

Hybrid Financial Representation Stacking Ensemble for Bankruptcy Prediction

Archana N, Dr. R N Kulkarni

¹Department of Management Studies, Department of Computer Science
Ballari Institute of Technology & Management, Ballari, Karnataka, India
email : archanan200018@gmail.com

Abstract—Corporate bankruptcy prediction is increasingly important because financial distress can affect creditors, owners, investors, managers, and lending institutions. This study proposes a Hybrid Financial Representation Stacking Ensemble (HFRSE) framework that combines original financial variables, traditional financial ratios, structural composition features, Autoencoder-based anomaly detection, a Temporal One-Dimensional Convolutional Neural Network (1D-CNN), CatBoost, XGBoost stacking, and SHAP-based explainability. The framework is developed using an open-source corporate bankruptcy dataset containing 78,682 firm-year observations and 18 financial variables covering the 1999–2018 reporting period. A company-based 80:20 train-test split is applied before three-year sequence construction to reduce data leakage. Because the dataset is highly imbalanced, Borderline-SMOTE is used on the training data, while an Autoencoder learns abnormal financial behaviour. The final HFRSE model achieves 72.12% accuracy, 11.88% precision, 55.29% recall, 19.56% F1-score, and 0.7015 ROC-AUC. The results indicate that the hybrid approach provides a useful balance between overall classification and detection of financially distressed firms. SHAP analysis further identifies the financial variables and anomaly information that most influence bankruptcy predictions.

Keywords—Corporate Bankruptcy, Financial Distress, Machine Learning, Deep Learning, Stacking Ensemble, SHAP

I. INTRODUCTION

Corporate bankruptcy is a major financial event because the failure of a company can create losses for creditors, owners, investors and other stakeholders. The project report identifies market volatility, economic problems and technological threats as conditions that can worsen business performance when they are not managed effectively. Bankruptcy prediction therefore sits at the intersection of finance and data analytics: financial statements provide the evidence of financial health, while analytical models can identify patterns that are difficult to observe using simple ratios alone.

Traditional bankruptcy prediction has often relied on statistical approaches such as Altman's Z-score and Ohlson's O-score. These approaches remain attractive because they are relatively easy to interpret, but they do

not fully capture complex relationships among many financial variables. The growth of digital financial information and artificial intelligence has made machine-learning and deep-learning methods increasingly useful for processing large financial datasets and detecting nonlinear patterns.

The present study develops a Hybrid Financial Representation Stacking Ensemble (HFRSE) for explainable corporate bankruptcy prediction. Instead of relying on a single representation or a single model, HFRSE combines original financial variables with eight traditional financial ratios and ten structural composition measures. An Autoencoder adds an anomaly score representing deviation from normal financial behaviour. A Temporal 1D-CNN learns patterns across three consecutive years, CatBoost models nonlinear relationships in structured financial data, and an XGBoost meta-learner combines the probability outputs of the two base models. SHAP is finally used to explain which financial factors influence the predictions.

The objective is not only to improve predictive performance but also to make the resulting predictions more useful for financial analysts, managers, investors and lending institutions. The framework is intended as an early-warning and risk-assessment tool that can help identify financially distressed firms before failure becomes unavoidable.

A. Problem Statement and Objectives

The study addresses the problem of designing and implementing an explainable HFRSE framework for corporate bankruptcy prediction. The specific objectives are to: (1) preprocess and validate corporate financial data; (2) generate traditional financial ratios and structural composition features; (3) combine original and engineered variables into a hybrid representation; (4) identify abnormal financial behaviour with an Autoencoder; (5) develop Temporal 1D-CNN and CatBoost prediction models; (6) combine the models through stacking; (7) interpret predictions using SHAP; and (8) evaluate the framework using Accuracy, Precision, Recall, F1-score and ROC-AUC.

B. Scope and Practical Relevance

The study uses historical financial information from an open-source bankruptcy dataset and focuses on early identification of corporate financial distress. The project is implemented in Python using open-source analytical and machine-learning libraries. The intended users include managers, investors, financial institutions, auditors and decision-makers who need a structured

indication of bankruptcy risk together with an explanation of the factors driving the prediction.

