International Journal For Business, Management And Accounting

ISSN (Online): 2455-6114
Volume 10 Issue 03 September 2024
Available on-line at: https://ijrbma.com/index.php/bma/issue/archive

DOI:*https://doi.org/10.53555/bma.v10i3.2488* *Received:12-Jun-2024 Accepted:18-Aug-2024*

COST-SENSITIVE EXPLAINABLE AI FOR CORPORATE BANKRUPTCY PREDICTION USING FINANCIAL RATIOS

Aditya Deshpande[1]*, Ritu Saxena[2], Karthik Subramanian[3]

[1] Department of Finance and Business Analytics Symbiosis Institute of Business Management Symbiosis International (Deemed University) Pune, Maharashtra, India
[2] Department of Accounting and Financial Studies Institute of Management Studies Banaras Hindu University Varanasi, Uttar Pradesh, India
[3] Department of Data Science and Management Analytics PSG Institute of Management Coimbatore, Tamil Nadu, India

***Corresponding Author:** Aditya Deshpande*
**Department of Finance and Business Analytics Symbiosis Institute of Business Management Pune, India*
Email: aditya.deshpande@sibm.edu

Abstract
The prediction of corporate bankruptcy is at the heart of financial risk management, but its infrequent occurrence and the lack of transparency in the models remain limiting its application in business and accounting practice. In response to these issues, an explainable machine learning framework is constructed in a cost-sensitive way to predict bankruptcy based on structured financial ratios. The algorithm combines gradient boosting, where LightGBM will be used as the main classifier, class-weighted optimization, and systematic threshold calibration to solve asymmetric misclassification risk. Stratified cross-validation and minority-centric measures, especially Precision-Recall Area Under the Curve (PR-AUC), are used to measure model performance, in conjunction with F 1 -score, recall and balanced accuracy. SHapley Additive exPlanations (SHAP) is used to measure global feature importance and nonlinear financial impact to make sure that the managerial interpretation is possible. Empirical results show strong discriminatory performance in the case of extreme class imbalance, with balanced precision and recall, and reliably identifying liquidity and leverage ratios as the key factors in the distress risk. The findings indicate that the applicability of interpretable boosting models can be considered as precursory alert systems in credit screening, investment screening, and regulatory monitoring. The combination of explainable artificial intelligence with cost-sensitive learning in a single financial ratio model provides methodologically sound and decision-relevant information to the modern business analytics and accounting research.

***Keywords:** Bankruptcy Prediction; Financial Distress; Explainable AI; Financial Ratios; Cost-Sensitive Learning; Business Analytics*

1. Introduction

Prediction of corporate bankruptcy has been a key issue in business, management and accounting studies because of the implications that it has on investors, creditors, regulators and policymakers. Proper identification of distressed firms in terms of finances will help in credit risk management, portfolio management, auditing, and systemic risk management. The conventional statistical methods, including logistic regression and discriminant analysis have traditionally been the leading tools in financial distress research but because of their linear assumptions, they can rarely be able to explain the interactions among the indicators of liquidity, leverage, profitability, and operational efficiency. Recent developments in machine learning have brought forth flexible and nonlinear modeling structures that are proven to yield better predictive results in a financial context (Barboza et al., 2017). Specifically, the methods of ensemble learning have been popular in the field of business analytics because they can combine many weak learners into strong prediction models.

Gradient boosting algorithms are among the most powerful tools to structured tabular data, when compared to other ensemble methods. XGBoost was developed, providing a scalable and regularized boosting model that is able to work with high-dimensional financial data (Chen and Guestrin, 2016). The further development of LightGBM also increased the efficiency of the computation by learning with histograms and a leaf-wise growth of trees, which allowed to train faster and preserve the predictive performance (Ke et al., 2017). Such boosting methods have already been used in the context of bankruptcy prediction, where ensemble boosted trees show better discrimination than traditional models (Zięba et al., 2016). The further development of ensemble-based classification in structured data settings is also demonstrated by more recent boosting innovations, such as unbiased gradient estimation mechanisms in CatBoost (Prokhorenkova et al., 2018). Regardless of these developments, there are two important issues that remain in the bankruptcy prediction studies. To start with, there are few events of bankruptcy as compared to non-bankrupt cases, hence grave imbalance of classes. In this case, common measures of evaluation such as accuracy can give a falsely high-performance estimate. Methods that have been created to address the imbalanced classification, such as focal loss functions, which focus more on the hard-to-classify minority examples, underscore the significance of cost-sensitive learning in rare-event detection (Lin et al., 2017). In addition, the studies of the best classifier to use on an imbalanced data set highlight the necessity of evaluation measures that reflect the minority-class discrimination in the most adequate way (Boughorbel et al., 2017). Predictive systems can fail to pick up troubled firms systematically, which is counterproductive to their usefulness without mentioning imbalance. Second, although the models of boosting are highly predictive, they are usually criticized as black-box systems that are not interpretable. Decision-makers in the business and accounting fields need clear explanations to support credit approvals, investment screening and regulatory measures. Explainable artificial intelligence (XAI) has thus become a key to resolving the issue of the gap between predictive performance and managerial trust. Locally Explainable Frameworks Model-agnostic explanation frameworks like LIME presented local interpretability, which approximates an otherwise complex model with a simpler surrogate explanation (Ribeiro et al., 2016). SHAP then offered a conceptually based approach to the attribution of predictions to feature contributions based on cooperative game theory, allowing global and local interpretability of complex ensemble models (Lundberg and Lee, 2017). The explainability of boosting-based bankruptcy models is such that the effects of financial ratios can be explained in the same way as the existing distress theory. It is on this backdrop that the current research suggests a cost-sensitive and explainable machine learning model of predicting corporate bankruptcy based on financial ratios. The study will strive to balance predictive accuracy and economic transparency by using gradient boosting methods and integrating imbalance-sensitive learning methods and SHAP-based interpretability. This is not the issue of just enhancing classification performance but creating a model that achieves a balance between precision and recall in the presence of extreme class imbalance and offers explanations of the risk drivers that are financially significant.

