Jurnal Ekonomi Kreatif dan Manajemen Bisnis Digital
https://transpublika.co.id/ojs/index.php/JEKOMBITAL
Online ISSN 2961-8428

https://doi.org/10.55047/jekombital.v4i2.1071

Original Article

Benchmarking Machine Learning Models for Corporate Bankruptcy Prediction using Financial Ratios

Sigit Mulyanto[1*], Dwika Lovitasari Yonia[2], Muhammad Arif[3], Bambang Sutejo[4]

[1,4]Department of Management, Faculty of Business, Universitas Darwan Ali, Sampit, Indonesia
[2]Department of Information Systems, Faculty of Computer Science, Universitas Darwan Ali, Sampit, Indonesia
[3]Department of Islamic Business Management, Faculty of Islamic Economics and Business, Universitas Islam Negeri Syekh Ali Hasan Ahmad Addary Padangsidimpuan, Indonesia
Email: [1)] sigitmul@gmail.com, [2)] dwika.lovitasari.yonia@gmail.com,
[3)] muhammadarif@uinsyahada.ac.id, [4)] tejosampit@gmail.com

Received : 26 October - 2025
Accepted : 27 November - 2025
Published online : 29 November - 2025

Abstract

Corporate bankruptcy prediction is a critical task in financial risk management, particularly under conditions of economic uncertainty and highly imbalanced datasets. This study presents a comprehensive benchmarking framework that evaluates multiple supervised learning models and a voting ensemble approach for corporate bankruptcy prediction. Using a publicly available dataset comprising 78,682 financial records from US-listed companies on NYSE and NASDAQ (1999-2018), we compare the performance of Random Forest, XGBoost, Gradient Boosting, Support Vector Machine, Decision Tree, and a Voting Classifier. Extensive preprocessing, including outlier removal, normalization, and feature selection, and cost-sensitive learning to mitigate severe class imbalance was conducted to ensure data quality. Model performance was assessed using multiple evaluation metrics such as accuracy, F1-score, and ROC AUC to account for class imbalance. Results demonstrate that the Voting Classifier, integrating Random Forest, XGBoost, and Gradient Boosting via hard voting, achieves superior overall performance with an accuracy of 93.6%, F1-score of 96.5%, and ROC AUC of 82.6%, outperforming individual models. The findings underscore the value of ensemble approaches in improving prediction robustness while addressing class imbalance challenges in financial distress forecasting. This study contributes a reproducible experimental design that can guide future research and practical implementation of learning models in corporate bankruptcy risk assessment.

Keywords: Bankruptcy Prediction, Ensemble Learning, Machine Learning, Model Evaluation, Voting Classifier.

1. Introduction

Corporate bankruptcy remains a pressing issue in financial risk management due to its cascading effects on stakeholders, including investors, regulators, creditors, and supply chains, especially in periods of heightened economic uncertainty such as the aftermath of the COVID-19 pandemic (Narvekar & Guha, 2021). Conventional statistical models, while historically important, often lack the capacity to capture complex, nonlinear financial patterns that precede bankruptcy. As a result, machine learning methods have increasingly gained traction for their ability to process multidimensional financial data and uncover subtle indicators of distress (Brygała & Korol, 2024).

Recent developments in bankruptcy prediction research have emphasized the role of data quality, model calibration, and evaluation strategies in influencing predictive outcomes.

Alam et al. tested several machine learning models on financial data from Polish firms and observed inconsistent performance among baseline models under different sampling and preprocessing strategies (Alam et al., 2021). Their findings underscore how factors like data imbalance and incomplete records may negatively impact model reliability.

Mattos and Shasha examined model performance in the context of Brazilian private companies and reported that the predictive capacity of decision trees and support vector machines was highly sensitive to inconsistencies in financial statement quality (da Silva Mattos & Shasha, 2024). These challenges are particularly acute in emerging markets, where standardized reporting is often lacking. Pawełek and Pociecha investigated whether hybrid classification models could outperform classical approaches, finding that even improved tree-based models failed to deliver robust performance without careful feature selection and data engineering (Pawełek & Pociecha, 2020).

Numerous studies have demonstrated that financial ratios such as return on assets, current ratio, and debt ratio can serve as effective predictors when integrated with algorithms like Random Forest or XGBoost (Shetty et al., 2022). However, many existing studies focus on limited model comparisons, employ proprietary or regional datasets, or rely heavily on hybrid input types such as sentiment data and macroeconomic variables that restrict replicability and model transparency (Arora & Singh, 2020). Furthermore, research often emphasizes accuracy as the primary evaluation metric, which can be misleading when class distribution is imbalanced, as is commonly the case in bankruptcy datasets (Dasilas & Rigani, 2024). There is a lack of consistent benchmarking across a wide range of algorithms using standardized, publicly available data with interpretable financial indicators (Gabrielli et al., 2023).

Another limitation in previous studies is the overreliance on accuracy as a single evaluation metric, which can be misleading in class-imbalanced settings. Dasilas & Rigani (2024) emphasized the need for broader metrics such as precision, recall, and F1 score to better assess model reliability in bankruptcy forecasting. Although ensemble and boosting methods have gained traction, few studies offer a systematic comparison of both traditional and advanced models within a unified and replicable evaluation setup. This study addresses that gap by incorporating a comprehensive benchmarking framework, including a voting ensemble that combines strong base learners. Recent studies support the use of evaluation metrics beyond accuracy, especially in class-imbalanced prediction tasks like bankruptcy forecasting (Mulyanto et al., 2025).

To address these gaps, this study presents a comprehensive and reproducible benchmark of five supervised learning models, such as Random Forest, Support Vector Machine, Decision Tree, XGBoost, and Gradient Boosting plus an ensemble Voting Classifier (Irmalasari & Dwiyanti, 2023). These models were chosen because they dominate prior work on financial distress, are well suited to structured ratio data, and collectively cover distinct algorithm families that offer complementary inductive biases (bagging trees for variance reduction, boosted trees for high accuracy on tabular data, and margin-based kernels for robust separation). They also provide calibrated or at least monotonic probability outputs required for ROC AUC analysis and threshold tuning, include options for class imbalance control, and span a practical trade-off between accuracy and interpretability, allowing both single model and ensemble effects to be evaluated fairly.

As part of our contribution, we benchmark all models on the US Company Bankruptcy Prediction dataset from Kaggle. The dataset contains 78,682 company year observations for firms listed on NYSE and NASDAQ between 1999 and 2018, with eighteen financial ratio features and a binary solvency label. Using this public and widely referenced resource ensures transparency, credibility, and straightforward replication of our workflow, since readers can obtain the same data, reproduce the preprocessing pipeline, and verify the reported results.

Ensemble methods such as Random Forest and XGBoost have demonstrated superior performance compared to individual classifiers in various financial prediction tasks, including credit scoring and fraud detection (Akinjole et al., 2024). All models were assessed with accuracy, F1 score, and the area under the receiver operating characteristic curve to capture both overall correctness and sensitivity to the minority bankruptcy class (Chaising & Srimaharaj, 2024). Using this trio of metrics prevents misleading conclusions that can arise when accuracy alone is reported on imbalanced data. The study thus offers a unified experimental framework that benchmarks traditional and contemporary algorithms while demonstrating their suitability for real-world financial distress prediction through interpretable features and fully replicable procedures.

