

Numerical Approximation Methods for Real-Time Employee Attrition Classification: A Systematic Review of Algorithm Performance, Computational Complexity, and Class-Imbalance Optimization

Abhimanyu Kumar
Lead Analyst, Corporate Systems

Abstract

Employee turnover is one of the most serious business problems that modern organisations face, and the cost of replacement for a lost employee can range from 50–200% of their annual salary (Allen et al., 2010). This synthesis is based on 16 peer-reviewed studies and conference proceedings that apply numerical and machine learning–based approximation methods to human resources (HR) data for predicting attrition in real time. Ensemble methods, especially gradient-boosted classifiers and stacking architectures, are identified as the most successful paradigm, with accuracy rates reaching 93.5% and Area Under the Receiver Operating Characteristic Curve (AUC-ROC) scores exceeding 0.96 on the IBM Watson HR Analytics benchmark dataset. Deep learning methods, such as Bidirectional Long Short-Term Memory (Bi-LSTM) networks, outperform other methods in the recall of minority-class attrition cases, particularly relevant given the inherent class imbalance problem in most enterprise HR datasets (where attrition rates range from 8.9% to 22.3% depending on the industry). Job satisfaction, overtime frequency, number of years at the company, and monthly income are consistently identified as the strongest predictors through feature-selection analyses. Proactive retention interventions can be achieved by reducing median decision latency to sub-200 millisecond thresholds via real-time deployment of predictive pipelines through Application Programming Interface (API) endpoints. Finally, this synthesis finds that no single algorithm is optimal for all cases; the choice of algorithm depends on the specific context, including data dimensionality, class-frequency distribution, and available computational resources.

Keywords: *employee attrition prediction, machine learning, ensemble methods, human resources analytics, Bi-LSTM, gradient boosting, random forest, feature selection, class imbalance, SMOTE, IBM Watson HR dataset, real-time prediction*

1. Introduction

One of the ongoing issues for human resource (HR) departments around the world is employee attrition—the phenomenon of employees leaving an organisation through resignation or dismissal.

10.48047/jocaaa.2023.31.02.52

The financial impact of workforce turnover is significant: direct replacement costs are commonly estimated at 50–200% of the departing employee's annual salary, encompassing recruitment, selection, onboarding, and lost productivity during the transition period (Garg et al., 2021). In large organisations, the aggregate cost of attrition frequently exceeds tens of millions of U.S. dollars annually, making the development of accurate predictive systems an urgent operational priority.

Machine learning (ML) and data-driven analytics have ushered in a paradigm shift in HR practice, moving from intuition-based decision-making towards information- and evidence-based talent management. Predictive attrition models enable HR personnel to identify potential leavers before they resign, allowing proactive retention strategies such as career development planning, compensation negotiation, and flexible working arrangements (Kang et al., 2020). These capabilities are particularly valuable in industries with high natural turnover, such as IT and technology (13.2–18.1%), healthcare (10.4–15.2%), and retail (17.3–22.3%), where each departure carries substantial cost implications (Khera & Divya, 2019).

Although the field has attracted significant scholarly attention, notable methodological fragmentation remains. Algorithms, feature engineering approaches, class-imbalance mitigation techniques, and evaluation frameworks vary widely across studies. The IBM Watson HR Analytics dataset—containing 1,470 synthetic employee records with 35 features—has become the current benchmark for cross-study comparability (Fallucchi et al., 2020; Mansor et al., 2021; Raza et al., 2022). The transferability of findings from this benchmark to real-world enterprise applications, however, warrants careful scrutiny.

1.1 Research Objectives

The primary goal of this synthesis is a comparative analysis of numerical approximation techniques applied to HR datasets for real-time employee attrition prediction. In particular, the study: (a) reviews the performance of classical, ensemble, and deep-learning algorithms; (b) examines feature selection and class-imbalance mitigation strategies; (c) assesses the computational feasibility of real-time deployment; and (d) identifies methodological gaps for further investigation as of 2022.

1.2 Scope and Delimitations

10.48047/jocaaa.2023.31.02.52

Only studies published from 2017 to 2022 employing supervised ML methods on structured tabular HR data are included. The synthesis excludes studies relying solely on qualitative methods, NLP-based employee survey analysis, or unsupervised clustering. Quantitative benchmarks reported herein are drawn directly from the cited studies and are presented as summarised.

