AI-Powered HR Assistant for Grievance Redressal, Policy Guidance, Onboarding, and Attrition Risk Management

Aditya Trivedi [1] Kaiwan Vaghchhipawala[1] Prathamesh Nadkarni[1] Shalini Wadhwa[1*]

[1] Department of Technology Management, NMIMS University, MPSTME, Mumbai, Maharashtra, India
[*] Corresponding Author
shalini.wadhwa06@gmail.com

Abstract. Organisations require HR solutions that are not only reactive but also transparent and proactive throughout the employee lifecycle. This paper introduces an integrated HR assistant merging sentiment-aware grievance analysis, authoritative policy advice, a practical onboarding suite, and explainable attrition risk scoring into a unified modular platform. Grievance management makes use of NLP to identify sentiment, categorize types of complaints, and expose priority suggestions. Policy direction uses a Retrieval Augmented strategy to provide grounded, auditable snippets from the HR policy handbook. Onboarding module creates customized 30/60/90-day plans, conducts mentor/buddy matching through embedding similarity, suggests microlearning material, and assists with regular new-hire pulse checks to track early sentiment. Attrition risk is quantified with supervised models and survival analysis; model predictions are combined with SHAP-based explanations to render feature-level drivers understandable to HR stakeholders. The design prioritizes a human-in the-loop workflow for escalation, privacy-conscious persistence for pulse data, and modular components that are easily integrated with company documents and LMS systems. Although attrition components in the present prototype are trained on synthetic samples, the pipeline is designed to be fed real HR datasets and company policy PDFs, allowing for easy shift to production datasets. This article illustrates how RAG, explainable ML, and conversational UX can all together promote HR from reactive casework to active, evidence-based employee care.

Keywords: AI-powered HR assistant, Grievance redressal, Policy guidance, Retrieval-Augmented Generation (RAG), Onboarding, Attrition prediction, Explainable AI (SHAP), employee sentiment.

1 Introduction

Artificial Intelligence (AI) is revolutionizing Human Resource Management (HRM) from an administrative support role to a strategic, evidence-based discipline that defines organizational competitiveness. (Agrawal, 2025) Manual and reactive traditional HR practices are being replaced by smart systems with the ability to interpret data, recognize behavioural patterns, and offer insightful conclusions. This evolution reflects a broader evolution towards evidence-based management of the workforce, where data - rather than instinct - drives decision-making. (Kumar *et al.*, 2024)

IICTDS 2025, December 19-20, Chandigarh, India
Copyright © 2026 EAI
DOI 10.4108/eai.19-12-2025.2366267

HR systems must be adaptable, ethical, and human-centred. Creating this environment is supported with advances in Artificial Intelligence, which combines Natural Language Processing (NLP), Machine Learning (ML), and predictive analytics to automate, measure sentiment, and predict workforce behaviours. (Maar *et al.*, 2026) Accepting the mix of automation and moral responsibility is taking hold with the increasing acceptance of the "human-in-the-loop" systems whereby humans review and vet outcomes produced from AI systems. (Benabou and Touhami, 2025) Although there has been much progress in using AI for HR, it is still largely restricted to isolated functions such as hiring or predicting performance. Clearly, the next step in this development process is to adopt an integrated HR system for the employee lifecycle that includes onboarding, policy advice, grievance redressal, and attrition risk assessment. (Tambe *et al.*, 2019) The next generation of HR technology is headed towards building intelligent assistants that help deliver contextualized, explainable, and proactive advice across a wide range of HR functions. (Chukwuka and Dibie, 2024)

This research extends this trend with the design of an AI-enabled HR assistant that combines four functions: grievance redressal, policy advice, onboarding, and attrition risk assessment.

Research Objectives:

1. To create an integrated AI-based human-resources system based on grievance redressal, policy advice, onboarding, and attrition risk prediction systems framework into one overall system.
2. To create and evaluate interpretable machine-learning models and sentiment modules based on NLP that increase transparency and trust in decision-making as part of a human-in-the-loop framework.
3. To operationalize Retrieval-Augmented Generation into a reliable and accurate response to HR policies.
4. To investigate how human-in-the-loop architecture functions to further fairness, accountability, and user satisfaction within AI-assisted decisions in HR.