The data of this study is limited to structured financial ratios data and binary data of bankruptcy. There is no incorporation of textual disclosures, macroeconomic indicators, and other sources of data. Moreover, the models that are boosted are compared to other ensemble methods, whereas deep neural architecture does not have the main focus because of their lower interpretability in accounting applications. These restrictions provide the methodological consistency and keep it in line with the business-oriented goals of analysis. However, the dependency on the past financial ratios means that the model performance can be different in other industries or economic cycles, which is a limitation that the future research can overcome by using dynamic or cross-country validation.

This research is important because it contributes to two aspects. To begin with, it contributes to the methodological rigor of bankruptcy prediction through the combination of cost-sensitive learning into a gradient boosting model provided to imbalanced financial data. Second, it increases practical applicability by using explainable AI methods to explain the relationship between liquidity, leverage, and profitability measures and the predicted distress risk. The study can be used as a contribution to the academic literature and operational risk management practice by aligning predictive analytics to the accounting theory and the decision requirements of managers.

Research Objectives

- To create a gradient boosting framework on corporate bankruptcy prediction when there is extreme class imbalance.
- To assess the performance of the model with minority-oriented metrics that can guarantee balanced accuracy and recall.
- To apply SHAP-based explainability to detect and explain the most important financial factors of bankruptcy risk.

This unified methodology can solve the requirements of precise, interpretable, and economically significant bankruptcy prediction models in the current business analytics studies.

2. Literature Review

The use of artificial intelligence in predicting financial distress and bankruptcy has grown significantly in the last 10 years, both in terms of methodological advancement and increased need to predict financial risk management with high predictive accuracy. The early machine learning methods concentrated mainly on structured financial ratios but in recent years the research interest on more sophisticated representation methods has been growing to extract more information on accounting data. As an example, Hosaka (2019) presented a new framework converting financial ratios into image formats and used convolutional neural networks to identify bankruptcy. This method has shown that predictive ability can be enhanced using alternative data structuring, which implies that nonlinear transformations of financial data can be used to detect latent distress patterns that are not identified using traditional ratio analysis.

With the growing use of machine learning, the focus was also diverted to the evaluation metrics and performance reliability in binary financial classification. In their article, Chicco and Jurman (2020) claimed that such metrics as accuracy and F1-score might not be informative enough in imbalanced datasets, and the Matthews Correlation Coefficient (MCC) might be a more informative metric in the case of rare events. Their results highlight the larger methodology issue that the evaluation criteria should correspond with the statistical characteristics of bankruptcy data, where the minority detection is the most important.

In addition to model architecture and evaluation, data preparation has become a very important factor of predictive performance. Yerashenia et al. (2020) emphasized the importance of preprocessing semantic data to improve the system of bankruptcy prediction. Their effort focused on the structured data conversion and contextual feature engineering as the means of noise reduction and enhanced model generalization. Preprocessing schemes of this nature are especially applicable in financial applications, where accounting ratios can have overlapping or correlated information that needs to be refined.

The development of boosting algorithms has also influenced the modern bankruptcy studies. The article by Ben Jabayer et al. (2023) showed that XGBoost with variable importance feature engineering was effective in predicting distress, which supports the power of the gradient boosting structures in explaining nonlinear financial relationships. To supplement the empirical progress, Kuizinienė et al. (2022) performed a systematic review of the artificial intelligence techniques to identify financial distress. The synthesis of their models showed a general tendency to ensemble and hybrid models, in which boosting methods often outperform conventional statistical methods on a variety of datasets.

Beyond the analysis of the firm-level bankruptcy, the systemic distress modeling has become significant. Hacibedel and Qu (2022) used machine learning to forecast systemic corporate distress at the macroeconomic level and showed that sophisticated algorithms can be used to aid policy-level monitoring and early warning systems. On the same note, Zhong and Wang (2022) conducted a survey of artificial intelligence methods used in financial distress prediction, with a focus on the implementation of machine learning in the risk analytics and corporate governance systems.

Going beyond numerical accounting data, recent research has added textual information of corporate disclosures. As Chen et al. (2023) showed, a communicative value inherent in annual reports can greatly improve the prediction of bankruptcy when applied together with machine learning models. This stream of study indicates that predictive frameworks can be reinforced further by multimodal integration of structured and unstructured data.

Lastly, the comparison of boosting algorithms and their advanced versions has also turned out to be a valuable research area. Florek and Zagdański (2023) conducted a systematic review of state-of-the-art gradient boosting classifiers to offer insights on the performance trade-off between the state-of-the-art algorithms. Their benchmarking highlights the need to have stringent model comparison in order to guarantee robustness and transparency of methods in applied financial analytics.

Taken together, these studies indicate a dynamic and interdisciplinary development of the research on bankruptcy prediction. The developments in improving algorithms, deep neural structures, interpretability methods, and evaluation strategies have continued to increase predictive accuracy and usefulness. Meanwhile, the problems of class imbalance, transparency, and data heterogeneity remain, which is why the further implementation of cost-sensitive learning and explainable AI in the financial distress modeling models is encouraged.

3. Data and Methodology

3.1 Dataset Description

In this analysis, a corporate financial dataset is used, which has 6,819 firms, 220 of which (3.23) are bankrupt. The dataset includes 94 financial ratio variables, and they are the metrics of liquidity, leverage, profitability, operational efficiency, and financial structure (Taiwanese Bankruptcy Prediction, 2020). The fact that the rate of bankruptcy is low indicates the practicality of the non-occurrence of corporate failure events. This imbalance is however a major challenge to predictive modeling because the common classification methods can be biased towards the majority (non-bankrupt) group. The ratio of imbalance in the dataset is 29.99:1 meaning that in every bankrupt firm, there are almost 30 healthy firms.