2. Literature Synthesis

2.1 Evolution of Predictive Attrition Research

In early applications of data mining to employee turnover prediction, Yiğit and Shourabizadeh (2017) applied decision tree and Naïve Bayes classifiers to HR data, achieving approximately 82.1% accuracy and an AUC-ROC of 0.85. This pioneering work demonstrated the viability of predictive attrition modelling while exposing the class-imbalance problem that subsequent studies addressed more systematically.

A clear shift towards advanced algorithmic techniques is evident in the 2019–2021 period. Khera and Divya (2019) demonstrated in the Indian IT industry that ensemble approaches outperform single-model classifiers on industry-specific datasets, with their best Random Forest (RF) configuration achieving 87.4% accuracy. Fallucchi et al. (2020) extended this line of inquiry by combining gradient-boosted classifiers with logistic regression (LR) on the IBM Watson HR dataset, observing that feature preprocessing—label encoding and min-max normalisation—significantly affected model convergence. Jain et al. (2020) introduced eXtreme Gradient Boosting (XGBoost) and AdaBoost within an explanatory framework, achieving 90.3% accuracy, and additionally applied SHapley Additive exPlanations (SHAP) to enhance interpretability for HR practitioners.

Kang et al. (2020) studied U.S. federal government workers—a population with distinct institutional incentives and tenure structures relative to private-sector employees. Evaluating Support Vector Machine (SVM), Naïve Bayes, and RF classifiers, they found that perceived comprehensive HR practices were a significant predictor of turnover intention, with direct policy implications for public-sector HRM.

Garg et al. (2021) conducted an extensive review of predictive HR analytics and proposed a conceptual taxonomic framework for HR analytics applications. Their synthesis identified attrition

prediction, performance forecasting, and recruitment optimisation as the three most common application domains—with attrition prediction dominating the published literature. The study also highlighted continued concerns about algorithmic fairness, noting that protected attributes such as age, gender, and ethnicity can function as implicit predictors in attrition models, potentially triggering compliance issues under anti-discrimination legislation.

Table 1: Summary of Key Studies on ML-Based Employee Attrition Prediction

Study	Algorithm(s)	Dataset	Accuracy (%)	AUC-ROC	Year
Yiğit & Shourabizadeh	Decision Tree, NB	IBM HR Watson	82.1	0.85	2017
Khera & Divya	LR, DT, RF, SVM	Indian IT Sector	87.4	0.89	2019
Fallucchi et al.	RF, Gradient Boost, LR	IBM HR Watson	89.1	0.92	2020
Jain et al.	RF, AdaBoost, XGBoost	IBM HR Watson	90.3	0.93	2020
Kang et al.	SVM, NB, RF	US Federal HR	84.7	0.88	2020
Najafi-Zangeneh et al.	RF + Feature Selection	IBM HR Watson	91.8	0.94	2021
Mansor et al.	LR, DT, RF, KNN	IBM HR Watson	88.6	0.91	2021
Saleh et al.	DNN (Deep Neural Net)	IBM HR Watson	88.2	0.92	2021
Wang & Zhi	XGBoost, LGBM	Proprietary HR	90.7	0.94	2021
Alsheref et al.	Ensemble (Stacking)	IBM HR Watson	93.5	0.97	2022
Raza et al.	RF, XGBoost, DNN	IBM HR Watson	91.6	0.95	2022
Guerranti & Dimitri	LR, SVM, RF, XGB, ANN	IBM HR Watson	91.9	0.95	2022
Musanga et al.	SVM, RF, GBM	HR Churn Dataset	89.4	0.93	2022

Note. LR = Logistic Regression; DT = Decision Tree; RF = Random Forest; SVM = Support Vector Machine; DNN = Deep Neural Network; Bi-LSTM = Bidirectional Long Short-Term Memory; GBM = Gradient Boosting Machine. Accuracy values reflect best-reported configurations per study.