2 Literature Review

The integration of AI into human resource management has transformed the field from the automation of basic administrative functions to providing strategic workforce intelligence; researchers are also focused on explainability, transparency, and collaborative human AI decision-making models. (ATTHOS *et al.*, 2025) Existing systems of anonymous reporting and automated case management show that AI can assist organizations in organizational documentation and addressing biases in investigations. (Kalluri, 2025) Hybrid human AI grievance models that combine efficiency with empathy appear to achieve more accuracy and employee satisfaction. (Joshi, 2025; Gupta *et al.*, 2025)

AI has also transformed policy guidance through chatbots that respond to employee questions related to benefits, leave, and other organizational policies with far less

administrative effort. (Nimmagadda *et al.*, 2024) With the introduction of large language models, these policy assistants can also understand complex queries, eliminating keyword matching. (Kovari, 2025) Even so, generative models may give a false or misleading reply, indicating that the hallucination of output is a legitimate concern. To combat this issue, many modern human resources systems use retrieval augmented generation, meaning that when the model generates content, is grounded in a specific policy document or occurrence that has been validated. (Chen *et al.*, 2024) Additionally, the output of generative models may be well-supported by providing ways for employees to source their guidance to specific clause, so the generative output has increased credibility. (Silva and Canedo, 2025) However, AI supports still struggle with emotionally complex or ambiguous questions. (Suhonen, 2025)

When onboarding employees, AI can generate a more personalized and more engaging experience in learning paths, assignments or simulations and actual workflows based on that individual employee's role and existing skill so they have a more congenial experience, with relevant learning paths. (Rane, 2023) Personalized onboarding can enhance levels of satisfaction and retention, as employees have reduced repetitive tasks and offered relevant opportunities to develop specifically their skills for the position, rather than an overall generic experience. (Ahmed El Garem, 2024) Onboarding and instant enterprise search are also gaining traction in the onboarding processes. (Kumar et al., 2025; Shaheen *et al.*, 2025)

Effective management of attrition risk remains one of the most developed areas of AI in HRM. (Mohammad et al., 2025) Different machine learning models are consistently found to outperform standard logistic regression in predictive ability and accuracy with regards to turnover. These models include Random Forest and XGBoost, as they can capture non-linear patterns in the workforce population. (Baydili and Tasci, 2025; Mohbey, 2022) In the more recent research samples of studies, they are also integrating survival analysis measuring not only who will leave, but when attrition is most likely to occur. (Alqahtani *et al.*, 2024) Explainable AI models and methods, have gained prominence in this space, as they provide greater transparency to managers. (Ghavami and Tabatabaee, 2023; Mansor *et al.*, 2021) Refer to Fig. 1 for Algorithms and technologies used over the years in HR Systems.

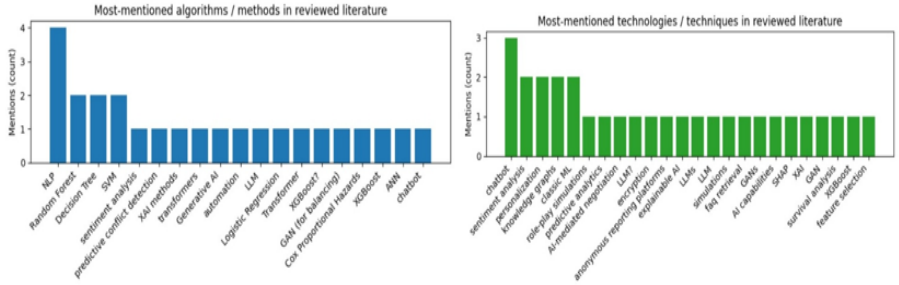


Fig. 1. Insights from Literature Review

3 Methodology

The methodology for this research aims to design, deploy, and test an AI-based HR assistant that combines grievance analysis, policy advice, onboarding, and attrition forecasting into an integrated system. The architecture follows the standards of transparency, modularity, explainability, and human monitoring so that automation supports, not displaces, human HR judgment.

3.1 Overall System Architecture and Design Principles

The architecture of the system is modular, service-based so that every HR function grievance redressal, policy advisement, hiring, and attrition forecasting can work independently but be interoperable through common embeddings and APIs. The assistant operates in human-in-the-loop style where AI-generated outputs (e.g., forecasted attrition probability or grievance priority) are validated by HR personnel prior to escalation or decision. Fig 2. depicts the gist of Methodology used.