2. Literature Review

Research on bankruptcy prediction continues to expand as scholars seek more reliable modelling approaches. Machine learning methods have introduced new solutions to persistent issues such as imbalance and evaluation bias. This review summarizes the key developments that frame the contribution of the present study.

2.1. Corporate Bankruptcy Prediction

Corporate bankruptcy forecasting seeks to identify firms at risk of failure before the event materializes. The stakes are high for investors, creditors, regulators, and supply chains, since failure events propagate losses and disrupt markets (Noh, 2023). Traditional accounting-based screening relies on ratio analysis and heuristics that are transparent, but they may underperform when patterns are nonlinear, high dimensional, or time varying. This motivates a shift to data-driven classifiers that can exploit interactions among profitability, leverage, liquidity, and efficiency signals.

Machine learning offers a pragmatic middle ground between interpretability and predictive power. Models can be calibrated to penalize missed bankruptcies more heavily than false alarms, enabling risk-sensitive early warning. In practice, robust evaluation requires attention to class imbalance and replicable workflows so that performance claims translate to operational value rather than overfitting on a single split.

2.2. Financial Ratios as Predictors

Financial ratios distil accounting statements into standardized indicators. Profitability measures capture the firm’s ability to generate earnings relative to assets or equity; liquidity ratios reflect near-term solvency and working capital health; leverage gauges exposure to fixed obligations; turnover and efficiency ratios reflect operating discipline. Individually, each group correlates with distress, but the strongest signals emerge from combinations that reflect joint pressures on cash flows and obligations (Giordani et al., 2014). Empirical studies consistently find that ratio sets outperform single indicators, particularly when interactions and nonlinearities are modelled. Yet ratios can be noisy due to fiscal seasonality, outliers, and differences in accounting policies. Consequently, preprocessing steps such as outlier filtering, along with scaling, often stabilize learning and improve generalization.

2.3. Landscape of Machine Learning for Bankruptcy

2.3.1. Decision Trees and Bagging

Decision trees partition the feature space with simple rules that are easy to interpret but can be unstable. Random forest addresses variance by aggregating many decorrelated trees trained on bootstrap samples and random feature subsets. The result is a strong baseline that handles mixed scales, is resilient to overfitting, and provides feature importance scores that help domain interpretation. In bankruptcy tasks, random forest often delivers competitive F1 scores and reliable ROC AUC without extensive tuning. Its limitations include less precise probability estimates and potential bias when minority class weighting is not applied. Nevertheless, as a production baseline, it balances accuracy, robustness, and interpretability (Gholampoor & Asadi, 2024).

2.3.2. Boosting on Tabular Data

Boosting builds learners sequentially, where each tree corrects residual errors from prior iterations. XGBoost and Gradient Boosting succeed on structured data by capturing nonlinear interactions and complex boundaries. When coupled with careful regularization and learning-rate control, boosted ensembles frequently set the state of the art on tabular classification. However, boosting is sensitive to class imbalance and hyperparameter selection (Imani et al., 2025). Without class weighting or threshold calibration, a model can optimize majority accuracy while failing to rank minority events, degrading ROC AUC. Proper tuning, early stopping, and cost-sensitive learning mitigate these effects and typically restore strong discrimination.

2.3.3. Margin Based Classifier

Support Vector Machine seeks a maximum margin boundary, which can be made nonlinear through kernels. With standardized features and tuned penalty and kernel parameters, SVM can form tight decision boundaries that generalize well even in high dimensions. In distressed-event prediction, SVM performance depends on scaling, imbalance handling, and kernel choice. It may achieve good accuracy but falter in probability ranking unless calibrated (X. Liu & He, 2022). Hence SVM often serves as a complementary comparator rather than the sole production model.

2.3.4. Ensemble Voting

Voting ensembles combine diverse base learners to reduce individual biases. Hard voting aggregates class labels, while soft voting averages predicted probabilities. When base models offer complementary strengths, ensembling improves stability across data splits and moderate's variance in threshold-dependent metrics (Amirshahi & Lahmiri, 2024). In bankruptcy prediction, soft voting often yields higher ROC AUC and comparable F1 score relative to single models, especially when at least one base learner captures minority patterns and another controls false alarms. This supports deployment needs where both sensitivity and precision matter.

2.4. Common Issues in the Literature

This section highlights three pervasive issues that most strongly affect the reliability of corporate bankruptcy prediction including class imbalance, data quality and preprocessing, and feature selection. These issues influence not only classification errors but also the stability of risk ranking across decision thresholds. The Figure 1. presents a concept map that synthesizes the three issue groups. On the right, class imbalance is addressed by reporting minority-sensitive metrics such as F1 score, ROC AUC, and PR AUC, and by applying remedies including class weighting, resampling, and cost-sensitive objectives. At the bottom, data quality and preprocessing emphasize handling missing values, extreme outliers, heterogeneous scales, imputation, outlier filtering, and standardization to stabilize training and avoid leakage. On the left, feature selection underscores correlation analysis and statistical screening such as the chi-square test, ideally nested within cross-validation to prevent information leakage, while improving interpretability and reducing computational cost without sacrificing ROC AUC. Taken together, the map guides researchers in choosing metrics, preprocessing procedures, and selection strategies that align with business goals and replication requirements.

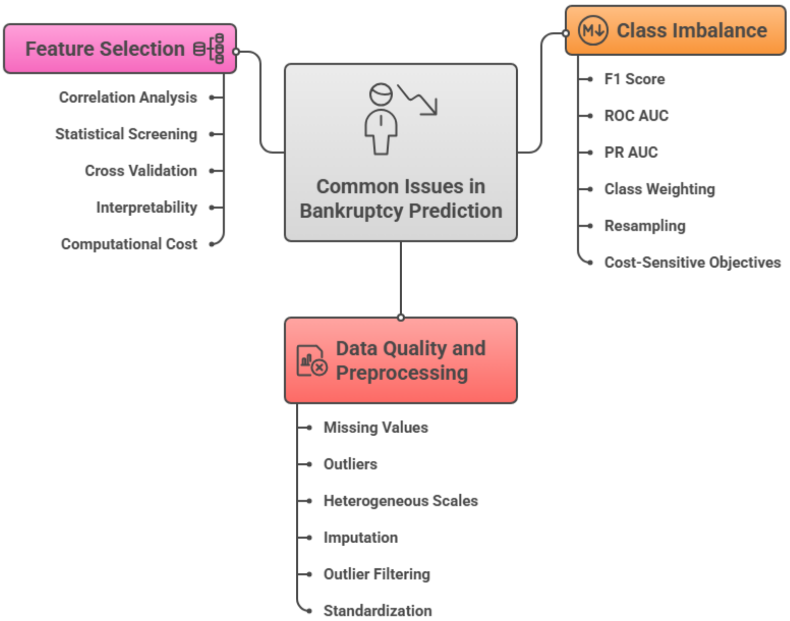

Figure 1. Common Issues in Bankruptcy Prediction

2.5. Summary of Related Studies

Comparative research increasingly shows that ensemble tree methods outperform linear baselines on tabular financial ratios, especially after addressing class imbalance. Random Forest remains a reliable baseline, while gradient boosting variants typically achieve higher ROC AUC when hyperparameters and class-cost settings are optimized. SVM remains competitive with proper scaling and probability calibration, though many studies combine it with ensembles to improve ranking performance. The literature is also converging toward multi-metric evaluation and transparent, reproducible pipelines using public datasets. Our study follows this direction by adopting an open workflow and reporting complementary threshold-dependent and threshold-independent metrics.

