

EXPLORING VOLUNTARY EMPLOYEE ATTRITION PREDICTION USING PERSONALITY TRAITS AND MACHINE LEARNING: AN HR ANALYTICS AND DECISION SUPPORT STUDY BASED ON A PUBLIC DATASET

DOI: 10.5937/JEMC2600008Y

UDC:
Original Scientific Paper

Nour El Houda YAHIAOUI¹

¹Djillali Liabes University, Faculty of Economics and Business Management, Sidi Bel Abbas, Algeria
Corresponding author. E-mail: nour_el_houda.yahiaoui@univ-sba.dz
ORCID ID (<https://orcid.org/0000-0002-5604-7073>)

Paper received: 10.04.2026.; Paper accepted: 11.07.2026.

Voluntary employee attrition imposes substantial organizational costs through workforce instability and knowledge depletion. Despite the rapid growth of people analytics, most predictive voluntary attrition models rely on observable variables, overlooking the psychological determinants of voluntary attrition. This study addresses this gap by developing and evaluating machine learning models that integrate Big Five personality traits with HR data, through a precision-optimized evaluation framework aligned with HR intervention costs. Three algorithms were trained on a dataset of 1,129 employees - Logistic Regression, Support Vector Machine (SVM), and Random Forest (RF). Precision was prioritized over Recall to minimize wasted retention resources. The RF achieved the highest performance (AUC = 0.69, Precision = 0.74), outperforming Logistic Regression (AUC = 0.64) and SVM (AUC ≈ 0.66). All five Big Five traits ranked among the top seven predictors alongside tenure and age, suggesting their contribution to attrition risk. The findings support personality-informed models for evidence-based retention decisions and workforce planning.

Keywords: Voluntary employee attrition; Attrition prediction; Big Five personality traits; Machine learning; HR analytics; Decision support.

INTRODUCTION

Voluntary employee attrition as distinct from involuntary separations initiated by the employer is the self-initiated, permanent departure of employees from an organization. It is a persistent and strategic significant challenge in contemporary human resource management (HRM). Far beyond direct replacement costs, its organizational impact extends to eroding institutional knowledge, disrupting team cohesion, undermining competitive capacity (Gümüşsoy, 2016; Ployhart et al., 2014) and negatively predicting performance across financial, operational, and service quality outcomes (Hancock et al., 2013). At the individual level, replacement costs are estimated at one to two times annual salary scaling monotonically with role complexity and organizational tenure (Cascio, 2000). In 2023, the

average annual attrition rate in the U.S. was 3.6% (U.S. Bureau of Labor Statistics, 2023), with voluntary departures accounting for approximately 13% of the workforce annually and nearly 40% of all attritions occurring within the first year (Hubstaff, 2024). These patterns underscore the strategic urgency of developing robust, forward-looking methods to identify and mitigate voluntary attrition risk before departure occurs.

The field of HR analytics — defined as the systematic application of data-driven methods to human resource phenomena in support of evidence-based organizational decisions (Marler & Boudreau, 2017) and its adjacent discipline of people analytics, which foregrounds the human, organizational, and ethical dimensions of workforce data—, have emerged as consequential

developments in contemporary HRM theory and practice. Predictive attrition modeling represents a primary axis of utility within this paradigm; its analytical outputs directly condition strategic workforce planning decisions and retention resource allocation (Tursunbayeva et al., 2018). AI-supported organizational decision-making positions such models as specialized decision-support architectures designed to augment human managerial heuristics with structured, evidence-backed insights (Bondarouk & Brewster, 2016; Tambe et al., 2019). Within this supportive taxonomy, machine learning (ML) models have demonstrated an escalating utility, delivering a level of discriminatory power and predictive precision that traditional, linear statistical methodologies fail to replicate (Fallucchi et al., 2020; Guerranti & Dimitri, 2022). Nevertheless, this analytical evolution remains structurally impeded by two persistent limitations.

First, existing ML attrition models predominantly rely on measurable demographic and organizational variables such as income, tenure, age, gender, and workload. This narrow operational scope overlooks the latent psychological drivers that govern an individual's attrition decisions (Wu et al., 2024; Zimmerman, 2008). Second, the Big Five personality model —comprising Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism provides a theoretically grounded, empirically robust, and cross-culturally validated framework for mapping these psychological dimensions (Barrick & Mount, 1991; Costa & McCrae, 1992). The robust predictive validity of these specific personality dimensions regarding long-term work attitudes and behavioral attrition intentions underscores their value as critical parameters in advanced predictive analytics (Wu et al., 2024; Zimmerman, 2008). Nevertheless, personality data have been largely absent from ML-based attrition prediction frameworks. This absence constitutes both a scientific limitation and a practical gap for HR managers seeking to understand the motivational antecedents of departure rather than its demographic correlates alone.

The consequence of this dual omission is strategically significant. Models that exclude personality data may correctly identify *who* is likely to leave based on demographic and organizational signals, but remain structurally unable to explain *why*, limiting HR managers to reactive, demographically-targeted responses rather than

psychologically-informed proactive retention interventions. From a workforce planning perspective, this distinction matters because interventions targeting dispositional antecedents of attrition such as role redesign, wellbeing support, and personalized development pathways require a fundamentally different evidence base than interventions targeting demographic risk profiles alone.

This study is therefore framed as a management-oriented HR analytics investigation that employs machine learning as its analytical instrument, rather than a machine learning study that incidentally uses HR data. It addresses the identified limitations by developing and evaluating ML models that integrate Big Five personality traits with demographic and organizational HR variables for voluntary employee attrition prediction, with the explicit objective of generating precision-optimized decision-support outputs aligned with workforce planning and retention management priorities. Three research hypotheses guide the study: (H1) whether non-linear ML models outperform the Logistic Regression baseline in attrition prediction, as measured by AUC and Precision; (H2) whether Big Five traits demonstrate non-redundant feature importance alongside established HR variables in the optimized model; and (H3) whether the resulting feature importance profiles provide an analytical basis for moving from *who*-focused to *why*-focused attrition management, enabling psychologically-informed retention interventions.

This study contributes to the HR analytics literature on three interconnected fronts. Primarily, it converges the Big Five taxonomic paradigm with structural demographic and organizational covariates inside a singular, integrated predictive framework. Furthermore, the investigation establishes an operationally aligned evaluative framework that privileges precision over baseline classification accuracy; this explicitly accounts for the economic and strategic asymmetry between resource-wasting false positives and omission-based false negatives within human capital management. Finally, the study closes the gap between algorithmic output and localized HR intervention by operationalizing these predictive signals into structured workforce planning methodologies and tailored, psychologically-grounded retention architectures.

The paper is organized as follows. Section 2 reviews the relevant literature across voluntary attrition

theory and organizational consequences, Big Five personality traits' organizational implications, machine learning for attrition prediction, and HR analytics and people analytics practice. Section 3 describes the dataset, preprocessing pipeline, and modeling approach. Section 4 reports the empirical results. Section 5 discusses results, managerial implications, and study limitations. Section 6 concludes with recommendations for research and practice.

LITERATURE REVIEW

Voluntary Employee Attrition: Theoretical Foundations and Organizational Consequences

Voluntary attrition materializes is driven by the intricate confluence of idiographic dispositions, job attitudes, organizational conditions, and macro-environmental vectors (Holtom et al., 2008; Hom et al., 2017). The resulting institutional fallout spans a multi-tiered taxonomy of strategic liabilities. On an individual scale, replacement costs scale with role complexity and seniority, frequently exhausting one to two times the departed incumbent's annualized compensation (Cascio, 2000). At the meso-level, voluntary exits destabilize collective operational equilibrium, arresting group performance via immediate productive capacity deficits while simultaneously draining the localized human capital reservoir of vital tacit knowledge (Hausknecht & Holwerda, 2012). At the enterprise level, these micro- and group-level disruptions aggregate to function as a reliable negative predictor of overarching financial metrics, operational velocity, and service quality parameters (Hancock et al., 2013).

From the resource-based perspective, human capital constitutes a foundational cornerstone of sustainable competitive advantage, thereby elevating governance protocols regarding talent acquisition, skill cultivation, and asset preservation to core strategic imperatives (Becker & Huselid, 1998; Ployhart et al., 2014). Within this conceptual framework, elevated voluntary attrition transitions from a localized operational friction to an existential strategic threat—one that actively dismantles the scarce, inimitable resource configurations necessary for sustaining market differentiation (Hancock et al., 2013).

