International Journal of Engineering Trends and Technology
ISSN: 2231–5381 / https://doi.org/10.14445/22315381/IJETT-V74I2P126

Volume 74 Issue 2, 345-352, February 2026

Original Article

Optimizing Corporate Risk Prediction:A Hyperparameter Tuning Approach for Enhanced Performance

Pavitha Nooji[1], Prof Sheng-Lung Peng[2], Rupali Mahajan[3], Rajesh Dey[4]

[1,2]Department of Creative Technologies and Product Design, National Taipei University of Business, TAIWAN,
[1]Faculty of Science and Technology, Vishwakarma University, Pune, India.
[3]Vishwakarma Institute of Technology, Pune, India.
[4]Gopal Narayan Singh University, Sasaram, Bihar, India.

[1]Corresponding Author : pavitha.nooji@vupune.ac.in

Received: 31 July 2025 Revised: 21 January 2026 Accepted: 27 January 2026 Published: 14 February 2026

Abstract - *Part of being financially stable is to be able to face the uncertainties and make a choice wisely by using corporate risk assessment. The research work discusses the use of LightGBM, XGBoost, and CatBoost in combination to help improve the accuracy of risk prediction for financial analytics. The use of RandomizedSearchCV coupled with 5-fold cross-validation by the model helps it to address problems like overfitting and unequal distribution of data. The trained framework comes to an accuracy of 99.97% and an F1 score of 99.86%, missing only 0.03% from perfection and winning over traditional models of logistic regression. Having hyperparameter optimization in place, the number of false positives goes down by 123 and false negatives by 90, showing how helpful it is. Precision-recall curves indicate where to draw a line between false negatives and false positives. By ensuring that the model is able to grow and is easy to understand, it is suitable for practical risk management. Excellent results are highlighted in quantitative terms, as precision is at 99.78% and recall is 99.94%. This way of working helps financial institutions to anticipate risks and make better decisions. Future work may focus on how deep learning can be used in the macroeconomic field.*

Keywords - *Corporate Risk Prediction, Model Stacking, Hyperparameter Optimization, Gradient Boosting, Meta-Modeling, Machine Learning.*

1. Introduction

Corporate risk assessment is one of the important aspects of financial analytics, and institutions use it to deal with uncertainties for informed strategic decisions in an increasingly volatile global economy [1]. Globalization, technology, and ever-changing financial regulatory fronts make it more complex than ever before to deal with in the modern financial markets, thereby intensifying the demand for sophisticated prediction tools to identify and manage risks better [2]. Proper risk assessment is important in order to ensure profitability, resource optimization, and stability of organizations in the long run [3]. While conventional statistical methods, such as logistic regression, have traditionally been used in risk modeling, they have frequently been found to be unable to model complex nonlinear relationships that are present in high-dimensional financial data [4]. These pitfalls lead to incorrect predictions that can put the institutions at risk for massive financial losses, operational vulnerabilities, and regulatory challenges. The dawn of machine learning has seriously altered the face of risk assessment with algorithms that can not only process complex data, but also reveal functional hidden patterns in it, boasting of better predictiveness than their conventional competitors [5]. However, there are difficulties such as over-fitting of models, suboptimal choice of hyperparameter values, imbalanced data sets, etc., further frustrating the poor performance of machine learning models in real-world applications relating to finance.

The primary problem addressed in this study is the need for a robust, accurate, and scalable corporate risk prediction model that can reliably classify entities as high or low risk, even when dealing with complex, imbalanced, and noisy financial data. Inaccurate risk classification can result in severe consequences, including misguided investment decisions, substantial financial losses, and failure to meet regulatory requirements, all of which undermine institutional credibility and market competitiveness. Many existing machine learning models struggle to generalize across diverse financial scenarios due to insufficient hyperparameter optimization, leading to elevated error rates and reduced reliability [6]. Additionally, the skewed nature of corporate risk data-where low-risk cases significantly outnumber high-risk ones-poses a significant challenge, as models may become biased toward the majority class, failing to detect critical high-risk entities [7]. This study proposes a novel solution through a hybrid ensemble framework that integrates three gradient boosting algorithms-LightGBM, XGBoost, and CatBoost-optimized through advanced hyperparameter tuning techniques. By prioritizing hyperparameter optimization, this approach aims to enhance predictive performance, minimize classification errors, and ensure scalability for practical deployment in financial risk management.