II. LITERATURE REVIEW

Recent research has progressed from traditional financial-ratio methods to machine learning, deep learning, ensemble methods, hybrid representations, temporal modelling and explainable AI:

- 1. Thota et al. (2025):** Examined bankruptcy prediction among Indian firms using Random Forest, Neural Networks and Gradient Boosting with SMOTE for class imbalance. The study supports AI-based prediction and explicit imbalance handling.
- 2. Islam et al. (2025):** Combined financial ratios with Management Discussion and Analysis disclosures using multimodal deep learning. The results support combining complementary financial representations for improved prediction.
- 3. Lakshmi, Martin, and Venkatesan (2016):** Used a genetic algorithm to select important financial ratios for bankruptcy prediction. Feature selection improved accuracy while reducing model complexity.
- 4. Ataş et al. (2025):** Combined Support Vector Machines with Meta Fuzzy Functions for bankruptcy prediction in Turkish commercial banks. The improved results support the use of complementary ensemble approaches.
- 5. Song and Peng (2019):** Studied imbalanced financial-risk classification and found SMOTE-based ensemble methods effective. The study also supports evaluating models using measures beyond accuracy..
- 6. Ansari et al. (2020):** Combined Magnetic Optimization Algorithm and Particle Swarm Optimization with Artificial Neural Networks for bankruptcy prediction. The hybrid optimization approach improved model performance and training efficiency.
- 7. Kadkhoda and Amiri (2024):** Combined network analysis with machine learning for financial-distress prediction. Network-based variables improved prediction, supporting feature construction beyond conventional ratios.
- 8. Gaurav et al. (2025):** Developed a CNN-based bankruptcy-prediction model using feature selection and SMOTE. CNN outperformed the reported logistic-regression and SVM models.
- 9. Magrini (2025):** Proposed an isometric log-ratio representation of financial statements to reduce redundancy and outlier effects. The study supports composition-based financial representations.
- 10. Shetty, Musa, and Bredart (2022):** Compared XGBoost, SVM and deep neural networks for Belgian SMEs and showed that a small set of informative indicators can provide useful prediction. This supports careful feature selection and ensemble modelling.
- 11. Nguyen, Viviani, and Ben Jabeur (2023):** Compared Random Forest, LightGBM, XGBoost and NGBoost and used SHAP for interpretation. The study directly supports combining ensemble learning with explainability.

12. Ogachi et al. (2020): Used financial ratios and Logistic Regression and identified key liquidity, leverage, efficiency and working-capital indicators. The findings confirm the continuing value of traditional financial ratios.

13. Gabrielli, Melioli, and Bertini (2026): Compared machine-learning methods for financial distress and found ensemble methods effective, while SMOTE improved robustness with imbalanced data. This supports HFRSE's ensemble and imbalance-handling components.

14. Hussain et al. (2026): Integrated NGBoost, CatBoost, LightGBM and SHAP across U.S. and Chinese markets. The study supports stable, interpretable ensemble models for financial decision-making.

15. Yoon, Jeong, and Kim (2020): Combined financial statements with news and social-media information for early financial-risk detection. Multiple information sources improved risk identification, supporting the broader hybrid-information principle.

16. Xiao and Liu (2026): Used a Dynamic Graph Convolutional Network with Temporal Aggregation to capture company relationships and financial changes over time. The findings support temporal learning for bankruptcy prediction.

17. Chen, Liu, and Wu (2025): Combined hybrid feature selection with an improved stacking ensemble for financial-distress prediction. The approach improved accuracy, robustness and sensitivity, supporting feature selection and stacking.

18. Yang et al. (2025): The 4C-FinNet framework used multi-channel financial representations and multi-year information to capture complex relationships. The results support richer financial representation and temporal deep learning.

19. Lombardo et al. (2024): The authors proposed an end-to-end deep-learning architecture that combined accounting data with textual information from financial reports. The integrated model performed better than models using either financial or textual information alone. The study demonstrates the value of hybrid information and points to future extensions of HFRSE.