Workforce planning depends on the ability to anticipate voluntary attrition before it occurs. Traditional approaches relying on historical attrition

rates and demographic trend analysis are inherently backward-looking (Holtom et al., 2008). The integration of predictive people analytics into workforce planning enables a qualitatively different, forward-looking capability; the identification of individual-level flight risk in advance of departure, allowing proactive rather than reactive intervention. Organizations with high-performance HR systems, including proactive retention management, outperform competitors on both financial and operational metrics, evidence that the capability to retain talent is itself a strategic competence, not merely a desirable outcome (Becker & Huselid, 1998).

Big Five Personality Traits and Voluntary Employee Attrition

The Big Five personality model is a taxonomic framework developed assess human personality across five dimensions—namely Openness to Experience, Conscientiousness, Extraversion, Agreeableness, and Neuroticism (Costa & McCrae, 1992). Known as the OCEAN model, this model is the most validated and empirically supported framework for assessing personality (Costa & McCrae, 1992; Jang et al., 1996). Its cross-cultural validity, empirical robustness, and demonstrated utility for predicting organizationally relevant behaviors make it an adequate personality framework in applied organizational research (Ching et al., 2014).

Psychometric traits represent foundational constructs within human psychology, exercising substantial explanatory power across a spectrum of strategic human capital domains—ranging from the predictive modeling of work performance (He et al., 2019) to the granular taxonomy of dispositional variances separating frontline personnel, supervisory cohorts, and entrepreneurial agents (Kang et al., 2023). Within this framework, the Big Five taxonomic paradigm maintains an empirically rigorous relationship with employee attitudinal structures and behavioral outcomes. Conscientiousness has been linked to higher job performance across diverse occupational contexts (Barrick & Mount, 1991; He et al., 2019; Wilmot & Ones, 2019), and to both intrinsic and extrinsic indicators of career success, a pattern attributable to characteristics such as persistence, dependability, and achievement orientation. Conversely, Neuroticism functions as a primary driver of attrition risk. It is characterized by susceptibility to negative affectivity, job dissatisfaction, and stress.

This dimension is highly reliable predictor of both intentions and actual attrition (Judge et al., 1999; Zimmerman, 2008). Validating these mechanics within an enterprise-scale people analytics framework via online text data, Wu et al. (2024) quantified this dynamic empirically, establishing that a marginal one-unit escalation in Neuroticism triggers a 0.6% increase in voluntary attrition probability, whereas Conscientiousness and Agreeableness operate as stabilizing psychological counterweights that insulate personnel against attrition pressures.

From an HR management standpoint, incorporating personality traits into attrition prediction models shifts the retention to a strategically important direction. A model that flags short tenure and young age as risk factors tells managers who is likely to leave. A model that also surfaces high Neuroticism and low Conscientiousness begins to explain why — enabling interventions that target the psychological roots of disengagement rather than demographic risk profiles alone. This distinction between who-focused and why-focused analytics represents a key managerial benefit of personality-augmented models for strategic workforce management. Beyond these organizational-level patterns, voluntary attrition is also driven by internal, psychological antecedents. Job dissatisfaction, perceived lack of development opportunity, poor supervisory relationships, and work-life imbalance consistently emerge as proximal predictors of attrition intention and behavior (Holtom et al., 2008; Zimmerman, 2008). Crucially, these attitudes are themselves shaped by stable personality traits. The attrition literature describes a mediation pathway running from dispositional traits through job attitudes to departure behavior (Hom et al., 2017; Zimmerman, 2008) – a pathway with an important implication: models that rely solely on demographic and organizational variables may capture the attitudinal correlates of departure while systematically overlooking the dispositional antecedents that give rise to those attitudes in the first place.

Machine Learning for Voluntary Employee Attrition Prediction

Traditional approaches to attrition modeling, often relying on linear regression and simple statistical tests, are largely constrained by assumptions of linearity and struggle with the complexity inherent in human resource data. Studies emphasize that attrition is not the result of a single factor; the

decision to leave is a complex, non-linear function of individual traits interacting with organizational stressors, which simple linear regression cannot adequately capture (Zimmerman, 2008). A growing body of research supports the superior performance of ML approaches, particularly ensemble and deep learning methods (Arqawi et al., 2022), in capturing the non-linear, interactive relationships that drive attrition (Fallucchi et al., 2020).

Extant personnel attrition literature remains heavily anchored in the systematic benchmarking of alternative machine learning (ML) architectures (Guerranti & Dimitri, 2022; Mansor et al., 2021; Raza et al., 2022; Shlash et al., 2025). Within these comparative frameworks, Random Forest and Gradient Boosting models are frequently deployed due to their robust discriminatory threshold and structural resilience when confronting common data issues including multicollinearity and high-dimensional data noise (Govindarajan et al., 2025). Logistic Regression remains a pervasive method serving as the traditional statistical baseline against which all modern ML models are benchmarked. Support Vector Machines (SVM) continue to serve as a standard non-linear method, frequently included in comparative studies, particularly when smaller or less complex datasets are involved (Govindarajan et al., 2025; Guerranti & Dimitri, 2022; Mansor et al., 2021; Raza et al., 2022; Shlash et al., 2025). Furthermore, the integration of advanced ensemble architectures—most notably XGBoost—has empirically demonstrated that the strategic convergence of heterogeneous weak learners yields significantly higher predictive stability and classification accuracy than isolated models (Alsheref et al., 2022; Alshiddy & Aljaber, 2023; Vaiyapuri & Sbair, 2025). A parallel line of body research has increasingly gravitated toward deep learning topologies to capture intricate, latent feature interactions, though the systematic interpretability deficits native to these black-box configurations present an enduring barrier to practical implementation (Al-Alawi & Ghanem, 2024). This evolutionary trajectory underscores a decisive shift away from rudimentary ML models toward highly sophisticated computational paradigms, such as Transformers and Graph Neural Networks (GNNs), to achieve deeper feature abstraction, albeit at the explicit expense of reduced interpretability.

A significant development in attrition modeling concerns the incorporation of employee psychological characteristics alongside organizational and

demographic variables. While variables such as tenure and age remain dominant predictors, personality traits are essential for explaining why an employee chooses to leave. Yu et al. (2025) demonstrate the contribution of Big Five personality traits to enhanced predictive accuracy, arguing that these traits enable models to capture the intrinsic motivational states and emotional dispositions associated with long-term organizational commitment. The most recent literature also reflects a growing concern for practical implementation, shifting attention from the maximization of raw predictive accuracy toward the generation of actionable results. The growing adoption of Explainable AI (XAI), including techniques such as SHAP, constitutes a primary emerging trend (Baydili & Taşci, 2025), reflecting the recognition that opaque predictive models are of limited utility to HR managers. Current research is therefore directed toward transparent, interpretable insights into why an individual employee is classified as an attrition risk, advancing the field beyond binary classification toward actionable insight for managers (Benabou & Touhami, 2025; Konar et al., 2025; Manafi Varkiani et al., 2025).

HR Analytics, People Analytics, and Predictive Voluntary Employee Attrition

HR analytics — understood as the systematic application of data analysis to human resource phenomena in support of evidence-based organizational decisions — has emerged as a significant development in contemporary HRM theory and practice (Marler & Boudreau, 2017), enabling organizations to anticipate future workforce behaviors and facilitate proactive managerial responses. People analytics, which emphasizes the human, organizational, and ethical dimensions of data-driven workforce management, has grown rapidly as an adjacent field alongside HR analytics (Tursunbayeva et al., 2018). Tursunbayeva et al. (2018) identify predictive attrition modeling as among the highest-value applications of people analytics, precisely because model outputs (individual risk scores) translate directly into retention resource allocation decisions. Despite this potential, Angrave et al. (2016) caution that HR analytics adoption frequently fails in practice due to insufficient data infrastructure, a shortage of analytically capable HR professionals, and critically, the disconnection between statistical model outputs and managerial decision logic.