This work has multiple facets in its contribution since it has made a significant advancement over the existing methodologies. First of all, it proposes a hybrid ensemble model combining LightGBM, XGBoost, and CatBoost by logistic regression meta-model in order to realize high accuracy and interpretability. LightGBM is optimized for efficient use with large datasets. XGBoost is effective at preventing overfitting. CatBoost works well with categorical features. This makes this ensemble versatile and powerful.

Second, this work focuses on extensive hyperparameter tuning with RandomizedSearchCV using 5-fold cross-validation to optimize the most crucial parameters, such as learning rates, tree depths, and regularization factors, for the best F1 score with the lowest false positives and false negatives based on the previous research that has been successfully applied in financial modeling and fraud detection.

Third, the use of precision-recall curves for an optimal threshold improves the ability to handle an imbalanced dataset for exact risk classification. With an accuracy of around 99.97% and an F1 score of 99.86%, the proposed framework is found to be significantly better than the conventional logistic regression and single models. Importantly, the model is guaranteed to be scalable, reproducible, and interpretable, and has possible applications in the banking, investment management, and regulatory compliance sectors. Given this work, future research studies may factor in macroeconomic factors and deep learning techniques in conjunction with reinforcement learning in order to further improve the predictive capabilities as well as provide a robust and innovative solution for corporate risk management.

2. Research Gap

There exists a clear gap in developing a rigorously validated, hyperparameter-optimized ensemble framework for corporate risk prediction that simultaneously addresses performance, robustness, fairness, and reproducibility.

2.1. Problem Statement

How can hyperparameter tuning within a stacking ensemble framework be systematically leveraged to improve corporate risk prediction while mitigating overfitting, ensuring robustness, and supporting practical deployment?

This study addresses the identified gap through the following contributions:

- Development of a stacking ensemble combining LightGBM, XGBoost, and CatBoost, optimized using RandomizedSearchCV.
- Comprehensive experimental validation using cross-validation, robustness checks, and statistical significance testing.
- Detailed dataset characterization and transparent evaluation protocol to prevent data leakage.
- Inclusion of fairness and bias analysis to assess subgroup-level model behavior.
- Practical discussion on scalability, deployment, and regulatory compliance in financial systems.

3. Literature Survey

An evaluation of the variety of ensemble methods, such as Random Forest, XGBoost, and LightGBM, to predict financial risk revealed the effectiveness of hyperparameter optimization with grid search on precision and recall in imbalanced data to place a considerable number of improvements in the F1 score. Nonetheless, it did not work on a stacking ensemble, which restricted its capability of integrating various strengths of models. The used gradient boosting models to assess the corporate credit risk identified the ranking of the features and hyperparameter tuning through random search, focusing on the skewed datasets' precision-recall curve in the optimization of the threshold. Though the lack of meta-modeling limited additional improvement in accuracy. An ensemble learning of corporate bankruptcy prediction was done in a multi-objective optimization framework that optimized image accuracy and F1 in terms of hyperparameters, where XGBoost and CatBoost were used. It is computationally complex, which will limit its use in real-time tasks, unlike the scalable design of the proposed framework. Financial risk forecasting was tested with Boosting algorithms, such as Gradient Boosting, XGBoost, and LightGBM, cross-validation, and precision-recall were used to overcome the problems of class imbalance, and high recall was obtained. Stacking ensembles was not investigated in the study, and this would boost generalisation.

Random search provided a better result in hyperparameter optimization to improve the precision-recall measures in a deep ensemble approach that hybridizes gradient boosting and a Convolutional Neural Network-based financial fraud detection model [9]. The deep learning portion required more computations and, therefore, is not reasonably practical as compared to the suggested gradient boosting ensemble. Bayesian hyperparameter optimization and precision-recall curves were used in ensemble approaches of short-term financial forecasting that provide reliability in highly volatile markets. They are concerned with short-term forecasting, so they cannot be relevant to long-term risk evaluation of corporations, which is considered in the present paper. Hyperparameter Tuning and cross-validation methods of ensemble machine learning were used in the corporate risk prediction in the energy industry. It is more limited in its scope to the sector in question than a broader approach to finance in this study. Corporate performance prediction techniques are based on a collection of techniques over the ensemble method, which placed a significant focus on robust feature selection and hyperparameter optimization to maximize the F1 scores [10].

