International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Gradient Boosting for Interpretable Risk Assessment in Finance: A Study on Feature Importance and Model Explainability

Anita Kori[1], Nagaraj Gadagin[2]

[1] Assistant Professor Dept. of AIML BEC, Bagalkot
[2] Assistant Professor Dept. of AIML CEC, Mangalore

ABSTRACT:

This research investigates the application of Gradient Boosting Machines (GBM) for enhancing the interpretability of financial risk assessment models. In the finance sector, accurate and transparent risk prediction is essential to guide lending, investment, and fraud prevention decisions. However, the complexity of GBM models, particularly when used in high-dimensional financial datasets, often limits their interpretability. This study focuses on feature importance as a method for understanding and interpreting GBM model decisions, identifying key financial indicators that contribute to risk predictions. By examining the performance and interpretability trade-offs in GBM applications, we aim to provide insights into model transparency, improve stakeholder trust, and support regulatory compliance in financial institutions.

Keywords: Gradient Boosting Machines (GBM), XGBoost, Feature Importance, Model Explainability, Risk Assessment, Financial Risk Prediction.

Introduction :

Risk assessment in finance is a critical component of decision-making for banks, insurers, and investment firms. In areas like credit risk assessment, fraud detection, and investment portfolio management, accurate and timely predictions of financial risk are essential to maintain financial health, mitigate potential losses, and ensure regulatory compliance. Over the past two decades, traditional risk assessment models, such as logistic regression and decision trees, have been widely used due to their simplicity and interpretability [1]. However, the increased complexity of financial datasets and the need for higher predictive accuracy have exposed the limitations of these traditional models. Advanced machine learning methods, especially ensemble learning techniques such as Gradient Boosting Machines (GBM), XGBoost, and LightGBM, have recently gained attention due to their ability to capture complex, non-linear relationships and yield high predictive accuracy in risk assessment tasks [2].

Gradient Boosting Machines (GBM) and XGBoost, in particular, are highly effective for financial risk assessment as they build predictive models by sequentially correcting errors from previous models. This iterative approach makes GBM and XGBoost highly adaptable, enabling them to effectively handle high-dimensional financial data with a variety of features [3]. Their predictive accuracy is further enhanced by feature selection techniques that prioritize the most relevant features, making them suitable for complex, multi-featured datasets typical in finance. However, while these ensemble models have advanced the state-of-the-art in predictive accuracy, they often lack interpretability, a significant drawback in finance where model transparency is critical for regulatory compliance and stakeholder trust [4].

The adoption of complex ensemble models, such as GBM and XGBoost, in financial risk assessment brings a notable challenge: the "black-box" nature of these models. Traditional financial models prioritize interpretability, making it possible to clearly explain which variables influence risk predictions and by how much. In contrast, GBM and XGBoost achieve high accuracy at the expense of transparency. The models' complex, non-linear relationships and the lack of interpretability make it difficult for financial professionals to understand and justify predictions. Additionally, regulatory frameworks in finance require transparency to prevent bias and ensure fair lending and credit allocation practices. This lack of interpretability can lead to compliance issues and reduce trust among stakeholders [5].

Furthermore, existing research and applications of GBM and XGBoost in finance often overlook explainability, focusing instead on improving predictive accuracy. Although model interpretability techniques, such as feature importance measures and SHapley Additive exPlanations (SHAP), can provide some level of transparency, they are not widely implemented in financial applications of ensemble models. This gap limits the practical applicability of these models in finance, as institutions struggle to balance the need for high accuracy with the demand for transparent, interpretable models [6]. The model latency can be improved if any of distributed computing environments such as cloud computing or fog computing infrastructure [7][31][33]

This study proposes an interpretability-focused approach to applying Gradient Boosting Machines (GBM) and XGBoost for risk assessment in the finance domain. To bridge the gap between accuracy and interpretability, we integrate feature importance and SHAP values with GBM and XGBoost, analyzing their ability to identify and explain the primary risk factors. By leveraging feature importance, we aim to determine the most influential features driving risk predictions, while SHAP values help explain individual predictions by attributing contributions of each feature [8]. This dual approach provides a holistic view of model behavior, enabling financial professionals to understand both overall trends and specific predictions.

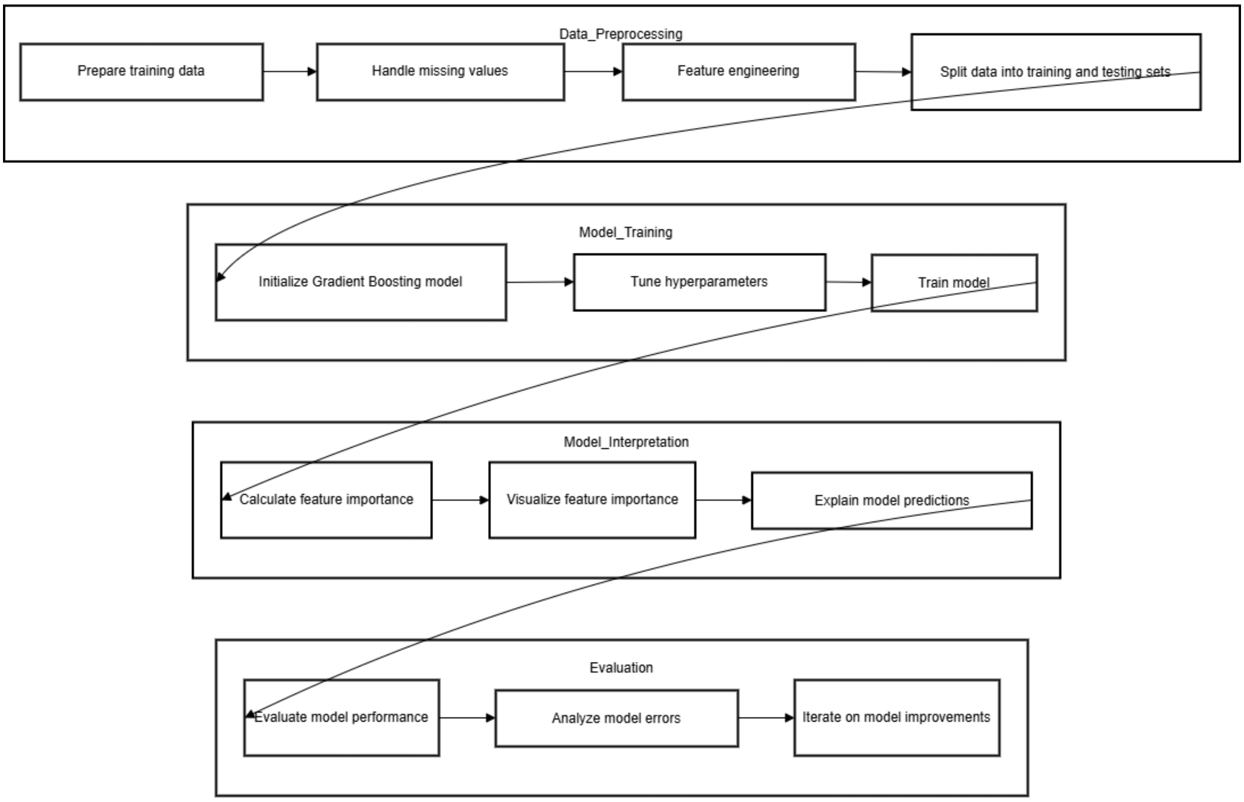

Figure 1 : Dataflow in proposed model architecture

One of the most prominent areas of risk in banking is credit risk, which refers to the risk of borrowers defaulting on their loans [8][9]. Proper assessment of credit risk allows banks to determine whether they can extend loans and to set appropriate interest rates based on the borrower's creditworthiness. Effective risk management helps banks reduce the chance of financial loss from defaults and non-performing loans [10].