As illustrated in Table 1, accuracy values across reviewed studies range from 82.1% (Yiğit & Shourabizadeh, 2017) to 93.5% (Alsheref et al., 2022), reflecting both algorithmic advancements and dataset-specific characteristics. AUC-ROC scores similarly improved from 0.85 in 2017 to

0.97 in 2021–2022, indicating substantially enhanced discrimination capability in more recent models.

2.2 Ensemble and Deep Learning Approaches

Substantial evidence from 2021–2022 confirms the superiority of ensemble methods over individual classifiers. In a systematic comparison of voting ensembles, stacking architectures, and single-model baselines, found that stacking—in which a meta-learner combines base-model predictions—achieved 92.8% accuracy and an AUC-ROC of 0.96, outperforming all individual classifiers evaluated in the same study. Alsheref et al. (2022) corroborated these findings using an automated ensemble pipeline with hyperparameter optimisation, achieving the highest accuracy (93.5%) reported in the synthesised literature.

A parallel track involves deep learning architectures. Saleh et al. (2021) applied a multi-layer Deep Neural Network (DNN) to attrition prediction, achieving 88.2% accuracy and an AUC-ROC of 0.92. While broadly comparable to RF baselines in performance, the DNN required substantially greater computational resources (training time ~187.4 s vs. ~12.3 s for RF) and offered less interpretability. This limitation employed a Bidirectional Long Short-Term Memory (Bi-LSTM) network (Wang & Zhi, 2021). The Bi-LSTM achieved the highest minority-class recall (89.7%) in the synthesised literature, making it particularly suitable for organisations with high sensitivity requirements in attrition identification.




Figure 1 — Comparative Performance of Machine Learning Algorithms for Employee Attrition Prediction (Accuracy, F1-Score, AUC-ROC $\times 100$) — to be inserted here. Synthesised from Yiğit & Shourabizadeh (2017), Fallucchi et al. (2020), Jain et al. (2020), Alsheref et al. (2022), Guerranti & Dimitri (2022).

Logistic regression and decision tree baselines perform substantially below gradient boosting and Bi-LSTM architectures across all three evaluation measures. The ensemble stacking approach approaches Bi-LSTM performance at a fraction of the computational cost, which has significant implications for deployment feasibility in resource-constrained HR technology environments.

3. Methodology and Technical Framework

3.1 Dataset Characteristics

Most synthesised studies use the IBM Watson HR Analytics benchmark dataset as the primary empirical tool. The dataset comprises 1,470 employee entries with 35 attributes spanning: Demographic (age, gender, marital status), Organisational (department, job role, job level), Compensation (monthly income, stock option level), and Experiential (years at company, number of companies worked, total working years) categories. The target variable—binary attrition classification (Yes/No)—exhibits inherent class imbalance, with positive cases (attrition) constituting 16.1% of total records (Mansor et al., 2021; Raza et al., 2022).

Other datasets used in synthesised studies include proprietary corporate HR databases (Wang & Zhi, 2021), sector-specific Indian IT industry data (Khera & Divya, 2019), the United States Office of Personnel Management workforce dataset (Kang et al., 2020), and the Zimbabwean HR churn dataset used by Musanga et al. (2022). Cross-dataset model transferability remains a methodological challenge, as models often perform poorly when applied to different organisational contexts.

3.2 The Machine Learning Pipeline

The generalised ML pipeline for attrition prediction encompasses seven sequential phases, as synthesised across reviewed studies: data collection, preprocessing, feature engineering, model training, validation and hyperparameter tuning, real-time deployment, and monitoring with feedback-driven retraining.




Figure 2 — End-to-End Machine Learning Pipeline for Real-Time Employee Attrition Prediction, annotated with representative techniques employed across synthesised studies — to be inserted here. Synthesised from Fallucchi et al. (2020), Najafi-Zangeneh et al. (2021), Raza et al. (2022).

The pipeline begins with structured HR data ingestion, followed by missing value imputation, one-hot and ordinal encoding of categorical variables, and feature normalisation. Feature engineering is a critical intermediate step; dimensionality reduction techniques such as Principal Component Analysis (PCA), Recursive Feature Elimination (RFE), and SHAP-based importance ranking are used to narrow the feature set supplied to the model. Model training encompasses both single-algorithm and ensemble configurations, with cross-validation and grid-search hyperparameter optimisation applied to mitigate overfitting. Real-time deployment via REST API endpoints enables sub-200 ms inference integration with enterprise HR management systems.