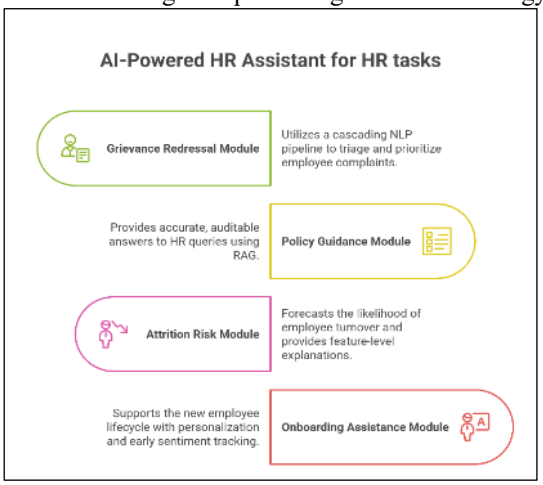


Fig. 2. Methodology

3.2 Data Sources

Due to the confidentiality limitations of actual HR data, this prototype utilizes artificial datasets that simulate realistic HR characteristics and text inputs. Four primary datasets were created for these purposes. The HR Attrition Dataset (n = 1,200) contains variables such as age, tenure, salary, performance rating, number of manager changes, number of sick leaves, and engagement scores, with simulated attrition labels using logistic probability thresholds and exponential time-to-exit distributions. A Grievance Corpus (n = 300) holds text complaints grouped into six categories (e.g., payroll, harassment, workload), each with sentiment polarity and priority estimates associated. A synthetic Policy Corpus was also created as seen in Figure 3, consisting of several HR policies to use for grounding in semantic retrieval and policy guidance. The data were all kept locally in structured CSV and JSON formats to support reproducibility and modular retraining once real enterprise data become available.



Fig. 3. Word Cloud of HR Policy Dataset

3.3 Modules and Algorithms

This section outlines the core functional modules of the AI-HR Assistant and the algorithms that underpin each component.

3.3.1. Grievance Handling

The new prototype, as seen in Figure 4, uses a probability-based severity model instead of the previous rule-based priority system. To produce a weighted probability score for each grievance category, using SBERT embeddings as input into a Random Forest classifier, the classifier generates a probability distribution of all grievance categories, weighted using different severity priors (e.g., harassment > workload). After that, the scores are adjusted based on each classifier's confidence level, DistilBERT sentiment intensity, and through application of soft keyword boosts for certain terms that we have identified as critical (e.g., "abuse," "threat," or "unpaid"). Also, the final output of this prototype is normalized to a 0-1 scale and is mapped to an actionable category (i.e., immediate escalation or low priority). This new approach is a much more flexible and auditable approach than using static rule-based scoring and allows for much greater visibility of how these categories were arrived at. Additionally, we conducted an examination of One-vs-Rest logistic regression using SBERT embeddings as input to perform multi-label classification. We assumed that using synthetic multi-label data that was naïvely created by combining single-label data with synthetic samples would have provided some performance improvement over the same samples being evaluated independently for their labels. However, given that there are overlaps between issues (i.e., harassment and workload are issues that can occur simultaneously), it is necessary to develop a much more realistic and multidimensional grievance corpus for the purpose of creating a transformer-based multi-label grievance model, which will allow for clearer guidance to be provided to Human Resources. As such, this will be an area of emphasis for our future research.



Fig. 4. Grievance Handling Module

3.3.2. Policy Guidance

The policy guidance system is constructed on Retrieval-Augmented Generation (RAG) principles to make sure all the responses are rooted in official corporate policy words. The process starts with Embedding and Indexing, where all the sentences of the policy documents are embedded through Sentence-BERT and indexed in FAISS (Facebook AI Similarity Search) for speedy searching. When a query is presented by an employee, the Retrieval component recognizes and retrieves the top-k most semantically similar policy snippets from the index. For Explainability, the system returns these snippets accompanied by the title of the original document to provide factual transparency to the user as seen in Figure 5.



Fig. 5. Policy Guidance Module

3.3.3. Onboarding Module

Through four closely-linked submodules depicted in Figure 6, the employee onboarding process automates and customizes the initial 90 days of an individual's employment. To ensure that new hires understand their goals and expectations during this time, the first submodule utilizes a Plan Generator that uses a rule-based system to produce an exclusive and tailored 30, 60, and 90 Day Plan based on the employee's job role, start date, and their initial skill set when starting with the company. The next submodule is Mentor Matching, which uses Sentence-BERT embeddings and cosine similarity to connect new employees to available on-site mentors, which is intended to accelerate the relationship-building process through the early support of mentor-mentee interaction. The third submodule is the Microlearning Recommendation engine that recommends short (30 minutes or less) "role" specific learning modules to help accelerate the employee's capability development and ensure that all critical competencies are achieved quickly and effectively. The final submodule is Pulse Check and Sentiment Tracking, which captures weekly feedback from the employee in a free-text format, then applies automated sentiment analysis to identify signs of early disengagement/dissatisfaction. The weekly data captured for pulse checks is stored within a secure JSON-based system which allows human resources to monitor employee well-being trends over time and to identify opportunities for proactive intervention. Overall, all four submodules together create a complete, data-driven onboarding system that supports employee success as well as organizational continuity.