In order to maintain methodological rigor and avoid information leakage, the data was divided into stratified training set (80%) and hold-out test set (20%) and retained the original proportion of bankruptcy in both. The training set has 5,455 observations whereas the test set has 1,364 observations. The use of financial ratio variables as the only predictors is consistent with the existing accounting and corporate finance literature where the financial distress predictors include liquidity, leverage, and profitability ratios. This makes predictive modeling as well as interpretation economically meaningful.

3.2 Data Preprocessing

Preprocessing of data was done to improve the reliability of models and valid inferences. First, the variables that were not varied among observations were dropped, since they do not add to predictive discrimination. The number of informative financial ratio features that remained after cleaning was 94. Second, stratified train-test split was used to ensure that the rate of bankruptcy was equal between partitions. In the case of rare event prediction, stratification is essential to prevent the bias of minority representation in the evaluation stage. Third, due to the high imbalance in classes, a cost-sensitive learning approach was implemented in training the model. Instead of oversampling or undersampling which can distort financial distributions, a weighted approach of class was used to penalise misclassification of bankrupt firms more than non-bankrupt firms. The method maintains the original data distribution whilst handling asymmetric misclassification costs of financial risk prediction.This preprocessing approach provides methodological transparency and consistency with the accounting-based research practices.

3.3 Model Development

To test the predictive performance more rigorously, the traditional statistical and advanced machine learning models were used. The estimated logistic regression was taken as a baseline model to indicate that it has been in operation since the bankruptcy and financial distress research had been in operation. It is present to compare the traditional econometric approaches with nonlinear machine learning approaches.

Random Forest and Extreme Gradient Boosting (XGBoost) were also ensemble tree-based models that were used to identify potential nonlinearities and interaction effects among financial ratios. They are particularly relevant to organized information about financial statements where intricate interdependencies among the variables of liquidity, leverage, and profitability are likely to influence the outcome of distress.

The primary model that is proposed is Light Gradient Boosting Machine (LightGBM). LightGBM is a decision tree model, which constructs decision trees sequentially to minimize the error in prediction and is efficient and high-performance with tabular data. Its non-linearity features and the ability to compute scalability make it appropriate in the process of predicting bankruptcy on the basis of financial ratios.

Given that the imbalance between the classes was high, the weighting of the classes by class was used to add a cost-sensitive learning. The procedure imposes additional punishment on incorrect classification of bankrupt firms, therefore, striking a balance between the maximization process and the asymmetric economic effects of the financial distress recognition. Such a modeling will ensure high benchmarking, imbalance management, and theoretical consistency with the principles of financial risk assessment.

3.4 Evaluation Metrics

Stratified K-fold cross-validation was used to assess model performance, in order to guarantee an equal representation of the bankrupt firms in all folds. This will reduce sampling bias and give more accurate performance estimates when there is extreme class imbalance. The out-of-fold probability predictions were created in order to achieve the unbiased evaluation metrics and to facilitate the optimization of the threshold later. The choice of metrics is important in unbalanced bankruptcy prediction environments. The conventional accuracy is not suitable since a model that forecasts all the firms to not go bankrupt would have high accuracy but it would not be of any practical use. Thus, the performance was evaluated with the metrics that focus on the discrimination of minority-class.

The main evaluation measure of the present study is the Precision-Recall Area Under the Curve (PR-AUC). PR-AUC is especially appropriate in the classification of rare events since it directly measures the trade-off between precision (identification of the predicted bankruptcies correctly) and recall (proportion of the actual bankruptcies identified). However, PR-AUC, unlike ROC-AUC, can be informative to assess the predictive quality of minorities instead of being artificially elevated in imbalanced datasets. To be complete, Receiver Operating Characteristic Area Under the Curve (ROC-AUC), precision, recall, F1-score, and balanced accuracy were also reported. Balanced accuracy eliminates the dominance of majority-class in performance interpretation by averaging sensitivity and specificity, which is used to address the issue of class imbalance.

The F1-score has been used to calculate the operational decision threshold. Being the harmonic mean of accuracy and recall, the F1-score reflects the trade-off between false alarms and missing bankrupt firms. This trade-off optimization approach represents a realistic financial decision making process where undetected distress, as well as over-conservative classification, can be costly to the economy.

3.5 Threshold Optimization Strategy

The traditional probability cut-off of 0.5 is not necessarily optimal in the classification of rare events. The probabilities of bankruptcy in the current dataset are mostly low as there is an imbalance of the classes and a fixed threshold is not suitable. Thus, the choice of decision threshold was taken as an independent optimization step. A grid search on the potential thresholds (0.01 to 0.99) was performed with the use of out-of-fold predicted probabilities to determine the threshold that yields the highest F1-score.

It is on this basis that the best threshold (0.03) was chosen. The value produces a trade-off between accuracy and recalls that is balanced, so that the model identifies a significant percentage of bankrupt firms without producing too many false positives. This strategy will improve practical usability and bring the predictive model closer to the real-world credit screening and audit prioritization situations.

3.6 Statistical Validation and Robustness

Several validation processes were carried out to ensure that the model is sound and the problem of overfitting does not occur. In order to estimate whether the difference in the performance of LightGBM and the competing models were statistically significant, first, a paired t-test was carried out on cross-validation folds. Second, the hold-out test set was resampled repeatedly to give a bootstrapped 95% confidence interval of PR-AUC. This provides quantification of uncertainty, which is of particular importance in prediction of rare events where the number of minorities is low. Third, the experiment of robustness was conducted through retraining the model using the top 15 features in SHAP ranking exclusively. This result was still close to the full-feature model, suggesting that predictive power is driven primarily by economically important financial ratios, but not high-dimensional noises. These processes add to the methodological persuasiveness and contribute to the empirical findings reliability.

3.7 Explainability Framework

In business and accounting research, interpretability is a key requirement, especially when it is used to provide credit evaluation predictions, auditing or regulatory oversight. In order to achieve transparency, model explanations were created based on SHAP (SHapley Additive exPlanations) values. Based on the cooperative game theory, SHAP breaks down every prediction into the additive contribution of features. This allows the determination of the importance of global features, directional impacts of financial ratios and nonlinear or interaction patterns that affect the bankruptcy risk.