One of the structural challenges in applying AI to HRM concerns the ethical risks of algorithmic

decision-making in personnel contexts. The value of AI-based HR tools lies not in supplanting managerial judgment but in augmenting it by providing structured, evidence-based inputs to decisions that managers ultimately make (Tambe et al., 2019). Rasmussen and Ulrich (2015) similarly argue that HR analytics avoids becoming a management fad only when it demonstrates a clear and direct line of sight from model output to managerial action, positioning people analytics as one element within broader organizational value-creation chains. HR Decision Support Systems (HR DSS) enhance decision-making by reducing cognitive biases, enabling evidence review, and targeting retention investments (Bondarouk & Brewster, 2016). Retention interventions including salary adjustments and mentoring involve high per-employee costs, broad application results in poor return on investment. Precision-optimized predictive models allow organizations to maximize financial efficiency and the probability of retention success by focusing retention efforts only on employees who require them (Marler & Boudreau, 2017; Tambe et al., 2019).

Research Hypotheses

The following three hypotheses are grounded in the literature reviewed above and operationalized to connect directly to the comparative model evaluation, the role of personality traits, and the managerial interpretation of performance outcomes.

- H1: Non-linear machine learning models (Random Forest, SVM) demonstrate higher predictive performance for employee attrition than the Logistic Regression baseline, as measured by AUC and Precision, consistent with the non-linear and interaction-driven nature of the voluntary attrition decision process. This hypothesis is grounded in the established consensus that voluntary attrition emerges from the complex, non-linear interplay of personality dispositions, job attitudes, and organizational conditions (Hom et al., 2017; Zimmerman, 2008). Linear models cannot capture the interaction effects, for example, between high Neuroticism and short tenure, that drive attrition risk, while ensemble methods such as Random Forest explicitly model these interdependencies.
- H2: Big Five personality traits demonstrate non-redundant feature importance in the optimized Random Forest model, ranking among the top predictors alongside established HR variables (tenure, age), consistent with the expectation

that psychological dispositions contribute non-linear predictive value to attrition modeling. Given the established empirical relationships between personality traits, particularly Neuroticism and Conscientiousness, and attrition intention (Barrick & Mount, 1991; Wu et al., 2024; Zimmerman, 2008), personality dimensions are expected to emerge as meaningful predictors even where their linear correlation with the outcome variable is weak or negligible.

- H3: The feature importance profiles produced by the personality-augmented Random Forest model provide a basis for moving from who-focused to why-focused attrition management: the ranking of personality dimensions alongside demographic variables enables HR managers to identify not only which employees are at risk but also which dispositional antecedents are associated with departure risk, supporting the design of psychologically-informed retention interventions (Hom et al., 2017; Zimmerman, 2008).

METHODOLOGY

Dataset Description

This study uses a publicly available employee attrition dataset originally shared from Edward Babushkin's HR analytics blog (Babushkin, 2017) and accessible via Kaggle (Employee Turnover Dataset, n.d.). The

dataset's design logic reflects a deliberate integration of two data layers: standard organizational and demographic HR variables, which are routinely available in organizational HR information systems, and continuous psychometric scores for each of the five Big Five personality dimensions measured on a 1–10 scale, consistent with the scoring conventions of standardized personality assessment instruments. This dual-layer structure enables direct comparison of demographic-only versus personality-augmented attrition prediction models.

The dataset comprises 1,129 employee records across 16 attributes. The binary target variable is voluntary attrition (Event: 1 = leaves, 0 = stays). Demographic variables include age and gender; organizational variables include industry sector (15 categories), professional role (15 categories), recruitment channel (8 categories), coaching during the trial period, supervisor gender, wage transparency status, and commuting mode. Personality is operationalized through five continuous scores — Extraversion, Agreeableness (labeled *independ*), Conscientiousness (labeled *selfcontrol*), Neuroticism (labeled *anxiety*), and Openness (labeled *novator*) — each scaled 1–10 and approximately normally distributed (range 2–8), consistent with psychometric instrument norms. The dataset contains no missing values across all 16 columns. Table 1 provides a complete feature-level description.

Table 1: Dataset Feature Description

Feature	Type	Description
stag	Numerical	Work experience / tenure in years (right-skewed distribution)
event	Binary	Target: Attrition = 1 (leaves), Non-attrition = 0 (stays)
gender	Categorical	Employee gender (male / female)
age	Numerical	Employee age in years (approx. normally distributed)
industry	Categorical	Employment sector: 15 categories (IT, Banks, Retail, etc.)
profession	Categorical	Professional role: 15 categories (HR, Sales, Finance, etc.)
traffic	Categorical	Recruitment channel: 8 categories (referral, advertising, etc.)
coach	Categorical	Coaching during trial period (yes / no / manager)
head_gender	Categorical	Immediate supervisor's gender
greywage	Categorical	Wage transparency: white (declared) or grey (undeclared)
way	Categorical	Commuting mode (bus / car / foot)
extraversion	Numerical (1-10)	Big Five: Extraversion
independ	Numerical (1-10)	Big Five: Agreeableness / Independence
selfcontrol	Numerical (1-10)	Big Five: Conscientiousness / Self-control
anxiety	Numerical (1-10)	Big Five: Neuroticism / Anxiety
novator	Numerical (1-10)	Big Five: Openness / Innovativeness

Source: (Employee Turnover Dataset, n.d.)

Three dataset-level limitations must be acknowledged before interpretation. First, the specific personality assessment instrument used in

the original data collection is not documented by Babushkin (2017), a limitation that introduces construct validity uncertainty: while the five

dimensions correspond nominally to the OCEAN model, the psychometric properties (reliability, norm group, cultural calibration) of the instrument cannot be verified. Second, the dataset's origin in a non-specified organizational context means that industry structures, labor market dynamics, and the cultural expression of Big Five traits may differ meaningfully across national contexts, limiting direct generalizability (Ching et al., 2014). Third, the near-equal class distribution (50.6% attrition, 49.4% non-attrition) is atypical of real organizational environments where attrition rates are generally 10–20%; metrics reported in this study should therefore be interpreted as upper-bound estimates under idealized balance conditions, with performance likely declining under realistic class imbalance.

Exploratory Data Analysis (EDA)

Exploratory data analysis (EDA) established clear associations between organizational variables and attrition. The largest occupational groups, HR and Retail, recorded the highest absolute numbers of attrition events. Specialized roles (Teaching, Commercial, Finance) and certain industries (Banks, Consult, Manufacturing) exhibited the highest proportionate risk of attrition. HoReCa (Hotels, Restaurants, Cafes) and Real Estate showed very low absolute counts of both employees and attrition events (see Table 7 in the appendix).

From industry and profession-level attrition patterns, the two largest professional groups HR and Retail recorded the highest absolute numbers of attrition events. Both fields are characterized by relatively low barriers to exit, high labor market mobility, and, in the case of Retail, significant proportions of short-tenure workers. At the proportionate risk level, specialized roles in Teaching, Commercial, and Finance exhibited the highest attrition rates relative to their headcount, a finding of particular relevance for workforce planning in knowledge-intensive organizations, where the departure of specialized employees carries disproportionate human capital costs (Hausknecht & Holwerda, 2012; Ployhart et al., 2014). At the industry level, Banks, Consulting, and Manufacturing exhibited the highest proportionate attrition risk, while HoReCa (Hotels, Restaurants, Cafes) and Real Estate showed very low absolute counts of both employees and attrition events (see Figure 5 in the appendix). Attrition patterns based on coaching during the trial period and commuting mode are detailed (see Figure 6 in the appendix).

From a gender composition and attrition rates perspective, female employees represented the majority of the workforce ($\approx 75.6\%$ of headcount) and recorded a slightly higher proportional attrition rate ($\approx 51.1\%$) compared to male employees ($\approx 48.9\%$; see Figure 5 in the appendix). While this difference is modest in absolute terms, it carries HR implications: organizations with female-majority workforces should examine whether their retention programs and career development pathways are adequately responsive to the work-life balance pressures and perceived development equity that the literature identifies as disproportionately affecting female employees' voluntary attrition intentions (Hom et al., 2017). This gender pattern also raises the algorithmic fairness consideration discussed in Section 5.3; a model trained on this data may encode gender-correlated risk signals that require explicit fairness auditing before organizational deployment.