20. Park et al. (2021): The study investigated explainability in bankruptcy prediction using LIME. LIME was effective for explaining individual predictions and produced interpretations consistent with traditional feature-importance approaches. The authors emphasized that interpretability is as important as predictive accuracy, supporting HFRSE's SHAP-based explanation layer.

A. Research Gap

The reviewed studies demonstrate substantial progress, but most concentrate on one or two methodological improvements. Some focus on financial ratios, some on feature selection, some on deep learning or ensemble learning, and others on explainability. A smaller group combines multiple information sources, but such

approaches may depend on corporate disclosures, news or other external text. The project identifies a further gap in the integrated treatment of financial representation, anomaly detection, imbalance handling, temporal learning, nonlinear structured-data modelling, stacking and explainability within one bankruptcy-prediction framework.

The HFRSE framework addresses this gap by integrating original financial variables, traditional ratios, composition-based features, Autoencoder anomaly information, Borderline-SMOTE, Temporal 1D-CNN, CatBoost, XGBoost stacking and SHAP. The objective is to improve prediction while retaining an explanation layer that can show why a firm was considered financially risky.

III. METHODOLOGY

A. Dataset and Data Preparation

The project uses an open-source corporate bankruptcy dataset containing 78,682 firm-year observations, company information, reporting year, bankruptcy status and 18 financial variables. The observations cover the 1999–2018 period. The target distribution is highly imbalanced: 73,462 observations (93.37%) represent healthy firms and 5,220 observations (6.63%) represent bankrupt firms. This imbalance makes accuracy alone inadequate for judging the model because a classifier can appear accurate while still missing a large number of bankrupt firms.

The project structures financial statements into three-year sequences ($t-2$, $t-1$ and t) so that the Temporal 1D-CNN can learn changes in financial condition rather than treating each year as an isolated observation. To reduce information leakage, companies are separated into training and testing groups using an 80:20 split before sequence creation. The training data is then processed for feature engineering, anomaly detection and imbalance handling, while the test distribution is retained for unbiased evaluation.

TABLE I. ORIGINAL FINANCIAL VARIABLES

Variable	Financial variable
X1	Current Assets
X2	Cost of Goods Sold
X3	Depreciation and Amortization
X4	EBITDA
X5	Inventory
X6	Net Income
X7	Total Receivables
X8	Market Value
X9	Net Sales
X10	Total Assets
X11	Total Long-Term Debt
X12	EBIT
X13	Gross Profit
X14	Total Current Liabilities
X15	Retained Earnings
X16	Total Revenue
X17	Total Liabilities
X18	Total Operating Expenses

In the article tables and formulas below, the descriptive financial-variable names are used instead of the X-

number labels so that the economic meaning of each calculation is immediately clear. Where the source project used X16 as Net Sales in a formula, the article uses the descriptive term Net Sales; the underlying dataset separately contains Total Revenue as X16.

B. Hybrid Financial Representation

Feature engineering expands the information available to the models. Eight traditional ratios capture liquidity, profitability, efficiency and leverage. Ten structural composition features express financial-statement items as proportions of total assets, total liabilities or revenue. These engineered features are combined with the original financial variables to create the Hybrid Financial Representation (HFR). Correlation analysis is then used to remove highly redundant variables and improve the stability of the final feature set.

TABLE II. TRADITIONAL FINANCIAL RATIOS

Category	Ratio	Formula	Purpose
Liquidity	Current Ratio	Current Assets / Total Current Liabilities	Measures ability to meet short-term obligations.
Liquidity	Working Capital Ratio	(Current Assets – Total Current Liabilities) / Total Assets	Measures working capital relative to total assets.
Profitability	Return on Assets (ROA)	Net Income / Total Assets	Measures how efficiently assets generate profit.
Efficiency	Asset Turnover	Net Sales / Total Assets	Measures how efficiently assets generate sales.
Efficiency	Inventory Turnover	Cost of Goods Sold / Inventory	Measures efficiency of inventory management.
Efficiency	Receivables Turnover	Net Sales / Total Receivables	Measures efficiency of collecting receivables.
Leverage	Debt Ratio	Total Liabilities / Total Assets	Measures the share of assets financed by liabilities.
Leverage	Long-Term Debt Ratio	Total Long-Term Debt / Total Assets	Measures the contribution of long-term debt to assets.