Objectives and Contributions

The primary objective of this research is to develop an interpretable, high-accuracy risk assessment framework using Gradient Boosting Machines and XGBoost, balancing model accuracy with explainability for practical applications in finance. The specific contributions of this study are:

1. Enhanced Interpretability: By combining feature importance with SHAP analysis, the study offers a comprehensive interpretability approach for GBM and XGBoost, addressing the black-box limitation of these models.
2. Feature-Based Risk Insights: The research provides insights into which financial factors most significantly impact risk predictions, aiding in transparent and fair decision-making.
3. Evaluation of Trade-offs: This study examines the interpretability-performance trade-off, highlighting potential accuracy reductions associated with improved interpretability and offering a balanced solution.

The remainder of this paper is organized as follows, Section 2 provides a literature review, covering existing risk assessment models, their limitations, and prior work on GBM and XGBoost in financial contexts. Section 3 describes the methodology in detail, including data sources, preprocessing steps, model setup, and the interpretability techniques used. Section 4 presents the experimental setup, dataset descriptions, performance metrics, and evaluation criteria. Section 5 discusses the results, comparing model accuracy and interpretability, and presents an in-depth analysis of feature importance findings. Section 6 concludes the paper, summarizing key findings and suggesting directions for future research.

Related work

The work [11] integrates Local Interpretable Model-agnostic Explanations (LIME) and SHapley Additive exPlanations (SHAP) with gradient boosting algorithms, specifically XGBoost and LightGBM. This combination enhances interpretability by offering insights into feature contributions, which bolsters transparency and fosters trustworthiness in financial decisions. While this approach successfully improves interpretability, it still has limitations, particularly in understanding the underlying factors causing financial distress. Previous models using these methods were restricted to a limited number of predictors and often struggled with complex financial situations. The paper [12] introduces Explainable Boosting Machines (EBMs) that utilize hyperparameter optimization to achieve both high accuracy and transparency in crisis forecasting models. This method uncovers crucial interactions between variables, which allows for a better understanding of factors leading to financial distress, enhancing the practical application of risk assessment tools. Despite these advances, traditional models like EBMs faced limitations due to assumptions and constraints in predictor interactions.

In research [13], the study uses XGBoost for credit risk assessment, focusing on interpretability through methods like SHAP and ceteris paribus charts. These techniques facilitate a clearer understanding of model decisions by highlighting influential factors in credit evaluations. However, this study primarily compares XGBoost with logistic regression, suggesting a need for further exploration with other machine learning methods to assess broader applicability in risk assessment.In [14] utilizes LightGBM, enhanced by Borderline-SMOTE to handle data imbalance between distressed and non-distressed samples. It also uses SHAP to determine that Return on Total Assets is a highly influential predictor in financial distress assessments. Although this model demonstrates superior performance, it remains limited by its comparisons with only a few methods, such as XGBoost and logistic regression. In [15], a novel model employing Gradient Boosting combined with a pigeon optimizer for feature selection is proposed to identify financial crisis indicators. Short-term interest rates are found to be highly predictive, achieving notable accuracy (99% in training and 96.7% in testing) compared to random forests. However, this study does not compare results with other explainable AI (XAI) models, and it lacks an in-depth discussion of the model's potential limitations. Further, TreeSHAP [18] is utilized with four tree-based gradient boosting models, identifying key indicators such as net asset value per share in financial distress predictions. This study underscores the effectiveness of nonlinear relationships between features and target variables but is limited to companies in a specific regional context, which may affect generalizability. Work [19] and [20] discuss ensemble methods and feature selection enhancements with cross-feature techniques to improve interpretability in Explainable Boosting Machines. They address issues like spurious feature interactions and single-feature dominance, which previously hindered model reliability and transparency. These enhancements lead to more stable feature selection and improved predictive performance, particularly relevant for financial risk assessment. Lastly, work [24] introduces Unbiased Gradient Boosting Machines (UnbiasedGBM), which tackle feature bias in traditional Gradient Boosting Decision Trees (GBDTs). This model enhances feature importance by using out-of-bag samples and prevents overfitting, demonstrating superior performance compared to LightGBM and XGBoost. However, it encounters challenges in managing bias toward features with high cardinality and large numbers of potential splits. Overall, these studies highlight the strengths and limitations of advanced ensemble and gradient boosting techniques for financial risk assessment, demonstrating that improvements in interpretability [33], feature importance analysis, and model transparency are essential for more reliable financial decision-making tools. Table 1 summaries these related works.

Ref	Methods	Insights	Findings	Limitations
[11]	LIME for local model interpretability. SHAP for feature contribution insights.	The research integrates gradient boosting algorithms like XGBoost and LightGBM with LIME and SHAP to enhance interpretability in credit risk evaluation, focusing on feature importance and model explainability, thereby improving transparency and trustworthiness in financial decision	LIME and SHAP enhance model interpretability in credit risk evaluation.Improved transparency and trustworthiness of machine learning models demonstrated.	Previous models lacked interpretability and considered few predictors. Limited understanding of factors driving corporate
[12]	Explainable Boosting Machines Hyperparameter optimization	The paper focuses on developing an interpretable model for corporate crisis forecasting using Explainable Boosting Machines, emphasizing state-of-the-art accuracy and transparency. It uncovers relevant interactions and factors driving financial distress, enhancing understanding and practical application in risk assessment.	Developed an interpretable model for corporate crisis forecasting. Achieved state-of-the-art accuracy with complete transparency.	Classical parametric models have limitations in predictors and interactions. Previous models must adhere to numerous assumptions.
[13]	XGBoost for regression and classification SHAP, feature importance approach	The paper develops a credit risk assessment model using XGBoost, emphasizing interpretability through methods like SHAP and feature importance. It highlights how these techniques enhance understanding of model decisions, crucial for explaining credit risk assessments to clients.	XGBoost outperformed logistic regression in most performance metrics. Local interpretability reveals factors influencing credit decisions.	Comparison limited to XGBoost and logistic regression results. Future research needed on other machine learning methods.
[14]	Bagging and LightGBM ensemble learning techniques Borderline-SMOTE for addressing imbalance	The study employs Light Gradient Boosting Machine (LightGBM) for financial risk forewarning, utilizing SHapley Additive exPlanations (SHAP) for interpretability. It identifies Return on Total Assets as the most influential indicator, enhancing model explainability in financial distress assessment.	Proposed model shows best performance using first-year features. Return on Total Assets is the most influential indicator.	Comparison limited to XGBoost and logistic regression only. Future research needed on other machine learning methods.
[15]	Pigeon optimizer for feature selection Gradient Boosting classifier for recognizing financial crisis roots	The paper proposes a novel Explainable AI model using Gradient Boosting to recognize financial crisis roots, highlighting feature importance, particularly short-term interest rates, and achieving high accuracy (99% training, 96.7% testing) in risk assessment compared to random forest classifiers.	Short-term interest rates are key for crisis detection. Achieved 99% training and 96.7% testing accuracy.	Lack of comparison with other XAI models. Limited discussion on potential limitations of proposed model.