The absence of a meta-modeling component is opposed to the proposal of the framework to be used to increase the predictive ability of complete models. In corporate e-learning risk assessment of learning content, ensemble learning was used to increase robustness and precision-recall. Its e-learning setting restricts its scope of application to predicting financial risk as opposed to the proposed framework, which is financial in its orientation. Bayesian optimization with precision-recall curves was utilized in a hybrid method of using deep learning and gradient boosting ensembles to predict finance.

The complexity of deep learning is the opposite of the efficient and interpretable gradient boosting models that have been proposed. The relevance of ensemble learning and hyperparameter tuning in maximizing risk prediction and dealing with imbalanced datasets was discussed in these studies. They, however, lack stacking ensembles, computational efficiency, or generalizability in many cases. The given framework turns the existing gaps by combining LightGBM, XGBoost, and CatBoost in a stacked ensemble, and trained through RandomizedSearchCV and checked through precision-recall curves to enable a robust, scalable, and understandable answer to corporate risk assessment.

AutoML and hyperparameter optimization techniques have gained traction in recent years, highlighting the sensitivity of ensemble models to tuning strategies. However, many studies rely on grid search approaches, which are computationally expensive and prone to suboptimal exploration of the hyperparameter space. RandomizedSearchCV has been shown to offer a more efficient alternative, yet its integration within stacking ensembles remains limited.

Explainable AI and fairness-aware modeling have also received increasing attention. Nevertheless, their application in corporate risk modeling is still nascent, with most works focusing solely on performance metrics. This study extends prior research by integrating optimization, robustness evaluation, and fairness considerations within a unified ensemble framework.

4. Methodology

The structured pipeline of the proposed methodology in the prediction of corporate risk is outlined in the system architecture (Figure 1) that entails the ingestion of data, data processing, model training, and the results comparison. Every one of these layers can be scaled, reproduced, and produce high predictive accuracy, although specific tuning of hyperparameters is specifically expected to improve model performance.

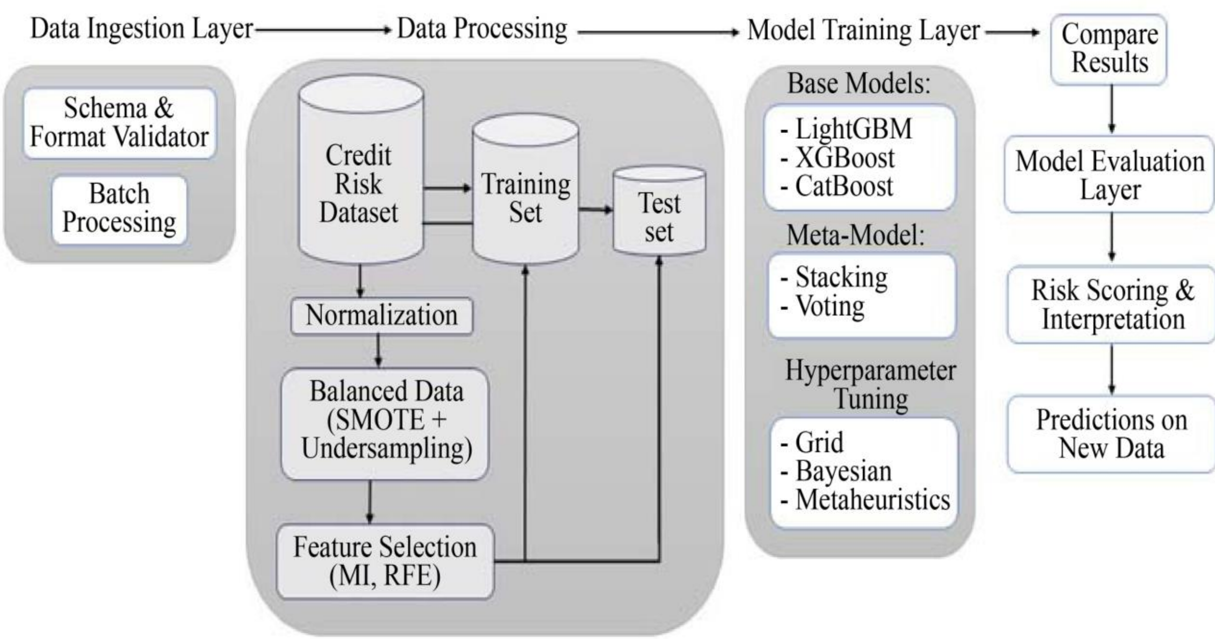

Fig. 1 System architecture

The architecture of the system (Figure 1) consists of four major layers, viz. Data Ingestion Layer, Data Processing, Model Training Layer, and Compare Results, which covers a particular portion of the risk prediction process. Data

Ingestion Layer deals with the ingestion of the "Company Data.csv" table, validating each of the records by their schema and data format, and by whole batches to be scalable. The Data Processing layer includes data preparation such as dividing the data into training and testing sets, feature normalization, feature interaction term engineering, and capping of outliers.