TABLE III. STRUCTURAL COMPOSITION FEATURES

Category	Composition feature	Formula	Purpose
Asset Structure	Current Assets to Assets	Current Assets / Total Assets	Share of total assets held as current assets.
Asset Structure	Inventory to Assets	Inventory / Total Assets	Share of assets invested in inventory.
Asset Structure	Receivables to Assets	Total Receivables / Total Assets	Share of assets tied up in receivables.
Asset Structure	Retained Earnings to Assets	Retained Earnings / Total Assets	Accumulated earnings relative to total assets.
Liability Structure	Long-Term Debt to Liabilities	Total Long-Term Debt / Total Liabilities	Share of liabilities financed by long-term debt.
Liability Structure	Current Liabilities to Liabilities	Total Current Liabilities / Total Liabilities	Share of short-term obligations.

Liabilities			
Income Structure	COGS to Revenue	Cost of Goods Sold / Total Revenue	Share of revenue consumed by COGS.
Income Structure	Gross Profit to Revenue	Gross Profit / Total Revenue	Gross profit margin from revenue.
Income Structure	Net Income to Revenue	Net Income / Total Revenue	Net profit margin from revenue.
Income Structure	Operating Expenses to Revenue	Total Operating Expenses / Total Revenue	Share of revenue spent on operating expenses.

C. Exploratory Analysis and Outlier Treatment

The exploratory analysis examines the target distribution, reporting-year distribution, financial-variable distributions, outliers and correlations. Most raw financial variables are positively skewed, with many observations concentrated at lower values and relatively few observations at very high values. Ratio variables such as Current Assets to Assets, Current Liabilities to Liabilities and Long-Term Debt to Liabilities show comparatively more stable distributions.

The project also identifies substantial outliers in raw financial variables. Examples include Market Value, Net Sales, Net Income, Debt Ratio and Inventory Turnover. These extreme observations are treated as meaningful characteristics of corporate financial data rather than automatically being discarded. Winsorization is used to reduce the influence of extreme values while retaining distressed firms that may carry important signals.




Fig. 1. Target distribution of healthy and bankrupt observations.




Fig. 2. Distribution of observations across reporting years.




Fig. 3. Distribution of original financial variables, first group.




Fig. 4. Distribution of original financial variables, second group.




Fig. 5. Distribution of traditional financial ratios and composition feature.




Fig. 6. Distribution of additional engineered financial variables.

D. Anomaly Detection and Class Imbalance

A deep Autoencoder is trained using financially healthy companies so that it learns a representation of normal financial behaviour. When a new company deviates from the learned pattern, the reconstruction error becomes larger. This error is converted into an AE_anomaly_score and added to the feature set. The

score does not replace the bankruptcy target; it provides an additional signal about abnormal financial behaviour that can help the downstream classifiers separate firms. Borderline-SMOTE is used to address class imbalance by generating synthetic minority-class observations near difficult class-boundary regions. The project applies the technique to the training process while leaving the test or validation distribution unchanged, so that final evaluation reflects the natural class distribution. This is important because the business cost of missing a financially distressed company can be greater than the cost of flagging a healthy company.

E. Temporal 1D-CNN

The Temporal 1D-CNN is designed to capture financial deterioration over time. It receives three consecutive years ($t-2$, $t-1$ and t) and learns temporal patterns in liquidity, profitability, leverage and efficiency. Instead of treating each annual statement independently, the model can therefore learn whether financial conditions are improving or deteriorating across consecutive reporting periods.

$$P(\text{Bankruptcy}) = \sigma(Wx + b)$$

Here, W represents learned model weights, x represents the extracted temporal feature representation, b is the bias term, and σ is the sigmoid function that converts the output into a probability between 0 and 1.

F. CatBoost Classifier

CatBoost is used because it is effective for structured financial data and can model nonlinear relationships among the hybrid financial variables. It analyses interactions among the original variables, financial ratios, composition features and Autoencoder anomaly score. Its bankruptcy probability is later passed to the stacking layer together with the Temporal 1D-CNN probability.