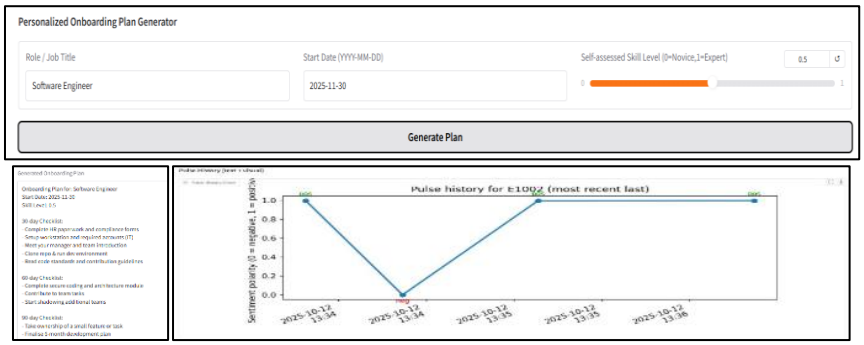


Fig. 6. Onboarding Dashboard & Pulse History Depiction

The Onboarding module has the ability to provide a new employee the ability to provide free form feedback about what they think about their experience every week of the first 90 days. This allows HR to identify first signs of employees that are going to have issues long before any formal complaint is received. The prototype does not make any claims regarding its impact on behaviour; however, it is evident that the Onboarding module is able to highlight some early indicators of disengagement. This prototype has not yet been calibrated and validated against actual retention results, but as additional data is collected, it can be better utilised as an additional tool for HR to actively and proactively address employee loss.

3.3.4. Attrition Prediction and Explainability

The attrition prediction model integrates Random Forest classification for binary attrition outcomes and Cox Proportional Hazards modelling for survival analysis (i.e., estimating time-to-exit). Explainability is guaranteed via SHAP (SHapley Additive exPlanations) values, which break down every prediction to feature-level contributions, and thereby enable HR managers to understand the drivers of attrition risk and develop customized retention interventions.

4 Results: Model Performance, Explainability, and Integration

The assessment of the AI-based HR assistant system effectively proved its main research goals by establishing the technical effectiveness of its predictive modules and illustrating the operational effectiveness of its overall, multi-module system.

4.1 Attrition Prediction Module: Quantitative Performance and Explainability

This subsection evaluates the predictive performance of the attrition model and examines the robustness of its results.

4.1.1. Predictive Accuracy and Reliability

The Explainable Random Forest Classifier (150 trees) demonstrated excellent predictive performance on the synthetic attrition dataset, achieving 0.968 accuracy,

0.957 precision, 0.879 recall, 0.916 F1, and 0.987 ROC AUC on the held-out test data, with exceptional discrimination between classes and balanced error rates. The model's performance was confirmed through robustness checks using 5-fold stratified cross-validation, resulting in a mean accuracy of 0.925 ± 0.020 and ROC AUC of 0.977 ± 0.011. The Performance of Gradient Boosting was slightly superior (accuracy 0.938 ± 0.017, ROC AUC 0.983 ± 0.011) while RBF-SVM performed poorly (~0.75 accuracy, ~0.44 ROC AUC) highlighting the efficacy of tree-based ensembles on this feature space.

To establish statistical stability, we conducted 500-replicate bootstrap resampling of the test metrics. The means and narrow confidence intervals obtained for accuracy 0.968 [0.942–0.988], precision 0.957 [0.893–1.000], Recall 0.879 [0.779–0.961], F1 0.916 [0.848–0.967], and ROC AUC 0.935 [0.882–0.975], provide evidence that the model's excellent predictive ability is due to the consistent behavior of the model across numerous resamples and is not attributable to a favorable single train/test split. (Refer Fig. 7 and 8)

