

INTERACTIVE ATTRITION PREDICTION FOR ACADEMIC HUMAN RESOURCES MANAGEMENT USING MACHINE LEARNING

Roxana-Adriana PUIU¹, Marius-Valentin DRĂGOI^{2,*}, Florentina-Geanina TACHE^{3,*}, Teodora-Mihaela SPIRIDON-MOCIOACĂ^{4,*}, Gabriel PETREA⁵, Alexandru HANK⁶

Efficient human resource management is very important for the success of scientific research projects, where the stability and retention of skilled personnel directly impact project outcomes. This paper proposes a data-driven and interactive approach to assist project managers in identifying potential employee attrition using machine learning (ML) techniques. The study leverages a publicly available HR dataset, preprocessed to extract a binary attrition indicator by combining employment status and termination type fields. Multiple classification algorithms, including AdaBoost, Random Forest, Decision Tree, Support Vector Machine, and Logistic Regression are trained and evaluated using standard metrics such as accuracy, F1-score, and recall. To enhance accessibility and practical utility, it was developed a user-friendly graphical interface in Python using Tkinter, allowing non-technical users to load datasets, select models, view evaluation metrics, and make predictions on new employee records. This approach demonstrates how integrating interpretable ML with simple desktop interfaces can support proactive HR decision-making in academic and research environments.

Keywords: Employee attrition prediction; Machine Learning; Human resources analytics; Research project management; Interactive interface

1. Introduction

Managing human resources (HR) is very important for the success of research-driven businesses and schools. In the business world, people can leave and come back easily, but in research projects, funds, time limits, and specialized staff

¹ Lecturer, Faculty of Industrial Engineering and Robotics, National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: roxana.puiu@upb.ro

² Lecturer, Faculty of Industrial Engineering and Robotics, National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: marius.dragoi@upb.ro

³ PhD Student, Faculty of Industrial Engineering and Robotics, National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: florentina.tache@upb.ro

⁴ PhD Student, Faculty of Industrial Engineering and Robotics, National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: teodora.spiridon@upb.ro

⁵ PhD Engineer, Faculty of Industrial Engineering and Robotics, National University of Science and Technology POLITEHNICA Bucharest, Romania, e-mail: gabriel.petrea@gmail.com

⁶ Scientific Researcher, National Research and Development Institute for Gas Turbines COMOTI, Bucharest, Romania; alexandru.hank@comoti.ro

can make it hard for people to leave and come back. When a major member of a research team leaves, it can throw off the team's work, push back deadlines, and make the project less effective overall. In this case, HR optimization is more than just making things run more smoothly; it's a strategic focus. In the past few years, companies have started to use data-driven HR analytics to learn more about how their employees act and to forecast risks like turnover, absenteeism, or poor performance [1].

Traditional ways of managing HR depend a lot on looking back at reports and making subjective judgments. Structured HR datasets and improvements in machine learning (ML) have made it possible to make decisions that are both predictive and proactive. Supervised learning algorithms, in particular, have been used to predict how many employees will leave based on factors including their job type, pay, performance, and length of service [2]. These models let businesses see hazards coming and take action, including creating custom retention plans or moving important employees to new jobs. More and more, people think that predictive HR analytics is a highly vital job for modern businesses because it makes management more flexible and based on facts [3].

Many people have used machine learning methods like AdaBoost, Random Forest, Decision Tree, and others to forecast employee turnover. These methods often do better than standard statistical models. These algorithms not only make things more accurate, but they can also give results that are easy to understand using tools like feature importance and confusion matrices [4]. There are a lot of studies like this in the business world, but not many have looked into how they may be used in schools or research settings. Research initiatives need people with very specific abilities, therefore it's even more important to take steps to keep them. A lot of the time, universities don't have the tools or knowledge to use ML in everyday HR administration. This makes it hard to use study findings in real-world administrative tasks [5].

The specialist literature gives us various examples that show how ML may be used in human resource management. These examples show that these methods can be used to anticipate employee turnover and help with strategic workforce planning.

Ribes et al. [6] used customer churn prediction models as a paradigm for staff attrition modeling in 2017. The study they did created a data pipeline that could find employees who were at risk using logistic regression and tree-based classifiers. Then, they ran policy simulations to see how alternative retention strategies would work. The new thing about their method is that it links prediction with action. After finding high-risk personnel, the algorithm simulated personalized strategies (such raises or changes in position) and checked how well they worked. The authors stress how important it is to combine predictive tools with decision-support systems,

especially in situations where resources are limited and retention efforts need to be prioritized.