Data Preprocessing and Feature Engineering

All preprocessing steps described below were applied to the training set only after the 70/30 train-test split; the resulting parameters (means, standard deviations, encoding schemes) were then applied to the test set without refitting, thereby preventing data leakage. Preprocessing proceeded through four stages. First, personality traits displayed near-zero linear correlation with the outcome variable (e.g., Openness: $\rho \approx 0.007$; Figure 1), indicating the absence of a meaningful linear relationship between personality dimensions and attrition. This near-zero Pearson correlation is equally consistent with the absence of any relationship and with the presence of non-linear or interaction effects that linear measures cannot detect. The theoretical basis for expecting non-linear contributions from personality traits derives from the attrition literature, which posits complex mediation pathways involving multiple interacting variables (Hom et al., 2017; Zimmerman, 2008). The empirical test of whether such non-linear signal exists is provided by the model performance comparison in Section 4.1, where Random Forest substantially outperforms Logistic Regression. Independently of this theoretical reasoning, the near-zero linear correlations justify the choice of non-parametric ensemble methods: when linear associations are absent, linear models will systematically underfit, while methods such as Random Forest, which make no linearity assumption, remain capable of detecting whatever structure exists in the data. Second, distributional analysis guided normalization: age

was approximately normally distributed; tenure was strongly right-skewed (necessitating log transformation); personality scores showed approximately normal distributions (range 2–8), consistent with scaled psychological instruments, and required no transformation (see Figure 7 in the Appendix). Third, all categorical variables were converted to machine-readable format via one-hot

encoding, generating 50 binary indicator columns from 8 original categorical features (Table 2). Fourth, all numerical features were standardized using Z-score normalization (StandardScaler) to ensure that no single variable exerts disproportionate influence on model learning due to scale differences — a requirement particularly critical for SVM.




Figure 1: Pearson Correlation Heat map: Feature Relationships and Personality Traits Multicollinearity

Table 2: One-Hot Encoding of Categorical Variables

Feature	Categories Generated	Example Values
Gender	2	Female, Male
Industry	15	IT, Banks, Retail, Manufacture, Consult, etc.
Profession	15	HR, Sales, Finance, Teaching, Engineering, etc.
Traffic	8	Referral, Advertising, Friends, KA, etc.
Coach	3	No, Yes, Manager
Head Gender	2	Female, Male
Greywage	2	White, Grey
Way	3	Bus, Car, Foot

Class Distribution and Sampling Strategy

The dataset exhibits a near class balance: 50.6% attrition events and 49.4% non-attrition, atypical of real organizational environments where attrition rates are generally much lower (10–20%). This balance eliminates the need for synthetic oversampling techniques such as SMOTE and improves metric reliability. However, it also introduces an important deployment caveat: models trained on balanced data may underperform when applied to real organizational settings with substantially lower attrition rates. Under realistic class imbalance, both Precision and Recall are expected to decrease relative to the figures reported here, and threshold recalibration or class-weighted loss functions would be required before deployment. This generalizability limitation is discussed further in Section 5.4. The 70/30 train-test

split used stratified sampling to preserve this ratio across both partitions using a fixed random seed of 42 for reproducibility.

Feature Selection: Recursive Feature Elimination with Cross-Validation

Recursive Feature Elimination with Cross-Validation (RFECV) was applied using a Random Forest estimator and 10-fold Stratified Cross-Validation, optimizing for the cross-validated accuracy score. This step is managerially relevant: a parsimonious model whose retained predictors are interpretable provides more actionable HR guidance than one absorbing all available variables indiscriminately. RFECV identified tenure and age as the two dominant predictors, consistent with the HRM literature on organizational tenure and early-career vulnerability (Hausknecht & Holwerda, 2012; Holtom et al., 2008). More importantly, all five Big Five personality trait scores were retained among the top seven most influential features.

Algorithm Selection, Training, and Evaluation Framework

Three algorithms were trained: Logistic Regression (linear statistical baseline), SVM with RBF kernel (non-linear comparator), and Random Forest (ensemble method). This progression from simple to complex model architectures enables rigorous testing of Hypothesis H1. The evaluation framework was constructed to reflect the cost structure of HR retention decisions rather than to optimize abstract accuracy metrics. In a retention intervention context, two error types carry fundamentally different organizational costs:

- False Positive: A stable employee incorrectly classified as a flight risk. This error results in the misallocation of retention resources (salary adjustments, coaching, development programs), potential adverse morale effects from unwarranted

managerial scrutiny, and poor return on the HR intervention budget.

- False Negative: A genuine leaver not identified by the model. This error results in a missed intervention opportunity; however, it represents a cost that organizations already incur in the absence of any predictive system. The marginal organizational harm associated with a false negative is therefore lower than that of a false positive in resource-constrained settings.

This asymmetry justifies prioritizing Precision over Recall. All analyses were implemented in Python 3.10 using the scikit-learn library (v1.3). A model with Precision = 0.74 delivers approximately 7.4 correct identifications per 10 intervention decisions — a meaningful return on retention expenditure. Generalization was assessed through 20-fold cross-validation. Hyperparameter optimization used Grid Search with 5-fold cross-validation optimizing for F1-score. The Grid Search explored the following parameter space: $n_estimators \in \{100, 200, 300\}$, $max_depth \in \{10, 20, None\}$, $min_samples_split \in \{2, 5, 10\}$, and $criterion \in \{gini, entropy\}$. The final optimized Random Forest architecture comprised 300 estimators, maximum depth of 20, minimum 2 samples per split, and Gini impurity as the splitting criterion.

RESULTS

Comparative Model Performance

Table 3 presents the initial performance metrics across the three models. The Random Forest outperformed both alternatives on all reported metrics — overall accuracy (0.69), Class 1 F1-score (0.69), and Class 1 Recall (0.65) — consistent with Hypothesis H1. Logistic Regression underperformed consistently, corroborating the non-linear, interaction-driven nature of the attrition prediction task. SVM achieved intermediate performance at substantially higher computational cost.

Table 3: Comparative Initial Model Performance

Model	Overall Accuracy	Class 1 F1-Score	Class 1 Recall	AUC
Logistic Regression	0.64	0.62	0.57	0.64
SVM (RBF kernel)	0.67	0.67	0.64	~0.66
Random Forest	0.69	0.69	0.65	0.69

The Random Forest's AUC of 0.69 reflects meaningful discriminatory capacity for a real-world attrition dataset incorporating genuine psychological trait measurements. The convergence of AUC (0.692) and test accuracy (0.690) — attributable to the near-equal

class distribution — is consistent with stable discriminatory performance across classification thresholds, supporting the model's reliability as an HR decision-support instrument (Figure 2).




Figure 2: ROC Curve Random Forest vs Logistic regression discriminatory performance

Feature Importance and the Predictive Value of Personality Traits

Feature importance rankings from the optimized Random Forest provide the most managerially relevant empirical finding of this study, and are consistent with Hypothesis H2 (Figure 3). All five Big Five personality traits ranked among the top seven predictive features alongside tenure and age. The top seven predictive features in descending order were: (1) tenure/stag, (2) age, (3) self-control/Conscientiousness, (4) independ/Agreeableness, (5) novator/Openness, (6) extraversion, and (7) anxiety/Neuroticism, outranking industry sector, professional role, and recruitment channel variables. This finding shows that the personality traits' predictive power is primarily non-linear and interaction-driven, undetectable by Logistic Regression but captured by the Random Forest's ensemble architecture. From a strategic HRM, this suggests that ethically-collected personality assessment data contributes non-redundant predictive value to attrition risk profiling, and that HR information systems should integrate such data rather than treating personality assessments as standalone, one-time tools disconnected from workforce management analytics.

Learning Algorithm Evaluation

Before finalizing the Random Forest as the primary model, a broader learning algorithm evaluation was

conducted using learning curve analysis — a diagnostic tool that reveals whether a model suffers from high bias (underfitting), high variance (overfitting), or neither, and how efficiently it scales with additional training data. Two additional algorithms were evaluated — Naive Bayes and SVM — precisely because their learning curve profiles provide the theoretical contrast that makes the case for Random Forest. The results of this evaluation, summarized in Table 4, are consistent with Hypothesis H1.

Naive Bayes — high bias (underfitting). The Naive Bayes learning curves showed the training score and cross-validation score converging at a low ceiling (≈ 0.58 – 0.60), with both curves lying very close together throughout. This pattern is the diagnostic signature of high bias: the model is too simple to detect the underlying patterns in the data, and increasing the training set size produces no improvement in performance. From a scalability perspective, Naive Bayes is computationally excellent — training time is negligible — but scalability is irrelevant when the model cannot learn. This result is consistent with the theoretical expectation: Naive Bayes assumes conditional independence among all features, which is violated in this dataset by the strong inter-correlations among personality traits and their non-linear interactions with organizational variables.