Recent studies on bankruptcy prediction highlight the growing dominance of ensemble and hybrid machine learning models over traditional methods. As shown in Table 1, most research demonstrates that approaches such as Random Forest, Gradient Boosting, and Voting Classifiers deliver superior accuracy and stability across imbalanced financial datasets. Moreover, the literature emphasizes the importance of multi-metric evaluation and model interpretability to ensure robust and transparent financial risk prediction.

In summary, recent studies consistently highlight the dominance of ensemble and boosting models in corporate bankruptcy prediction, supported by their strong discrimination and robustness across sectors and dataset scales. Classical models such as SVM remain viable when supported by appropriate preprocessing, yet the empirical evidence increasingly favours tree-based ensembles, particularly under imbalanced conditions and multi-metric evaluation frameworks. Table 1 provides a consolidated overview of these state-of-the-art findings, illustrating the methodological patterns that inform the design of the present study.

Table 1. State of the Art

Study (Year)	Dataset and Size	Methods Compared	Evaluation Metrics	Key Findings and Notes
Irmalasari & Dwiyanti (2023)	Public corporate financial ratios; multi-sector	Decision Tree, Random Forest, Gradient Boosting, Voting ensemble	Confusion matrix, Accuracy, ROC AUC	Ensemble voting improved stability across splits; random forest strong baseline with competitive F1 score.
Dasilas & Rigani (2024)	Financial distress panel; multi-year	Classical ML vs. ensembles	Precision, Recall, F1 score	Emphasized multi-metric evaluation for imbalanced labels; accuracy alone was misleading.
Nguyen et al. (2025)	International firm-level financial ratios	ML + SHAP Interpretability	F1 Score, ROC AUC	Gradient boosting performed best; SHAP improved transparency and ratio-level interpretation.
Noh (2023)	Imbalanced financial data from multiple industries	ANN, SVM, Tree Models, Ensembles	Accuracy, Recall, ROC AUC	Ensemble methods outperformed linear and shallow models under heavy imbalance conditions.
Shetty et al. (2022)	U.S. financial-ratio dataset	Random Forest, XGBoost, Logistic Regression	Accuracy, F1 Score	Machine learning models outperformed traditional ratio-based predictions; XGBoost strongest overall.
Gnip et al. (2025)	Large-scale imbalanced bankruptcy dataset	Range of imbalanced-learning algorithms	ROC AUC, PR AUC, G-mean	Cost-sensitive boosting and balanced resampling yielded highest stability across folds.
Ainan et al. (2024)	Polish corporate dataset (imbalanced)	Hybrid ML Techniques	Accuracy, F1 Score, ROC AUC	Hybrid models improved minority detection; emphasized impact of preprocessing quality.

3. Research Methodology

The research workflow of corporate bankruptcy prediction consists of three main phases: data preparation, model development, and validation. It begins with data preprocessing and feature selection to ensure that the input features are clean and relevant. The overall process, including model training using individual classifiers, ensemble integration through a voting classifier, and final evaluation to produce predicted labels, is depicted in Figure 2.

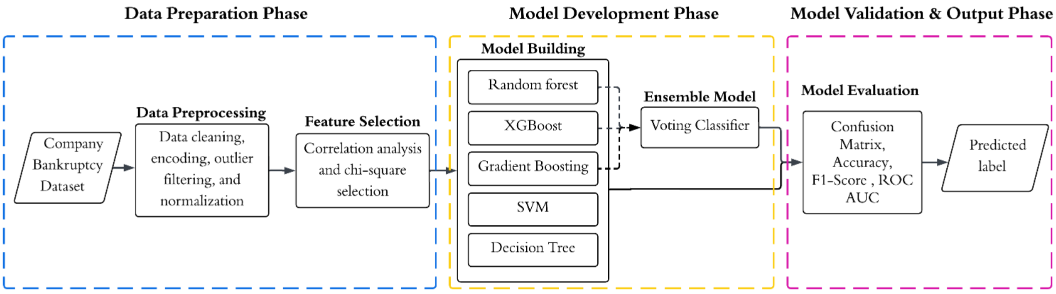

Figure 2. Workflow of Corporate Bankruptcy Prediction

3.1. Dataset Description

The analysis draws on the “US Company Bankruptcy Prediction” dataset released on Kaggle by Fedesoriano and licensed under Creative Commons Attribution-ShareAlike 4.0, which allows unrestricted use provided proper attribution is maintained. The collection contains 78 682 annual-statement observations for mid- and large-capitalisation firms listed on the NYSE and NASDAQ from 1999 to 2018, each described by 18 financial-ratio predictors that capture profitability, liquidity, leverage and operating efficiency as summarized in Table 2. Key indicators such as EBITDA (earnings before interest, taxes, depreciation and amortization) and EBIT (earnings before interest and taxes) measure operating profit with and without non-cash charges (Brygała, 2022; Jandaghi et al., 2021).

Despite the two-decade coverage, these ratios remain compatible with the post-2017 International Financial Reporting Standards, especially IFRS 9 Financial Instruments and IFRS 15 Revenue from Contracts with Customers because the newer rules modify timing and classification rather than the basic measurement of assets, liabilities and retained earnings. The target label identifies bankrupt firms (6.6 %) versus solvent ones (93.4 %), producing a markedly imbalanced class distribution that mirrors real-world rarity of failure (Gnip et al., 2025). The dataset spans multiple industries, thereby supporting the creation of broadly generalizable bankruptcy-prediction models (Ainan et al., 2024) while ensuring that the findings remain relevant under the current IFRS framework.

Table 2. Dataset Attributes

Variable Name	Description	Variable Name	Description
X1	Current assets	X11	Long-term debt
X2	Cost of goods sold	X12	EBIT
X3	Depreciation & amortization	X13	Gross profit
X4	EBITDA	X14	Current liabilities
X5	Inventory	X15	Retained earnings
X6	Net income	X16	Total revenue
X7	Receivables	X17	Total liabilities
X8	Market value	X18	Operating expenses
X9	Net sales	year	Year of record
X10	Total assets	status_label	Bankruptcy status: Failed or Alive

3.2. Data Preparation Phase

In the early phase in Figure 2, the raw data underwent a series of preprocessing steps to ensure its quality and suitability for model development. These steps included handling missing values using mean imputation for numerical features and encoding the binary target variable (alive or failed) into numeric format (Hassan & Yousaf, 2022). Outliers were detected and removed using the Isolation Forest algorithm to reduce the influence of extreme financial ratios. Finally, all features were normalized using z-score standardization to ensure that each variable contributes proportionally during model training. These operations were conducted in parallel with exploratory data analysis (EDA) to observe distributions and validate imputation strategies (Chakraborty & Ranjan, 2024).