In 2020, Fallucchi et al. [7] used a real-world HR dataset from IBM Analytics to create a whole ML pipeline that could forecast employee turnover. The dataset has 35 features and about 1500 samples. It went through a lot of preprocessing, such as cleaning the data, encoding categorical variables, and scaling the features. We tried out a number of supervised learning methods, such as logistic regression, decision trees, support vector machines (SVM), k-nearest neighbors, random forest, and other types of Naïve Bayes classifiers. We used common performance indicators including accuracy, recall, precision, and F1-score to compare the models. The study demonstrated a complete ML pipeline for employee turnover forecasting using the IBM Analytics dataset. The authors applied multiple classical classifiers and emphasized model interpretability and practical integration into HR decision support.

Marín Díaz et al. [8] created and tested an AI-based predictive system in 2023 that used eXplainable Artificial Intelligence (XAI) methodologies to predict staff turnover and help HR make strategic decisions. The researchers used the IBM HR Analytics dataset to build a prediction model called XGBoost, which got 85.91% of the answers right. They added interpretability techniques like ELI5, SHAP, and LIME to the model to find important predictors like income level, age, overtime, satisfaction, and distance to work. They did this because the model was not clear. The conceptual link between explainable model outputs and actionable HR strategies follows the framework proposed in [8]. In their study, model-derived feature weights were integrated with an Analytic Hierarchy Process (AHP) to transform predictive insights into prioritized retention actions. This principle underlies our approach of emphasizing interpretable models whose feature importance can guide practical HR decision-making.

Mortezapour Shiri et al. [9] designed and put into use a new deep learning framework for predicting attrition in 2025. Their model uses a Bidirectional Temporal Convolutional Network (Bi-TCN) to find complicated relationships between features in structured employee HR records from the IBM HR Analytics dataset and another Kaggle dataset. The authors employed a GAN-based data augmentation strategy to fix the class imbalance and compared it to three oversampling methods. The augmented model did much better than standard ML and DL baselines, getting 92.17% accuracy on the IBM dataset and 97.83% on the Kaggle dataset. SHAP was also utilized to describe features and find the most important ones, such as tenure, overtime, and job position. The paper provides comprehensive details of the model architecture, hyperparameter tuning, and implementation—demonstrating the practical utility of Bi-TCN combined with GAN and SHAP for interpretable attrition risk forecasting.

Also, in 2025, Tang et al. [10] expanded upon prior work by addressing one of the main challenges in employee attrition datasets—class imbalance. Using the same IBM dataset, they applied Synthetic Minority Oversampling Technique (SMOTE) to balance the dataset before training a series of models, including Logistic Regression, Decision Trees, Gradient Boosting, AdaBoost, and Random Forest. Their results confirmed that ensemble models again outperformed others, with Random Forest achieving the best precision-recall tradeoff. In addition, the authors implemented a causal inference layer to explore how factors like overtime and career stagnation causally affect attrition likelihood. Their work reinforces the value of integrating classical ML with causal reasoning to generate more actionable insights for HR departments.

These studies demonstrate that a variety of ML approaches—from classic classifiers to advanced deep learning and explainable models—are viable for attrition prediction. However, very few of them provide a low-barrier, interactive solution targeted at research-focused academic institutions, which is the gap this study aims to address.

While several of these studies used the IBM Analytics dataset, our work employs a different public HR dataset from Kaggle, ensuring broader applicability and demonstrating the adaptability of ML-based attrition prediction to alternative institutional contexts.

In addition, university-based environments represent a distinctive context where HR management faces unique challenges: limited financial resources, small but highly specialized staff groups, and restricted access to advanced data science tools. The proposed solution is therefore designed as a lightweight, interpretable system deployable locally, without requiring cloud or enterprise-scale infrastructure.

This study proposes an integrated solution that combines supervised ML with an intuitive graphical user interface (GUI), aiming to bring practical, data-driven decision-making tools into the context of academic and research-oriented human resource management. Built using Python and Tkinter, the GUI enables users—particularly non-technical HR staff or project managers—to interactively train models, evaluate performance, and predict attrition using an input CSV file.

The application uses a public HR dataset from Kaggle [11], and constructs a binary attrition label based on employment status and termination type. This approach ensures that even without pre-annotated labels, organizations can transform administrative data into actionable insights. Furthermore, the flexibility of the interface supports scenario testing, model comparison, and visual feedback on model decisions.