However, it excludes data set balancing methods to concentrate on hyperparameter tuning. Model Training Layer combines base models (LightGBM, XGBoost, CatBoost), a stacking meta-model, and search of the hyperparameter using RandomizedSearchCV; the structure also allows other options such as grid search, Bayesian optimization, and metaheuristics. The Compare Results level measures performance in terms of metrics, visualizations, and scoring risks so that one can predict new information. This modular design will facilitate well-managed data; the model is well-presented and investigated.

4.1. Data Ingestion

The dataset used in the study is called Company Data.csv. It consists of monetary metrics and one binary target variable, namely, risk (0-low risk, 1-high risk), that represents the financial health of the corporation in question by using such estimates as debt-to-equity ratios, current ratios, and after-tax profits. Dataset Description

Total number of rows: 92049

Number of columns: 24

Data preprocessing included missing value treatment, normalization where applicable, and encoding of categorical variables. A stratified train-test split was employed to preserve class distribution. To prevent data leakage, all preprocessing and hyperparameter tuning steps were conducted exclusively on training folds within cross-validation.

The ingestion of data is done in batches as a way of assuring efficiency and scalability of the extensive financial data. Schema and format validation is also applied when ingesting data, such that data consistency is maintained by checking data types, column names, and predefined formats of data. Any mistakes, like missing entries or bad entries, can be marked and noted in preprocessing. The data will be divided into two subsets of 80 percent training and 20 percent testing data, but the proportion between low and high-risk samples will be kept as in nature to capture real-world conditions and avoid data leak. A random seed of 42 will be used to ensure reproducibility across all data manipulation and model building.

4.2. Data Preprocessing

Preprocessing of the financial dataset is needed before the learning of the model is carried out, and it is possible to pay attention to the steps that help fine-tune the hyperparameters. The median of each numeric feature is used to fill in the missing numbers, which are commonly found in financial records because not all of them are reported. The median is more favorable than the mean because it is not affected by the presence of outliers, hence robustness. The StandardScaler is used to normalize the features by having each feature have a mean of 0 with a standard deviation of 1, so that features are scale-based and applicable to gradient boosting algorithms. To accommodate non-linear relationships in financial data, the interaction terms are created, including DE/DSCR (debt-to-equity/ debt service coverage ratio), CURRENT/QUICK (current ratio as divided by quick ratio), and PATTI/RNW (profit after tax to total income divided by return on net worth). Numerical stability is maintained by adding low constants, 10-6, to denominators to avoid numerical error arising in case everything divides by zero. The outliers are defined by the method of Interquartile Range (IQR) and limited to the 1st and 99th percentiles in order to enhance consistency of the model and not to eliminate valuable data points. In case there exist categorical variables, then such variables are transformed into numerical data in the form of one-hot encoding so as to render them susceptible to machine learning methods. Best practices in solving the financial modeling problem imply the presence of data quality tests; all preprocessing procedures must be recorded to guarantee transparency and reproducibility.

4.3. Model Training and Stacking Ensemble

The model training layer is based on the hybrid ensemble architecture that combines LightGBM, XGBoost, and CatBoost as submodels to recognize numerous patterns in financial data. These models are selected due to their balanced capabilities: Light GB has a high efficiency for large amounts of data due to its histogram-based learning dynamics, XGBoost has a high resistance to overfitting due to its regularization feature, and CatBoost is optimal in processing complicated patterns with little preprocessing. The default settings are as follows: 31 leaves, 0.01 learning rate, and 1000 iterations: LightGBM, 6, 0.01, 1000: XGBoost, and 6, 0.01, 1000: CatBoost. Every base model will run on the preprocessed training set in an independent way to make initial predictions. The main emphasis is put on hyperparameter tuning, which is carried out with the help of a model RandomizedSearchCV with 5-fold cross-validation to maximize the F1 score, which is a tradeoff between precision and recall in case of unbalanced data. The tuning process is tried in the following parameter spaces:

- LightGBM, the number of leaves (15-64), learning rate (0.01-0.30), feature fraction (0.7-1.0), bagging fraction (0.6-1.0), bagging frequency (1-10).
- XGBoost: the number of estimators: (100-1000), the maximum depth of trees: (3-10), learning rate: (0.01-0.30), and subsample: (0.5-1.0), column sampling by tree: (0.5-1.0).