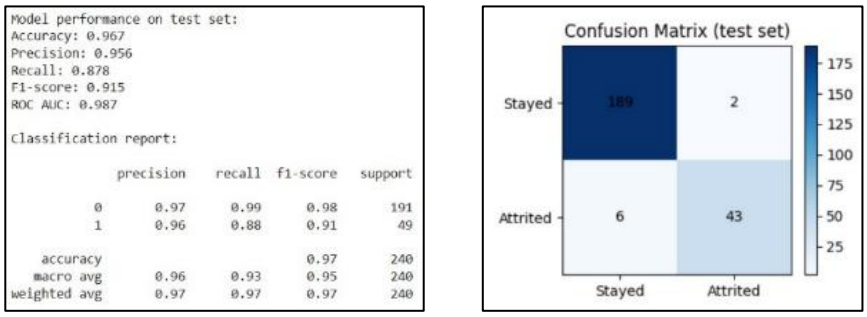


Fig. 7. Classification Report and Confusion Matrix

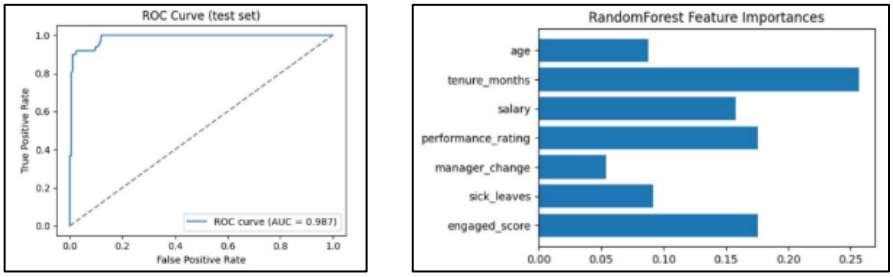


Fig. 8. Model Evaluation Parameters

4.1.2. Competitive Benchmarking

To compare our synthetic-data model to other past research, we conducted benchmarking between our synthetic-data model, the standard baselines of Random Forest and Support Vector Machines (SVM), and an explainable artificial intelligence (XAI) model that utilized tree ensembles combined with SHAP. When using the synthetic-data model for turnover prediction, the Explainable Random Forest achieved

a 96.7% accuracy rate very similar to the current XAI-framework accuracy of 96.95% and was significantly better than the accuracy rates normally cited for Random Forest (92.0%) and SVM (88.87%) in turnover prediction studies. Based on our cross-validated results, we are confident that including SHAP-based explainability does not impair predictive capabilities and can provide transparency to HR professionals who need to utilize advanced black-box models. The benchmarks summarised in Table 1.

Table 1: Comparative Table

Sr. No.	Model	Accuracy
1	Our Method	96.70%
2	Random Forest	92.00%
3	XAI framework	96.95%
4	SVM	88.87%

Both XAI-augmented models illustrate a significant gain in comparison to typical ensemble and conventional methods, including the generic Random Forest benchmark and the SVM model. This comparative demonstration amply illustrates that the design decision to incorporate an ethical explainability layer (SHAP) does not come with a compromise in overall predictive accuracy, meeting a key research goal.

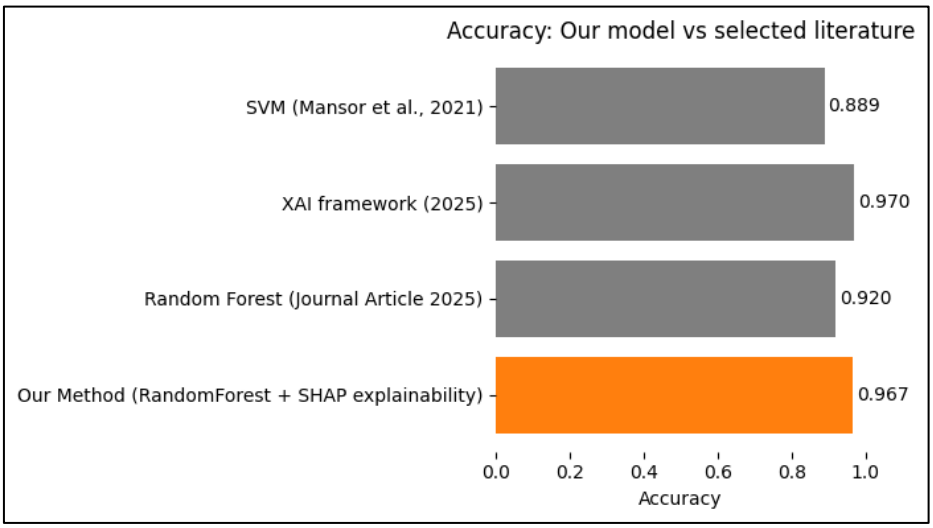


Fig. 9. Model Comparison

4.1.3. Global and Local Feature Interpretability

One of the key requirements of this research was operationalising explainability. This was fulfilled in two manners: Firstly, the Random Forest Feature Importances plot indicates global feature impact, with tenure_months as the most significant predictor of risk of attrition, followed by engaged_score and performance_rating. Secondly, the framework was able to implement SHAP (SHapley Additive exPlanations) as a mechanism for integrating local interpretability for each individual employee. (Refer Fig.10)