G. XGBoost Stacking Meta-Learner

The stacking layer combines the probability outputs of the Temporal 1D-CNN and CatBoost using an XGBoost meta-learner. The purpose is to exploit complementary strengths: the CNN captures temporal patterns, while CatBoost captures nonlinear relationships in structured financial features. The final decision threshold is selected by maximizing F2-score, giving greater importance to Recall because missing a bankrupt firm is considered more costly than incorrectly flagging a healthy firm.

$$\text{Final Prediction} = \text{XGBoost}(P_{\text{CNN}}, P_{\text{CatBoost}})$$

H. SHAP Explainability

SHAP (Shapley Additive Explanations) is applied after prediction to quantify the contribution of financial features to the model output. This creates an explanation layer that can identify the variables most responsible for increasing or decreasing bankruptcy risk. Such explanations are intended to make the model more useful to financial analysts and decision-makers rather than treating the prediction as an unexplained black-box result.

I. Evaluation Metrics

The framework is evaluated using Accuracy, Precision, Recall, F1-score and ROC-AUC, together with confusion matrices. Accuracy measures the overall proportion of correctly classified observations. Precision measures how many predicted bankrupt firms are actually bankrupt. Recall measures the ability to identify bankrupt firms. F1-score balances Precision and Recall, while ROC-AUC evaluates discrimination between healthy and bankrupt firms across classification thresholds.

IV. RESULTS

A. Exploratory Data Analysis Findings

The target-distribution analysis confirms that bankruptcy is a rare event in the dataset. Of 78,682 observations, 73,462 are healthy and 5,220 are bankrupt. The strong imbalance explains why the study emphasizes Recall, F1-score and ROC-AUC in addition to Accuracy. The year-distribution analysis covers twenty reporting years from 1999 to 2018 and shows sufficient observations across the period, supporting the use of sequential financial information.

The financial-variable histograms show strong positive skewness in many raw monetary variables. This is consistent with large differences in firm size, revenue, assets and liabilities. Several ratio and composition variables are comparatively stable because they are normalized by financial totals. The analysis supports the use of feature engineering and outlier treatment before model training.




Fig. 7. Boxplots of financial and engineered variables, first group.




Fig. 8. Boxplots of financial and engineered variables, second group.




Fig. 9. Boxplots of financial and engineered variables, third group.




Fig. 10. Boxplots of financial and engineered variables, fourth group.

The boxplots show observations beyond the upper and lower whiskers for several variables. These outliers are attributed to natural differences among firms rather than assumed data-entry errors. The project therefore retains the observations and uses Winsorization to reduce their influence. Retaining them helps the model learn from both financially stable and financially distressed firms.

B. Correlation Analysis and Feature Selection




Fig. 11. Correlation heatmap before feature selection. The pre-selection correlation heatmap shows several highly correlated variables, indicating redundancy in the financial representation. Feature selection is therefore

necessary to reduce multicollinearity and allow the models to focus on more informative variables.




Fig. 12. Correlation heatmap after feature selection.

After removing highly correlated variables, the remaining feature set has lower levels of correlation while retaining important financial information. This reduces redundancy and improves the stability and efficiency of the predictive modelling stage.

C. HFRSE Performance

TABLE IV. PERFORMANCE OF THE PROPOSED HFRSE FRAMEWORK

Metric	HFRSE value
Accuracy	0.721215 (72.12%)
Precision	0.118839 (11.88%)
Recall	0.552866 (55.28%)
F1-score	0.195628 (19.56%)
ROC-AUC	0.701495 (70.15%)

On the unknown test dataset, the HFRSE framework achieved 72.12% Accuracy, 11.88% Precision, 55.28% Recall, 19.56% F1-score and 0.7015 ROC-AUC. The relatively low Precision reflects the difficulty of predicting the minority bankruptcy class, while the higher Recall is valuable because the study prioritizes the detection of financially distressed firms. The ROC-AUC indicates satisfactory discrimination between healthy and bankrupt companies.