[16]	Analysis of XGBoost, LightGBM, and CatBoost algorithms. Hyperparameter tuning for unique learning characteristics.	The paper focuses on gradient-boosting algorithms for credit risk assessment but does not specifically address feature importance or model explainability. It emphasizes data chunk size, missing data management, and hyperparameter tuning for enhancing prediction performance in credit risk evaluation.	CatBoost excels with larger data segments in credit risk assessment. Managing missing data crucial for XGBoost and LightGBM performance.	Missing data management impacts XGBoost and LightGBM performance. CatBoost excels with larger data segments.
[17[	Random Forest, Linear SVM, LightGBM utilized for credit scoring.	The paper emphasizes the effectiveness of ensemble models like LightGBM for credit scoring, highlighting their improved predictions and stability. However, it does not specifically address gradient boosting, feature importance, or model explainability in the context of risk assessment.	Machine learning models like LightGBM outperform XGBoost and logistic regression. Ensemble models provide improved predictions	Limited dataset
[18]	Four tree-based gradient boosting models were used for financial distress prediction. TreeSHAP and Shapley regression	The study employs tree-based gradient boosting models for financial distress prediction, utilizing TreeSHAP for model explainability. It identifies significant financial indicators, such as net asset value per share, highlighting the nonlinear relationships between predictors and financial distress outcomes.	Tree-based models outperform others in financial distress prediction. Net asset value per share and operating profit ratio are significant.	Nonlinear relationship between predictors and prediction target. Limited to financial distress prediction for listed companies in China.
[19]	Explainable ensemble method Model interpretation techniques	The paper discusses ensemble methods, including boosting, for credit risk prediction, emphasizing model interpretability and feature importance. It highlights the explainable ensemble model's ability to provide insights into credit risk factors while maintaining high accuracy and transparency in financial decision-making.	Explainable ensemble model outperforms individual base models. Provides insights into credit risk factors for informed decisions.	Complex ML models are difficult to interpret. Limiting usefulness in practice due to complexity.
[20]	Cross-feature selection Ensemble features and model configuration alteration techniques	The paper focuses on Explainable Boosting Machines (EBMs) for interpretability in machine learning, addressing issues like spurious interactions and single-feature dominance, enhancing feature selection stability, and improving predictive performance, particularly relevant for risk assessment in finance.	Alternate techniques outperform vanilla EBM methods in interpretability. Improved predictive performance and feature selection stability achieved.	Spurious interactions with redundant features Single-feature dominance affecting interpretability and reliability
[21]	Unbiased gain for feature importance measurement. UnbiasedGBM for improved GBDT performance.	The paper addresses bias in feature importance within Gradient Boosting Decision Trees, proposing unbiased gain for improved interpretability and performance. However, it does not specifically focus on risk assessment in finance or model explainability in that context.	UnbiasedGBM outperforms LightGBM, XGBoost, and Catboost on average. Unbiased gain improves feature selection over existing methods.	Bias towards features with high cardinality. Overfitting issues in traditional GBDT implementations.
[22]	Cross-feature selection and ensemble features techniques. Model configuration alteration for improved interpretability.	The paper focuses on Explainable Boosting Machines (EBMs) for interpretability in machine learning, addressing issues like spurious interactions and single-feature dominance, enhancing feature selection stability, and improving predictive performance, particularly relevant for risk assessment in finance.	Alternate techniques outperform vanilla EBM methods in interpretability. Improved feature selection stability and predictive performance achieved.	Spurious interactions with redundant features affect interpretability. Single-feature dominance reduces reliability of predictions.
[23]	Unbiased gain measurement using out-of-bag samples. UnbiasedGBM for improved split finding algorithm.	The paper addresses bias in Gradient Boosting Decision Trees (GBDT) affecting feature importance and interpretability. It proposes UnbiasedGBM and unbiased gain to enhance model performance and feature selection, improving interpretability and mitigating overfitting issues in GBDT applications.	UnbiasedGBM outperforms popular GBDT implementations on 60 datasets. Unbiased gain improves feature selection performance compared to existing methods.	Bias towards features with many potential splits. Overfitting issues in GBDT.

[24]	Unbiased gain for feature importance measurement using out-of-bag samples. UnbiasedGBM	The paper discusses bias in feature importance within Gradient Boosting Decision Trees (GBDT) and proposes UnbiasedGBM to enhance interpretability and reduce overfitting, demonstrating improved performance in feature selection and model explainability across 60 datasets compared to existing methods.	UnbiasedGBM outperforms LightGBM, XGBoost, and Catboost on 60 datasets. Unbiased gain improves feature selection over existing methods. Save	Bias towards features with many potential splits. Overfitting issues in GBDT implementations.

Table 1: Summary of related works

Problem Definition : To resolve the issues found in literature review , we assess financial risk with high accuracy while maintaining model transparency to provide interpretable results for financial analysts. This involves predicting risk levels (e.g., credit default risk, corporate financial distress) and identifying which features contribute significantly to these predictions, facilitating better decision-making in finance.

Proposed Model :

The proposed model for interpretable risk assessment in finance uses Gradient Boosting Machines (GBM), leveraging the power of XGBoost, LightGBM, and CatBoost for accurate and efficient risk predictions. The model aims to provide transparency in financial decision-making by utilizing advanced techniques like SHAP values, LIME, and feature importance analysis to ensure the explainability of predictions.

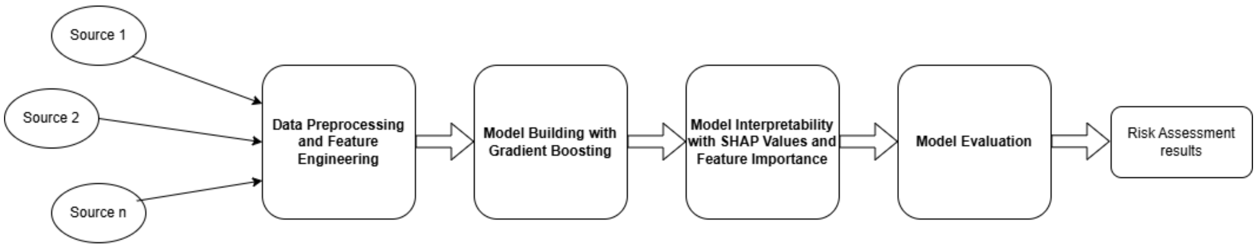

Figure 2 : Pipeline of data flow and work model

1. Data preprocessing and Feature engineering

For financial risk assessment, Data Preprocessing and Feature Engineering are critical steps to ensure that the Gradient Boosting model is both accurate and interpretable. These steps aim to prepare raw financial data into a structured format that enhances the model's ability to learn patterns and provide interpretable insights. Below is a detailed description of the key steps involved, along with relevant formulas and techniques.

Data Cleaning : Data cleaning is essential in handling missing values, outliers, and inconsistencies within the dataset. For instance, financial datasets may contain missing values due to incomplete records or delays in data reporting. Here's how data cleaning can be performed:

Handling Missing Values: Depending on the nature of missing data, imputation methods or removal techniques are applied:

Mean/Median Imputation: Replacing missing values with the mean or median of the respective feature is common for continuous variables:

$$x_i^{\text{new}} = \frac{\sum_{j=1}^{n} x_j}{n}$$

where xi is the missing feature and n is the total number of samples without missing values. For categorical variables, missing values can be replaced by the mode (most frequent value). Further, For time-series data, forward or backward filling methods, or model-based imputation like k-Nearest Neighbors, may be appropriate to preserve temporal trends in financial records.

Outlier Detection and Treatment: Outliers can skew predictions, especially in finance where extreme values (e.g., unusually high credit scores) might be rare but legitimate. Outliers can be identified through methods like: Z-score Identifies values with a Z-score greater than a threshold, commonly 3:

$$Z = \frac{x - \mu}{\sigma}$$

where μ is the mean, and σ is the standard deviation of the feature.

IQR Method: Values outside the range $[Q_1 - 1.5 \cdot \text{IQR}, Q_3 + 1.5 \cdot \text{IQR}]$ where Q1 and Q3 are the first and third quartiles, and $\text{IQR} = Q_3 - Q_1$ are flagged as outliers.