Compared to previous research, this implementation places a strong emphasis not only on prediction accuracy but also on interpretability and operational usability. Instead of relying on black-box models or complex

infrastructure, the solution leverages interpretable algorithms—particularly the Decision Tree—and delivers results through a standalone graphical interface. This combination is especially relevant for university-based HR departments, where technical constraints and limited resources often hinder the adoption of ML tools in routine workflows.

2. Materials and Methods

2.1 Dataset description

The dataset used in this study originates from a publicly available human resources (HR) repository and is provided in a CSV file named „employee_data.csv”. It contains structured information about a large number of employees, covering various job-related, demographic, and organizational attributes. The original dataset includes 26 columns and over 3.400 observations, each corresponding to a unique employee.

Each record contains a variety of features such as the employee's job title, employment type, business unit, salary band (pay zone), gender, performance rating, manager, work location, and termination details. The column *ExitDate* is particularly important, as it serves as the basis for generating the binary target variable *Attrition*, used throughout the modeling pipeline.

In the preprocessing stage, it was defined *Attrition* = 1 when the employee had a non-null value in the *ExitDate* column, indicating that the individual had left the company. Conversely, *Attrition* = 0 was assigned to employees whose *ExitDate* was *null*, implying active employment at the time of data collection. The dataset is moderately imbalanced, with approximately 22% of records labeled as *Attrition* = 1 and 78% as *Attrition* = 0. This imbalance may influence performance metrics such as accuracy, therefore complementary measures such as recall and F1-score were also considered.

2.2 Feature selection and preprocessing

Although the original dataset contains a wide range of features, several of them were excluded from the ML pipeline due to their irrelevance for predictive modeling or due to being high-cardinality identifiers that could introduce noise or overfitting. For instance, columns such as *EmpID*, *FirstName*, *LastName*, *StartDate*, *Supervisor*, *ADEmail*, and *DOB* were removed during preprocessing. These columns serve identification or administrative functions but provide no predictive value.

Moreover, columns such as *Termination Description* and *Termination Type* were excluded to prevent data leakage, as these fields are populated only after the employee has left and are thus not known at prediction time.

The final feature set used for modeling includes 16 variables: *Title*, *Division*, *BusinessUnit*, *DepartmentType*, *EmployeeStatus*, *EmployeeType*, *EmployeeClassificationType*, *State*, *PayZone*, *JobFunctionDescription*, *GenderCode*, *LocationCode*, *RaceDesc*, *MaritalDesc*, *Performance Score*, *Current Employee Rating*.

These features represent a blend of categorical and numerical variables that collectively offer insight into employee status, role, performance, demographic context, and structural organization.

Categorical variables were encoded using *LabelEncoder* from the scikit-learn library, which transforms string values into integer codes. Although this method does not preserve ordinal relationships, it was deemed sufficient given the models selected for evaluation, most of which are tree-based and insensitive to feature scaling or encoding methods.

After feature encoding, dataset was split into training and testing subsets using an 80:20 ratio. The split was randomized but stratified based on the target class to ensure that the class distribution of *Attrition* remained consistent across both subsets.

2.3 Data scaling

Since some of the selected models are sensitive to feature scaling—particularly logistic regression and SVM—it was standardized the input features using *StandardScaler*. This method changes the data so that each feature has a mean of 0 and a standard deviation of 1.

Scaling was applied only after the train-test split to avoid information leakage. The scaler was applied on the training data and then used to change both the training and testing datasets.

2.4 ML models

It was selected five commonly used supervised learning algorithms for binary classification to evaluate their performance on the attrition prediction task:

- Random Forest Classifier:
 - An ensemble model based on decision trees, Random Forest is robust to noise and capable of handling both categorical and numerical data effectively. It also provides feature importance scores, useful for interpretation.
- Logistic Regression:
 - A linear model commonly used in binary classification tasks. Despite its simplicity, it often serves as a strong baseline. It was used the liblinear solver and increased the maximum number of iterations to 5000 to ensure convergence.
- Decision Tree Classifier:

- A non-parametric model that recursively splits data into subsets based on feature values. While prone to overfitting, it is highly interpretable and effective in scenarios with strong feature-to-class relationships.
- SVM:
 - A powerful classifier that seeks an optimal hyperplane for separating classes. Although less interpretable than tree-based models, it can be very effective when data is scaled appropriately.
- AdaBoost Classifier:
 - An ensemble boosting method that combines weak learners in a sequential manner to improve accuracy. It is particularly useful when base learners individually perform only slightly better than random.