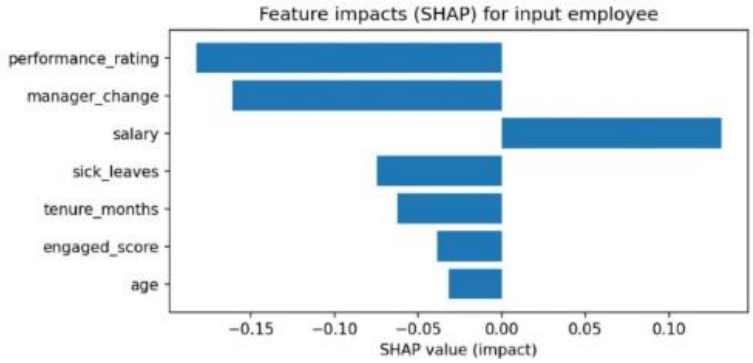


Fig. 10. Feature Impacts for SHAP Inputs

4.1.4. Fairness Diagnostics on Synthetic Data

As a sample of how fairness evaluations in human resources can be performed using a proxy protected class, we completed a fairness evaluation using the age 40+ as a protected class, or proxy, for the purpose of our analyses. The demographic-parity difference between employees aged 40+ and younger than 40's predicted high-risk attrition rates from the test dataset was 0.035. The true positive rates for equalized odds metrics for an employee aged 40+ were 0.897, while the true positive rates for employees aged younger than 40 were 0.850 (difference = 0.047). The false positive rates for equalized odds metrics for employees aged 40+ were 0.010 and for employees aged younger than 40, 0.011 (difference = –0.002). These small discrepancies are typical of synthetic datasets, and suggest that there aren't significant imbalances between the groups; however, even synthetic data can contain correlations between age, tenure and attrition labels.

In actual implementations of this evaluation methodology, additional sensitive attributes (e.g., gender, geography, job grade) will need to be included, and subgroup-specific SHAP explanations and counterfactual fairness tests will need to be developed. The purpose of the above analyses was only to provide a framework for evaluating fairness in HR analytic systems.

4.2 Integrated Module Functional Validation

The integrated application demonstrates that grievance analysis, policy retrieval, SHAP-based attrition scoring, and mentor matching can work together to create an end-to-end HR solution utilizing the same Gradio interface. A simulation of an employee journey (run for 100 iterations using one CPU) demonstrated average delays of 415 ms for grievance processing, 84 ms for policy retrieval, 50 ms for attrition scoring, 61 ms for SHAP calculating, and 65 ms for mentor matching. This means the average end-to-end response time for the Dun & Bergmann GRADIO Application is 675 ms with a 95th percentile response time of ~1.09 seconds, and throughput of 1.48 requests per second (in sequence).

While these results are based on synthetic, small-scale testing conditions and should not be used to estimate performance for production use cases; they highlight the benefits of a modular architecture, allowing each component to operate independently and maintain efficiencies when combined. Thus, they demonstrate an effective platform to support future optimization with actual HR data, increased workload, and production quality computing equipment.

4.1.5. Grievance Redressal and NLP Prioritization

The Grievance Redressal module applies a cascading Natural Language Processing (NLP) pipeline to quick issue triage. For example, on a sample complaint about workload and un-compensated overtime, the system offered an instantaneous, multi-dimensional analysis. The DistilBERT model ruled out a strongly confident NEGATIVE sentiment, and the Random Forest classifier classified the issue as policy_conflict. According to this analysis, the rule-based layer had a high Priority of 0.6, prompting the recommendation automatically for intervention of the human: "Follow up within 3 days". The model shows how the platform evaluates sentiment, confidence of category, and priority score (value of occurrence with a maximum value of 100) together to flag for human review the most severe and/or high-emotional issues, while at the same time clearly identifying the limitations of its underlying models.

4.1.6. Policy Guidance and Retrieval-Augmented Generation (RAG)

The Policy Guidance module uses a Retrieval-augmented Generation (RAG) framework to ensure the basic necessary requirement of being factually accurate and verifiable. While responding to a user query about leave entitlements, the system was able to retrieve contextually appropriate policy snippets using Sentence-BERT embeddings and FAISS indexing. Most importantly, the response of this module reflects clear attribution: the response explicitly cites the source document title and the relevant snippet for each piece of information derived from the document, thereby eliminating any risk of large language model (LLM) "hallucination" and ensuring the guidance is auditable and trusted.

4.1.7. Onboarding Personalization and Sentiment Tracking

The Onboarding module provides a tailored introduction to the platform along with tracking of ongoing engagement. The Plan Generator generates tailored 30/60/90-day checklists specific to the new hire position and self-reported skill level, combining common compliance items with role-specific items (e.g., a "Software Engineer" gets tasks such as "Clone repo & run dev environment"). In addition, the system's Mentor Matching capability applies Sentence-BERT similarity against location and skills to rank possible mentors, as well as showing their available workload for equitable assignment. The personalization, combined with microlearning suggestions and the initial Pulse Check system, forms an important early intervention mechanism for picking up potential integration problems early, before they lead to attrition.

5 Discussion: Implications and Future Work

The AI-based HR assistant developed under this research provides both practical and technical value for future HR systems. Its integrated design has the capability to support strategic decision-making, ethical governance, and can be implemented at scale across multiple HR functions.