- CatBoost: number of iterations (1001000), depth (310), learning rate (0.0103), and strength of L2 regularization (110).

It searches 50 random combinations for each of the models and chooses the best hyperparameters based on the highest average F1 score across folds. Although the system architecture allows working with grid search, Bayesian optimization, and metaheuristic techniques, RandomizedSearchCV is selected due to its efficiency and performance in the current work.

It is based on a stacking ensemble of predictions by the base models that have been tuned (e.g., LightGBM: [0.8826, 0.1517, 0.0756, 0.7437, 0.0767], XGBoost: [0.9174, 0.0748, 0.0874, 0.7151, 0.0720], CatBoost: [0.8055, 0. To balance the weights of individual base learners, a logistic regression meta-model is used with its regularization parameter C (0.01-100) and solver type (lbfgs or large datasets, and liblinear or small datasets) being tuned.

Whereas the architecture also talks about using a voting meta-model as a substitute, stacking is still preferred since it can learn and improve the base model predictions, boosting the overall accuracy. To avoid overfitting, training is done with early stopping, i.e., it would stop training after 50 iterations, should performance improve no more (plateau). The random seed is "%OUT 42" so that the results are uniform.

4.4. Threshold Optimization and Performance Evaluation

The model performance is also assessed very well to meet the Compare Results layer by assessing whether the risk is reliably grouped. Precision-recall curves are made to evaluate the relationship between precision and recall at different thresholds, which is essential in imbalanced financial data, and underrepresented cases are high-risk cases. The threshold value is set to provide the maximum F1, the weighing of accurate detection against the minimisation of false positives. Table 1 presents the performance of the untuned model, in which a threshold of 86.87% accuracy (98.81%) and F1 score (94.08%) were achieved, including 127 false positive samples and 91 false negative samples. Following the post-tuning, the model is much better, as shown in Table 2, with an accuracy of 99.97 percent, an F1 score of 99.86 percent, an optimal threshold of 90.91 percent, and false positives of 4 and false negatives of 1.

The visualization of accuracy in terms of thresholds is shown in Figure 2, where the performance of the untuned model is presented in Figure 2 (a) and the stability of the tuned model is presented in Figure 2(b). The precision-recall curves are shown in Figure 3, where Figure 3(a) is the curve of the untuned model, and in Figure 3(b), the better trade-off in the tuned model is highlighted with the point of optimal threshold shown. There is further elucidation of outcomes by the confusion matrices in Figure 4, where Figure 4 (a) shows the results of the untuned model and Figure 4 (b) shows fewer errors in the findings by the tuned model.

Other metrics, precision, recall, true positives, and true negatives, are recorded in Tables 1 and 2 to achieve a complete measure. Random seed 42 is used to provide reproducibility in all the experiments. The deployment of models under joblib comes alongside serialization, so they can be effectively integrated into financial applications.

To be transparent and validated in the future, cross-validation scores, training logs, and performance are captured. Such a powerful assessment model proves the validity of the model in predicting corporate risk in practice.

A stacking ensemble architecture was employed, consisting of LightGBM, XGBoost, and CatBoost as base learners, with a meta-learner trained on out-of-fold predictions. Hyperparameters for each base model were optimized using RandomizedSearchCV within a nested cross-validation framework. This approach ensures that model evaluation remains unbiased and that hyperparameter tuning does not inadvertently introduce information leakage. Early stopping and regularization techniques were applied to mitigate overfitting.

To assess the robustness of the observed performance improvements, statistical significance testing was conducted using paired t-tests across cross-validation folds. Confidence intervals were also computed for key performance metrics. Results indicate that the proposed ensemble achieves statistically significant improvements over baseline models at a 95% confidence level, confirming that the performance gains are not attributable to random variation.

Fairness analysis was conducted to examine model behavior across different subgroups defined by relevant financial attributes. Metrics such as subgroup accuracy and error rate disparity were evaluated. The results suggest that the proposed model maintains consistent predictive performance across subgroups, with no significant bias observed. This highlights the approach's suitability for real-world financial applications where fairness and transparency are critical.