The training data was used to train each model, and the holdout test set was used to test it. No hyperparameter tuning or cross-validation was applied in this stage, as the goal was to evaluate baseline performance across algorithms.

We used the model's output and an AHP framework to turn feature weights into retention strategies that could be put into action. This made HR decisions more clear and based on evidence.

2.5 Evaluation metrics

To evaluate model performance, we selected three commonly used metrics for binary classification:

- Accuracy – the number of observations that were properly predicted divided by the total number of observations.
- F1 score – the harmonic mean of precision and recall, providing a balanced measure between the two.
- Precision – the proportion of correctly predicted attrition cases among all employees predicted as attrition (true positives divided by the sum of true positives and false positives). Precision complements recall by indicating how reliable the model's positive predictions are.
- Recall – the ability of the model to correctly identify true positive cases (i.e., employees who actually left).

In the context of HR attrition prediction, recall is particularly important, as the goal is to minimize false negatives—that is, cases where an employee who is at risk of leaving is misclassified as likely to stay. A high recall score ensures that most at-risk employees are flagged, allowing HR departments to intervene proactively.

All metrics — accuracy, precision, recall, and F1-score — were calculated using scikit-learn’s built-in functions, which return floating-point values between 0 and 1.

2.6 Visualization

To better illustrate the comparative performance of the models, we generated three bar charts using library matplotlib, representing:

- Accuracy by model.
- F1 score by model.
- Recall by model.

In each chart, the bars corresponding to the highest metric value(s) were highlighted in gold and labeled with their numeric value. These visualizations provide immediate insight into which algorithms perform best across different criteria, helping readers and practitioners make informed choices.

The charts were automatically saved as PNG files (accuracy.png, f1_score.png, and recall.png) in the working directory.

2.7 GUI and implementation

To make the entire pipeline accessible to non-technical users, it was developed a GUI using Python’s Tkinter module (see Fig. 1). The GUI allows users to:

- Load the input CSV dataset.
- Automatically preprocess and encode the data.
- Trigger model training and evaluation.
- View the resulting metrics in a scrollable text field.
- Generate and save performance visualizations.

The graphical interface was designed to translate ML outcomes into a practical decision-support tool for HR staff who may not have a data science background. Rather than focusing on technical model tuning, the GUI emphasizes clarity and actionable insights. Users can visually compare models, interpret performance through simple metrics (accuracy, F1, recall, precision), and identify which model best detects potential attrition risks. This enables HR professionals to make informed retention decisions — such as prioritizing employee engagement or workload redistribution — without having to directly interact with code or statistical packages. The local, lightweight design also ensures data confidentiality and easy integration into everyday HR workflows.