3.3 Feature Engineering and Selection

Features in HR attrition datasets contribute unequally to model performance; redundant and correlated features can obscure discriminative signal. Najafi-Zangeneh et al. (2021) developed an a hybrid feature selection approach combining correlation analysis with wrapper-based RFE, reducing the IBM Watson HR dataset feature space from 35 to 18 variables and increasing classification accuracy from 88.4% to 91.8%.




Figure 3 — Top-10 Feature Importance Scores for Employee Attrition Prediction (normalised importance values averaged across Random Forest and XGBoost classifiers) — to be inserted here. Synthesised from Najafi-Zangeneh et al. (2021) and Jain et al. (2020).

Table 2: Feature Importance Scores Across Multiple Classification Algorithms

Feature	RF Score	XGBoost Score	LR Score	DNN Score	Rank
Job Satisfaction	0.187	0.194	0.162	0.178	1
Overtime	0.162	0.171	0.148	0.156	2
Years at Company	0.143	0.138	0.127	0.141	3
Monthly Income	0.121	0.118	0.134	0.129	4
Work-Life Balance	0.109	0.102	0.118	0.113	5
Age	0.094	0.097	0.088	0.091	6
Distance from Home	0.078	0.081	0.072	0.076	7

Note. RF = Random Forest; XGBoost = eXtreme Gradient Boosting; LR = Logistic Regression; DNN = Deep Neural Network. Scores represent normalised SHAP-derived importance values. Rank is determined by mean importance across all four algorithms.

As presented in Table 2 and visualised in Figure 3, job satisfaction is consistently the most important predictor across all four algorithm configurations, with importance scores ranging from

0.162 (LR) to 0.194 (XGBoost). Overtime frequency ranks second, followed by tenure (years at company) and monthly income. These results are consistent with organisational behaviour research identifying job satisfaction, workload intensity, and compensation equity as primary drivers of turnover intention (Jain et al., 2020; Wang & Zhi, 2021).

3.4 Class Imbalance Mitigation

The treatment of class imbalance is among the most technically significant methodological decisions in attrition model development. When data are imbalanced, classifiers frequently achieve high overall accuracy but low sensitivity, correctly classifying the majority of non-attrition cases while repeatedly failing to identify true attrition cases (Guerranti & Dimitri, 2022; Musanga et al., 2022).

Table 3: Class Imbalance Mitigation Techniques: Effectiveness and Key References

Technique	Description	Improvement in Minority Recall (%)	Key Reference
SMOTE	Synthetic Minority Over-sampling Technique	+14.2%	Fallucchi et al. (2020); Raza et al. (2022)
ADASYN	Adaptive Synthetic Sampling	+12.8%	Najafi-Zangeneh et al. (2021)
Cost-Sensitive Learning	Class-weight penalisation in loss function	+9.6%	Wang & Zhi (2021)
Random Undersampling	Majority-class record removal	+7.1%	Mansor et al. (2021)
Ensemble Resampling	Combined SMOTE + Bagging	+18.3%	Alsheref et al. (2022)
Threshold Optimisation	Adjusting classification decision boundary	+11.4%	Guerranti & Dimitri (2022)

Note. Improvement percentages reflect the increase in minority-class recall relative to unmitigated baseline classifiers on the IBM Watson HR dataset. SMOTE = Synthetic Minority Over-sampling Technique; ADASYN = Adaptive Synthetic Sampling.

10.48047/jocaaa.2023.31.02.52

As shown in Table 3, ensemble resampling combining SMOTE with bagging-based architectures yields the greatest minority-class recall gain (+18.3%) (Alsheref et al., 2022). SMOTE alone achieves +14.2% recall improvement by interpolating existing minority-class feature vectors to generate synthetic instances. Threshold optimisation—shifting the classification boundary below the default 0.5—is a computationally lightweight alternative achieving +11.4% minority recall improvement without modifying the training distribution (Guerranti & Dimitri, 2022).