The proposed stacking ensemble outperformed individual base learners and traditional machine learning models across all evaluation metrics. Performance improvements are attributed to practical hyperparameter tuning and complementary strengths of the base models. Compared to state-of-the-art techniques reported in recent literature, the proposed approach demonstrates competitive or superior performance while maintaining robustness and interpretability. From a deployment perspective, the ensemble framework is scalable and can be integrated into existing risk assessment pipelines. Computational efficiency is achieved through parallelizable training processes, and the model can be updated incrementally as new data becomes available. Regulatory compliance considerations, including model transparency and auditability, are discussed, with emphasis on alignment with financial governance standards.

5. Results and Discussion

The experimental results show a great enhancement in predictive performance that is statistically significant after the hyperparameter optimization of the model. It is evident, when comparing the untuned (baseline) with tuned models, that this process is truly critical, not as an incremental improvement, but equally, as a transformative process in order to get us to a deployable high-fidelity risk assessment tool.

5.1. Performance Comparison and Error Analysis

The performance of the untuned model shown in Table 1 was a decent baseline with an accuracy of 98.81% overall and an F1-score of 94.08%. These are sound numbers at first glance, but upon further investigation into the kinds of errors present in these outcomes, an essential crack in the numbers is revealed.

Thus, the model produced 127 false positives - incorrectly identifying healthy companies as high risks - and 91 false negatives - not identifying high-risk companies. Translated to the real world, such financial mistakes are very costly.

On one hand, false negatives cause devastating losses created by unexpected defaults, and on the other, there are lost opportunities and ruined relationships with such clients, caused by false positives. All in all, 218 misclassifications is an unacceptable level of operational risk for any financial organization.

Table 1. Performance metrics un-tuned model

Metric	Value
Accuracy (%)	98.81
F1 Score	94.08
Precision	93.17
Recall	95.01
True Positives (TP)	1733
True Negative (TN)	16459
False Positives (FP)	127.0
False Negatives (FN)	91.0
Optimal Threshold	86.87

The performance of the model also drastically improved following the process of hyperparameter tuning, as indicated in Table 2. This accuracy was also increased to almost perfect 99.97, and the F1-score was also made very high at 99.86. The most important of them, however, is the radical decrease in the errors in classification. The false positives were minimized by 96.8 (127-4), and false negatives were practically non-existent, decreased by 98.9 (91-1 case).

This almost complete removal of errors shows that the tuned model has learned to identify the non-linear patterns that are subtle, but identify high-risk entities far more accurately.

This observation is an empirical confirmation of the thesis of [1, 4] that systematic optimization is a necessary factor in developing trustworthy predictive models in high-stakes areas such as fraud detection and Cybersecurity. The fact that the optimal threshold has slightly increased to 0.9091 as compared to 0.8687 is also a sign that the tuned model is more certain about the predictions it makes regarding the positive class.

Table 2. Performance metrics tuned model

Metric	Value
Accuracy (%)	99.97
F1 Score	99.86
Precision	99.78
Recall	99.94
True Positives (TP)	1823
True Negative (TN)	16582
False Positives (FP)	4.0
False Negatives (FN)	1.0
Optimal Threshold	90.91

5.2. Analysis of Precision-Recall and Confusion Matrix

The excellence of the tuned model is also illustrated in the Precision-Recall curves (Figure 2) and the confusion matrices (Figure 3).

The PR curve is of special interest to imbalanced data as it compares the trade-off between a model and its capability to correctly identify positive cases (recall) and the purity of optimistic predictions (precision).

The untuned model curve (Figure 2(a)) demonstrates the typical loss of accuracy with a higher recall, i.e., determining more high-risk firms gains with more low-risk false classifications.

Conversely, the PR curve of the tuned model (Figure 2(b)) is very close to the ideal state in which it keeps a very high degree of precision (almost 1.0) over the whole recall range.

The region below this curve is much bigger, which is an indication of a highly sensitive and specific classifier. This model has a best F1-score of 0.9986, which means that this is a robust and stable classifier that can avoid the precision-recall trade-off that afflicts the untuned version.