To further refine the input space, feature selection was applied to identify the most relevant predictors. Correlation analysis helped eliminate highly collinear variables, reducing redundancy (Premalatha et al., 2023). Additionally, SelectKBest with chi-square (χ^2) tests was used to rank features based on their statistical association with the bankruptcy label (Ainan et al., 2024). This two-step process ensured that the final feature set was both informative and non-redundant, helping prevent overfitting and reduce computational cost.

3.3. Model Development Phase

To develop a robust predictive model for corporate bankruptcy, several machine learning algorithms were evaluated and compared to determine the most effective approach. As illustrated in Figure 2, we implemented five base classifiers: Random Forest, XGBoost, Gradient Boosting, Support Vector Machine (SVM), and Decision Tree. Each model was tuned using grid search and trained with cross-validation to balance bias-variance trade-offs and reduce the risk of overfitting (Mulyanto et al., 2025). Their performance was measured based on accuracy, recall, precision, and F1-score to identify strengths and weaknesses under imbalanced class conditions. The model specifications are as follows:

3.3.1. Random Forest

Random Forest is an ensemble method that constructs N decision trees $\{T1, T2, ..., TN\}$, each trained on a bootstrapped subset of data and random feature subsets. The final prediction for classification is obtained via majority voting (Equation 1):

$$\hat{y} = mode(T1(x), T2(x), ..., TN(x)) \tag{1}$$

where x is the input feature vector (Nguyen et al., 2025). Random feature sampling helps reduce variance and mitigate overfitting compared to single decision trees.

3.3.2. XGBoost (Extreme Gradient Boosting)

XGBoost (Extreme Gradient Boosting) is a boosting algorithm that sequentially adds new trees to minimize the following objective function (Equation 2):

$$\boldsymbol{\mathcal{L}} = \sum_{i=1}^{n} \boldsymbol{l}(\boldsymbol{y}_{i,}\hat{y}_i^{(t-1)} + \boldsymbol{ft}(\boldsymbol{xi})) + \Omega(ft) \tag{2}$$

where l is a differentiable loss function (e.g., log-loss for classification), ft is the new tree added at iteration t, and $\Omega(ft)$ is a regularization term that controls model complexity (Yotsawat et al., 2023). XGBoost includes shrinkage, subsampling, and column sampling to further prevent overfitting.

3.3.3. Gradient Boosting

Gradient Boosting is similar to XGBoost but typically lacks advanced regularization. The core update at each stage adds a new learner $h_m(x)$ to correct the residual errors (Equation 3):

$$\hat{y}^{(m)} = \hat{y}^{(m-1)} + \gamma_m h_m(x) \tag{3}$$

where γ_m is the learning rate (Babu et al., 2024).

3.3.4. Support Vector Machine (SVM)

SVM constructs a decision boundary by maximizing the margin between two classes. A margin-based classifier with an RBF kernel to capture non-linear patterns; it was included for its strength in high-dimensional decision boundaries (Rahman et al., 2024). For non-linear problems, the RBF (Radial Basis Function) kernel is often used (see Equation 4):

$$K(xi, xj) = exp(-\gamma \,||\, xi - xj \,||^2) \tag{4}$$

where xi and xj are input feature vectors, $||\, xi - xj \,||^2$ is the squared Euclidean distance between them, and γ is a hyperparameter controlling the kernel width.

3.3.5. Decision Tree

A simple, interpretable classifier used as a baseline model, prone to overfitting but useful for understanding feature interactions (Abir & Salam, 2024). A Decision Tree recursively partitions the feature space by selecting the feature j and threshold s that minimizes impurity (Equation 5):

$$Gini(S) = 1 - \sum_{k-1}^{K} pk^2 \tag{5}$$

where pk is the proportion of class k in subset S.

3.3.6. Ensemble Voting Classifier

The ensemble voting classifier aggregates predictions from multiple base learners. In hard voting (Equation 6):

$$\hat{y} = mode\{\hat{y}1, \hat{y}2, \ldots, \hat{y}K\} \tag{6}$$

where $\hat{y}K$ is the predicted class from model k (Gohil et al., 2023). This strategy leverages the complementary strengths of individual models, often improving stability and generalization, particularly in imbalanced datasets.

Ensemble Voting Classifier was constructed by combining top-performing models (Random Forest, XGBoost, Gradient Boosting) using a hard voting scheme (Z. Liu et al., 2024). This ensemble strategy allowed us to leverage the complementary strengths of multiple models and provided more stable and balanced predictions. While the ensemble offered only a marginal performance boost, it contributed to a more reliable classification, particularly for the minority (bankrupt) class (Tien et al., 2022). All models were implemented using Python's Scikit-learn, XGBoost, and LightGBM libraries, forming a comprehensive development phase grounded in best practices of machine learning experimentation.

Table 3. Hyperparameters and Constraints

Model	Hyperparameters
Random Forest	n_estimators ∈ [100, 300], max_depth ∈ [5, 20], min_samples_split ∈ [2, 10]
XGBoost	learning_rate ∈ [0.01, 0.3], n_estimators ∈ [100, 500], max_depth ∈ [3, 10]
Gradient Boosting	learning_rate ∈ [0.01, 0.3], n_estimators ∈ [100, 500], num_leaves ∈ [31, 128]
SVM	C ∈ [0.1, 10], kernel ∈ {'linear', 'rbf'}, gamma ∈ [0.001, 1]
Decision Tree	max_depth ∈ [3, 10], min_samples_split ∈ [2, 10]
Voting Classifier	estimators ∈ {RandomForest, XGBoost, GradientBoosting}, voting ∈ {'hard', 'soft'}

As detailed in Table 3, the hyper-parameter search spanned a broad yet practical range: for Random Forest we varied the number of trees (*n_estimators*) from 100 to 300, tree depth (*max_depth*) from 5 to 20, and the minimum samples to split a node (*min_samples_split*) from 2 to 10 to balance bias and variance. XGBoost and Gradient Boosting shared grids that adjusted learning rate between 0.01 and 0.3 and boosting rounds from 100 to 500, with XGBoost exploring tree depths of 3–10 and Gradient Boosting manipulating leaf counts from 31 to 128 to capture deeper interactions. The SVM grid tuned the penalty parameter C from 0.1 to 10, switched between linear and RBF kernels, and tested *gamma* values from 0.001 to 1 to optimise the decision boundary. For the Decision Tree benchmark, depths ranged from 3 to 10 and *min_samples_split* from 2 to 10, prioritizing interpretability while preventing overfitting. The Voting Classifier ensemble combined Random Forest, XGBoost, and Gradient Boosting, comparing hard versus soft voting to assess whether majority labels or averaged probabilities yielded more robust final predictions.

3.4. Model Validation and Output Phase

Model evaluation was conducted using a combination of an 80/20 stratified train-test split and 5-fold cross-validation, as illustrated in Figure 2, to ensure robust and unbiased performance measurement (Islam et al., 2024). Let D denote the entire dataset, then D was partitioned into the training set D_{train} (80%) and test set D_{tes} (20%), maintaining the original class distribution via stratified sampling to preserve representativeness, especially under class imbalance conditions (Chen et al., 2024). In this study, the target variable is a binary label indicating bankruptcy status, where "alive" refers to solvent or financially healthy companies, while "failed" indicates bankrupt or insolvent companies.