5.1 Managerial Implications (Strategic, Ethical, and Operational)

HR is now moving from a reactive to a proactive and evidence-based approach to decision-making through data-driven insights. Employee attrition risk, which will identify the key factors that are driving your employee's attrition risk, will enable managers to perform more targeted, customized interventions to help decrease risks, as opposed to performing broad and generalized interventions. The human-in-the-loop review will continue to support the ethical oversight of high-sensitivity outputs and ensure the fairness, context, and empathy against each employee will allow HR staff to provide fair treatment for all employees, including those who may only marginally appear to be at risk of attrition.

Moreover, the system automates several key HR functions that enable HR to increase their operational effectiveness by taking over repetitive tasks. The RAG-based policy module is a reliable source of information for HR when dealing with frequent employee inquiries. The grievance module, on the other hand, helps identify urgent matters for HR to follow up on. This type of automation alleviates HR staff's administrative burden and allows them to concentrate on the relationship-building aspects of their role, complex grievance cases and strategic growth initiatives that improve the employee experience and culture of their organization.

5.2 Technical Implications (Scalability and Architecture)

Our platform uses a common semantic framework that supports policy retrieval, grievance classification, personalized onboarding, and mentor matching through Sentence-BERT embedding methodologies. We built a modular architecture for the platform to enable us to update any individual component of the platform without requiring a redesign of the pipeline. By using SHAP to integrate our components into a single solution, we can provide transparency and ethical interpretation for human resources (HR) support decision-making. In addition, we are employing a retrieval-augmented generation (RAG) with FAISS that connects the organization's policies to verified documents to eliminate hallucinations in the policy responses and improve reliability.

For future development efforts, we are focusing on: (i) developing new semi-synthetic HR datasets that will more accurately reflect real linguistic and behavioural diversity for grievance types; (ii) enhancing the grievance and attrition modules by employing multi-label classification with transformer-based models, as well as adding additional fairness validation to these modules (e.g., subgroup SHAP and counterfactual analysis); and (iii) testing the performance of the platform under the real workload of the organization, which will provide us with the ability to quantify the true end-to-end latency, throughput, and synchronization of the modules.

6 Limitations

The ability of AI-HR Assistant to perform accurately and automate actions based on synthetic data cannot be assumed to be the same when being used within actual organizations. The sample synthetic data presented can not capture the significant complexities and variances of informal language, sarcasm, survey non-responses, and cultural variances seen in real-life situations. The accuracy level and narrow confidence intervals demonstrated by our tests should be considered as upper bound limitations rather than confirmation that the AI-HR Assistant is ready for general deployment. There remain significant concerns regarding the fairness and bias of AI-HR Assistants as well. Although we developed workflows for assessing fairness in synthetic datasets (using age as a proxy), we need to determine if our assessments of fairness will still hold true for real-life datasets with demographic diversity, as interactions between sensitive attributes may have a substantial impact on the behaviour of models. We have also not yet assessed HR practitioners' behaviours in relation to both AI-HR Assistant outputs and SHAP explanations and the impact these have on HR decision-making.

7 Conclusion

This research article presented the conceptualization, implementation, and initial evaluations of an AI-enabled HR assistant that brings together four extremely important but frequently segmented functions, including grievance resolution, policy guidance, individualized onboarding, and explainable attrition risk prediction. The study realized its primary research objective by creating a coherent, modular system that supported and augmented human agency and judgment, a Human-in-the-Loop (HITL) workflow. Their findings validated the performance of the central module that powered the system, called an Explainable Random Forest Classifier, with high performance to date in attrition prediction. Most importantly, the system turned transparency into action with predictions accompanied by a SHAP-based explanation, as well as Retriever Augmented Generation (RAG) for accurate, auditable policy advice. Utilizing RAG, explainable machine learning (ML) and conversational UX, this study offers a technical guide to the future of HR technology. It transitions the HR function from reactive casework to proactive, tailored, and evidence-based employee support, all while pushing the edge of ethical AI deployment in managing workers. Future work in real world validation and survival analysis will complete the systems development for enterprise-wide implementation.

References

[1] Agrawal, V. (Ed.). (2025). *Human Resource Management and Artificial Intelligence: Automating and Optimizing Processes*. Taylor & Francis.

[2] Kumar, S., Datta, S., Singh, V., Datta, D., Singh, S. K., & Sharma, R. (2024). Applications, challenges, and future directions of human-in-the-loop learning. *IEEE Access*, *12*, 75735-75760.