D. Model-Level Confusion Matrices




Fig. 13. Confusion matrix of the standalone Temporal 1D-CNN. The Temporal 1D-CNN correctly classified 8,019 healthy companies and 456 bankrupt companies, while 3,998 healthy companies were classified as bankrupt and 329 bankrupt companies were classified as healthy.

Its high Recall makes it useful for detecting distressed firms, although the relatively large number of false positives reduces overall accuracy.




Fig. 14. Confusion matrix of the standalone CatBoost classifier. CatBoost correctly classified 9,324 healthy companies and 381 bankrupt companies, while 2,693 healthy companies were classified as bankrupt and 404 bankrupt companies were classified as healthy. CatBoost provides stronger overall classification performance, but its Recall is lower than that of the Temporal 1D-CNN.




Fig. 15. Confusion matrix of the proposed HFRSE stacking ensemble.

The HFRSE ensemble correctly classified 8,799 healthy companies and 434 bankrupt companies, while 3,218 healthy companies were classified as bankrupt and 351 bankrupt companies were classified as healthy. The ensemble reduces some of the weaknesses of the standalone models and provides a balanced combination of overall classification and bankruptcy detection.

TABLE V. COMPARATIVE PERFORMANCE OF THE EVALUATED MODELS

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
CatBoost (Standalone)	75.81%	12.39%	48.54%	19.75%	0.7017
Temporal 1D-CNN (Standalone)	66.20%	10.24%	58.09%	17.41%	0.6494
HFRSE Stacking Ensemble	72.12%	11.88%	55.29%	19.56%	0.7015

CatBoost records the highest Accuracy, Precision, F1-score and ROC-AUC among the three evaluated models, whereas the Temporal 1D-CNN records the highest Recall. HFRSE combines their complementary characteristics: it retains a high Recall relative to CatBoost while improving on the standalone CNN's

overall Accuracy and providing an integrated explainable framework.

E. ROC and Precision-Recall Analysis




Fig. 16. ROC curve of the proposed HFRSE framework. The HFRSE ROC curve gives an ROC-AUC of 0.7015. This indicates that the framework has a useful ability to distinguish healthy firms from bankrupt firms across classification thresholds.




Fig. 17. Precision-Recall curve of the proposed HFRSE framework.

The Precision-Recall curve highlights the trade-off between identifying more bankrupt firms and maintaining precision. The project emphasizes Recall because bankruptcy is the minority class and missed bankruptcies can be more costly for financial risk assessment.

F. SHAP Explainability Analysis



Fig. 18. SHAP feature-importance bar plot.

The SHAP feature-importance analysis identifies Market Value at t-2, Retained Earnings at t-2, Net Income to Revenue, Debt Ratio, Return on Assets and the Autoencoder Anomaly Score among the most influential variables. This indicates that both conventional financial health indicators and learned abnormality information contribute to the final prediction.




Fig. 19. SHAP summary plot showing feature importance and directional influence.

The SHAP summary plot shows that feature values influence bankruptcy risk in different directions and with different magnitudes. The result demonstrates the practical advantage of the explainability layer: instead of receiving only a bankruptcy label, a decision-maker can examine which financial conditions contributed to the prediction.

V. DISCUSSION

The results show that corporate bankruptcy prediction is strongly affected by class imbalance. The model's Accuracy of 72.12% should therefore not be interpreted independently of its Recall and F1-score. The 55.29% Recall indicates that the framework identifies more than half of the bankrupt observations in the test set, which is important for an early-warning application. At the same time, the 11.88% Precision indicates that a substantial number of healthy firms are also flagged, so the model should be used as a risk-screening tool rather than as an automatic final decision mechanism.

The comparison of the individual learners explains why stacking is useful. The Temporal 1D-CNN is strong at recognizing temporal deterioration and produces the highest Recall, but it also creates more false positives. CatBoost is stronger at modelling nonlinear relationships in structured financial variables and produces the highest Accuracy, Precision, F1-score and ROC-AUC among the compared models, but it misses more bankrupt firms than the CNN. HFRSE combines the probability outputs of both models and therefore provides a more balanced operating point.