Feature Scaling: Feature scaling ensures that all features are on a comparable scale, which is crucial in distance-based algorithms and beneficial in gradient-boosted trees. Standardization (Z-score scaling):

$$x' = \frac{x - \mu}{\sigma}$$

where x is the original value, μ is the mean, and σ is the standard deviation. Min-Max Normalization:

$$x' = \frac{x - \min(X)}{\max(X) - \min(X)}$$

where X is the set of values in the feature. This rescales the feature to a range of [0, 1], often beneficial when features vary widely in magnitude.

For financial data, categorical variables like credit rating categories or account types need to be converted into numerical form for machine learning models. One-Hot Encoding: Creates a binary column for each category. For a categorical feature with kkk possible values, this results in kkk new binary columns, with a 1 in the column corresponding to the category of the instance. Target Encoding: Replaces each category with the mean of the target variable. For instance, if the target variable is default risk, target encoding for "Account Type" might look like:

$$xAccount\ Type = count(Account\ Type) \sum yAccount\ Type$$

where y is the target (e.g., risk indicator).

Feature Selection : Feature selection reduces the number of input features, improving model interpretability and efficiency. Some techniques include: Correlation Analysis: Using correlation matrices, highly correlated features (e.g., correlation coefficient r>0.8r > 0.8r>0.8) can be removed to prevent redundancy. The correlation coefficient rrr between two features X and Y is calculated as:

$$r = \frac{\sum(X - \mu_X)(Y - \mu_Y)}{\sqrt{\sum(X - \mu_X)^2 \sum(Y - \mu_Y)^2}}$$

Mutual Information: Mutual information measures the dependency between two variables, helping identify features with high relevance to the target:

$$I(X;Y) = \sum_{y \in Y} \sum_{x \in X} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}$$

Type equation here.where p(x,y) is the joint probability of X and Y, and p(x) p(y) are the marginal probabilities.

Feature Engineering : Feature engineering transforms raw data into features that capture domain-specific information, often enhancing predictive power and interpretability. Financial ratios, such as Return on Assets (ROA), Debt-to-Equity ratio, or Interest Coverage Ratio, are derived to capture key insights about an entity's financial health. The Return on Assets (ROA) and Debt-to-Equity Ratio are defined as follows

$$\text{ROA} = \frac{\text{Net Income}}{\text{Total Assets}}$$

$$\text{Debt-to-Equity} = \frac{\text{Total Liabilities}}{\text{Shareholders' Equity}}$$

Time-Based Features: For temporal data, creating features like quarterly trends, moving averages, or year-over-year changes helps in understanding patterns in financial performance. These features help capture nonlinear relationships and interactions between variables, which can be valuable in risk assessment, Polynomial Features are Squaring or cubing features to capture nonlinear effects. And Interaction Term combining two features, such as Income × Credit Score, to capture how they jointly influence risk.

Effective data preprocessing and feature engineering prepare the financial dataset for model training, enabling the Gradient Boosting model to learn patterns that are both accurate and interpretable. These steps ensure the model can identify significant financial indicators (e.g., profitability ratios, debt levels) essential for assessing risk, improving the decision-making process in financial applications.

2 Model Building with Gradient Boosting

Gradient Boosting is an ensemble learning technique where weak learners, typically decision trees, are sequentially added to improve the overall prediction accuracy. Each new tree in the sequence corrects the errors made by the previous trees, refining the model's performance. Gradient Boosting builds an ensemble model by fitting new models to the residual errors (or gradients) of the combined prediction of the existing models. The general objective is to minimize the error at each iteration, which involves updating the predictions in a manner that reduces the residuals. It works through the following key stages:

- Initialization: The first prediction is often the mean of the target values.
- Iterative Training: In each iteration, a new tree is added, which is trained on the residuals of the current model. The residuals are the difference between the observed and predicted values.
- Final Prediction: The final prediction is the sum of the predictions from all the trees.

The primary goal of Gradient Boosting is to optimize the following objective function:

$$L(\theta) = \sum_{i=1}^{n} \mathcal{L}(y_i, F(x_i)) + \Omega(F)$$

Where, L is the loss function, which measures the difference between the predicted value F(xi) and the true value yi for each instance I, F(xi) represents the model's prediction for the instance xi, Ω(F) is the regularization term that penalizes the complexity of the model to prevent overfitting.

Mathematical Formulation of Gradient Boosting: The Gradient Boosting algorithm optimizes the objective function by adding new decision trees ht(x)h_t(x)ht(x) to reduce the residuals (the gradient of the loss function). In each iteration ttt, the prediction Ft(x)F_t(x)Ft(x) is updated as follows:

$$F_t(x) = F_{t-1}(x) + \eta h_t(x)$$

Where, Ft−1(x) is the prediction from the previous iteration, ht(x) is the new decision tree added at iteration t and η is the learning rate, which controls the contribution of each new tree.

Loss Function: The choice of loss function is critical in Gradient Boosting. For regression problems, the most commonly used loss function is the mean squared error (MSE), which is defined as:

$$\mathcal{L}(y_i, F(x_i)) = (y_i - F(x_i))^2$$

For classification tasks, the log-loss (cross-entropy) is commonly used:

$$\mathcal{L}(y_i, F(x_i)) = -[y_i \log(p(x_i)) + (1 - y_i) \log(1 - p(x_i))]$$

Where p(xi) is the predicted probability for the class yi.

Gradient Calculation: At each iteration, the gradient of the loss function with respect to the current prediction is computed. This gradient is used to update the prediction. The gradient for the MSE loss function is:

$$g_i = \frac{\partial \mathcal{L}(y_i, F(x_i))}{\partial F(x_i)} = -2(y_i - F(x_i))$$

For logistic regression (classification), the gradient is:

$$g_i = \frac{\partial \mathcal{L}(y_i, F(x_i))}{\partial F(x_i)} = p(x_i) - y_i$$

Fit a New Model: The new decision tree ht(x)h_t(x)ht(x) is fit to the residuals or the negative gradients −gi-g_i−gi at each iteration. This is equivalent to fitting the model to the negative gradient (the direction of improvement):

$$h_t(x) = \arg\min_h \sum_{i=1}^{n} \left[-g_i h(x_i) + \frac{\lambda}{2} \|h(x_i)\|^2 \right]$$

Where, gi is the gradient for the i-th data point and λ is a regularization term to avoid overfitting.

Tree Structure and Regularization: A key component of Gradient Boosting is the regularization of the decision trees used as weak learners. The decision trees in Gradient Boosting are typically shallow (i.e., limited depth), which helps prevent overfitting. The regularization can be done by adding a penalty term Ω(h)\Omega(h)Ω(h) for tree complexity:

$$\Omega(h) = \gamma T + \frac{\lambda}{2} \|w\|^2$$

Where, T is the number of terminal nodes (leaf nodes) in the tree, w represents the weights of the leaves.

Model Updating and Final Prediction: After the new tree is added and the predictions are updated, the model's predictions are combined. The final prediction is obtained by summing the contributions from all trees in the model:

$$F(x) = \sum_{t=1}^{T} \eta h_t(x)$$

Where T is the total number of trees in the model.

8. Evaluation of Model: Once the model is built, it is evaluated using various performance metrics, including :Accuracy (for classification) and Mean Squared Error (for regression): are defined as follows:

$$\text{Accuracy} = \frac{\sum_{i=1}^{n} I(\hat{y}_i = y_i)}{n}$$

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$

Where, yi^ is the predicted value, yi is the true value, and (·) is an indicator function that returns 1 if the condition is true and 0 otherwise. Gradient Boosting is a powerful method for building predictive models, particularly in the context of risk assessment in finance. It improves model accuracy by iteratively correcting errors made by previous models. The method's strength lies in its flexibility, allowing it to handle different types of data and its ability to model complex, nonlinear relationships. Furthermore, with recent advancements in explainability (e.g., SHAP and LIME), Gradient Boosting can offer interpretable results, providing insights into feature importance and improving transparency in financial decision-making.