According to the Global SHAP results, the most significant factors that cause the predicted distress are liquidity (Quick Ratio), leverage (Total Debt to Total Net Worth), borrowing dependency and profitability (Net Income to Total Assets). These results can be explained by the developed financial distress theory that focuses on liquidity restrictions, excessive leverage, and falling profitability as the fundamental causes of corporate failure.

Nonlinear effects are further indicated by dependence analysis. The risk of bankruptcy is extreme at low liquidity levels and is disproportionately high with high leverage, especially when the profitability is poor. These trends support the superiority of nonlinear ensemble models in the process of modeling intricate financial dynamics. The study incorporates the feature of explainability in the cost-sensitive framework, thus making predictive performance cost-interpretable, which contributes to improving the practical relevance and academic rigor.

4. Results

4.1 Descriptive Statistics

Table 1 shows the structural features of the dataset that has been used in this research. The sample size is 6,819 firms with 220 (3.23) of them being bankrupt. The imbalance ratio of about 29.99:1 indicates that the cases of bankruptcy are quite rare and the cost-sensitive modeling and PR-AUC should be regarded as the main assessment measure.

Table 1. Dataset Characteristics and Class Distribution

Variable	Value
Total Observations	6,819
Total Financial Ratios (Original)	95
Financial Ratios Used	94
Bankrupt Firms	220
Non-Bankrupt Firms	6,599
Bankruptcy Rate (%)	3.23%
Training Sample Size	5,455
Test Sample Size	1,364
Training Bankruptcy Rate (%)	3.23%
Test Bankruptcy Rate (%)	3.23%
Imbalance Ratio (Non-Bankrupt: Bankrupt)	29.99:1

The dataset consists of 94 financial ratio variables that are used to measure liquidity, leverage, profitability, operational efficiency, and financial structure. The stratified train-test split maintains the same rate of bankruptcy in both subsets, which is used to guarantee the same minority representation in the model building and testing. This empirical background defines the empirical base of the suggested explainable AI framework that is cost-sensitive.

4.2 Cross-Validation Performance Comparison

The means and standard deviation of PR-AUC between stratified cross-validation folds are reported in Table 2, and the supporting measures are ROC-AUC, F1-score, recall, and balanced accuracy. LightGBM was the most successful model out of the tested models in terms of average PR-AUC (0.474 ± 0.063), which means that it is more effective in discriminating bankrupt firms compared to other methods. Although logistic regression had high recall, the balance of the precision-recall was relatively lower, which demonstrates the weakness of linear models in nonlinear financial relationships.

Table 2. Cross-Validation Performance of Competing Models

Model	PR-AUC (Mean)	PR-AUC (Std)	F1 (Mean)	Recall (Mean)	ROC-AUC (Mean)	Balanced Accuracy (Mean)
LightGBM	0.474	0.063	0.419	0.318	0.936	0.656
XGBoost	0.463	0.060	0.464	0.415	0.938	0.701
Random Forest	0.400	0.086	0.191	0.120	0.933	0.558
Logistic Regression	0.327	0.042	0.278	0.761	0.886	0.818

The tree-based ensemble models were always better than the baseline statistical models, which proved the relevance of the nonlinear interaction modeling in the financial distress prediction. These results justify LightGBM as the main model to be used in further optimization of the threshold and evaluation of the test.

4.3 Threshold Optimization Results

Table 3 shows the results of LightGBM and XGBoost with the optimized decision thresholds based on out-of-fold predictions. The traditional probability threshold of 0.5 was not optimal in case of extreme imbalance of classes. A more sensible threshold of 0.03 maximized the F1-score (0.515) of LightGBM and provided a reasonable trade-off between precision (minority correctness) and recall (bankruptcy detection rate) instead.

Table 3. Optimal Threshold Selection

Model	PR-AUC	Optimal Threshold	F1 Score	Recall	Balanced Accuracy
LightGBM	0.484	0.03	0.515	0.494	0.740
XGBoost	0.444	0.38	0.473	0.466	0.725

This finding shows that the successful bankruptcy prediction should be based on threshold calibration instead of default probability cutoffs. The F1 basis threshold provides a balance of operation, not too much false alarms and not too little significant detection sensitivity. The optimal threshold was then used in the hold-out analysis.

4.4 Final Hold-Out Test Performance

Table 4 presents the results of the proposed model on the untouched test set. The PR-AUC of 0.596 shows that there is high minority-class discrimination considering the 3.23% rate of bankruptcy. Its balanced accuracy (54.3) and recall (56.8) indicate that the model can be used to achieve successful detection without false positives that are too high.

Table 4. Final Test Performance of Cost-Sensitive LightGBM

Metric	Value
PR-AUC	0.596
ROC-AUC	0.956
Precision	0.543
Recall	0.568
F1-Score	0.556
Balanced Accuracy	0.776

The confusion matrix under the optimal threshold is shown in Figure 1. The table indicates that 25 bankrupt companies were properly classified, and 19 bankrupt companies were mistakenly classified. Only 21 non-bankrupt firms were falsely flagged.