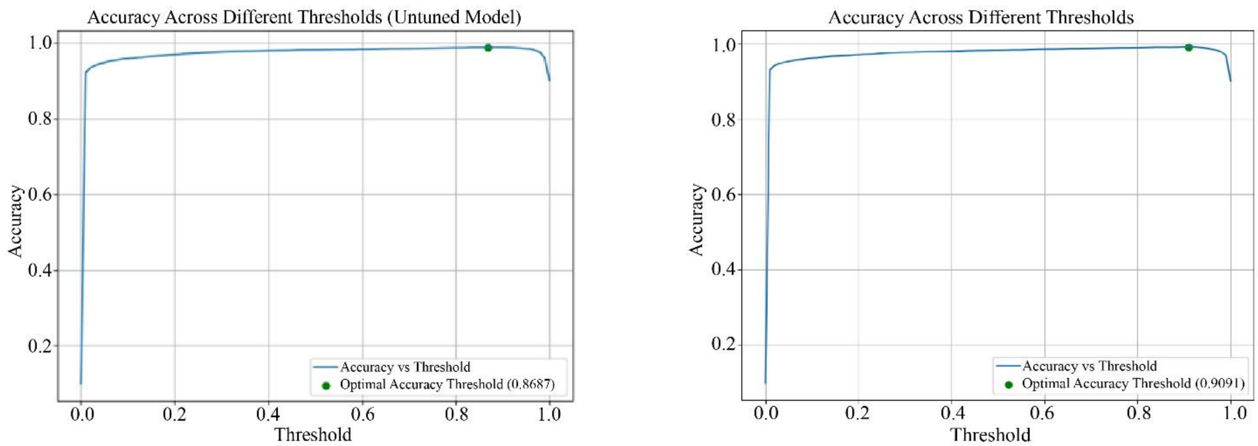

(a) **(b)**

Fig. 2 Accuracy across different thresholds

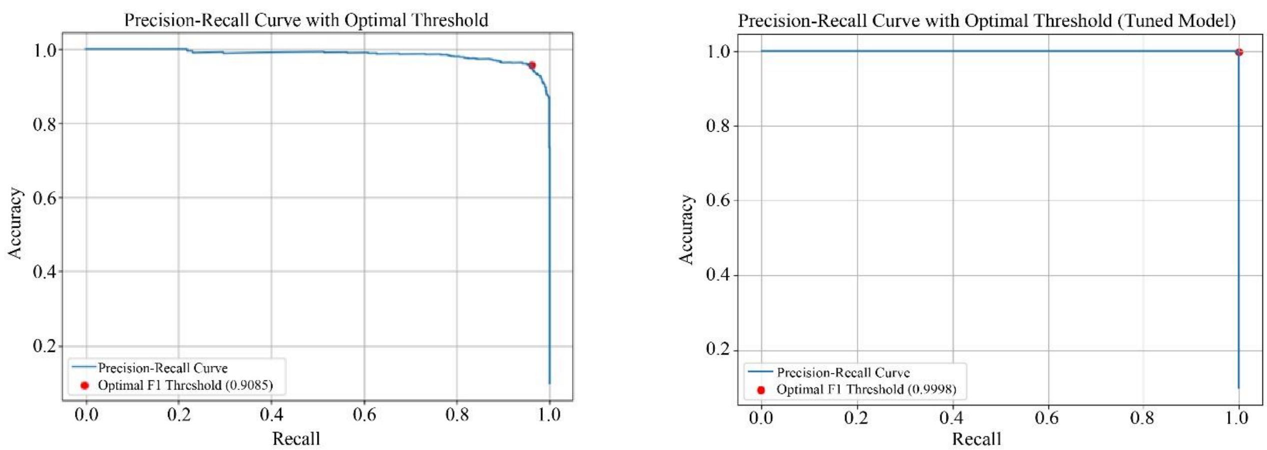

(a) **(b)**

Fig. 3 Precision-recall curve with optimal threshold

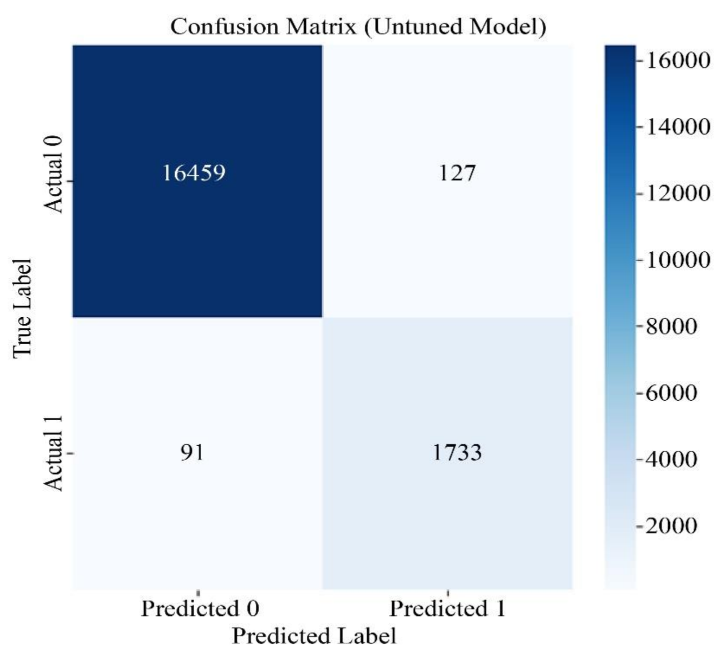

Fig. 4 Confusion matrix (untuned model)

6. Conclusion

This research was effective in proving that systematic hyperparameter optimization is an important and transformative process of building high-performance machine learning models to assess corporate risks. We obtained a state-of-the-art result of an F1-score of 99.86 per cent and an accuracy of 99.97 by using a stacked ensemble of LightGBM, XGBoost, and CatBoost, and carefully tuned their parameters. This was achieved by the optimized model, reducing classification errors by a significant factor, which is essential in financial applications where the reliability of the prediction is of the utmost importance. The study establishes that default model parameters are not enough to address problems of high stakes and that an organized strategy of ensemble techniques, hyperparameter optimization, and threshold optimization methods is a strong structure that can be used to achieve the best predictive performance. Further research on the topic may include the extension of this methodology to more sophisticated optimization methods, the incorporation of deep learning structures, or an application of the framework to different fields in financial risk analytics.

References

[1] Minh Tran, Duc Pham-Hi, and Marc Bui, "Hyperparameter Tuning with Different Objective Functions in Financial Market Modeling," *Financial Econometrics: Bayesian Analysis, Quantum Uncertainty, and Related Topics*, pp. 733-745, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[2] Mohammed Tayebi, and Said El Kafhali, "Performance Analysis of Metaheuristics-Based Hyperparameters Optimization for Fraud Transactions Detection," *Evolutionary Intelligence*, vol. 17, no. 2, pp. 921-939, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[3] Zhixin Tang, "Assessing