The feature-engineering layer also has an important role. Traditional ratios preserve familiar financial concepts such as liquidity, profitability, efficiency and leverage, while composition features normalize financial-statement items relative to assets, liabilities or revenue. The Autoencoder adds a different perspective by measuring deviation from normal financial behaviour. The combination gives the models more information than any single feature family.

From a management perspective, the framework can support early risk identification, investment screening, lending-risk assessment and financial review. SHAP is especially important because managers and financial analysts often need to understand the reason behind a

high-risk classification before taking action. The project therefore treats explainability as a decision support layer rather than as an optional visual output.

The main limitation evident from the results is the low Precision of the minority-class prediction. The framework therefore should not be interpreted as a substitute for financial due diligence, audit procedures or professional credit assessment. Further work can investigate richer external information, alternative temporal architectures, more extensive calibration, cost-sensitive learning, and additional financial or macroeconomic variables while preserving the explainability objective.

A. Implementation Environment

The project was implemented using Windows 10 or later and Python. The software stack includes pandas for data loading and preprocessing, NumPy for numerical operations, Matplotlib and Seaborn for visualisation, scikit-learn for preprocessing and evaluation, TensorFlow/Keras or PyTorch for deep learning, CatBoost for structured classification, XGBoost for the stacking meta-learner, imbalanced-learn for Borderline-SMOTE, SHAP for explainability, and OpenPyXL for spreadsheet handling. The reported recommended hardware includes an Intel Core i5-class processor, at least 16 GB RAM, an SSD with at least 25 GB free space, and a 64-bit operating system.

TABLE VI. IMPLEMENTATION MODULES OF THE HFRSE FRAMEWORK

Module	Main task	Output
Data Preparation	Validate and preprocess bankruptcy data.	Cleaned financial dataset
Feature Engineering	Create ratios and composition features.	Hybrid feature dataset
Exploratory Data Analysis	Study distributions and class structure.	EDA reports and visualisations
Company-Based Data Split	Separate firms to reduce data leakage.	Training and testing datasets
Feature Selection	Remove redundant/highly correlated variables.	Selected feature dataset
Anomaly Detection	Learn normal behaviour and generate anomaly scores.	Enhanced dataset
Temporal Learning	Learn three-year financial sequences with 1D-CNN.	CNN probabilities
CatBoost Model	Model nonlinear structured relationships.	CatBoost probabilities
Stacking Ensemble	Combine base-model probabilities with XGBoost.	Final bankruptcy prediction
Performance Evaluation	Measure Accuracy, Precision, Recall, F1, ROC-AUC and confusion matrix.	Performance results
Explainability Analysis	Identify influential features using SHAP.	Feature-importance and explanation plots

VI. CONCLUSION

This study presents a Hybrid Financial Representation Stacking Ensemble for explainable corporate bankruptcy prediction. The framework integrates original financial variables with traditional financial ratios and structural composition features, then adds an Autoencoder anomaly score, a Temporal 1D-CNN, CatBoost, an XGBoost stacking meta-learner and SHAP explanations. The approach is designed to address the

main difficulties identified in the project: complex financial relationships, temporal deterioration, highly imbalanced bankruptcy data, redundant variables and the need for interpretable predictions.

Using 78,682 firm-year observations from the 1999–2018 period, the proposed HFRSE framework achieved 72.12% Accuracy, 11.88% Precision, 55.29% Recall, 19.56% F1-score and 0.7015 ROC-AUC. The Temporal 1D-CNN achieved the highest Recall, while CatBoost achieved the highest Accuracy, Precision, F1-score and ROC-AUC. The stacking framework combined these strengths and produced a balanced result. SHAP further showed that Market Value, Retained Earnings, Net Income to Revenue, Debt Ratio, ROA and the Autoencoder Anomaly Score are among the most influential factors.

The results suggest that HFRSE can serve as an early-warning and financial-risk assessment framework, particularly where decision-makers need both a prediction and an explanation. Future extensions can incorporate additional economic and business characteristics, richer disclosure or textual information, improved temporal architectures, calibrated probabilities and cost-sensitive approaches to further improve minority-class precision and practical deployment.