Figure 3: Most Important Features: Recursive Feature Elimination

SVM (RBF kernel) — moderate variance, poor scalability. The SVM learning curves revealed a different profile: for smaller training samples, the training score substantially exceeded the cross-validation score, indicating moderate overfitting (variance) — Figure 4. As the training sample grew, the two curves converged — consistent with the known behavior of kernel SVMs, which improve with additional data but slowly. Cross-validation scores stabilized between 0.60 and 0.62, below the Random Forest's performance level. Critically, training time for the SVM increased non-linearly with sample size, reflecting the $O(n^2)$ to $O(n^3)$ computational complexity characteristic of kernel methods. This represents an unfavorable trade-off for organizational deployment: marginally better performance than Naive Bayes at a substantially higher computational cost, and with performance still well below Random Forest.

The learning curve analysis therefore provides the theoretical grounding for the model selection decision. Both Naive Bayes (≤ 0.60) and SVM (≤ 0.62) achieved substantially lower performance than Random Forest (0.69), and both exhibited fundamental structural limitations (underfitting and poor scalability respectively) that make them unsuitable for the complex, non-linear, interaction-rich feature space that characterizes the personality-augmented attrition prediction problem. Random Forest avoids both failure modes: it is designed to capture non-linear feature interactions — the type of structure that personality trait predictors introduce into the feature space — without the variance instability of individual decision trees, and it scales acceptably for organizational-scale HR datasets. These results are consistent with Hypothesis H1.

Table 4: Learning Curves, Scalability, and Performance Comparison Summary

Model	Bias / Variance Profile	Performance Score	Scalability	HR Deployment Suitability
Naive Bayes	High Bias - Underfitting	$\approx 0.58 - 0.60$	Excellent (Fast)	Not suitable - cannot capture non-linear patterns
SVM (RBF kernel)	Moderate Variance/ Overfitting risk	$\approx 0.60 - 0.62$	Poor (Slow - $O(n^2)$)	Not suitable - poor scale, moderate variance
Random Forest	Low Bias, Low Variance/ Optimal	0.69 (Test Accuracy)	Good (Parallelizable)	Selected - best performance and fit for HR context




Figure 4: Learning Curves, Stability and Performance

Hyperparameter Tuning and the Precision-Optimized Final Model

Hyperparameter tuning via Grid Search produced the performance shift summarized in Table 5. The tuned model achieved a Class 1 Precision of 0.74,

surpassing the 0.70 threshold, demonstrating that the precision-first optimization strategy calibrated the model toward HR decision-support requirements. Table 6 presents the confusion matrix for the optimized Random Forest at the default classification threshold of 0.50. Threshold adjustment could trade

Precision for Recall depending on the organizational context and intervention budget.

The practical interpretation of the final model's performance is as follows: when the model classifies an employee as a potential flight risk, that classification is correct in approximately 74% of cases. For every 10 employees identified for targeted retention intervention, approximately 7.4 are genuine departure risks – a return rate that justifies the allocation of structured, individualized retention

resources and compares favorably with the documented accuracy of unstructured managerial judgment. The Recall of 0.61 indicates that approximately 39% of actual leavers remain unidentified, a limitation addressed in Section 5.4. Nevertheless, the model's marginal value relative to unstructured targeting remains important, particularly in organizations that currently lack any systematic attrition risk assessment.

Table 5: Model Performance Before and After Hyperparameter Tuning

Metric	Before Tuning	After Tuning	Interpretation
Overall Test Accuracy	0.69	0.68	Marginal decline: acceptable
F1-Score (Class 1: Attrition)	0.69	0.67	Slight decline: expected trade-off
Precision (Class 1: Attrition)	0.73	0.74	Primary HR decision target achieved
Recall (Class 1: Attrition)	0.65	0.61	Accepted cost of precision-first optimization

Table 6: Confusion Matrix: Optimized Random Forest (Classification Threshold = 0.50, Test Set n = 339)

	Predicted: Stays (0)	Predicted: Leaves (1)	Total Actual
Actual: Stays (0)	TN = 130	FP = 37	167
Actual: Leaves (1)	FN = 67	TP = 105	172
Total Predicted	197	142	339

Of 172 employees who actually departed, the model correctly identified 105 (TP) while missing 67 (FN). Of 167 employees who stayed, 130 (TN) were correctly classified as low-risk while 37 (FP) were incorrectly flagged for intervention. The FP count of 37 is the operationally critical figure: these represent the wasted retention resources that the Precision metric directly controls.

DISCUSSION

Implications for HR Decision-Making

The findings of this study carry direct implications in HRM. The connection between model output and managerial action satisfies the criterion that Rasmussen and Ulrich (2015) identify as the defining test of whether people analytics generates genuine organizational value. The optimized model produces a rank-ordered register of at-risk employees that HR managers may use to concentrate retention investments, including career development discussions, mentoring programs, compensation benchmarking, and wellbeing support, on those individuals who present the

greatest departure risk. With Precision = 0.74, 74% of employees targeted for intervention are genuine departure risks, substantially improving return on investment relative to undifferentiated retention programs. This precision-targeted approach operationalizes the strategic HRM argument that selective investment in identified high-risk employees maximizes human capital retention efficiency (Becker & Huselid, 1998). The ranking of personality traits among the top predictors enables HR managers to transition from who-focused to why-focused attrition management, the fact that supports Hypothesis H3. Employees with elevated anxiety and low self-control profiles constitute a distinct risk category that may benefit from psychologically-informed interventions including stress management support and wellbeing programs. As the model identifies statistical associations rather than causal mechanisms, intervention design should draw on these findings in conjunction with individual managerial judgment. This is consistent with Zimmerman's (2008) finding that personality traits are associated with attrition through satisfaction and commitment pathways that generic retention packages do not address. Tenure's

position as the most influential predictor, with short-tenure employees presenting the highest risk, is consistent with the case for investment in structured onboarding programs and first-year mentoring as retention instruments. This aligns with Holtom et al.'s (2008) finding that organizational embeddedness established early in tenure is a powerful protective factor against voluntary departure, and with the empirical fact that 40% of all departures occur within the first year (Hubstaff, 2024).

The practical value of the model's Precision = 0.74 becomes clearest when benchmarked against the realistic alternative: unstructured or intuition-based identification of flight risk by line managers. Research on managerial judgment accuracy in HR contexts indicates that unaided identification of at-risk employees performs only marginally above chance (Highhouse, 2008; Kuncel et al., 2013), meaning approximately 50–55% of employees flagged for intervention by managers without analytical support are not genuine departure risks. At Precision = 0.74, the model reduces wasted intervention decisions from approximately 5 per 10 to approximately 2.6 per 10, a reduction of nearly 50% in misdirected retention expenditure. To translate this into financial terms: if a structured retention intervention (career development discussion, compensation benchmarking, mentoring assignment) carries an average direct cost of \$500–\$1,500 per employee, an organization targeting 100 employees annually saves approximately \$12,000–\$36,000 in misdirected expenditure relative to unstructured targeting, before accounting for the indirect costs of managerial time and employee morale effects of unwarranted scrutiny. This calculation establishes a quantifiable return-on-investment rationale for precision-optimized attrition modeling that moves beyond abstract metric reporting toward the organizational value chain that Rasmussen and Ulrich (2015) identify as the defining test of HR analytics.

Contribution to HR Analytics and People Analytics Practice

This study contributes to the people analytics literature in a way that is distinct from prior machine learning attrition work. Rather than maximizing predictive accuracy as a standalone criterion, this study advances the argument that evaluation metrics must be aligned with the organizational cost structure of the decisions they inform. The

precision-first evaluation framework developed here is a transferable methodological contribution: any organization deploying an attrition model to guide costly retention interventions should calibrate its evaluation toward Precision rather than overall accuracy or Recall.

More broadly, the exploratory inclusion of Big Five personality traits contributes to the evidence base for psychological people analytics — the integration of personality, motivation, and affect data into HR decision support systems (Tambe et al., 2019; Tursunbayeva et al., 2018). The finding that personality traits demonstrated predictive signal not detectable through linear correlation analysis has direct implications for how HR teams should structure their data infrastructure. Specifically, personality assessments administered during recruitment or onboarding should be retained in the employee data record and incorporated into retention risk models, rather than being treated as standalone instruments. This recommendation must be implemented within a robust ethical and legal framework discussed in the next section.

