025.
- [26] E. Martinez and L. Huang, "Mitigating Algorithmic Bias in Employee Turnover Prediction Models," *Journal of Business Ethics*, vol. 195, pp. 45–62, 2025.
- [27] P. K. Das and S. Mohanty, "Enterprise Deployment of Machine Learning Web Applications: A Framework for Interactive DSS," *Software: Practice and Experience*, vol. 55, no. 3, 2024.
- [28] K. T. Lee, "Hyperparameter Optimization for XGBoost in Highly Imbalanced Human Resource Datasets," *Pattern Recognition Letters*, vol. 178, pp. 112–119, 2024.
- [29] F. Ricci, M. Rossi, and A. Bianchi, "From Reactive to Proactive: A Cost-Benefit Analysis of AI-Driven Employee Retention Interventions," *European Journal of Operational Research*, vol. 320, no. 1, pp. 288–301, 2025.
- [30] S. Verma and T. O. Ahsan, "MLOps for Human Resource Analytics: Continuous Monitoring and Model Retraining Strategies," *ACM Computing Surveys*, vol. 58, 2026.