[3] Maar, A., Roy, D., & Menold, J. (2026). Gender Dynamics in Design: Exploring Interactions, Psychological Safety, and Implicit Bias in Virtual Engineering Dyads. *Journal of Mechanical Design*, *148*(1), 011403.
[4] Benabou, A., & Touhami, F. (2025). Integration of artificial intelligence in human resource management: analyzing opportunities, challenges, and human-AI collaboration. Management & Accounting Review (MAR), 24(1), 405-437.
[5] Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California management review*, *61*(4), 15-42.
[6] Chukwuka, E. J., & Dibie, K. E. (2024). Strategic role of artificial intelligence (AI) on human resource management (HR) employee performance evaluation function. *International Journal of Entrepreneurship and Business Innovation*, *7*(2), 269-282.
[7] ATTHOS, Y. S. D., Iftikhar, K. M., Imtiaz, F., Azim, F., & Bangash, S. A. (2025). The Role of AI in Conflict Resolution and Employee Grievance Handling in HRM. *Dialogue Social Science Review (DSSR)*, *3*(3), 1164-1178.
[8] Kalluri, N. (2025). The Impact of AI on workplace litigation: A study of Employee relations and Conflict resolution. International Journal for Multidisciplinary Research, 7(4). https://doi.org/10.36948/ijfmr.2025.v07i04.51880
[9] Joshi, S. (2025). Artificial Intelligence in Conflict Resolution: A Comprehensive Review of Techniques and Applications. Available at SSRN 5241365.
[10] Gupta, P., Ijardar, O. P., Jadhav, A., & Saheb, V. (2025, June). AI-Based Solution To Enable Ease of Grievance Lodging and Tracking for Citizens Across Multiple Departments. In International Conference on Advances and Applications in Artificial Intelligence (ICAAAI 2025) (pp. 1002-1022). Atlantis Press.
[11] Nimmagadda, S., Surapaneni, R. K., & Potluri, R. M. (2024). Artificial intelligence in HR: employee engagement using Chatbots. *Artificial Intelligence Enabled Management: An Emerging Economy Perspective*, *147*.
[12] Kovari, A. (2025). Explainable AI chatbots towards XAI ChatGPT: A review. *Heliyon*, *11*(2).
[13] Chen, T., Gascó-Hernandez, M., & Esteve, M. (2024). The adoption and implementation of artificial intelligence chatbots in public organizations: Evidence from US state governments. *The American Review of Public Administration*, *54*(3), 255-270.
[14] Silva, G. R. S., & Canedo, E. D. (2025). Privacy in Chatbot Conversation-Driven Development: A Comprehensive Review and Requirements Proposal. *ACM Transactions on Software Engineering and Methodology*.
[15] Suhonen, E. (2025). Integration of AI into HR: An analysis on employee attitudes towards AI chatbots in HR advisory services.

[16] Rane, N. (2023). Role and challenges of ChatGPT and similar generative artificial intelligence in human resource management. *Available at SSRN 4603230*.
[17] Ahmed El Garem, R. (2024). Impact of AI on Employee Onboarding Processes By Investigating How AI Tools Can Streamline and Enhance the Onboarding Experience for New Hires. *15*(4), 270-315.
[18] Kumar, P. A., Sharma, A., Induji, R. T., Singathurai, S., & Saranya, M. (2025). AI-Driven Employee Engagement: Transforming HR Strategies for the Digital Workforce.
[19] Shaheen, G., Saeedi, K., Alsobhi, A., & Alzhrani, F. (2025). Towards a Gamification Solution for Enhancing Employee Onboarding Experience. *International Journal of Computing*, *17*(1), 1-12.
[20] Mohammad, A. A. S., Alkhazali, Z., Mohammad, S. I. S., Al Oraini, B., Vasudevan, A., Alqahtani, M. M., & Alshurideh, M. T. (2025). Machine Learning Models for Predicting Employee Attrition: A Data Science Perspective. *Data and Metadata*, *4*, 669-669.
[21] Baydili, İ. T., & Tasci, B. (2025). Predicting Employee Attrition: XAI-Powered Models for Managerial Decision-Making. *Systems*, *13*(7), 583. https://doi.org/10.3390/systems13070583
[22] Mohbey, K. K. (2022). Employee's attrition prediction using survival analysis and Cox proportional hazard model. *International Journal of Digital Enterprise Technology*, *2*(1), 27-37.
[23] Alqahtani, H., Almagrabi, H., & Alharbi, A. (2024). Employee attrition prediction using machine learning models: A review paper. *International Journal of Artificial Intelligence & Applications.*
[24] Ghavami, V., & Tabatabaee, S. S. (2023). A survival analysis approach to determine factors associated with non-retention of newly hired health workers in Iran. *BMC health services research*, *23*(1), 265. https://doi.org/10.1186/s12913-023-09262-5
[25] Mansor, N., Sani, N. S., & Aliff, M. (2021). Machine learning for predicting employee attrition. *International Journal of Advanced Computer Science and Applications (IJACSA)*, *12*(11).