4. Results and Comparative Analysis

4.1 Algorithm Performance Benchmarking

A clear hierarchical structure is observable when comparing algorithm performance across synthesised studies. Precision scores range from 79.3% (Yiğit & Shourabizadeh, 2017) to 92.8% (Alsheref et al., 2022), and recall scores from 76.8% to 91.4%, reflecting advances in both preprocessing and algorithmic maturation over the review period. The AUC-ROC score—the most robust single-metric indicator of discriminative capability under class imbalance—ranges from 0.85 to 0.97 across synthesised studies.

Table 4: Longitudinal Performance Comparison Across Key Studies

Study (Year)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC	Dataset Size
Yiğit & Shourabizadeh (2017)	79.3	76.8	78.0	0.85	1,470
Khera & Divya (2019)	83.1	80.4	81.7	0.87	2,380
Fallucchi et al. (2020)	86.4	83.7	85.0	0.92	1,470
Najafi-Zangeneh et al. (2021)	89.7	86.2	87.9	0.94	1,470
Alsheref et al. (2022)	92.8	91.4	92.1	0.97	1,470
Raza et al. (2022)	90.4	88.9	89.6	0.95	1,470

Note. AUC-ROC = Area Under the Receiver Operating Characteristic Curve. Dataset size refers to the number of employee records utilised. All metrics reflect best-configuration performance as reported by original authors.




Figure 4 — Performance Heatmap: Single vs. Ensemble Models for Attrition Prediction. Cell values represent percentage scores for Accuracy, Precision, Recall, F1-Score, and AUC-ROC ($\times 100$). Darker cells indicate superior performance — to be inserted here. Synthesised from Alsheref et al. (2022), and Guerranti & Dimitri (2022).

As shown in Table 4, top-performing ensemble configurations dominate every performance metric. The AUC-ROC gap between stacking ensembles and Bi-LSTM models is less than 0.1, and the accuracy difference is within 0.7 percentage points, confirming that ensemble methods offer near-optimal performance at substantially lower computational cost (Alsheref et al., 2022).

4.2 Computational Efficiency Analysis

Beyond predictive accuracy, enterprise HR deployment requires consideration of training time, memory requirements, inference speed, and horizontal scalability. Table 5 summarises the computational complexity profile for eight algorithm types, as synthesised from the reviewed literature.

Table 5: Computational Efficiency Metrics Across Algorithm Categories

Algorithm	Training Time (s)	Memory (MB)	Latency (ms)	Scalability
Logistic Regression	0.8	12	3	Very High
Decision Tree	1.4	18	5	High
Random Forest (100 trees)	12.3	142	18	High
SVM (RBF Kernel)	28.7	96	22	Moderate
Gradient Boosting (XGBoost)	34.2	210	25	High
Deep Neural Network	187.4	580	48	Moderate
Bi-LSTM	312.6	840	76	Low-Moderate
Stacking Ensemble	48.9	320	32	High

Note. Training time and memory estimates are derived from IBM Watson HR dataset (1,470 records, 35 features) experiments reported in synthesised studies. Latency refers to single-record inference time. Scalability assessment is qualitative, reflecting author-reported assessments.

Logistic regression offers the fastest training (0.8 s) and inference (3 ms), making it well-suited for high-throughput, real-time systems. In contrast, Bi-LSTM requires 312.6 seconds of training time and 840 MB of memory, rendering it computationally intensive for organisations lacking GPU infrastructure (Saleh et al., 2021). Stacking ensemble architectures occupy a reasonable middle ground, with 48.9 s training time, 320 MB memory, and 32 ms inference latency—comfortably within the sub-200 ms threshold for real-time HR system integration (Wang & Zhi, 2021).

4.3 Industry Attrition Rate Analysis

Significant variation in employee attrition prevalence exists across industries and time periods.



Figure 5 — Employee Attrition Rate Trends Across Industry Sectors (2017–2022). Rates represent percentage of workforce separating annually. Shaded areas highlight IT and Retail sector trajectories — to be inserted here. Synthesised from Khera & Divya (2019), Kang et al. (2020), and Garg et al. (2021).

The retail sector consistently exhibits the highest attrition rates—reaching 22.3% in 2022—attributable to seasonal work patterns, narrow pay bands, and low barriers to occupational exit. IT and technology represent the most significant growth segment across the review period, increasing from 13.2% in 2017 to 18.1% in 2022, driven by escalating inter-firm competition for specialised skills in cloud computing, artificial intelligence, and cybersecurity (Garg et al., 2021; Khera & Divya, 2019). Government sector attrition remains consistently low at 5.1–6.4%, consistent with findings by Kang et al. (2020) on the retention effects of defined-benefit pension schemes and public-sector job security measures.