During evaluation, the confusion matrix C is constructed as (Equation 7):

$$C = \begin{pmatrix} TP & FP \\ FN & TN \end{pmatrix} \quad (7)$$

where:

TP (True Positive): bankrupt correctly predicted as bankrupt,

FP (False Positive): solvent incorrectly predicted as bankrupt,

FN (False Negative): bankrupt incorrectly predicted as solvent,

TN (True Negative): solvent correctly predicted as solvent (Doroshenko & Savchuk, 2024)

From the confusion matrix in Equation (7), multiple performance metrics are derived:

3.4.1. Accuracy

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (8)$$

representing the proportion of correctly classified samples over the total population (Equation 8) (Mulyanto et al., 2025).

Given the class imbalance, greater emphasis was placed on the F1-score, which balances precision and recall to better reflect the model's performance on the minority class (bankrupt firms). The precision measuring the proportion of correctly identified bankrupt companies among all predicted bankrupt cases (Equation 9) and recall indicating the proportion of actual bankrupt firms correctly detected (Equation 10) were calculated respectively as:

$$Precision = \frac{TP}{TP + FP} \tag{9}$$

$$Recall = \frac{TP}{TP + FN} \tag{10}$$

and the F1-score is computed as the harmonic mean of precision and recall (Equation 11):

3.4.2. F1-Score

$$F1 = 2 \bullet \frac{Precision \bullet Recall}{Precision + Recall} \tag{11}$$

which combines precision and recall via their harmonic mean to balance the trade-off, especially relevant for imbalanced datasets (Smiti et al., 2024).

3.4.3. ROC Curve and AUC (Area Under the Curve)

The ROC curve plots True Positive Rate (TPR) against False Positive Rate (FPR) (Equation 12):

$$TPR = \frac{TP}{TP + FN}, \qquad FPR = \frac{FP}{FP + TN} \tag{12}$$

The AUC summarizes the overall discriminative capacity of the model across all classification thresholds. A perfect classifier achieves AUC = 1, while random classification yields AUC = 0.5 (Sharmily et al., 2024). To mitigate the risk of overfitting and ensure model generalization, 5-fold cross-validation was incorporated during the training phase. The dataset was partitioned into $k = 5$ folds at each iteration, one-fold was retained as the validation set and the remaining folds served for training, rotating across folds. The averaged metrics across all folds provided an unbiased estimation of model performance (Pretnar Žagar & Demšar, 2022). Final model performance was then reported on the independent hold-out test set D_{tes}.

4. Result and Discussion

This section presents the results of the bankruptcy prediction models evaluated in the previous methodology phase. Each model's performance is assessed based on predefined evaluation metrics, including confusion matrix, accuracy, F1-score, and ROC–AUC. The discussion highlights the strengths and limitations of each model in handling imbalanced data and identifying bankrupt companies effectively.

4.1. Research Results

The findings highlight how each model performs under conditions of severe class imbalance, emphasizing differences in discrimination ability, error patterns, and overall predictive reliability. The results also provide a basis for assessing which methods offer the most robust and operationally relevant early-warning signals for financial distress.

4.1.1. Exploratory Data Analysis

The exploratory data analysis revealed a significant class imbalance, where the number of companies labelled as “alive” vastly outnumbered those marked as “failed,” confirming the need for evaluation metrics beyond simple accuracy. Several numerical features exhibited skewed distributions, particularly those related to profitability and debt ratios, indicating the presence of extreme values or outliers. A correlation matrix was generated to examine the linear relationships among features, as presented in Figure 3. The matrix shows multiple instances of high positive correlation, such as between X1 and X18 (0.99), X2 and X9 (0.98), and X4 and X12 (0.97), suggesting strong multicollinearity. These patterns prompted the removal of redundant features during the feature selection phase to reduce overfitting risk and improve model interpretability. Overall, the EDA findings guided key preprocessing decisions, including outlier handling, normalization, and the refinement of the feature set used for model training.

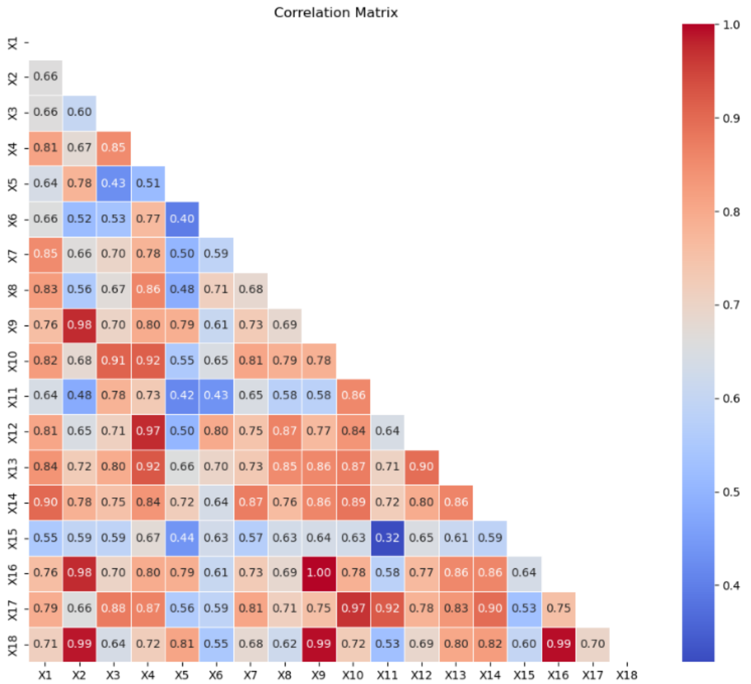

Figure 3. Correlation Matrix of Financial Features Used in Bankruptcy Prediction

4.1.2. Model Development, Evaluation, and Deployment

The performance of the five candidate models was assessed with four evaluations: confusion matrix, accuracy, F1-score, and ROC-AUC, that both overall correctness and minority-class sensitivity could be judged objectively. The confusion matrix in Figure 4 makes clear that Random Forest offers the most reliable bankruptcy detection: it correctly labels 14.659 firms as bankrupt and only misses 10 actual failures, while generating 1.011 false alarms out of 14.666 solvent cases. Gradient Boosting is a close second, matching the 14.659 true positives and 10 false negatives but at the cost of 1.055 false positives and 13 true negatives. XGBoost also performs strongly, identifying 14.636 bankruptcies with 33 misses and 1002 false alarms, correctly classifying 66 non-bankrupt firms. The Voting Classifier achieves 14.392 true positives, but its higher miss rate of 444 and 300 false positives illustrates the trade-off inherent in ensembling soft-voted probabilities. In contrast, Decision Tree struggles, misclassifying 866 bankruptcies and 791 solvent firms, and SVM fails outright by detecting no bankrupt firms (0 true positives) despite perfect recognition of healthy companies (14.669 true negatives). These results establish Random Forest as the top performer, Gradient Boosting as the runner-up, and SVM as the least effective model for this imbalanced classification task.