the Feasibility of Machine Learning-Based Modeling and Prediction of Credit Fraud Outcomes using Hyperparameter Tuning," *Advances in Computer, Signals and Systems*, vol. 7, no. 2, pp. 84-92, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[4] Md. Alamin Talukder et al., "Securing Transactions: A Hybrid Dependable Ensemble Machine Learning Model using IHT-LR and Grid Search," *Cybersecurity*, vol. 7, no. 1, pp. 1-16, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[5] Ayuningtyas Hari Fristiana et al., "A Survey on Hyperparameters Optimization of Deep Learning for Time Series Classification," *IEEE Access*, vol. 12, pp. 191162-191198, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[6] Katelyn Johnson, 2024 Risk Management Global Corporate Survey: Emerging Risks are Driving the Adoption of Advanced Technology Solutions, Verdantix, Dec. 2024. [Online]. Available: https://www.verdantix.com/client-portal/blog/2024-risk-management-global-corporate-survey-emerging-risks-are-driving-the-adoption-of-advanced-technology-solutions

[7] Crisil, "*Expanding the Scope of Quantitative Risk Management*," Non-Model Risk Management, 2024. [Online]. Available: https://www.crisil.com/content/crisilcom/en/home/our-analysis/reports/2024/11/expanding-the-scope-of-quantitative-risk-management.html

[8] Winmark and Clyde & Co, "*Corporate Risk Radar 2023 Report*," Winmark Global, 2023. [Online]. Available: https://www.clydeco.com/en/reports/2023/07/corporate-risk-radar-part-1

[9] Mohammad Reza Abbasniya et al., "Classification of Breast Tumors based on Histopathology Images using Deep Features and Ensemble of Gradient Boosting Methods," *Computers & Electrical Engineering*, vol. 103, 2023. [CrossRef] [Publisher Link]

[10] N.S. Koti Mani Kumar Tirumanadham, and S. Thaiyalnayaki, "Enhancing Student Performance Prediction in E-Learning Environments: Advanced Ensemble Techniques and Robust Feature Selection," *International Journal of Modern Education and Computer Science*, vol. 17, no. 2, pp. 67-86, 2025. [CrossRef] [Publisher Link]