4.4 Comparative Framework: Analytical Paradigm Benchmarking

Table 6 presents a detailed comparison of four analytical paradigms—traditional HR analytics, basic single-model ML, advanced ensemble ML, and deep learning—across seven evaluative dimensions.

Table 6: Comparative Framework: Analytical Paradigm Benchmarking for Attrition Prediction

Metric	Traditional HR Analytics	Basic ML (Single Model)	Advanced Ensemble	Deep Learning (Bi-LSTM)
Prediction Accuracy (%)	61.2	85.4	93.5	93.1
F1-Score (%)	N/A	81.6	92.1	91.4
AUC-ROC	0.64	0.88	0.97	0.97
Inference Latency (ms)	>1000	18–25	32–48	76–120
Interpretability	High	Moderate	Moderate (SHAP)	Low
Implementation Cost (USD)	Low (~\$5K)	Moderate (~\$30K)	High (~\$75K)	Very High (~\$150K)
Retraining Frequency	Annual	Quarterly	Monthly	Monthly

Note. Implementation cost estimates reflect approximate total cost of ownership for deployment in a 1,000-employee organisation, including infrastructure, licensing, and personnel costs. SHAP = SHapley Additive exPlanations. N/A = Not applicable for traditional systems lacking probabilistic output.

The paradigm comparison confirms that advanced ensemble methods offer the best combined profile of predictive performance (93.5% accuracy, AUC-ROC 0.97), explainability via SHAP frameworks, and inference latency (32–48 ms), at half the implementation cost of deep learning deployments (\$75,000 vs. \$150,000 for a 1,000-employee organisation). Traditional HR analytics, while inexpensive (~\$5,000), achieves only 61.2% accuracy—insufficient for reliable identification of at-risk employees—and lacks probabilistic output, precluding uncertainty-quantified retention risk scoring (Garg et al., 2021; Wang & Zhi, 2021).

5. Discussion

5.1 Principal Findings

The comparative analysis yields four principal findings. First, ensemble methods—particularly stacking architectures with gradient-boosted base learners—represent the current state of the art for tabular HR attrition data, matching deep learning performance at significantly lower computational cost. Second, feature selection makes a measurable positive contribution to model

10.48047/jocaaa.2023.31.02.52

performance: SHAP-based importance ranking and RFE can reduce the feature set by up to 49% (from 35 to 18 features) and increase accuracy by 3–4 percentage points (Najafi-Zangeneh et al., 2021). Third, class imbalance constitutes the most critical preprocessing challenge in attrition prediction, and SMOTE-ensemble hybridisation yields the greatest gains in operationally relevant recall metrics. Fourth, ensemble and single-model classifiers can be deployed within sub-200 ms latency thresholds, enabling live integration with HR information systems for continuous attrition risk monitoring.

5.2 Challenges and Limitations

Several limitations constrain the generalisability of synthesised findings. Reported performance metrics are heavily conditioned on the IBM Watson HR Analytics dataset, which is synthetic with only 1,470 records. Real enterprise HR datasets may differ substantially in record volume, feature distribution, and temporal dynamics of attrition drivers, all of which can affect algorithm performance relative to controlled benchmark conditions (Raza et al., 2022; Musanga et al., 2022). Table 7 summarises identified challenges and proposed mitigations.

Table 7: Key Challenges in HR Attrition Prediction and Proposed Mitigations

Category	Challenge	Proposed Mitigation	Supporting Reference
Data Quality	Missing values and inconsistent records	Multiple imputation; mean/mode substitution	Mansor et al. (2021)
Class Imbalance	Attrition cases as low as 8.9% of dataset	SMOTE, ensemble resampling, threshold tuning	Alsheref et al. (2022); Fallucchi et al. (2020)
Feature Dimensionality	Up to 35 input variables causing overfitting	PCA, SHAP-based selection, RFE	Najafi-Zangeneh et al. (2021)
Interpretability	Black-box model outputs limit HR adoption	SHAP values; LIME explanations	Jain et al. (2020); Wang & Zhi (2021)
Computational Cost	Deep models require GPU infrastructure	Cloud-based inference; model distillation	Saleh et al. (2021)
Bias & Fairness	Protected attribute leakage (age, gender)	Fairness-aware preprocessing; disparate impact removal	Garg et al. (2021); Kang et al. (2020)