Ethical Considerations in Personality-Informed HR Decision-Making

Integrating psychometric indicators into algorithmic HR decision-making raises a triad of ethical, jurisprudential, and privacy liabilities that demand rigorous structural mitigation prior to organizational deployment. The first is systemic algorithmic bias. Empirical personality constructs correlate with demographic variables such as gender and cultural background (Ching et al., 2014; Kang et al., 2023), and a model incorporating personality scores may indirectly encode and amplify pre-existing demographic disparities. Pre-deployment fairness auditing, entailing the systematic examination of model predictions across protected demographic subgroups, is a prerequisite for responsible implementation, mentioned by Newman et al. (2020) in their analysis of algorithmic reductionism in HR decision processes. Second, self-fulfilling prophecy risk. If HR managers are informed that a particular employee has been classified as a flight risk, they may unconsciously reduce developmental investment in that individual, a response that would itself increase the probability of the predicted departure. Deployment protocols need therefore to frame model outputs explicitly as decision-support inputs subject to managerial review, rather than as deterministic classifications that autonomously

determine HR actions, and should incorporate managerial training on the appropriate interpretation and use of algorithmic risk scores. Third, transparency and power asymmetry. The deployment of personality-informed attrition models creates a structural information asymmetry; the organization gains predictive insight into individual employees' psychological profiles and departure probabilities that employees themselves may not be aware of, cannot access, and cannot contest. This asymmetry raises questions of procedural fairness that extend beyond legal compliance. Organizations deploying such models should consider whether employees are informed that personality assessments administered at onboarding will be used in ongoing attrition risk profiling, and whether employees have a right to view, query, or request correction of their risk classification. Transparency in algorithmic HR processes is increasingly recognized as both an ethical obligation and a legal requirement under data protection frameworks that include the right to explanation (Newman et al., 2020; Tambe et al., 2019).

Another distinct managerial risk concerns automation bias: the documented tendency of human decision-makers to over-rely on algorithmic outputs and to reduce the critical scrutiny they would otherwise apply to the same decision (Highhouse, 2008). In the HR attrition context, automation bias manifests as treating every model-generated risk flag as actionable fact rather than as one probabilistic input among several. Given that the model misclassifies approximately 26% of flagged employees, deployment protocols need to actively counteract automation bias through structured override procedures, mandatory managerial review of all flagged cases, and regular feedback loops that expose managers to the model's error rate. A related risk is the accountability gap: when model-guided retention decisions produce adverse outcomes, the diffusion of responsibility between algorithmic output and human judgment can prevent organizational learning and impede the attribution of responsibility. Clear governance protocols specifying the respective roles of model, HR analyst, and line manager in each retention decision are a prerequisite for responsible deployment (Bondarouk & Brewster, 2016).

Data retention and misuse risk. The recommendation to retain personality assessment data in employee records carries specific misuse risks that organizations need actively to govern.

Psychological profile data collected for voluntary attrition prediction could be repurposed — deliberately or inadvertently — to inform promotion decisions, redundancy selection, performance evaluations, or insurance-related determinations, all of which lie outside the declared purpose of the attrition model. This scope extension is a violation of the purpose limitation principle and represents a form of psychological profiling that employees have not consented to. Organizations need to implement technical and procedural safeguards that constrain personality data to the specific HR decision context for which it was collected and for which consent was obtained.

Study Limitations and Future Horizons

Although this study contributes to the HR analytics literature, a number of limitations warrant acknowledgement and should be addressed in future research. First, the dataset is a publicly available secondary source (Babushkin, 2017) from a single organizational context whose provenance is not fully documented and has not been validated against an independent organizational dataset, limiting the external validity of the personality-attrition relationships identified. Second, the near-equal class distribution (50.6%/49.4%) is atypical of real organizational environments, where attrition rates are generally considerably lower. Models trained on balanced data may underperform when deployed in settings characterized by class imbalance, and the Precision and Recall metrics reported here may therefore overstate performance under realistic operational conditions. Thus, deployment in real organizational contexts requires threshold recalibration and likely retraining on organization-specific data before use in environments with lower attrition rates. Third, the dataset does not specify the national or cultural context of its employees, which limits cultural and organizational generalizability. Given that Big Five trait profiles and their relationships with work-related attitudes vary meaningfully across cultural contexts (Ching et al., 2014), the model requires recalibration before deployment in different geographic or sectoral environments. Fourth, the Random Forest algorithm provides feature importance rankings but does not reveal the functional form or directional nature of predictor-outcome relationships. The integration of XAI techniques (such as SHAP or LIME) in future research would provide directional, individual-level explanations of attrition risk, substantially enhancing the model's managerial utility by yielding specific, actionable explanations (e.g.,

"Employee X is a risk due to low Self-control and short Tenure") to line managers, thereby advancing beyond aggregate feature importance rankings. Future research can also examine interaction terms between personality traits and organizational factors (e.g., the interaction of Independence and Manager Gender) as additional features to further boost non-linear model performance. Investigating the performance of deep learning architectures, such as Recurrent Neural Networks (RNNs) or deep feedforward networks, may uncover more intricate non-linear patterns within the data, potentially yielding higher discriminatory performance. Fifth, the study employs a static, cross-sectional design; the model treats attrition as a binary classification problem at a single point in time, thereby neglecting the temporal dynamics of voluntary attrition decisions. Time-to-event modeling (Shabbir et al., 2025) represents a natural and important methodological extension. A final structural boundary condition concerns the external validity of the empirical architecture, which remains inherently circumscribed by its reliance on a singular organizational dataset. Consequently, the specific psychometric-attrition dynamics uncovered within this localized corporate ecosystem resist uncritical extrapolation across divergent industrial sectors, national labor markets, or distinct socio-cultural environments absent systematic replication utilizing independent, heterogeneous data samples.

Apart from dataset and algorithmic limitations, several organizational constraints affect the practical validity of personality-augmented attrition models. First, models trained on a static historical snapshot cannot adapt to organizational changes reshaped by mergers, restructurings, leadership transitions, and cultural shifts alter the attrition dynamics, requiring periodic retraining. Second, deployment in a real HR information system requires integration with existing HRIS infrastructure and data governance frameworks. Third, a feedback loop problem arises once the model is deployed: if retention interventions successfully reduce departure rates among high-risk employees, the outcome distribution shifts and the model's training data become progressively less representative of the current organizational reality. Fourth, organizational resistance from HR professionals, line managers, and employees themselves may limit adoption; the literature on HR analytics implementation identifies human and cultural resistance as a more common cause of analytics failure than technical inadequacy (Angrave et al., 2016). Fifth, the model's outputs are

only as actionable as the retention interventions available to the organization; a model correctly identifying high-Neuroticism, short-tenure employees as flight risks delivers limited value in organizations that lack the management capability, budget, or program infrastructure to respond with differentiated, psychologically-informed retention actions.

CONCLUSIONS

This study developed, evaluated, and interpreted a machine learning-based voluntary employee attrition prediction framework integrating Big Five personality traits with organizational and demographic data, contributing to both the HR analytics literature and the practical domain of workforce management. Four principal conclusions emerge.

First, non-linear ensemble models outperformed linear approaches in this study, consistent with the theoretical position that voluntary attrition arises from complex, interacting antecedents that linear models are structurally unable to capture. Second, Big Five personality traits demonstrated non-redundant predictive signal, ranking among the top seven features despite near-zero linear correlation with the outcome variable, supporting the case for their further investigation within HR analytics frameworks and HR information systems. Third, model evaluation in HR decision contexts needs to be governed by the cost structure of the interventions being planned. When false positives impose greater organizational costs than false negatives, Precision — rather than overall accuracy — constitutes the appropriate primary evaluation criterion. Fourth, personality-augmented attrition models have the potential to transform the character of the retention management conversation from who-focused to why-focused, enabling interventions that address the psychological antecedents of dissatisfaction rather than demographic risk profiles alone.