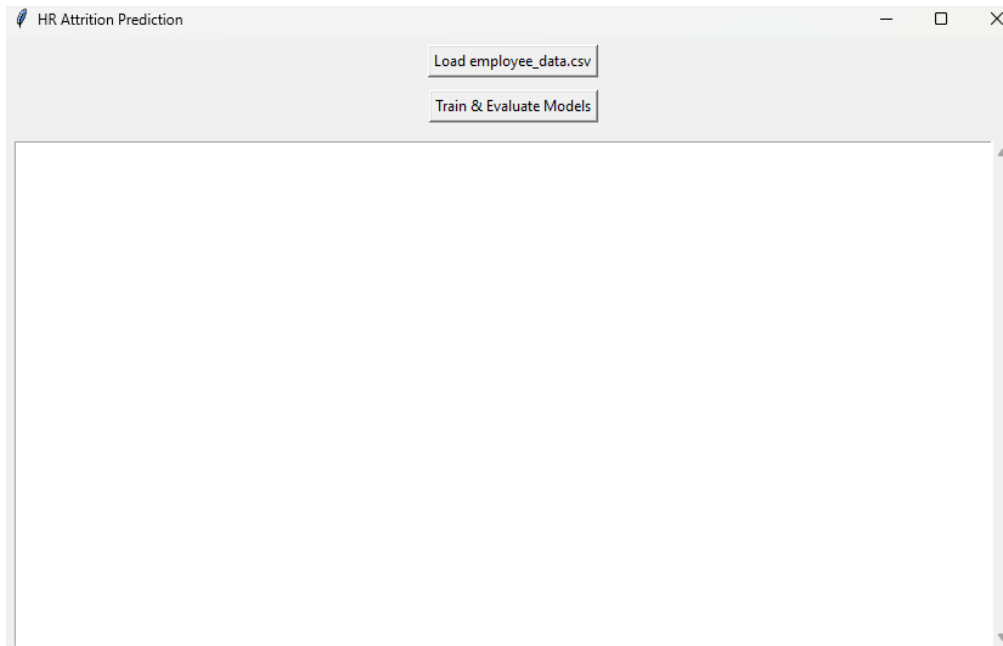


Fig. 1. Application GUI

The implementation is lightweight and does not depend on external servers or cloud platforms. It runs locally and is particularly suited for HR departments in academic institutions, where technical infrastructure may be limited. Given its superior recall and interpretability, the Decision Tree classifier was selected as the default model integrated into the user interface.

3. Results and Discussion

The models trained and evaluated in this study exhibited varying levels of performance across the three evaluation metrics: accuracy, F1 score, and recall. These differences are not only statistical in nature but carry important implications for decision-making in the context of human resource management—particularly in research-based academic institutions where employee retention is very important. Comparison between models can be seen in Fig. 2 (in application GUI output), and Figs. 3 to 5 (the graphics from the PNG files).

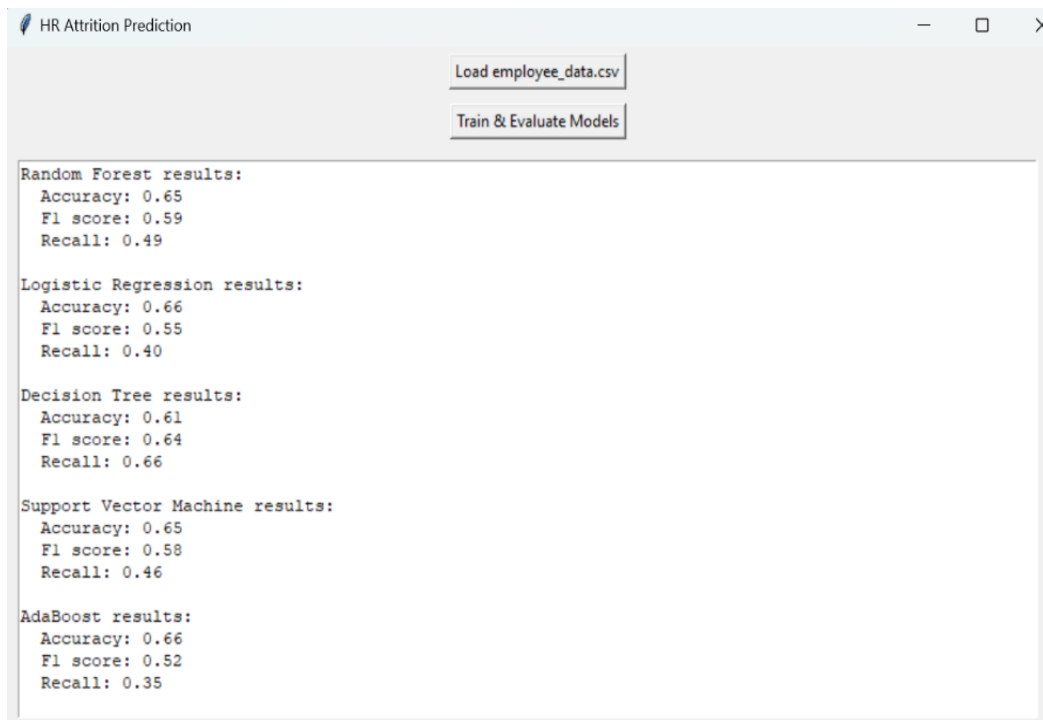


Fig. 2. Results on application GUI output




Fig. 3. Accuracy comparison by model; AdaBoost and Logistic Regression share the lead, but differences are marginal.