Note. SMOTE = Synthetic Minority Over-sampling Technique; PCA = Principal Component Analysis; SHAP = SHapley Additive exPlanations; LIME = Local Interpretable Model-agnostic Explanations; RFE = Recursive Feature Elimination.

Algorithmic bias represents a distinct challenge. When variables correlated with protected attributes such as age or distance from home are included, prediction patterns may produce discriminatory outcomes inconsistent with equal employment opportunity principles (Garg et al., 2021; Kang et al., 2020). Fairness-aware preprocessing techniques—including disparate impact removal and instance re-weighting—have been proposed but remain underexplored in the attrition prediction literature as of 2022. Similarly, organisations lacking dedicated ML infrastructure are largely unable to access DNN and Bi-LSTM architectures due to their high computational demands (Saleh et al., 2021).

5.3 Real-Time Deployment Considerations

Deploying attrition models for real-time inference introduces engineering considerations beyond batch prediction. Wang and Zhi (2021) proposed embedding predictive models behind REST API endpoints, enabling real-time scoring of employee records as HR events—salary updates, performance reviews, project assignments—are logged into enterprise systems. Their management analytics framework demonstrated that monthly model retraining achieves superior performance stability by capturing evolving organisational dynamics. Musanga et al. (2022) similarly advocated for frequent model revalidation to reflect workforce composition changes and macroeconomic shocks affecting attrition propensity distributions.

Concept drift—the progressive divergence between input feature distributions and attrition outcomes due to organisational, economic, or societal changes—is a systemic threat to deployed model validity. This risk is illustrated by the COVID-19 context in the 2020 data (Kang et al., 2020), when the government sector experienced an attrition rate increase (5.6% to 6.1%) attributable to labour market disruptions not captured by pre-pandemic training data.

6. Conclusion

This synthesis has conducted a systematic comparative analysis of numerical approximation techniques applied to HR datasets for real-time employee attrition prediction, drawing on 16 peer-reviewed studies published from 2017 to 2022. The principal conclusions are as follows. Ensemble

10.48047/jocaaa.2023.31.02.52

methods, including stacking architectures, represent the most effective paradigm for tabular HR data, achieving up to 93.5% accuracy and an AUC-ROC of 0.97 on the IBM Watson HR Analytics benchmark. Bi-LSTM deep learning architectures provide an advantage in minority-class recall (89.7%) over architectures with comparable overall accuracy, making them preferable in contexts where false negatives carry greater organisational cost than computational parsimony.

Feature selection is a critical performance lever: SHAP-based dimensionality reduction from 35 to 18 variables yielded a 3.4 percentage point accuracy gain without increasing overfitting risk. Four variables—Job Satisfaction, Overtime Frequency, Years at Company, and Monthly Income—emerge as the most predictive across all organisations examined, consistent with established organisational behaviour theory. SMOTE-ensemble hybridisation achieves the largest gains in operationally relevant recall metrics (+18.3%), substantially offsetting the systematic under-prediction of attrition events that results from training on unmitigated imbalanced datasets.

On deployment, ensemble models demonstrate the best performance-to-resource ratio for real-time inference, sustaining sub-200 ms latency at prediction accuracy exceeding 93%. The implementation cost advantage of ensemble over deep learning approaches (\$75,000 vs. \$150,000 for a 1,000-employee organisation) further reinforces the ensemble paradigm for most enterprise scenarios.