The performance comparison in Figure 5 reveals that the Voting Classifier leads with an accuracy of 0.936, an F1 score of 0.965, and a ROC AUC of 0.826, demonstrating its strength at correctly classifying both majority and minority cases and at ranking bankruptcy risk across all thresholds. Random Forest follows closely with an accuracy of 0.935, an F1 score of 0.965, and a ROC AUC of 0.821, reflecting its reliable bagging of decision trees. XGBoost achieves an accuracy of 0.934 and the highest F1 score of 0.966, but its ROC AUC of 0.778 shows slightly less robustness in probability ranking. Support Vector Machine records an accuracy of 0.932 and an F1 score of 0.964, yet its ROC AUC of 0.581 indicates limited separation between failed and solvent firms. Gradient Boosting posts an accuracy of 0.932 and an F1 score of 0.965 but a ROC AUC of only 0.505, suggesting that despite strong threshold-based metrics it struggles to rank cases effectively. Decision Tree lags with an accuracy of 0.894, an F1 score of 0.964, and a ROC AUC of 0.508, highlighting overfitting and misclassification issues. Overall, the Voting Classifier is the top model, Random Forest ranks second, and Gradient Boosting performs the weakest when all three measures are considered.

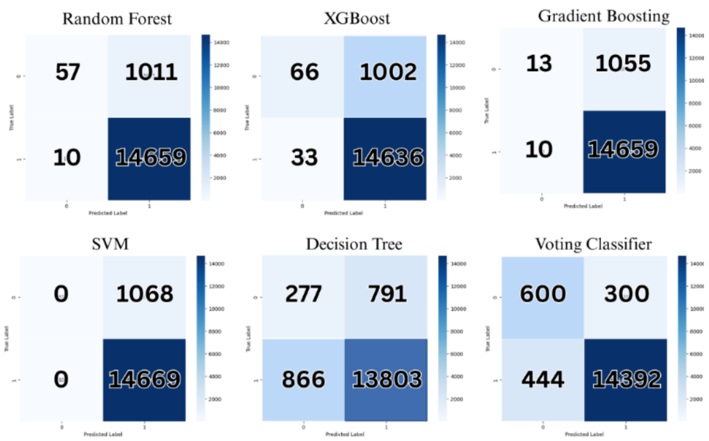

Figure 4. Confusion matrix comparing bankruptcy prediction performance

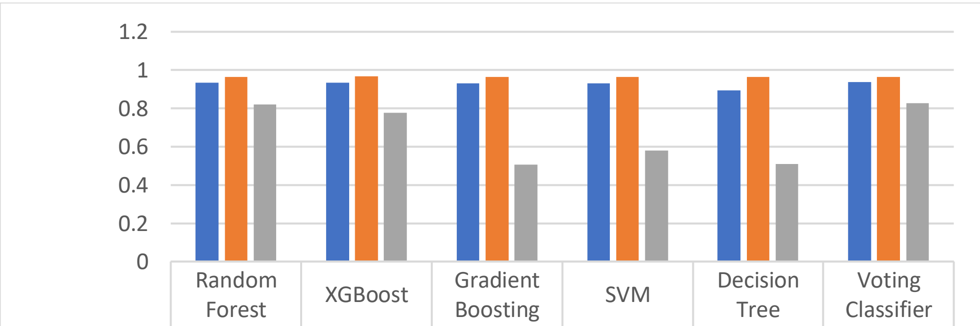

	Random Forest	XGBoost	Gradient Boosting	SVM	Decision Tree	Voting Classifier
Accuracy	0.935	0.934	0.932	0.932	0.894	0.936
F1-Score	0.965	0.966	0.965	0.964	0.964	0.965
ROC-AUC	0.821	0.778	0.505	0.581	0.508	0.826

Figure 5. Accuracy, F1-Score and ROC-AUC comparison

The comparative evaluation of six machine learning models for corporate bankruptcy prediction confirms the ensemble Voting Classifier as the most effective solution, with the Random Forest a close second and Gradient Boosting trailing. The Voting Classifier recorded an accuracy of 0.936, an F1 score of 0.965, and a ROC AUC of 0.826. Its confusion matrix shows 600 true-positive bankruptcies, 444 false negatives, and 300 false alarms. Random Forest achieved 0.935 accuracy, an identical F1 score of 0.965, and a ROC AUC of 0.821 while identifying 14.659 bankrupt firms with only 10 missed cases and 1.011 false positives. Gradient Boosting matched the Random Forest in true positives and false negatives yet produced 1 055 false positives and a ROC AUC of 0.505, revealing weak discrimination between failed and solvent firms despite apparently strong threshold-based metrics.

4.2. Discussion

This section interprets the empirical results presented in the previous subsection and explains their implications for bankruptcy prediction under severe class imbalance. The discussion highlights how each model's behavior reflects its underlying learning mechanism, with particular attention to discrimination ability, error structure, and suitability for early-warning applications. The analysis also connects these findings to insights from prior studies, clarifying where the present results reinforce or challenge existing evidence.

4.2.1. Comparative Interpretation of Model Performance

The comparative analysis demonstrates clear differences in how the models respond to the severe class imbalance characteristic of bankruptcy data. The Voting Classifier consistently delivers the most balanced predictive behaviour, benefiting from the complementary strengths of its constituent models and maintaining strong performance across both threshold-based and threshold-independent metrics. As illustrated in Figure 5, Random Forest stands out as the most dependable single model, offering high recall with minimal missed bankruptcies and stable discrimination driven by its capacity to model nonlinear interactions. By contrast, Gradient Boosting, despite achieving competitive F1 scores, exhibits substantially weaker ranking capability, indicating vulnerability to imbalance effects. These results reinforce the broader evidence that ensemble approaches provide superior stability and discriminative power for structured financial-ratio tasks.

The quantitative comparison further confirms the superiority of the Voting Classifier, which achieves 0.936 accuracy, 0.965 F1 score, and 0.826 ROC AUC. Random Forest closely matches this performance with 0.935 accuracy, 0.965 F1 score, and 0.821 ROC AUC, supported by a markedly low number of missed bankruptcies. Gradient Boosting, although comparable in true positives and false negatives, records 1,055 false alarms and a ROC AUC of 0.505, demonstrating limited ability to separate distressed from solvent firms. Taken together, the results empirically establish ensemble-based models as the most reliable framework for bankruptcy risk assessment and confirm Random Forest as the strongest stand-alone alternative for practical early-warning systems.

These results illustrate the importance of evaluating models with multiple metrics in imbalanced datasets. Accuracy alone may look impressive, but ROC AUC exposes whether a model truly separates classes across decision thresholds. Gradient Boosting's low ROC AUC underscores this point, as it effectively defaulted to majority-class predictions. Conversely, the Voting ensemble consistently delivered high values across all key measures, balancing sensitivity and specificity in detecting at-risk companies. Therefore, the Voting Classifier stands as the best model for reliable early bankruptcy warning, the Random Forest is the second-best, and Gradient Boosting is the weakest performer. This finding fulfils the research objective of identifying a dependable machine learning framework to support stakeholders in proactively managing corporate financial distress.