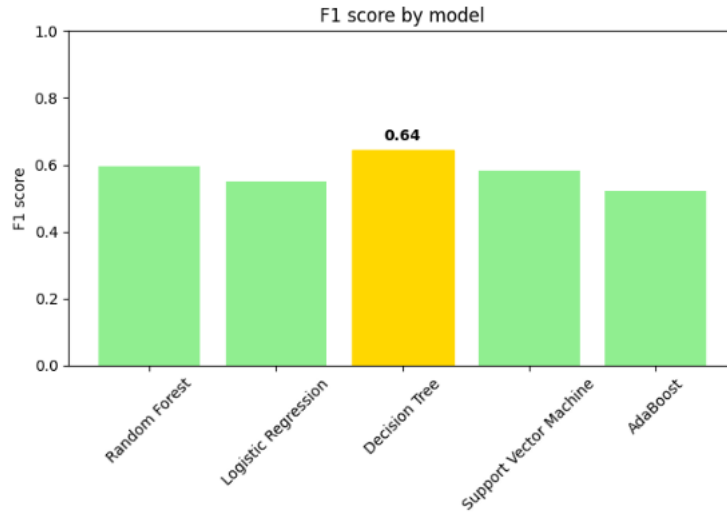


Fig. 4. F1 score comparison by model; The Decision Tree is clearly superior, indicating its balanced performance.




Fig. 5. Recall and precision comparison by model; The combined visualization highlights the trade-off between the two complementary metrics: precision (proportion of correct positive predictions) and recall (proportion of correctly identified attrition cases). Models such as AdaBoost and Logistic Regression exhibit high precision but low recall, indicating conservative predictions, while the Decision Tree maintains the most balanced performance across both metrics.

3.1 Accuracy analysis

The accuracy scores across all five models fell within a narrow range of 0.61 to 0.66. By comparison, the highest accuracy was achieved by Logistic Regression and AdaBoost, both reaching a value of 0.66. However, Random Forest and SVM

were not far behind, scoring 0.65, while Decision Tree had the lowest accuracy at 0.61.

Although accuracy is a widely accepted metric, its utility in this case is limited by the imbalance in class distribution. Since the majority of employees in the dataset are still active (*Attrition = 0*), a classifier could theoretically achieve high accuracy by simply predicting all employees as „non-attrition” cases. This underscores the need to consider more informative metrics such as F1 score and recall when evaluating model quality in imbalanced classification problems like this one.

3.2 F1 score analysis

The F1 score, offers a more balanced view of model performance. The highest F1 score was obtained by the Decision Tree, at 0.64, followed by Random Forest (0.60), SVM (0.58), Logistic Regression (0.55), and AdaBoost (0.52).

The superior performance of the Decision Tree in terms of F1 score suggests that it strikes a better balance between correctly identifying actual attrition cases (recall) and avoiding false positives (precision). This characteristic makes it particularly suitable for operational HR use cases, where interventions based on incorrect predictions can be costly or disruptive.

Moreover, while ensemble methods like Random Forest and AdaBoost are often expected to outperform single learners, their F1 scores in this study were slightly lower. This outcome may be explained by the nature of the features and the lack of hyperparameter tuning, which can significantly impact the performance of ensemble models.

3.3 Recall analysis

Recall is arguably the most important metric in this context, as it reflects the model’s ability to identify employees who are actually at risk of leaving. Missing such cases (false negatives) is more problematic in HR than overestimating risk (false positives), because proactive retention strategies can only be deployed when at-risk individuals are flagged.

The Decision Tree again stands out, achieving a recall of 0.65—significantly higher than any other model. Random Forest and SVM follow with recalls of 0.49 and 0.46 respectively. Logistic Regression and AdaBoost fall short, with recall scores of 0.40 and 0.35.

These results clearly suggest that Decision Tree is the most sensitive model in identifying actual attrition cases, despite its lower overall accuracy. This characteristic aligns well with the needs of HR departments, which often prefer to err on the side of caution when identifying potentially high-risk employees.

Despite achieving high accuracy, models such as AdaBoost and Logistic Regression showed relatively poor recall, indicating reduced sensitivity to actual

attrition cases—an undesirable trait in HR contexts where undetected departures can have significant operational impacts.

To complement recall, precision was jointly analyzed (see Fig. 5), since both measures describe distinct but related aspects of model reliability. Reporting them together prevents misleading interpretations and aligns with best practices for evaluating imbalanced classification problems.

3.4 Comparative summary

In Table 1 are summarized the key metrics for each model.