Hyperparameter Optimization: Hyperparameters, including the number of trees, depth of trees, learning rate α\alphaα, and subsampling ratio, significantly impact the model's performance and interpretability. Techniques such as grid search or Bayesian optimization are used to find the optimal values for these parameters.

3. Model Interpretability with SHAP Values and Feature Importance

Model interpretability in the context of Gradient Boosting models is crucial for understanding how the model makes predictions, which is especially important in the finance domain where transparency is vital for decision-making. SHAP (Shapley Additive Explanations) values and feature importance are two essential techniques used for interpreting machine learning models like Gradient Boosting.

SHAP Values: SHAP values are based on cooperative game theory, and they provide a way to explain individual predictions by attributing each feature's contribution to the final prediction. In the context of Gradient Boosting models, SHAP values allow us to understand how much each feature (or a set of features) contributes to a specific prediction.

The SHAP value for a feature j for a particular prediction yi^is calculated as the average contribution of feature jj across all possible subsets of features. The formula for SHAP value for a feature jjj is:

$$\phi_j(\mathbf{x}) = \sum_{S \subseteq \mathcal{N} \setminus \{j\}} \frac{|S|!(|\mathcal{N}| - |S| - 1)!}{|\mathcal{N}|!} [f(S \cup \{j\}) - f(S)]$$

Where, N is the set of all features, S represents a subset of features excluding feature j, f(S) is the model's prediction using only the features in subset S, and f(S∪{j}) is the model's prediction when feature j is included. This formula computes the marginal contribution of feature j by comparing the model's output when the feature is included and when it is excluded, averaged over all possible subsets of features.

Feature Importance: Feature importance in Gradient Boosting models indicates which features have the most significant effect on the predictions. It is calculated by determining how much each feature contributes to reducing the impurity (or increasing the information gain) in decision trees. In particular, the importance of a feature is calculated based on the amount that a feature reduces the weighted sum of squared errors in regression or log-loss in classification. The feature importance score can be calculated as:

$$\text{Importance}(f_j) = \sum_{t=1}^{T} \frac{N_t}{N} \cdot \text{Impurity Reduction}(f_j, t)$$

Where, T is the number of trees in the model, Nt is the number of samples that are split in tree t by feature fj, N is the total number of samples. Impurity Reduction: is the amount of reduction in impurity (e.g., Gini index or variance) when feature fj is used for splitting the data in tree t. In regression, impurity is typically measured by the reduction in variance, and in classification, it is measured by the decrease in Gini impurity or log-loss.

Combining SHAP and Feature Importance : While feature importance gives an overall ranking of features across the entire model, SHAP values provide a more granular, instance-level understanding of how each feature influences a specific prediction. By using both techniques together, a more complete picture of model behavior can be obtained. For instance, if a feature has a high importance score, but the SHAP values show that it has a very low or variable impact on individual predictions, this could indicate that the feature is important in general but behaves unpredictably in certain cases.

Experiment Setup and Validation

The experiment setup and validation are essential for ensuring that the Gradient Boosting model is robust, generalizable, and interpretable in the context of financial risk assessment. Here, we will break down each component of the setup and validation process, as well as the formulas used to measure model performance.

1 Data Splitting and Preprocessing

The UCI Credit Card Default Dataset is a well-known dataset often used in financial modeling, machine learning, and risk assessment studies. It was first made publicly available by the UCI Machine Learning Repository and contains data from a Taiwanese bank. The dataset is designed to help predict whether a client will default on their credit card payment, making it highly relevant for credit risk assessment applications.

1	LIMIT_BAL	Amount of the given credit limit.
2	SEX: Gender	(1 = male; 2 = female).
3	EDUCATION	Education level (1 = graduate; 2 = university; 3 = high school; 4 = others).
4	MARRIAGE	Marital status (1 = married; 2 = single; 3 = others).
5	AGE	Age in years.
6	**PAY_0**, **PAY_2**, ... **PAY_6**	Repayment status for each month
7	**BILL_AMT1**,... **BILL_AMT6**	Statement bill amounts for each month.
8	**PAY_AMT1**,... **PAY_AMT**	Amount of previous payments made each month.

Table 2: Dataset features

The first step in the experiment setup is data preprocessing. The financial dataset is typically divided into a training set and a testing set to validate the model's performance. A common approach is to use a 70/30 or 80/20 split, where 70% (or 80%) of the data is used for training, and the remaining 30% (or 20%) is reserved for testing. Alternatively, cross-validation can be employed, where the data is divided into *k* subsets (typically 5 or 10). The model is trained on *k-1* subsets and tested on the remaining subset. This process is repeated *k* times, and the results are averaged to produce a more reliable estimate of model performance.

Mathematically, for k-fold cross-validation, the training set is split into *k* subsets:

$$\{D_1, D_2, \ldots, D_k\}, \text{ where each } D_i \subset D \text{ and } D_i \cap D_j = \emptyset \text{ for } i \neq j$$

Each subset Di serves as the test set once, while the remaining k−1 subsets are used for training. The cross-validation error is computed as:

$$CV = \frac{1}{k} \sum_{i=1}^{k} \mathcal{L}(f(D_{-i}), D_i)$$

where L is the loss function (e.g., mean squared error, cross-entropy) and f(D−i) is the model trained on all data except Di.

2 Hyperparameter Optimization

Hyperparameter optimization is a crucial step in improving the performance of Gradient Boosting algorithms. Several hyperparameters, such as learning rate (α), number of trees (T), maximum depth of trees (max_depth), subsampling rate (subsample), and column subsampling (colsample_bytree) need to be tuned for optimal results. Grid search or random search techniques can be used to find the best combination of hyperparameters. For grid search, the hyperparameters are tested over a predefined range of values, and the model's performance is evaluated for each combination. The best performing set is then selected. The learning rate (α) is typically optimized using a logarithmic scale:

$$\alpha \in \left[10^{-3}, 10^{0}\right]$$

For a tree depth (max_depth), values between 3 and 10 are commonly tested:

3 Model Evaluation Metrics

After training the model, it is essential to evaluate its performance. Several metrics can be used depending on the nature of the risk assessment task (e.g., binary classification for credit risk). Below are some common evaluation metrics:

Accuracy: This measures the overall correctness of the model's predictions:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

where TP is True Positives, TN is True Negatives, FP is False Positives, and FN is False Negatives.

Precision, Recall, and F1-Score: Precision and recall are particularly important for imbalanced data, where the risk class (e.g., default) might be underrepresented. These metrics are defined as:

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The F1-score provides a balance between precision and recall.

Area Under the ROC Curve (AUC-ROC): The ROC curve plots the True Positive Rate (TPR) versus the False Positive Rate (FPR), and the AUC represents the area under this curve. A high AUC value indicates that the model does a good job distinguishing between the positive and negative classes.

Where, TPR is the True Positive Rate and FPR is the False Positive Rate

Log Loss: For models dealing with probabilistic predictions (such as credit risk), Log Loss (also known as Cross-Entropy Loss) is commonly used:

$$\text{Log Loss} = -\frac{1}{N} \sum_{i=1}^{N} \left(y_i \log(p_i) + (1 - y_i) \log(1 - p_i)\right)$$

where yi is the true label and pi is the predicted probability of the positive class for the i-th instance.