Although ROC-AUC and F1-score demonstrate the technical discriminative power of each model, lenders ultimately care about the economic cost of mis-classification. A false-positive bankruptcy flag can discourage credit approval or trigger higher interest spreads, inflating a solvent firm's cost of capital; conversely, a false negative exposes the lender to default losses. Therefore, model selection and decision-threshold tuning should balance recall and specificity in line with the institution's risk appetite and pricing framework. A cost-sensitive analysis can translate classification outcomes into monetary terms and highlight the break-even threshold where marginal gains in recall outweigh added false-positive costs.

4.2.2. Implications, Limitations, and Directions for Future Research

The results carry important implications for both researchers and practitioners. For researchers, the findings demonstrate the necessity of multi-metric evaluation and reproducible workflows when studying imbalanced financial-risk datasets, as reliance on accuracy alone can produce misleading conclusions. For practitioners, the clear superiority of the Voting Classifier and Random Forest suggests that ensemble-centric solutions offer dependable early-warning systems with strong minority-class detection and manageable false-positive rates. Nonetheless, the study is constrained by its use of a single publicly available dataset, which, while supporting reproducibility, limits cross-market generalization. Future research should focus on improving probability calibration, optimizing decision thresholds, and exploring cost-sensitive objectives to align model behaviour with the economic consequences of errors. Incorporating interpretability techniques such as SHAP values would also strengthen stakeholder trust by clarifying the financial ratios that drive predictions. These enhancements can improve the operational value of bankruptcy-prediction models without compromising transparency or replicability.

5. Conclusion

The primary aim of this study was to benchmark six widely used machine learning models for corporate bankruptcy prediction and to close key methodological gaps in prior work, particularly the lack of multi-metric evaluation, reproducible workflows, and fair assessment under imbalanced data. Using a large dataset of 78,682 U.S. firms and a consistent pipeline from preprocessing to evaluation, the study demonstrates that the Voting Classifier delivers the most balanced performance, effectively detecting minority bankruptcy cases with moderate false alarms, while Random Forest emerges as the strongest single-model alternative. Gradient Boosting, despite matching F1 scores, shows markedly lower ROC AUC, indicating weak ranking ability under imbalance. Quantitatively, the Voting Classifier achieved 0.936 accuracy, 0.965 F1, and 0.826 ROC AUC, whereas Random Forest reached 0.935 accuracy, 0.965 F1, and 0.821 ROC AUC; in contrast, Gradient Boosting recorded only 0.505 ROC AUC. For researchers, the main implication is the clear superiority of ensemble methods for stable discrimination, while practitioners can adopt Voting Classifier or Random Forest as reliable components of early-warning systems that balance recall of at-risk firms with manageable false positives.

Future work should retain the current dataset to ensure replicability while prioritizing methodological improvements that address class imbalance and enhance probability calibration. Cost sensitive learning, threshold optimization informed by precision and recall curves, and deeper tuning of Gradient Boosting can improve ranking performance and minority class detection. Incorporating interpretability methods such as SHAP values would strengthen transparency and stakeholder trust by clarifying the financial ratios that influence predictions. A training objective guided by a cost matrix may also raise recall for at-risk firms without creating an unacceptable increase in false positives, making future models more suitable for practical early warning applications.

6. References

Abir, M. I. H., & Salam, T. (2024). Comparative Analysis and Prediction of Machine Learning Algorithms for MRI-Based Alzheimer's Detection Using Multi-modal Data. *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)*, 1–5. https://doi.org/10.1109/COMPAS60761.2024.10797119

Ainan, U. H., Por, L. Y., Chen, Y.-L., Yang, J., & Ku, C. S. (2024). Advancing Bankruptcy Forecasting With Hybrid Machine Learning Techniques: Insights From an Unbalanced Polish Dataset. *IEEE Access*, *12*, 9369–9381. https://doi.org/10.1109/ACCESS.2024.3354173

Akinjole, A., Shobayo, O., Popoola, J., Okoyeigbo, O., & Ogunleye, B. (2024). Ensemble-Based Machine Learning Algorithm for Loan Default Risk Prediction. *Mathematics*, *12*(21), 3423. https://doi.org/10.3390/math12213423

Alam, T. M., Shaukat, K., Mushtaq, M., Ali, Y., Khushi, M., Luo, S., & Wahab, A. (2021). Corporate Bankruptcy Prediction: An Approach Towards Better Corporate World. *The Computer Journal*, *64*(11), 1731–1746. https://doi.org/10.1093/comjnl/bxaa056

Amirshahi, B., & Lahmiri, S. (2024). Bankruptcy prediction using optimal ensemble models under balanced and imbalanced data. *Expert Systems*, *41*(8). https://doi.org/10.1111/exsy.13599

Arora, I., & Singh, N. (2020). Prediction of Corporate Bankruptcy using Financial Ratios and News. *International Journal of Engineering and Management Research*, *10*(5), 82–87. https://doi.org/10.31033/ijemr.10.5.15

Brygała, M. (2022). Consumer Bankruptcy Prediction Using Balanced and Imbalanced Data. *Risks*, *10*(2), 24. https://doi.org/10.3390/risks10020024

Brygała, M., & Korol, T. (2024). Personal bankruptcy prediction using machine learning techniques. *Economics and Business Review*, *10*(2), 118–142. https://doi.org/10.18559/ebr.2024.2.1149

Chaising, S., & Srimaharaj, W. (2024). Ensemble Framework for Bankruptcy Prediction. *2024 Joint International Conference on Digital Arts, Media and Technology with ECTI Northern Section Conference on Electrical, Electronics, Computer and Telecommunications Engineering (ECTI DAMT & NCON)*, 545–550. https://doi.org/10.1109/ECTIDAMTNCON60518.2024.10480004

Chakraborty, D. B., & Ranjan, R. (2024). Missing Data Imputation With Contextual Granules and AI-driven Bankruptcy Prediction. *2024 14th International Conference on Pattern Recognition Systems (ICPRS)*, 1–8. https://doi.org/10.1109/ICPRS62101.2024.10677843

Chen, M., Fan, W., Tang, W., Liu, T., Li, D., & Dib, O. (2024). *Review of Machine Learning Algorithms for Breast Cancer Diagnosis* (pp. 229–243). https://doi.org/10.1007/978-981-97-0844-4_17

da Silva Mattos, E., & Shasha, D. (2024). Bankruptcy prediction with low-quality financial information. *Expert Systems with Applications*, *237*, 121418. https://doi.org/10.1016/j.eswa.2023.121418

Dasilas, A., & Rigani, A. (2024). Machine learning techniques in bankruptcy prediction: A systematic literature review. *Expert Systems with Applications*, *255*, 124761. https://doi.org/10.1016/j.eswa.2024.124761

Doroshenko, A. V., & Savchuk, D. Y. (2024). Research of data mining methods for classification of imbalanced data sets. *Ukrainian Journal of Information Technology*, *6*(1), 48–57. https://doi.org/10.23939/ujit2024.01.048

Gabrielli, G., Melioli, A., & Bertini, F. (2023). High-dimensional Data from Financial Statements for a Bankruptcy Prediction Model. *2023 IEEE 39th International Conference on Data Engineering Workshops (ICDEW)*, 1–7. https://doi.org/10.1109/ICDEW58674.2023.00005