For HR practice, the findings support the integration of personality assessment data into HR information systems and the adoption of precision-optimized attrition models as decision-support instruments for retention resource allocation and early-tenure risk management. For research, four directions are recommended: (1) integrating SHAP analysis to provide directional, employee-level explanations of attrition risk; (2) replicating the personality-augmented framework on field-collected

organizational datasets with real-world class imbalance; (3) exploring time-to-event modeling to predict when departure will occur rather than merely whether it will; and (4) examining the organizational and cultural boundary conditions under which personality-informed retention interventions produce the strongest positive outcomes, a critical gap between prediction and prescription in the people analytics pipeline.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this manuscript, the authors used ChatGPT (OpenAI, GPT-5) solely as an assistive tool for language editing and improving the clarity of the text. The author(s) reviewed and verified all AI-assisted content and accept full responsibility for the final version of the manuscript.

REFERENCES

- Al-Alawi, A. I., & Ghanem, Y. A. (2024). Predicting Employee Attrition Using Machine Learning: A Systematic Literature Review. In *2024 ASU International Conference in Emerging Technologies for Sustainability and Intelligent Systems, ICETIS 2024*, 526-530. <https://doi.org/10.1109/icetis61505.2024.10459451>
- Alsheref, F. K., Fattoh, I. E., & Mead, W. (2022). Automated Prediction of Employee Attrition Using Ensemble Model Based on Machine Learning Algorithms. *Computational Intelligence and Neuroscience*, 2022(1), 7728668. <https://doi.org/10.1155/2022/7728668>
- Alshiddy, S. M., & Aljaber, B. N. (2023). Employee Attrition Prediction using Nested Ensemble Learning Techniques. *International Journal of Advanced Computer Science and Applications*, 14(7), 932-938. <https://doi.org/10.14569/ijacsa.2023.01407101>
- Angrave, D., Charlwood, A., Kirkpatrick, I., Lawrence, M., & Stuart, M. (2016). HR and analytics: Why HR is set to fail the big data challenge. *Human Resource Management Journal*, 26(1), 1-11. <https://doi.org/10.1111/1748-8583.12090>
- Arqawi, S. M., Abu Rumman, M. A., Zitawi, E. A., Rabaya, A. H., Sadaqa, A. S., Abunasser, B. S., & Abu-Naser, S. S. (2022). Predicting Employee Attrition and Performance Using Deep Learning. *Journal of Theoretical and Applied Information Technology*, 100(21), 6526-6536.
- Babushkin, E. (2017). Employee Turnover: How to Predict Individual Risks of Quitting. *HR Analytics Blog*. Retrieved from <https://edwvb.blogspot.com/2017/10/employee-turnover-how-to-predict-individual-risks-of-quitting.html>
- Barrick, M. R., & Mount, M. K. (1991). The Big Five personality dimensions and job performance: A meta-analysis. *Personnel Psychology*, 44, 1-26. <https://doi.org/10.1111/j.1744-6570.1991.tb00688.x>
- Baydili, I. T., & Taşci, B. (2025). Predicting Employee Attrition: XAI-Powered Models for Managerial Decision-Making. *Systems*, 13(7), 583. <https://doi.org/10.3390/systems13070583>
- Becker, B. E., & Huselid, M. A. (1998). High-performance work systems and firm performance: A synthesis of research and managerial implications. *Research in Personnel and Human Resources Management*, 16, 53-101. JAI Press Inc.
- Benabou, A., & Touhami, F. (2025). Optimising HRM practices in call centres: predicting and explaining employee turnover intention using classification and regression trees. *International Journal of Organizational Analysis*, 34(3), 870-891. <https://doi.org/10.1108/ijoa-12-2024-5117>
- Bondarouk, T., & Brewster, C. (2016). Conceptualising the future of HRM and technology research. *The International Journal of Human Resource Management*, 27(21), 2652-2671. <https://doi.org/10.1080/09585192.2016.1232296>
- Cascio, W. F. (2000). Costing Human Resources: The Financial Impact of Behavior in Organizations (4th ed.). *South-Western College Publishing*.
- Ching, C. M., Church, A. T., Katigbak, M. S., Reyes, J. A. S., Tanaka-Matsumi, J., Takaoka, S., Zhang, H., Shen, J., Arias, R. M., Rincon, B. C., & Ortiz, F. A. (2014). The manifestation of traits in everyday behavior and affect: A five-culture study. *Journal of Research in Personality*, 48(1), 1-16. <https://doi.org/10.1016/j.jrp.2013.10.002>
- Costa, P. T., & McCrae, R. R. (1992). The five-factor model of personality and its relevance to personality disorders. *Journal of Personality Disorders*, 6(4), 343-359. <https://doi.org/10.1521/pedi.1992.6.4.343>
- Employee Turnover Dataset. (n.d.). Kaggle. Retrieved September 23, 2025, from <https://www.kaggle.com/datasets/davinwijaya/employee-turnover>
- 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>
- Govindarajan, R., Kumar, N. K., Sudhakar Reddy, P., Sai Pravalika, E., Dhatri, B., & Pavan Kumar, G. (2025). Predicting Employee Attrition: A Comparative Analysis of Machine Learning Models Using the IBM Human Resource Analytics Dataset. *Procedia Computer Science*, 258, 4084-4093. <https://doi.org/10.1016/j.procs.2025.04.659>
- 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>

- Gümüşsoy, C. A. (2016). The Effect of Five-Factor Model of Personality Traits on Turnover Intention among Information Technology Professionals. *Online Academic Journal of Information Technology*, 7(22), 7–28. <https://doi.org/10.5824/1309-1581.2016.1.001.x>
- Hancock, J. I., Allen, D. G., Bosco, F. A., McDaniel, K. R., & Pierce, C. A. (2013). Meta-analytic review of employee turnover as a predictor of firm performance. *Journal of Management*, 39(3), 573–603. <https://doi.org/10.1177/0149206311424943>
- Hausknecht, J. P., & Holwerda, J. A. (2012). When does employee turnover matter? Dynamic member configurations, productive capacity, and collective performance. *Organization Science*, 24(1), 210–225. <https://doi.org/10.1287/orsc.1110.0720>
- He, Y., Donnellan, M. B., & Mendoza, A. M. (2019). Five-factor personality domains and job performance: A second order meta-analysis. *Journal of Research in Personality*, 82, 103848. <https://doi.org/10.1016/j.jrp.2019.103848>
- Highhouse, S. (2008). Stubborn reliance on intuition and subjectivity in employee selection. *Industrial and Organizational Psychology*, 1(3), 333–342. <https://doi.org/10.1111/j.1754-9434.2008.00058.x>
- Holtom, B. C., Mitchell, T. R., Lee, T. W., & Eberly, M. B. (2008). 5 Turnover and retention research: A glance at the past, a closer review of the present, and a venture into the future. *Academy of Management Annals*, 2(1), 231–274. <https://doi.org/10.5465/19416520802211552>
- Hom, P. W., Lee, T. W., Shaw, J. D., & Hausknecht, J. P. (2017). One hundred years of employee turnover theory and research. *Journal of Applied Psychology*, 102(3), 530–545. <https://doi.org/10.1037/apl0000103>
- Hubstaff. (2024). Employee Turnover Statistics and Trends for 2024. Retrieved September 23, 2025, from <https://hubstaff.com/blog/employee-turnover-statistics/>
- Jang, K. L., Livesley, W. J., & Vernon, P. A. (1996). Heritability of the Big Five Personality Dimensions and Their Facets: A Twin Study. *Journal of Personality*, 64(3), 577–592. <https://doi.org/10.1111/j.1467-6494.1996.tb00522.x>
- Judge, T. A., Higgins, C. A., Thoresen, C. J., & Barrick, M. R. (1999). The Big Five personality traits, general mental ability, and career success across the life span. *Personnel Psychology*, 52(3), 621–652. <https://doi.org/10.1111/j.1744-6570.1999.tb00174.x>
- Kang, W., Guzman, K. L., & Malvaso, A. (2023). Big Five personality traits in the workplace: Investigating personality differences between employees, supervisors, managers, and entrepreneurs. *Frontiers in Psychology*, 14, 976022. <https://doi.org/10.3389/fpsyg.2023.976022>
- Konar, K., Das, S., Das, S., & Misra, S. (2025). Employee Attrition Prediction Using Bayesian Optimized Stacked Ensemble Learning and Explainable AI. *SN Computer Science*, 6, 672. <https://doi.org/10.1007/s42979-025-04204-w>
- Kuncel, N. R., Klieger, D. M., Connelly, B. S., & Ones, D. S. (2013). Mechanical versus clinical data combination in selection and admissions decisions: a meta-analysis. *Journal of Applied Psychology*, 98(6), 1060–1072. <https://doi.org/10.1037/a0034156>
- Manafi Varkiani, S., Pattarin, F., Fabbri, T., & Fantoni, G. (2025). Predicting employee attrition and explaining its determinants. *Expert Systems with Applications*, 272, 126575. <https://doi.org/10.1016/j.eswa.2025.126575>
- Mansor, N., Sani, N. S., & Aliff, M. (2021). Machine Learning for Predicting Employee Attrition. *International Journal of Advanced Computer Science and Applications*, 12(11), 435–445. <https://doi.org/10.14569/ijacsa.2021.0121149>
- Marler, J. H., & Boudreau, J. W. (2017). An evidence-based review of HR Analytics. *The International Journal of Human Resource Management*, 28(1), 3–26. <https://doi.org/10.1080/09585192.2016.1244699>
- Newman, D. T., Fast, N. J., & Harmon, D. J. (2020). When eliminating bias isn't fair: Algorithmic reductionism and procedural justice in human resource decisions. *Organizational Behavior and Human Decision Processes*, 160, 149–167. <https://doi.org/10.1016/j.obhdp.2020.03.008>
- Ployhart, R. E., Nyberg, A. J., Reilly, G., & Maltarich, M. A. (2014). Human capital is dead; long live human capital resources. *Journal of Management*, 40(2), 371–398. <https://doi.org/10.1177/0149206313512152>
- Rasmussen, T., & Ulrich, D. (2015). Learning from practice: How HR analytics avoids being a management fad. *Organizational Dynamics*, 44(3), 236–242. <https://doi.org/10.1016/j.orgdyn.2015.05.008>
- 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>
- Shabbir, H. H., Farooq, M. H., Zafar, A., Akram, B. A., Waheed, T., & Aslam, M. (2025). Enhancing employee churn prediction with Weibull time-to-event modeling. *Journal of Engineering Research*, 13(4), 3231–3248. <https://doi.org/10.1016/j.jer.2025.03.009>
- Shlash, M. A. A., Alkhazali, Z., Shelash Mohammad, S. I., 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*, AG Editor, 4: 669. <https://doi.org/10.56294/dm2025669>
- 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. <https://doi.org/10.1177/0008125619867910>
- Tursunbayeva, A., Di Lauro, S., & Pagliari, C. (2018). People analytics: A scoping review of conceptual boundaries and value propositions. *International*