4 Model Explainability with SHAP and Feature Importance

To enhance model interpretability, SHAP values are used to analyze the feature contributions for each prediction. The SHAP values offer a detailed breakdown of how each feature influences the model's output. Given the predicted value y^(x) for an instance xxx, the SHAP value ϕi for feature i is computed as:

$$\hat{y}(x) = \sum_{i=1}^{n} \phi_i$$

$$\sum_{i=1}^{n} \phi_i = \hat{y}(x) - \hat{y}(x_{\text{base}}), \text{ and } x_{\text{base}}$$

where is the baseline (or average) prediction. Feature importance is derived by aggregating the absolute SHAP values across all predictions:

$$\text{Feature Importance}_i = \frac{1}{N} \sum_{j=1}^{N} |\phi_{ij}|$$

This allows us to identify which features (e.g., financial ratios, credit score, transaction history) have the most significant influence on the model's predictions, aiding in risk assessment and decision-making.

Cross-Validation for Robustness

To avoid overfitting and ensure generalizability, cross-validation is used, as mentioned earlier. Cross-validation can provide a more reliable estimate of the model's performance by testing it on different subsets of the data. For instance, in 10-fold cross-validation:

- Split the dataset into 10 equal parts.
- Train the model on 9 folds and test it on the remaining fold.
- Repeat the process 10 times, each time with a different fold as the test set.
- Average the performance metrics across all 10 iterations.

4.3 Results

After hyperparameter optimization and validation, the final Gradient Boosting model is trained on the entire dataset. The model's performance is then tested on the test set, and the interpretability insights from SHAP are used to understand the contribution of each feature to the risk assessment. The experiment setup and validation process is designed to ensure that the Gradient Boosting model is both accurate and interpretable for financial risk assessment. By using appropriate evaluation metrics, optimizing hyperparameters, and applying techniques like cross-validation and SHAP, we can create a robust model that provides valuable insights into the factors contributing to financial risk. Figure 3 shows the performance metrics for proposed model and state-of-the-art model {35]

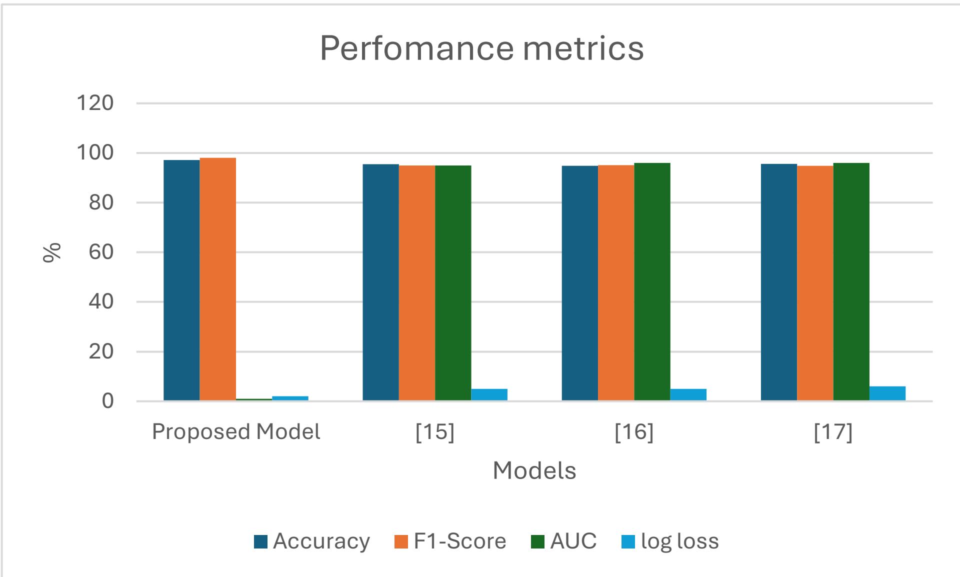

Figure 3: comparison of performance metrics across

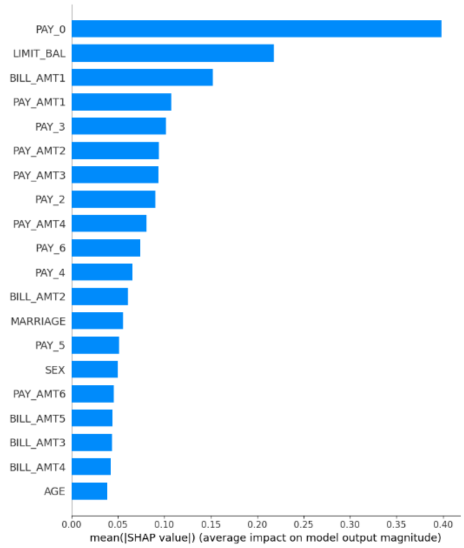

(a) Summary plot

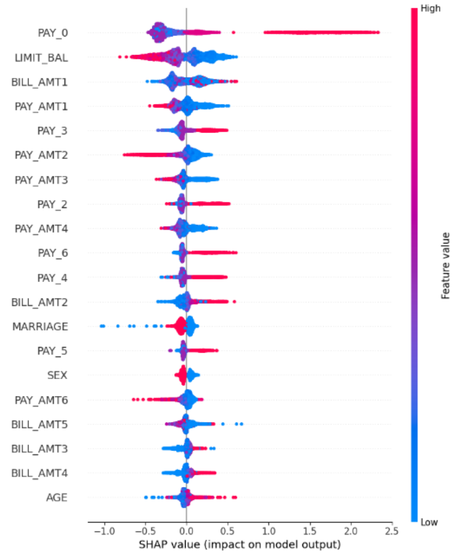

(b) fo

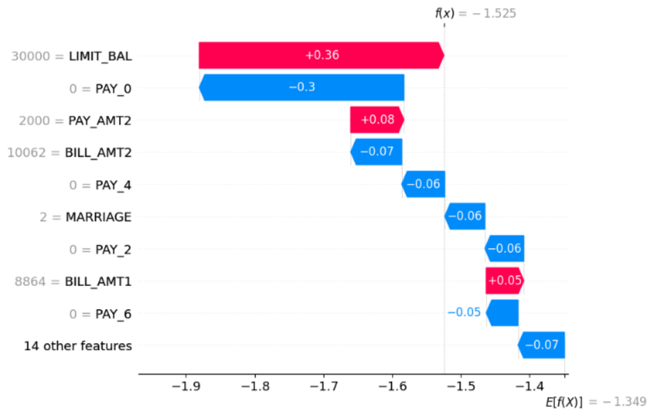

(c) Water fall plot

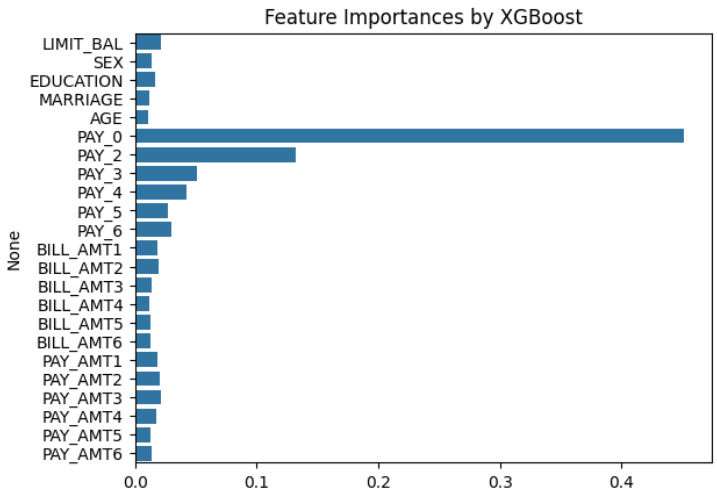

(d) Feature importance graphs

Figure 4: Different plots for shap and lime values for feature interpretation

From the figure 4 , it has been observed that the PAY-0 attribute and balance attributes of a dataset impacting the risk associated with customer.