Future research directions identified at the 2022 knowledge horizon include: (a) external validity evaluation of benchmark findings on large-scale, longitudinal enterprise HR datasets; (b) development of fairness-constrained training objectives that simultaneously optimise predictive performance and demographic parity; (c) investigation of concept drift detection and retraining protocols to sustain model accuracy in dynamic organisational settings; and (d) integration of natural language processing techniques to incorporate unstructured data sources—performance review text, engagement survey responses—as supplementary attrition predictors alongside structured HR data. The frameworks and comparisons presented herein offer a unified reference for practitioners and researchers in the emerging field of data science-driven strategic HRM

References

10.48047/jocaaa.2023.31.02.52

- Allen, D. G., Bryant, P. C., & Vardaman, J. M. (2010). Retaining talent: Replacing misconceptions with evidence-based strategies. *Academy of Management Perspectives*, 24(2), 48–64. <https://doi.org/10.5465/amp.24.2.48>
- Alsheref, F. K., Fattoh, I. E., & Ead, W. M. (2022). Automated prediction of employee attrition using ensemble model based on machine learning algorithms. *Computational Intelligence and Neuroscience*, 2022, 1–9. <https://doi.org/10.1155/2022/7728668>
- Fallucchi, F., Coladangelo, M., Giuliano, R., & De Luca, E. W. (2020). Predicting employee attrition using machine learning techniques. *Computers*, 9(4), 86. <https://doi.org/10.3390/computers9040086>
- Garg, S., Sinha, S., Kar, A. K., & Mani, M. (2021). A review of machine learning applications in human resource management. *International Journal of Productivity and Performance Management*, 71, 1590–1610. <https://doi.org/10.1108/IJPPM-08-2020-0427>
- Guerranti, F., & Dimitri, G. M. (2022). A comparison of machine learning approaches for predicting employee attrition. *Applied Sciences*, 13(1), 267. <https://doi.org/10.3390/app13010267>
- Jain, P. K., Jain, M., & Pamula, R. (2020). Explaining and predicting employees' attrition: A machine learning approach. *SN Applied Sciences*, 2. <https://doi.org/10.1007/s42452-020-2519-4>
- Kang, I., Croft, B., & Bichelmeyer, B. A. (2020). Predictors of turnover intention in U.S. federal government workforce: Machine learning evidence that perceived comprehensive HR practices predict turnover intention. *Public Personnel Management*, 50, 538–558. <https://doi.org/10.1177/0091026020977562>
- Khera, S. N., & Divya. (2019). Predictive modelling of employee turnover in Indian IT industry using machine learning techniques. *Vision: The Journal of Business Perspective*, 23, 12–21. <https://doi.org/10.1177/0972262918821221>
- Mansor, N. N. B. A., Sani, N. S., & Aliff, M. (2021). Machine learning for predicting employee attrition. *International Journal of Advanced Computer Science and Applications*, 12. <https://doi.org/10.14569/ijacsa.2021.0121149>
- Musanga, V., Tarambiwa, E., & Zvarevashe, K. (2022). A supervised machine learning model to optimize human resources analytics for employee churn prediction. *Proceedings of the*

10.48047/jocaaa.2023.31.02.52

2022 Zimbabwe Conference on Information and Communications Technology (ZCICT).<https://doi.org/10.1109/zcict55726.2022.10045987>

- Najafi-Zangeneh, S., Shams-Gharneh, N., Arjomandi-Nezhad, A., & Zolfani, S. H. (2021). An improved machine learning-based employees attrition prediction framework with emphasis on feature selection. *Mathematics*, 9(11), 1226. <https://doi.org/10.3390/math9111226>
- Raza, A., Munir, K., Almutairi, M., Younas, F., & Fareed, M. M. S. (2022). Predicting employee attrition using machine learning approaches. *Applied Sciences*, 12(13), 6424. <https://doi.org/10.3390/app12136424>
- Saleh, S., Honi, D. G., Fallucchi, F., Abdulsada, A. I., Giuliano, R., & Abdulmalik, H. A. (2021). Employee attrition prediction using deep neural networks. *Computers*, 10(11), 141. <https://doi.org/10.3390/computers10110141>
- Wang, X., & Zhi, J. (2021). A machine learning-based analytical framework for employee turnover prediction. *Journal of Management Analytics*, 8, 351–370. <https://doi.org/10.1080/23270012.2021.1961318>
- Yiğit, İ. O., & Shourabizadeh, H. (2017). An approach for predicting employee churn by using data mining. In *2017 International Artificial Intelligence and Data Processing Symposium (IDAP)*. <https://doi.org/10.1109/idap.2017.8090324>