- Journal of Information Management*, 43, 224-247.
<https://doi.org/10.1016/j.ijinfomgt.2018.08.002>
- U.S. Bureau of Labor Statistics, Department of Labor. (2023). Chart 2. Hires and total separations rates, seasonally adjusted. Retrieved September 24, 2025, from <https://www.bls.gov/jlt>
- Vaiyapuri, T., & Shai, Z. (2025). Bayesian Optimized Boosted Ensemble Models for HR Analytics- Adopting Green Human Resource Management Practices. *International Journal of Technology*, 16(2), 561-572.
<https://doi.org/10.14716/ijtech.v16i2.7277>
- Wilmot, M. P., & Ones, D. S. (2019). A century of research on conscientiousness at work. *Proceedings of the National Academy of Sciences*, 116(46), 23004-23010.
<https://doi.org/10.1073/pnas.1908430116>
- Wu, R., Zhao, X., Li, Z., & Xie, Y. (2024). The role of employee personality in employee satisfaction and turnover: insights from online employee reviews. *Personnel Review*, 53(7), 1581-1611.
<https://doi.org/10.1108/pr-04-2023-0309>
- Yu, P., Gao, Y., & Fan, Y. (2025). Predicting Employee Attrition with Big Five Personality Involved Using Machine Learning. *Advances in Economics, Management and Political Sciences*, 215, 20-29.
<https://doi.org/10.54254/2754-1169/2024.26743>
- Zimmerman, R. D. (2008). Understanding the Impact of Personality Traits on Individuals' Turnover Decisions: A Meta-Analytic Path Model. *Personnel Psychology*, 61(2), 309-348.
<https://doi.org/10.1111/j.1744-6570.2008.00115.x>

APPENDICES

Table 7: Dataset Descriptive Statistics

	stag	event	gender	age	industry	profession	traffic	coach
count	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000
mean	36.627526	0.505757	0.244464	31.066965	8.614703	6.560673	4.178919	0.838795
std	34.096597	0.500188	0.429959	6.996147	4.232955	2.514161	2.252630	0.607766
min	0.394251	0.00000	0.00000	18.00000	0.00000	0.00000	0.00000	0.00000
25%	11.728953	0.00000	0.00000	26.00000	5.00000	6.00000	2.00000	0.00000
50%	24.344969	1.00000	0.00000	30.00000	10.00000	6.00000	4.00000	1.00000
75%	51.318275	1.00000	0.00000	36.00000	12.00000	6.00000	7.00000	1.00000
max	179.449692	1.00000	1.00000	58.00000	15.00000	14.00000	7.00000	2.00000
	head_gender	greywage	way	extraversion	independ	selfcontrol	anxiety	novator
count	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000	1129.000
mean	0.517272	0.887511	0.500443	5.592383	5.478034	5.597254	5.665633	5.879628
std	0.499923	0.316107	0.676512	1.851637	1.703312	1.980101	1.709176	1.904016
min	0.00000	0.00000	0.00000	1.00000	1.00000	1.00000	1.700000	1.00000
25%	0.00000	1.00000	0.00000	4.60000	4.10000	4.10000	4.800000	4.40000
50%	1.00000	1.00000	0.00000	5.40000	5.50000	5.70000	5.600000	6.00000
75%	1.00000	1.00000	1.00000	7.00000	6.90000	7.20000	7.100000	7.50000
max	1.00000	1.00000	2.00000	10.00000	10.00000	10.00000	10.00000	10.00000





Figure 5: Dataset Descriptive Statistics

Coaching During Trial Period and Attrition Rate



Communing Mode Distribution and Attrition Rate




Figure 6: Visualization of Categorical Variables Based on Attrition Status




Figure 7: Feature Distributions Analysis

ISTRAŽIVANJE PREDVIĐANJA DOBROVOLJNOG ODLIVA ZAPOSLENIH KORIŠĆENJEM OSOBINA LIČNOSTI I MAŠINSKOG UČENJA: STUDIJA HR ANALITIKE I PODRŠKE ODLUČIVANJU ZASNOVANA NA JAVNO DOSTUPNOM SKUPU PODATAKA

Dobrovoljni odliv zaposlenih izaziva značajne organizacione troškove usled nestabilnosti radne snage i gubitka znanja. Uprkos brzom razvoju analitike zaposlenih (*people analytics*), većina modela za predviđanje dobrovoljnog odliva zasniva se na uočljivim varijablama, zanemarujući psihološke determinante dobrovoljnog odliva. Ovo istraživanje nastoji da prevaziđe navedeni nedostatak razvojem i evaluacijom modela mašinskog učenja koji integrišu osobine ličnosti prema modelu Big Five sa podacima iz oblasti upravljanja ljudskim resursima, primenom evaluacionog okvira optimizovanog prema kriterijumu preciznosti i usklađenog sa troškovima HR intervencija. Na skupu podataka koji obuhvata 1.129 zaposlenih obučena su tri algoritma – logistička regresija (*Logistic Regression*), mašina potpornih vektora (*Support Vector Machine* – SVM) i slučajna šuma (*Random Forest* – RF). Pri evaluaciji modela prednost je data meri Precision u odnosu na Recall, kako bi se smanjilo neefikasno korišćenje resursa namenjenih zadržavanju zaposlenih. Model slučajne šume ostvario je najbolje rezultate (AUC = 0,69; Precision = 0,74), nadmašivši logističku regresiju (AUC = 0,64) i SVM (AUC ≈ 0,66). Svih pet dimenzija modela Big Five našlo se među sedam najznačajnijih prediktora, zajedno sa dužinom radnog staža i starošću zaposlenih, što ukazuje na njihov doprinos riziku od odliva zaposlenih. Rezultati podržavaju primenu modela koji uključuju osobine ličnosti kao podršku donošenju odluka o zadržavanju zaposlenih zasnovanih na dokazima i planiranju radne snage.

Ključne reči: Dobrovoljni odliv zaposlenih; Predviđanje odliva zaposlenih; Osobine ličnosti Big Five; Mašinsko učenje; HR analitika; Podrška odlučivanju.