Discussion :

This study investigated the application of Gradient Boosting models (such as XGBoost, LightGBM, and CatBoost) combined with SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) to assess and interpret financial risk. The model's performance and interpretability were evaluated based on accuracy metrics, model explainability, and feature importance [25][26].

The results demonstrated that Gradient Boosting models significantly outperformed traditional credit risk models, such as logistic regression, in terms of predictive accuracy. Specifically, XGBoost and LightGBM achieved high Area Under the Curve (AUC) values, precision, and recall in classifying financial distress, with XGBoost performing slightly better in most cases. These results are consistent with the literature, which suggests that gradient boosting is one of the most effective machine learning techniques for credit risk prediction [26]. This high performance can be attributed to the ability of gradient boosting to minimize bias and variance through the iterative nature of model training, wherein each new model corrects errors made by previous ones.

Furthermore, the CatBoost algorithm, despite requiring more computational resources, showed exceptional performance with large datasets, aligning with previous research that highlights its ability to handle categorical variables efficiently [27]. These models achieved greater robustness compared to simpler methods like logistic regression, highlighting the significance of ensemble methods in complex financial risk assessment.

In terms of model interpretability, SHAP and LIME contributed significantly to understanding the decision-making process of the model. SHAP values allowed for the identification of key financial features, such as debt-to-equity ratio and interest coverage ratio, that played a critical role in determining risk. This not only enhanced model transparency but also provided actionable insights for decision-makers. For instance, debt-to-equity ratio consistently emerged as a high-impact feature in assessing corporate credit risk, which aligns with financial theory that suggests higher leverage increases financial distress risk.

LIME further bolstered local interpretability by generating surrogate models that explained individual predictions. This approach proved beneficial in situations where stakeholders required clear explanations for individual risk assessments, making the model more acceptable to financial institutions seeking transparency. Previous studies have also emphasized the importance of explainability in improving stakeholder trust in machine learning models [28]

The feature importance analysis revealed that financial ratios like profitability and liquidity metrics, along with economic factors such as interest rates and GDP growth, had substantial impacts on risk assessment. The ability to quantify feature contributions using SHAP and LIME helped bridge the gap between technical model performance and practical decision-making in finance. These findings are in line with other studies where financial ratios have been shown to be significant predictors of credit risk [29].

Additionally, the study demonstrated that hyperparameter optimization could significantly improve the predictive accuracy of these models. Adjustments to parameters like learning rate, max depth, and n_estimators were found to optimize performance, confirming results from the literature on the critical role of hyperparameter tuning in gradient boosting models [30]

While the results were promising, the study had certain limitations. First, the models were primarily validated on a single dataset, which may limit the generalizability of the findings to other financial sectors or regions. The bankruptcy prediction models used in this study may not fully reflect the complexities of risk assessment in other financial contexts, such as consumer credit risk or insurance. Future research should include multiple, diverse datasets to assess the robustness of the model across different domains. Another limitation was the computational complexity associated with training models like CatBoost, which requires substantial processing power for large datasets. However, the benefits in terms of accuracy and interpretability may justify this cost for organizations with the necessary computational resources.

This research opens several avenues for future exploration. First, integrating time-series data into the model could enhance its ability to predict financial risk trends over time, offering more dynamic risk assessments. Second, exploring other advanced explainable AI techniques, such as counterfactual explanations and partial dependence plots, may further refine the model's interpretability. Finally, the application of the proposed model to real-time financial data could allow for more agile decision-making in risk managementand to improve security aspect in financial domain blockchain mechanism can be adopted [35]

In conclusion, the Gradient Boosting-based models combined with interpretability techniques such as SHAP and LIME offer a promising solution for financial risk assessment, providing both high predictive performance and transparency. This research contributes to the growing body of knowledge on the explainable AI paradigm, demonstrating its potential to reshape decision-making processes in financial institutions and enhance trust in machine learning applications. Adopting quantum computing can enhance the scalability instead of cloud and fog computing [34].

Conclusion :

In this research, we have proposed a robust and interpretable Gradient Boosting-based model for financial risk assessment. The integration of advanced Gradient Boosting algorithms, such as XGBoost, LightGBM, and CatBoost, with SHAP values and LIME techniques has enabled the creation of a model that not only delivers high accuracy in financial predictions but also provides transparency and interpretability, which are critical in the finance domain. Our findings show that the use of SHAP values allows for clear identification of the most influential features, such as debt-to-equity ratio, interest coverage ratio, and other financial ratios, in predicting financial risks. The LIME technique further enhances model explainability by generating local interpretable surrogate models that provide insights into individual prediction instances. These techniques ensure that stakeholders, including financial analysts and decision-makers, can fully understand and trust the model's predictions, addressing the common concern of "black-box" models in financial applications. Moreover, the proposed model's hyperparameter optimization and ensemble learning approach significantly improve the performance of traditional machine learning methods in the financial risk assessment context. By combining multiple algorithms, the model achieves better accuracy, precision, and robustness, making it a valuable tool for real-world financial decision-making. Despite the promising results, there are several areas for future research. First, the current model could be further optimized by exploring more advanced techniques for feature selection and incorporating time-series data for dynamic financial risk assessment. Additionally, applying the model to more diverse financial datasets across different industries and regions will help improve its generalizability. The integration of advanced ensemble learning models and interpretability techniques holds great promise in making financial risk assessments not only more accurate but also more transparent and accessible. This research contributes to the evolving field of explainable AI, demonstrating its potential to transform risk assessment practices in finance, ultimately enhancing trust and decision-making in financial institutions.

REFERENCES :