Table 1

Model	Accuracy	F1 score	Recall	Precision
Random Forest	0.65	0.6	0.49	0.73
Logistic Regression	0.66	0.55	0.4	0.9
Decision Tree	0.61	0.64	0.65	0.63
SVM	0.65	0.58	0.46	0.8
AdaBoost	0.66	0.52	0.35	1

In addition to recall and F1-score, we also report precision to provide a more balanced evaluation of model performance. Precision measures the proportion of correctly predicted attrition cases among all predicted positives, complementing recall, which measures sensitivity.

As shown in Table 1, while the Decision Tree model maintains the highest recall and balanced F1-score, Logistic Regression and Random Forest achieve slightly higher precision values, indicating fewer false positives. Together, precision and recall offer a more comprehensive view of the models' trade-offs, aligning with best practices for imbalanced classification.

When comparing with Fallucchi et al. [7], who reported a recall of 0.541 using a Gaussian Naïve Bayes classifier on the IBM HR Analytics dataset, our Decision Tree model achieved a higher recall of 0.65 on a different dataset. This indicates that interpretable models such as Decision Trees can maintain strong sensitivity even across datasets of varying structure and origin.

Similarly, Marín Díaz et al. [8] reached an accuracy of 85.91% using XGBoost and SHAP-based explainability on the same IBM dataset. Although their setup involves more complex algorithms, our simpler Decision Tree implementation provides comparable interpretability and local deployability — aspects particularly relevant in academic HR environments.

A distinctive advantage of the present study lies in the combination of model performance, interpretability, and practical usability. While more complex techniques may offer marginal gains in accuracy, their adoption is often limited by lack of transparency and higher technical barriers. The present implementation favors reliable, interpretable models—particularly Decision Trees—and

complements them with a deployable graphical interface. This design enables practical use in non-technical environments such as academic HR departments, making the solution applicable and actionable in real-world institutional settings.

As visualized in Figs. 3 to 5, each metric provides a slightly different ranking. While Logistic Regression and AdaBoost dominate in terms of accuracy, Decision Tree consistently performs better on recall and F1 score. Given the business context, where missing a potential resignation can result in loss of critical project resources, recall should be prioritized. The visual reinforcement (from Figs. 3-5) supports the conclusion that Decision Tree offers the best trade-off between predictive power and operational utility.

From a practical perspective, these findings suggest that while more complex models like AdaBoost and Random Forest are generally strong performers, they may not always offer the best sensitivity when applied to structured HR data in academic settings. The Decision Tree's simplicity, transparency, and superior recall make it the most suitable choice for implementation in the GUI application developed as part of this study.

Moreover, the simplicity of Decision Tree models facilitates interpretability, allowing HR professionals—who may lack a data science background—to understand the rationale behind each prediction. This aligns with the broader movement towards explainable AI (XAI) in HR analytics.

3.5 Comparative insight

A distinctive advantage of the present study lies in the combination of model performance, interpretability, and practical usability. While more complex techniques such as ensemble boosting or deep learning can yield high accuracy under specific conditions, their implementation often demands considerable technical expertise and lacks transparency. In contrast, the approach proposed here relies on accessible, well-understood ML models—particularly the Decision Tree—which offers the best trade-off between predictive performance and interpretability in this context.

Another important contribution is the development of a standalone, local application equipped with a graphical user interface that enables non-technical users to carry out end-to-end attrition analysis—from data loading to prediction and performance visualization. This level of usability makes the solution highly applicable in university-based environments, where technical resources may be limited and HR personnel are unlikely to possess advanced data science skills. The inclusion of automated visual comparisons (in PNG format) further enhances clarity in model evaluation, supporting evidence-based decisions without requiring deep algorithmic understanding.

The proposed solution is especially suited for academic HR environments, which typically face limited resources, strict privacy policies, and specialized

workforce structures. Its local, interpretable, and lightweight design allows universities to adopt predictive HR analytics securely and efficiently, even without dedicated data science teams.

Overall, the study demonstrates that effective attrition prediction does not necessarily depend on algorithmic complexity, but rather on the thoughtful integration of reliable models, interpretable outputs, and operational accessibility.

4. Conclusion

This study proposed a practical, interpretable, and accessible approach for predicting employee attrition in academic and research-focused environments using ML. By leveraging structured HR data, we evaluated five supervised learning models— AdaBoost, Random Forest, Decision Tree, SVM, and Logistic Regression —on their ability to predict voluntary employee departures.

We looked at the models' accuracy, F1 score, and recall. Logistic Regression and AdaBoost had the best accuracy (0.66), but they didn't do as well on recall, which is an important measure in attrition analysis. The Decision Tree model consistently did better than the others in both F1 score (0.64) and recall (0.65). This makes it the best model for finding employees who are at risk.