REFERENCES

- Ataş, P. K., Tak, N., Özögür-Akyüz, S., & Erdogan, B. E. (2025). An innovative approach to ensemble learning in bankruptcy prediction using support vector machines and meta fuzzy functions. *Information Sciences*, 719, Article 122450.
- Ansari, A., Ahmad, I. S., Bakar, A. A., & Yaakub, M. R. (2020). A hybrid metaheuristic method in training artificial neural network for bankruptcy prediction. *IEEE Access*, 8, 176640–176655.
- Chen, X., Liu, J., & Wu, C. (2025). Multi-class financial distress prediction based on hybrid feature selection and improved stacking ensemble model. *Expert Systems With Applications*, 282, Article 127832.
- Gabrielli, G., Melioli, A., & Bertini, F. (2026). Corporate financial distress prediction: A machine learning approach in the era of big data. *Journal of Accounting & Organizational Change*, 22(7), 31–65.
- Gaurav, A., Gupta, B. B., Bansal, S., & Psannis, K. E. (2025). Bankruptcy forecasting in enterprises and its security using hybrid deep learning models. *Cyber Security and Applications*, 3, Article 100070.
- Hussain, A. B., Jingchun, S., Hussain, B., Shaique, M., & Rehman, A. U. (2026). Cross-market bankruptcy prediction: An interpretable ensemble learning framework using SHAP analysis. *Computational Economics*.
- Kadkhoda, S. T., & Amiri, B. (2024). A hybrid network analysis and machine learning model for enhanced financial distress prediction. *IEEE Access*, 12, 52759–52777.
- Lakshmi, T. M., Martin, A., & Venkatesan, V. P. (2016). A genetic bankrupt ratio analysis tool using a genetic algorithm to identify influencing financial ratios. *IEEE Transactions on Evolutionary Computation*, 20(1), 38–51.
- Lombardo, G., Bertogalli, A., Consoli, S., & Reforgiato Recupero, D. (2024). Natural language processing and deep learning for bankruptcy prediction: An end-to-end architecture. *IEEE Access*, 12, 151075–151091.
- Magrini, A. (2025). Bankruptcy risk prediction: A new approach based on compositional analysis of financial statements. *Big Data Research*, 41, Article 100537.
- 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, 107–148.
- Ogachi, D., Ndege, R., Gaturu, P., & Zoltan, Z. (2020). Corporate bankruptcy prediction model, a special focus on listed companies in Kenya. *Journal of Risk and Financial Management*, 13(3), Article 47.
- Park, M. S., Son, H., Hyun, C., & Hwang, H. J. (2021). Explainability of machine learning models for bankruptcy prediction. *IEEE Access*, 9, 124887–124899.
- Shetty, S., Musa, M., & Brédart, X. (2022). Bankruptcy prediction using machine learning techniques. *Journal of Risk and Financial Management*, 15(1), Article 35.
- Song, Y., & Peng, Y. (2019). A MCDM-based evaluation approach for imbalanced classification methods in financial risk prediction. *IEEE Access*, 7, 84897–84908.
- Thota, N., Puli, S., Subrahmanyam, A. C. V., & Yarala, S. (2025). The influence of public listing on bankruptcy prediction in India: An AI-ML approach. *Borsa Istanbul Review*, 25, 1463–1475.
- Islam, M. T., Maua, J., Islam, R., Ahmed, M., Mridha, M. F., & Hossen, M. J. (2025). A dual-encoder approach to bankruptcy prediction with financial ratios and corporate disclosures. *IEEE Open Journal of the Computer Society*, 6, 1896–1910.
- Xiao, Y., & Liu, X. (2026). Dynamic graph convolutional network with temporal aggregation for corporate bankruptcy prediction. *IEEE Access*, 14, 14985–14995.
- Yang, S., Li, C., Li, Z., Li, T., & Wang, Z. (2025). 4C-FinNet: A multi-channel feature map and spatio-temporal feature fusion method for financial risk prediction. *Expert Systems With Applications*, 284, Article 127789.
- Yoon, B., Jeong, Y., & Kim, S. (2020). Detecting a risk signal in stock investment through opinion mining and graph-based semi-supervised learning. *IEEE Access*, 8, 161943–161957.