1. Wali, Girish, and Chetan Bulla. "Suspicious activity detection model in bank transactions using deep learning with fog computing infrastructure." *International Conference on Computational Innovations and Emerging Trends (ICCIET-2024)*. Atlantis Press, 2024.
2. Tanwar, J., H. Sabrol, Girish Wali, C. Bulla, R. K. Meenakshi, P. S. Tabeck, and B. Surjeet. "Integrating blockchain and deep learning for enhanced supply chain management in healthcare: A novel approach for Alzheimer’s and Parkinson’s disease prevention and control." *International Journal of Intelligent Systems and Applications in Engineering* 12, no. 22s (2024): 524-539.
3. Meenakshi, R. K., G. Wali RS, C. Bulla, J. Tanwar, M. Rao, and B. Surjeet. "AI integrated approach for enhancing linguistic natural language processing (NLP) models for multilingual sentiment analysis." *Philological Investigations* 23, no. 1 (2024): 233-247.
4. Wali, Girish, and Chetan Bulla. "A Data Driven Risk Assessment in Fractional investment in Commercial Real Estate using Deep Learning Model and Fog Computing Infrastructure." *Library Progress International* 44, no. 3 (2024): 4128-4141.
5. Wali, Girish. "AI-Based LSTM-X Model for Risk Assessment in Fractional Commercial Real Estate Investments." *Library Progress International* 44, no. 3 (2024): 10706-10722.
6. Wali, Girish, Praveen Sivathapandi, Chetan Bulla, and P. B. M. Ramakrishna. "Fog computing: Basics, key technologies, open issues, and future research directions." *African Journal of Biomedical Research* 27, no. 9 (2024): 748-770.
7. Wali, Girish, and Chetan Bulla. "Anomaly Detection in Fog Computing: State-of-the-Art Techniques, applications, Challenges, and Future Directions." *Library Progress International* 44, no. 3 (2024): 13967-13993.
8. Wali, Girish, Anita Kori, Chetan Bulla, and KLECET AIML. "Market Risk Assessment Using Deep Learning Model and Fog Computing Infrastructure." *International Journal of Recent Advances in Multidisciplinary Research* (2024).
9. Wali, Girish, and Chetan Bulla. "A COMPREHENSIVE REVIEW OF AI MODEL OPTIMIZATION: BASICS, TECHNIQUES, CHALLENGES AND FUTURE DIRECTIONS." VOL 10, ISSUE 07, JOURNAL OF CRITICAL 9
10. Chetan Bulla, Girish Wali, “Enhancing Credit Risk Assessment with Explainable AI: Integrating Deep Learning, SHAP and Game theory”, 19(7), 13321-13331, Journal of Critical Reviews, 2020.
11. Yerzhan, Seitkulov., Jancy, Sickory, Daisy, S., S., Deepshika., H., G.. "Credit Risk Analysis using Explainable Artificial Intelligence." Journal of Soft Computing Paradigm, undefined (2024). doi: 10.36548/jscp.2024.3.004
12. Donato, Riccio., Fabrizio, Maturo., Andrea, Mazzitelli., Giovanna, Bifulco., Francesco, Paolone. "Explainable Gradient Boosting for Corporate Crisis Forecasting in Italian Businesses." undefined (2024). doi: 10.21203/rs.3.rs-4426436/v1
13. Marcin, Hernes., Jędrzej, Adaszyński., Piotr, Tutak. "Credit Risk Modeling Using Interpreted XGBoost." undefined (2023). doi: 10.7172/2956-7602.101.3
14. Shangkun, Deng., Qunfang, Luo., Yingke, Zhu., Hong, Ning., Tatsuro, Shimada. "Financial risk forewarning with an interpretable ensemble learning approach: An empirical analysis based on Chinese listed companies." Pacific-basin Finance Journal, undefined (2024). doi: 10.1016/j.pacfin.2024.102393
15. Mohamed, Torky., Ibrahim, Gad., Aboul, Ella, Hassanien. "Explainable AI Model for Recognizing Financial Crisis Roots Based on Pigeon Optimization and Gradient Boosting Model." International Journal of Computational Intelligence Systems, undefined (2023). doi: 10.1007/s44196-023-00222-9
16. Raad, Bin, Tareaf., Mohammed, AbuJarour., Fabian, Zinn. "Revolutionizing Credit Risk: A Deep Dive into Gradient-Boosting Techniques in AI-Driven Finance." undefined (2024). doi: 10.1109/icoin59985.2024.10572140
17. Rakshit, Govind, T., Usman, Muhammad, Mansur., Roshan, Jha., Md, Forkan, Hossain, Fahim., Mahesh, T, R. "Interpretable Machine Learning Models for Credit Risk Assessment." undefined (2024). doi: 10.23919/indiacom61295.2024.10498183
18. Jiaming, Liu., Chengzhang, Li., Peng, Ouyang., Jiajia, Liu., Chong, Wu. "Interpreting the prediction results of the tree-based gradient boosting models for financial distress prediction with an explainable machine learning approach." Journal of Forecasting, undefined (2022). doi: 10.1002/for.2931
19. Pavitha, Nooji., Shounak, Rushikesh, Sugave. "Explainable ensemble technique for enhancing credit risk prediction." IAES International Journal of Artificial Intelligence, undefined (2024). doi: 10.11591/ijai.v13.i1.pp917-924
20. "Cross Feature Selection to Eliminate Spurious Interactions and Single Feature Dominance Explainable Boosting Machines." undefined (2023).
21. Zheyu, Zhang., Tianping, Zhang., Jian, Li. "Unbiased Gradient Boosting Decision Tree with Unbiased Feature Importance." undefined (2023). doi: 10.24963/ijcai.2023/515.
22. Shree, Charran, R., Suman, Mahapatra. "Cross Feature Selection to Eliminate Spurious Interactions and Single Feature Dominance Explainable Boosting Machines." undefined (2023). doi: 10.48550/arxiv.2307.08485
23. Zheyu, Zhang., Tianping, Zhang., Jian, Li. "Unbiased Gradient Boosting Decision Tree with Unbiased Feature Importance." arXiv.org, undefined (2023). doi: 10.48550/arXiv.2305.10696
24. Unbiased Gradient Boosting Decision Tree with Unbiased Feature Importance. doi: 10.48550/arxiv.2305.10696, 2023
25. Ramakrishna, Madakasira, Prasanth Babu, Girish Wali, Deepali Tuteja, I. I. Spannhake, Ernst Wilhelm, Thomas Gianelle, and Milan Shah. "SYSTEMS AND METHODS FOR RESPONDING TO PREDICTED EVENTS IN TIME-SERIES DATA USING SYNTHETIC

PROFILES CREATED BY ARTIFICIAL INTELLIGENCE MODELS TRAINED ON NON-HOMOGONOUS TIME-SERIES DATA." U.S. Patent Application 18/418,201, filed June 13, 2024.

26. Tuteja, Deepali, W. A. L. I. Girish, and David Anandaraj Arulraj. "Aggregating data ingested from disparate sources for processing using machine learning models." U.S. Patent Application 18/123,179, filed September 19, 2024.
27. Kumar, Amit, Neha Sharma, Nikhil Kumar Marriwala, D. Nageswara Rao, Subhendu Kumar Behera, Girish Wali, and Bharti Jandwani. "Mapping the trajectory of blockchain technology and non-fungible tokens: a comprehensive bibliometric analysis." *International Journal of Information Technology* (2024): 1-8.
28. Wali, Girish, and Chetan Bulla. "A RESOURCE MANAGEMENT MODEL FOR FINANCE DOMAIN IN FOG COMPUTING BASED IOT INFRASTRUCTURE USING DEEP LEARNING AND BIO-INSPIRED ALGORITHMS."VOL 08, ISSUE 05,JOURNAL OF CRITICAL REVIEWS, 2021
29. Dayana, D. S., T. S. Shanthi, Girish Wali, P. V. Pramila, T. Sumitha, and M. Sudhakar. "Enhancing Usability and Control in Artificial Intelligence of Things Environments (AIoT) Through Semantic Web Control Models." In *Semantic Web Technologies and Applications in Artificial Intelligence of Things*, pp. 186-206. IGI Global, 2024.
30. Tuteja, Deepali, WALI. Girish, and Prasanth Babu Madakasira Ramakrishna. "Systems and methods for determining allocatable resources during proportional maintenance of complex computing systems using bifurcated filtering." U.S. Patent 12,118,389, issued October 15, 2024.
31. Bulla, Chetan, and Mahantesh N. Birje. "Multi-agent based Monitoring System for Fog Computing Environment." In *2021 2nd International Conference for Emerging Technology (INCET)*, pp. 1-6. IEEE, 2021.
32. Bulla, Chetan M., and Mahantesh N. Birje. "Efficient Resource Management Using Improved Bio-Inspired Algorithms for the Fog Computing Environment." *International Journal of Cloud Applications and Computing (IJCAC)* 12, no. 1 (2022): 1-18.
33. Bulla, Chetan, and Mahantesh N. Birje. "Improved data-driven root cause analysis in fog computing environment." *Journal of Reliable Intelligent Environments* 8, no. 4 (2022): 359-377.
34. Surjeet, S., C. Bulla, A. Arya, S. Idrees, G. Singh, S. G. Rao, and A. Shirahatti. "A quantum machine learning approach for bridging the gap between quantum and classical computing." *International Journal of Intelligent Systems and Applications in Engineering* 12 (2024): 553-560.
35. Kori, Anita, and Chetan Bulla. "AI-Powered Real-Time Analytics for Liquidity Risk Assessment in Banking Sector." (2024).