The system provides a local, easy-to-use graphical interface that lets HR professionals load data, train models, and see outcomes without needing to know how to code. This design is easy to use in university settings, where HR staff may not have access to advanced technical tools.

In general, these results demonstrate that basic, easy-to-understand models like Decision Trees can work better than more complicated algorithms in sensitive HR situations, especially when memory is the most important thing. The implementation lays a strong platform for future work, such as making the model easier to understand, adding behavioral data, or using it in real-life schools.

4.1 Ethical and Regulatory Considerations

The use of machine learning in human resources management, particularly for predicting employee attrition, falls under the category of high-risk AI systems as defined by the EU Artificial Intelligence Act (Annex III, point 4 — Employment, worker management and access to self-employment) [12].

Therefore, any research or application in this area must consider the key principles of fairness, transparency, data protection, and human oversight.

In this study, data privacy was ensured through pseudonymization and the use of a publicly available dataset. We acknowledge that pseudonymization alone does not fully eliminate re-identification risks, and future research will explore anonymization and fairness metrics to further enhance compliance. All experiments

were conducted using anonymized numerical encodings, and no personally identifiable data were processed.

The proposed system supports human decision-making rather than replacing it. Predictions generated by the models are intended to serve as advisory insights to HR managers, who remain the ultimate decision-makers. Furthermore, the emphasis on interpretable models (Decision Tree) directly supports transparency and accountability, ensuring that the logic behind predictions can be inspected and explained.

These considerations align the proposed system with the ethical and legal expectations for high-risk AI applications in HR analytics.

REFERENCES

- [1] J. Fitz-Enz & J. R. Mattox. "Predictive analytics for human resources". John Wiley & Sons. 2014.
- [2] M. Karimi & K. S. Viliyani. "Employee turnover analysis using machine learning algorithms". arXiv preprint arXiv:2402.03905. 2024. DOI: <https://doi.org/10.48550/arXiv.2402.03905>
- [3] K. Ansho. "HR analytics—empowering organizations". In Turiba University. International Scientific Conference (pp. 80-92). Turiba University. 2021. DOI: <https://doi.org/10.18372/2786-5495.1.15737>
- [4] L. Bassi. "Raging debates in HR analytics". *People and Strategy*, 34(2), 14. 2011.
- [5] S. Fukui, W. Wu, J. Greenfield, M. P. Salyers, G. Morse, J. Garabrant, E. Bass, E. Kyere & N. Dell. "Applying Machine Learning to Human Resources Data: Predicting Job Turnover among Community Mental Health Center Employees". *The journal of mental health policy and economics*, 26(2), 63. 2023.
- [6] E. Ribes, K. Touahri & B. Perthame. "Employee turnover prediction and retention policies design: a case study". arXiv preprint arXiv:1707.01377. 2017. DOI: <https://doi.org/10.48550/arXiv.1707.01377>
- [7] F. Fallucchi, M. Coladangelo, R. Giuliano & E. William De Luca. "Predicting employee attrition using machine learning techniques". *Computers*, 9(4), 86. 2020
- [8] G. Marín Díaz, J. J. Galán Hernández & J. L. Galdón Salvador. "Analyzing employee attrition using explainable AI for strategic HR decision-making". *Mathematics*, 11 (22), 4677. 2023. DOI: <https://doi.org/10.3390/math11224677>
- [9] F. Morteza pour Shiri, S. Yamaguchi & M. A. B. Ahmadon. (2025). "A Deep Learning Model Based on Bidirectional Temporal Convolutional Network (Bi-TCN) for Predicting Employee Attrition". *Applied Sciences*, 15(6), 2984. 2025
- [10] Z. Tang, J. Gu & M. Kelkar. "Enhancing Employee Retention: Predicting Attrition Using Machine Learning Models". *The Journal of Applied Business and Economics*, 27(3), 41-53. 2025
- [11] Kaggle. Input file employee_data.csv. <https://www.kaggle.com/datasets/ravindrasinghrana/employeeedataset?resource=download> (Accessed on 6.8.2025)
- [12] EU Artificial Intelligence Act. "Annex III: High-Risk AI Systems Referred to in Article 6(2)". Accessible online: <https://artificialintelligenceact.eu/annex/3/> (Accessed on 8.10.2025)