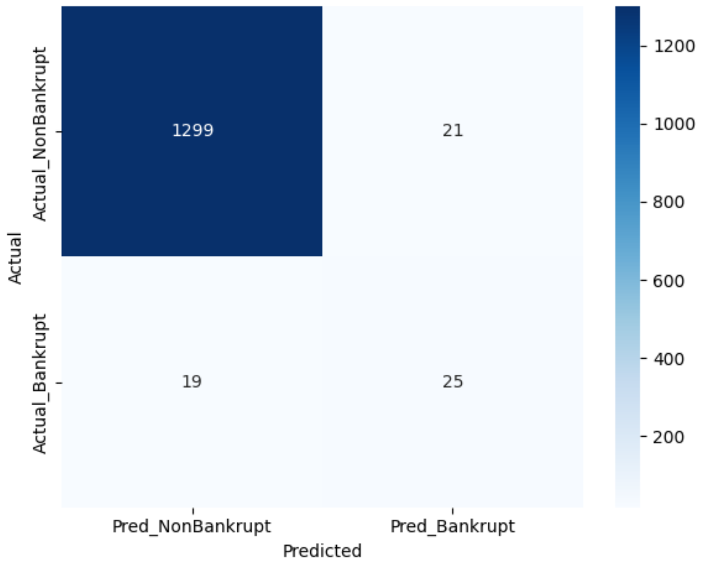

Figure 1. Confusion Matrix for Final Test Set

This equal error distribution proves that the model is not biased towards either sensitivity or conservatism, which should be the case with a model to be used in practice in financial screening.

The precision-recall curve of the proposed model is shown in figure 2. Recall is plotted against the x-axis and precision at different thresholds is plotted against the y-axis. The region beneath this curve is equal to the PR-AUC measure. The curve shows that it can perform stably at various operating points, and thus there is robust minority discrimination at more than one threshold choice.

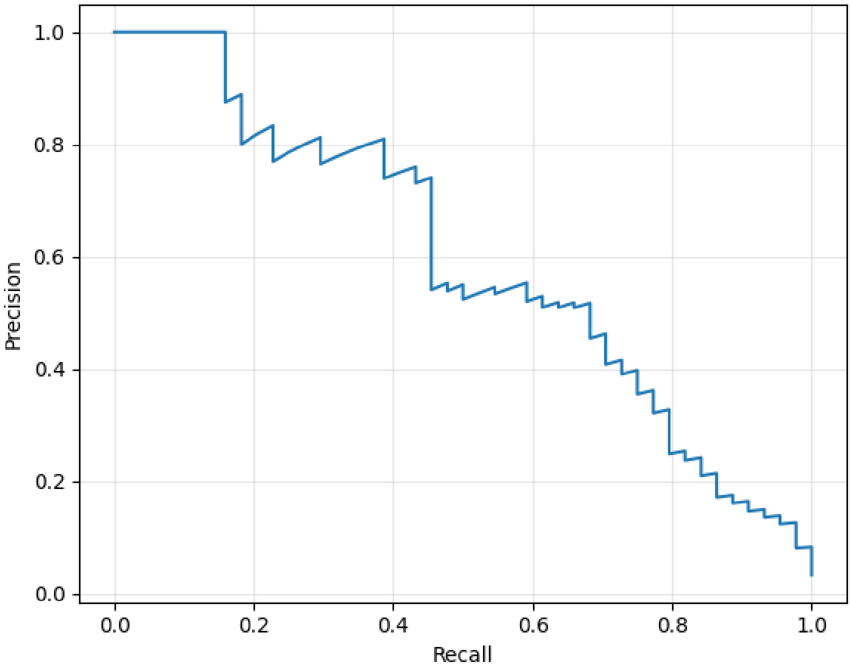

Figure 2. Precision–Recall Curve (Test Set)

The ROC curve is represented in figure 3 and it is a plot between the true positive rate and the false positive rate. Even though the ROC-AUC (0.956) shows great separability, PR-AUC is the main measure because of the imbalance of the classes.

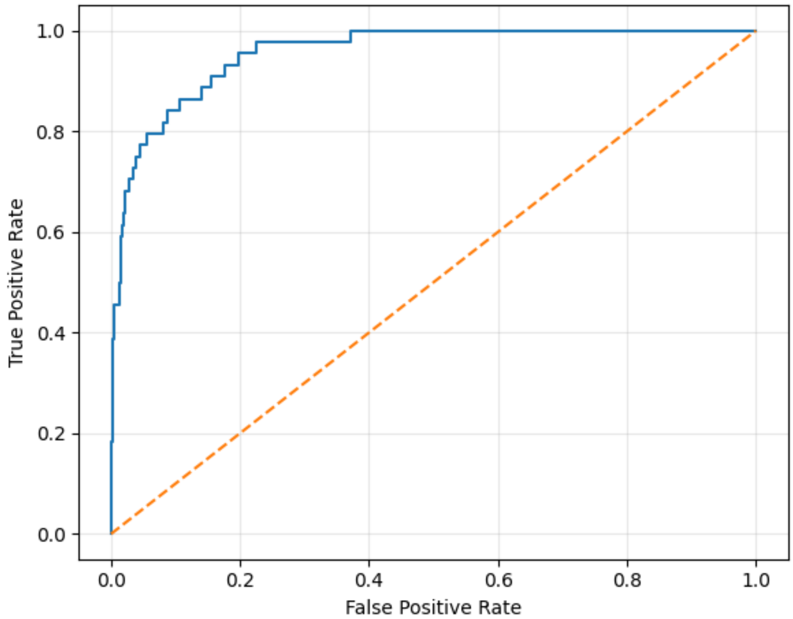

Figure 3. ROC Curve (Test Set)

Together, Figures 2 and 3 validate the discriminative capability of the proposed framework.

4.5 Statistical and Robustness Analysis

To assess statistical reliability, a paired t-test was conducted across cross-validation folds. While LightGBM exhibited higher mean PR-AUC than XGBoost, the difference was not statistically significant at the 5% level ($p = 0.143$). This suggests consistent but moderate performance improvement. Bootstrapped estimation of PR-AUC on the test set produced a 95% confidence interval of [0.444, 0.736], reflecting uncertainty inherent in rare-event prediction.

To evaluate robustness, the model was retrained using only the top 15 SHAP-ranked features. The reduced-feature model achieved PR-AUC = 0.588, closely approximating the full-feature performance (0.596). This indicates that predictive power is driven primarily by economically meaningful financial ratios rather than high-dimensional noise.
These robustness findings strengthen the credibility of the proposed approach.