Gholampoor, H., & Asadi, M. (2024). Risk Analysis of Bankruptcy in the U.S. Healthcare Industries Based on Financial Ratios: A Machine Learning Analysis. *Journal of Theoretical and Applied Electronic Commerce Research*, *19*(2), 1303–1320. https://doi.org/10.3390/jtaer19020066

Giordani, P., Jacobson, T., Schedvin, E. von, & Villani, M. (2014). Taking the Twists into Account: Predicting Firm Bankruptcy Risk with Splines of Financial Ratios. *Journal of Financial and Quantitative Analysis*, *49*(4), 1071–1099. https://doi.org/10.1017/S0022109014000623

Gnip, P., Kanász, R., Zoričak, M., & Drotár, P. (2025). An experimental survey of imbalanced learning algorithms for bankruptcy prediction. *Artificial Intelligence Review*, *58*(4), 104. https://doi.org/10.1007/s10462-025-11107-y

Gohil, R., S, D., M, V., & J, J. (2023). Election Forecasting with Machine Learning and Sentiment Analysis: Karnataka 2023. *2023 International Conference on Ambient Intelligence, Knowledge Informatics and Industrial Electronics (AIKIIE)*, 1–6. https://doi.org/10.1109/AIKIIE60097.2023.10390333

Hassan, A., & Yousaf, N. (2022). Bankruptcy Prediction using Diverse Machine Learning Algorithms. *2022 International Conference on Frontiers of Information Technology (FIT)*, 106–111. https://doi.org/10.1109/FIT57066.2022.00029

Imani, M., Beikmohammadi, A., & Arabnia, H. R. (2025). Comprehensive Analysis of Random Forest and XGBoost Performance with SMOTE, ADASYN, and GNUS Under Varying Imbalance Levels. *Technologies*, *13*(3), 88. https://doi.org/10.3390/technologies13030088

Irmalasari, I., & Dwiyanti, L. (2023). Algorithm Analysis of Decision Tree, Gradient Boosting Decision Tree, and Random Forest for Classification (Case Study: West Java House of Representatives Election 2019). *2023 International Conference on Electrical Engineering and Informatics (ICEEI)*, 1–5. https://doi.org/10.1109/ICEEI59426.2023.10346727

Islam, J., Saha, S., Hasan, M., Mahmud, A., & Jannat, M. (2024). Cognitive Modelling of Bankruptcy Risk: A Comparative Analysis of Machine Learning Models to Predict the Bankruptcy. *2024 12th International Symposium on Digital Forensics and Security (ISDFS)*, 1–6. https://doi.org/10.1109/ISDFS60797.2024.10527269

Jandaghi, G., Saranj, A., Rajaei, R., Ghasemi, A., & Tehrani, R. (2021). Identification of the Most Critical Factors in Bankruptcy Prediction and Credit Classification of Companies. *Interdisciplinary Journal of Management Studies*, *14*(4). https://doi.org/10.22059/ijms.2021.285398.673712

Liu, X., & He, W. (2022). Adaptive kernel scaling support vector machine with application to a prostate cancer image study. *Journal of Applied Statistics*, *49*(6), 1465–1484. https://doi.org/10.1080/02664763.2020.1870669

Liu, Z., Zhang, H., Xiong, F., Huang, X., Yu, S., Sun, Q., Diao, L., Li, Z., Wu, Y., Zeng, Y., & Huang, C. (2024). Prediction of clinical pregnancy outcome after single fresh blastocyst transfer during in vitro fertilization: an ensemble learning perspective. *Human Fertility*, *27*(1). https://doi.org/10.1080/14647273.2024.2422918

Mulyanto, S., Yonia, D. L., & Sutejo, B. (2025). Finance Loan Risk Assessment Using Machine Learning for Credit Eligibility Prediction and Model Optimization. *IJISTECH (International Journal of Information System and Technology)*, *8*(5). https://doi.org/10.30645/ijistech.v8i5.376

Narvekar, A., & Guha, D. (2021). Bankruptcy prediction using machine learning and an application to the case of the COVID-19 recession. *Data Science in Finance and Economics*, *1*(2), 180–195. https://doi.org/10.3934/DSFE.2021010

Nguyen, H. H., Viviani, J.-L., & Ben Jabeur, S. (2025). Bankruptcy prediction using machine learning and Shapley additive explanations. *Review of Quantitative Finance and Accounting*, *65*(1), 107–148. https://doi.org/10.1007/s11156-023-01192-x

Noh, S.-H. (2023). Comparing the Performance of Corporate Bankruptcy Prediction Models Based on Imbalanced Financial Data. *Sustainability*, *15*(6), 4794. https://doi.org/10.3390/su15064794

Pawełek, B., & Pociecha, J. (2020). Corporate Bankruptcy Prediction with the Use of the Logit Leaf Model. In *Studies in Classification, Data Analysis, and Knowledge Organization* (pp. 129–146). Springer. https://doi.org/10.1007/978-3-030-52348-0_9

Premalatha, G., Priyanka, R., & Chaitya, K. (2023). Feature selection for predicting bankruptcy: Comparative analysis. *2023 Fifth International Conference on Electrical, Computer and Communication Technologies (ICECCT)*, 1–5. https://doi.org/10.1109/ICECCT56650.2023.10179633

Pretnar Žagar, A., & Demšar, J. (2022). Model Evaluation. In *Tourism on the Verge* (pp. 253–274). Springer. https://doi.org/10.1007/978-3-030-88389-8_13

Rahman, M. M., Kobir, K. H., Akther, S., & Kallol, M. A. H. (2024). Ensemble Machine Learning for Enhanced Breast Cancer Prediction: A Comparative Study. *International Journal of Advanced Computer Science and Applications*, *15*(7). https://doi.org/10.14569/IJACSA.2024.0150792

Sharmily, R. R., Karthik, B., & Vijayan, T. (2024). Improved MRI based Automatic Brain Tumor Categorization Employing Deep Learning Techniques. *2024 4th Asian Conference on Innovation in Technology (ASIANCON)*, 1–5. https://doi.org/10.1109/ASIANCON62057.2024.10837709

Shetty, S., Musa, M., & Brédart, X. (2022). Bankruptcy Prediction Using Machine Learning Techniques. *Journal of Risk and Financial Management*, *15*(1), 35. https://doi.org/10.3390/jrfm15010035

Smiti, S., Soui, M., & Ghedira, K. (2024). Tri-XGBoost model improved by BLSmote-ENN: an interpretable semi-supervised approach for addressing bankruptcy prediction. *Knowledge and Information Systems*, *66*(7), 3883–3920. https://doi.org/10.1007/s10115-024-02067-w

Tien, H. L. Q., Quang Tran, L., & Hop Do, T. (2022). An Empirical Study on Bankruptcy Prediction using Ensemble Learning. *2022 RIVF International Conference on Computing and Communication Technologies (RIVF)*, 173–178. https://doi.org/10.1109/RIVF55975.2022.10013848

Yotsawat, W., Phodong, K., Promrat, T., & Wattuya, P. (2023). Bankruptcy prediction model using cost-sensitive extreme gradient boosting in the context of imbalanced datasets. *International Journal of Electrical and Computer Engineering (IJECE)*, *13*(4), 4683. https://doi.org/10.11591/ijece.v13i4.pp4683-4691