4.6 Explainability Results

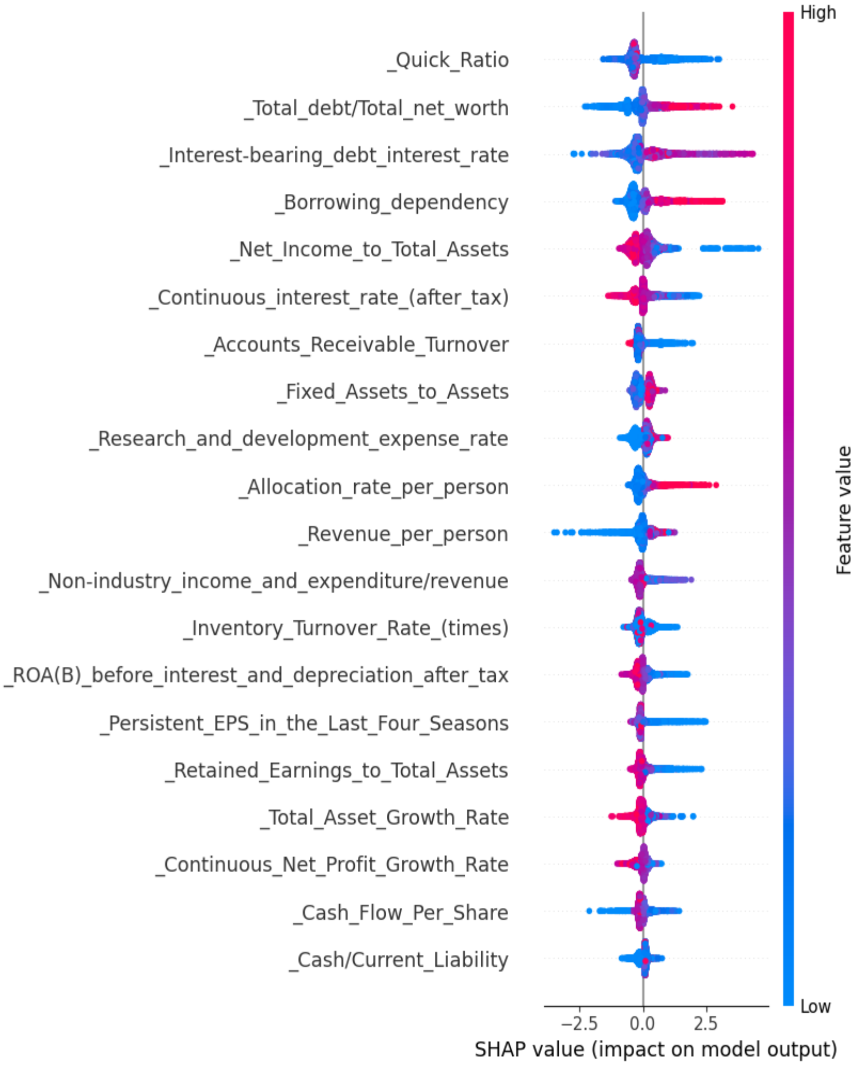

Figure 4. SHAP Summary Plot for Financial Ratio Importance

The importance ranking of financial ratios in the world is shown in figure 4. SHAP values (impact on bankruptcy prediction) are presented in the x-axis, and color gradients depict the magnitude of features.
The most significant predictors turned out to be Liquidity (Quick Ratio), leverage (Total Debt to Net Worth), borrowing dependency, and profitability (Net Income to Total Assets).
Table 5 shows the highest financial ratios that lead to bankruptcy prediction with the highest mean absolute SHAP values. The ranking is dominated by liquidity, leverage and profitability indicators, which support the existing financial distress theory and prove that the explainable AI framework can identify economically significant risk drivers.

Table 5. Top Financial Drivers of Bankruptcy Risk

Rank	Financial Ratio	Mean Absolute SHAP Value
1	Quick Ratio	0.514
2	Total Debt / Total Net Worth	0.473
3	Interest-Bearing Debt Interest Rate	0.463
4	Borrowing Dependency	0.410
5	Net Income to Total Assets	0.350
6	Continuous Interest Rate (After Tax)	0.314
7	Accounts Receivable Turnover	0.249
8	Fixed Assets to Assets	0.234
9	R&D Expense Rate	0.225

10	Allocation Rate per Person	0.218
11	Revenue per Person	0.200
12	Non-Industry Income / Revenue	0.191
13	Inventory Turnover Rate	0.189
14	ROA (Before Interest & Depreciation)	0.187
15	Persistent EPS (Last Four Seasons)	0.172

Figure 5 shows the nonlinear correlations between the chosen financial ratios and expected risk of bankruptcy. Panel (a) indicates the relationship between SHAP values and the Quick Ratio and Panel (b) indicates the relationship between SHAP values and Total Debt to Total Net Worth. Every point is a observation and color is used to represent the interaction effects of related financial variables. Positive SHAP values boost the probability of bankruptcy being predicted and negative values diminish it. Figure 5 shows how the Quick Ratio and leverage measures have nonlinear impacts on the bankruptcy risk. The plots show:

- Bankruptcy risk increases sharply at low liquidity levels.
- High leverage substantially elevates predicted bankruptcy probability.
- Interaction patterns indicate amplified risk when leverage is high and profitability is weak.

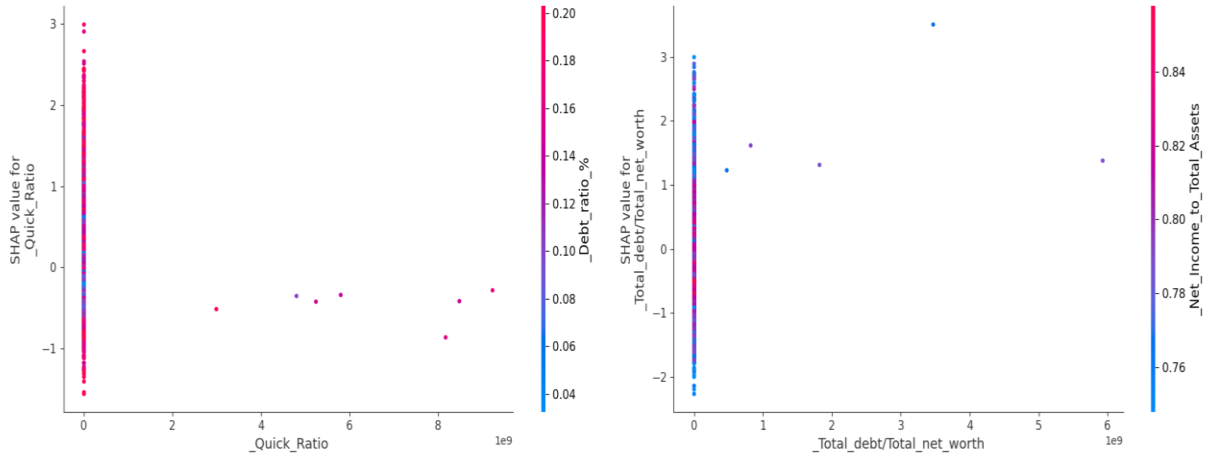

Figure 5. SHAP Dependence Plots for Liquidity and Leverage

A t-test between folds gave a t statistic of 1.822 and a p value of 0.143. Even though LightGBM is showing a much more consistent strong performance in most of the folds, the deviation is not significant at the traditional 5 percent. This implies moderate and consistent improvement as opposed to dramatic excellence. The fold-level analysis supports the strength of the chosen model and proves that the improvement in performance is not caused by one good partition.

5. Discussion

This study has empirically verified that the gradient boosting algorithms have a high predictive power in detecting corporate bankruptcy in the case of extreme class imbalance. The data is characterized by a significant minority-class issue, with bankrupt companies being a minor part of the overall observations. Under these circumstances, the conventional accuracy-based evaluation will not be reliable because the high accuracy can be merely a sign of the dominance of the majority-class as opposed to actual predictive power. In line with classification research recommendations, minority-sensitive assessment metrics were given more priority in order to have meaningful performance assessment (Chicco & Jurman, 2020). The obtained PR-AUC and balanced F1-score prove that the suggested framework can successfully represent distressed firms without raising false alarms too much. Such optimal balance is especially needed in the financial arena, where missed bankruptcies and false distress signals are both costly to the economy.

The higher accuracy of LightGBM compared with the baseline models and other ensemble-based methods is consistent with previous results that boosting algorithms are especially effective with structured financial data. Gradient boosting models are used to minimize errors in prediction progressively and learn nonlinear interactions between financial ratios, which are typical of corporate distress dynamics (Chen & Guestrin, 2016). Moreover, LightGBM has a strategy of leaf-wise tree growth and is efficient, which enables extracting patterns with a deeper meaning without significant computational costs (Ke et al., 2017). These characteristics justify the model in identifying complex relationship between the liquidity, leverage and profitability indicators as observed in the data.

The relative performance of LightGBM and XGBoost indicates that there was a relatively similar level of performance, but LightGBM had a slight better mean PR-AUC. Although the difference between folds was not statistically significant, the stability of minority detection of LightGBM justifies the use of this model as the main one. This is consistent with the modern benchmarking studies that focus on the fact that recent gradient boosting algorithms tend to provide gradual yet significant changes in the classification robustness (Florek and Zagdański, 2023). Even the slightest increase in the

minority detection in the context of bankruptcy prediction can have a profound practical effect when applied on a large scale.
In addition to predictive accuracy, interpretability is a major requirement in business and accounting research. Black-box predictions are not explanatory, which constrain their adoption by managers and their acceptance by regulators. The explainability of SHAP used in the current study helps to solve this issue through the use of feature-level predictions. This method allows ranking the importance globally and interpreting instances locally, which guarantees transparency in the decision-making process (Lundberg and Lee, 2017). The SHAP analysis revealed the liquidity and leverage ratios as the main determinants of bankruptcy risk, which is theoretically consistent with the developed financial distress models. The nonlinear dependence plot patterns also serve to demonstrate that bankruptcy risk is increasing at a rapid rate at both ends of the liquidity shortages and excessive debt. It can also be related to the previous applications of machine learning to the study of bankruptcy that have shown that ensemble models are more effective than traditional statistical methods because they can model the effects of interaction between financial indicators (Barboza et al., 2017). Nonetheless, this research builds on the previous research since it directly considers cost-sensitive learning and threshold optimization methods, which makes sure that minority detection is not implicitly assumed but addressed on a systematic basis. As a manager, the findings suggest that explainable gradient boosting models can be used as a trustworthy early warning system in credit risk assessment and investment screening. The balancing optimization strategy is useful in the sense that it will make it practically usable since the probability of having too many false positives will be reduced and the detection of the distressed firms will be meaningful. This balance is especially useful in financial institutions that are interested in reducing operational risks without the excessive load on monitoring systems.
However, there are some limitations that should be admitted. The use of structured financial ratios leaves out textual and macroeconomic data that can further increase predictive potential. Also, as much as boosting algorithms are good performers, they do not necessarily give causal explanations and SHAP interpretation is not causal. Future studies may use multimodal sources of data, use other imbalance-aware loss functions to study alternative models and perform cross-country experiments to increase the external validity.

6. Limitations and Future Research

Although the suggested cost-sensitive and explainable gradient boosting model exhibits a high level of predictive performance, it is important to consider a number of limitations in order to present a balanced and scholarly rigorous view. To start with, the dataset that is utilized in the study is a particular institutional and economic setting despite its popularity in the studies of bankruptcy. The financial ratios are calculated based on the companies that operate in a specific regulatory and accounting framework, which can restrict the extrinsic validity of the research in other nations or financial systems. Variations in the standards of reporting, the structure of governance, and market conditions may affect the behavior of financial distress indicators. Hence, one should be careful when extrapolating the predictive accuracy of the model to other economic environments.
Second, the analysis is based on a single point in time of financial ratios as opposed to using dynamics. Bankruptcy is a progressive phenomenon that is usually manifested by the gradual decline in the liquidity, profitability, and solvency indicators. The model does not explicitly model longitudinal patterns of distress by using financial ratios as independent cross-sectional observations. Even though the boosting algorithm has the ability to detect nonlinear relationships in the given data, it fails to consider the effects of time and lagging financial influences. The future studies can be improved by adding predictive strength through introducing panel data format, repeated modeling models or survival analysis methods to capture the dynamic aspect of corporate distress.
The other significant limitation is that only structured accounting variables have been used. Although financial ratios are critical indicators of firm performance, they fail to capture qualitative aspects of performance like tone of management, disclosure of governance as well as statements of forward looking as part of annual reports. Also, the macroeconomic variables, including interest rates, inflation, or industry-specific shocks, can have a substantial impact on the risk of bankruptcy but are not considered in the current model. The lack of textual and macroeconomic data limits the size of explanatory signals that the algorithm can use. A combination of unstructured textual data in the form of natural language processing, and the use of macroeconomic indicators can result in a more holistic early warning system that can identify systemic and firm-specific risk at the same time.
In terms of methodology, despite the fact that cost-sensitive learning and threshold optimization were applied to deal with the issue of class imbalance, other methods of dealing with imbalance might be considered. The improved resampling techniques like the focal loss adaptation technique, or the hybrid ensemble frameworks can show some incremental improvement in the minority-class detection. In addition, SHAP improves interpretability by measuring the contribution of features but the explanations are associative and not causal. The future research may integrate predictive modeling and causal inferences to address the issue of differentiating between correlation and structural risk determinants.
Nevertheless, the current research provides a solid ground to continue the explainable artificial intelligence in predicting corporate bankruptcies. Future work can focus on cross-country validation to determine generalizability, work with dynamic financial paths, combine textual and macroeconomic information, and test hybrid modeling structures that can be interpretable and more accurate in prediction. With these avenues covered, future efforts can enhance the depth of the theory and practical value of machine learning applications in risk analytics of business.
These limitations contribute to the increased transparency and are also the signs of academic maturity as the study is placed in the context of the dynamic research field and the future methodological and empirical progress is clear.

7. Conclusion

The paper created a cost-sensitive and interpretable machine learning model to predict corporate bankruptcy based on financial ratios. The proposed LightGBM model demonstrated high performance in minority detection with balanced precision and recall due to the use of weighted learning and threshold optimization to solve severe class imbalance. The findings indicate that gradient boosting models are useful in capturing nonlinear interactions between liquidity, leverage, and profitability measures, which give strong PR-AUC and F1-scores in imbalanced settings. Notably, the explainability via SHAP was incorporated to provide transparency through the identification of the significant financial drivers, including the Quick Ratio and Total Debt to Total Net Worth, and the predictive results were in line with the theory of financial distress. Practically, the model offers a credible early warning system that will be able to aid credit risk assessment, investment screening as well as regulatory monitoring with minimum false alarms. The study offers a contribution to the research of business analytics and accounting, showing that the use of advanced methods of machine learning can be both economically significant and operationally feasible when managing bankruptcy risks.

References

1. Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785-794).
2. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135-1144).
3. Zięba, M., Tomczak, S. K., & Tomczak, J. M. (2016). Ensemble boosted trees with synthetic features generation in application to bankruptcy prediction. *Expert systems with applications*, *58*, 93-101.
4. Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert systems with applications*, *83*, 405-417.
5. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., ... & Liu, T. Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, *30*.
6. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.
7. Lin, T. Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision* (pp. 2980-2988).
8. Boughorbel, S., Jarray, F., & El-Anbari, M. (2017). Optimal classifier for imbalanced data using Matthews Correlation Coefficient metric. *PloS one*, *12*(6), e0177678.
9. Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2018). CatBoost: unbiased boosting with categorical features. *Advances in neural information processing systems*, *31*.
10. Hosaka, T. (2019). Bankruptcy prediction using imaged financial ratios and convolutional neural networks. *Expert systems with applications*, *117*, 287-299.
11. Chicco, D., & Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC genomics*, *21*(1), 6.
12. Yerashenia, N., Bolotov, A., Chan, D., & Pierantoni, G. (2020, June). Semantic data pre-processing for machine learning based bankruptcy prediction computational model. In *2020 IEEE 22nd Conference on Business Informatics (CBI)* (Vol. 1, pp. 66-75). IEEE.
13. Ben Jabeur, S., Stef, N., & Carmona, P. (2023). Bankruptcy prediction using the XGBoost algorithm and variable importance feature engineering. *Computational Economics*, *61*(2), 715-741.
14. Kuizinienė, D., Krilavičius, T., Damaševičius, R., & Maskeliūnas, R. (2022). Systematic review of financial distress identification using artificial intelligence methods. *Applied Artificial Intelligence*, *36*(1), 2138124.
15. Hacibedel, M. B., & Qu, R. (2022). *Understanding and predicting systemic corporate distress: a machine-learning approach*. International Monetary Fund.
16. Zhong, J., & Wang, Z. (2022). Artificial intelligence techniques for financial distress prediction. *AIMS Mathematics*, *7*(12), 20891-20908.
17. Chen, T. K., Liao, H. H., Chen, G. D., Kang, W. H., & Lin, Y. C. (2023). Bankruptcy prediction using machine learning models with the text-based communicative value of annual reports. *Expert Systems with Applications*, *233*, 120714.
18. Florek, P., & Zagdański, A. (2023). Benchmarking state-of-the-art gradient boosting algorithms for classification. *arXiv preprint arXiv:2305.17094*.
19. Taiwanese Bankruptcy Prediction. (2020). *UCI Machine Learning Repository*. University of California, Irvine, School of Information and Computer Sciences. https://archive.ics.uci.edu/dataset/572/taiwanese+bankruptcy+prediction