Bol. Soc. Paran. Mat.
©SPM – E-ISSN-2175-1188
SPM: www.spm.uem.br/bspm

(3s.) **v. 2026 (44) 8** : 1–15.
ISSN-0037-8712
doi:10.5269/bspm.83916

Enhancing Loan Approval Decisions through Machine Learning Algorithms

M. Aruna Prasanna Kumari and Shaiku Shahida Saheb

ABSTRACT: This study presents academic state-of-the-art current research on using Machine Learning to improve the predictability of bank loan approval procedures for financial inclusion. This study provides solutions to several issues persisting in traditional methods of evaluating credit with standard credit assessment procedures, including embedded bias, imprecision, and limited scalability across levels of socioeconomic status. The methodology used in this article utilised a significant workflow to analyse the machine learning process includes preparing the data, exploring it, creating new features, and system training and testing of nine classification models, including both classical and ensemble learning methods. The outcomes specify that ensemble learning models (Random Forest, Gradient Boosting, XGBoost and LightGBM) are capable of providing very high accuracy with 98% accuracy, precision, recall, and F1 scores. By including additional data sources and using explainable AI, ensemble models can demonstrate compliance with regulatory requirements through their integrity, transparency, and consistency. The findings of this research show how ensemble models can reduce classification errors, lower operational costs, and improve decision-making. Therefore, they have the capability of decreasing operational costs and improving decision making may enable people in underserved communities to have more access to financial services. The research highlights new ways that machine learning can support greater financial inclusion and represents a scalable, trustworthy, and responsible approach for leading banking institutions to broaden their ability to reach out to those who are not currently benefiting from the availability of credit.

Keywords: Financial inclusion, machine learning model, ensemble learning model, performance metrics, confusion matrix.

Contents

1 Introduction 1

2 Methodology 3

3 Machine Learning Models 6

4 Results and Discussions 12

5 Conclusion 13

1. Introduction

Financial Inclusion (FI) is the new priority for all governments' policies for finance and banking. FI refers to the equal opportunity for everyone to have access to banking products, especially through credit [6]. Banking products support the ability of individual consumers and small businesses to engage in economic activities to be resilient against future financial overwhelming and develop their quality of life through the use of banking services [5]. After many years of government policy reform and advancements in technology, a large number of potential borrowers are still unable to access the formal banking sector. Many reasons exist why borrowers are still unable to access the banks: some of the issues relate to inefficiency, discrimination, and an inability to evaluate the risks and uncertainties associated with lending to potential borrowers [14]. With the increased demand for data-driven and systematic evaluations of borrowers' creditworthiness, banking institutions are using intelligent computational tools and techniques for assessing borrower creditworthiness to promote and implement a more inclusive banking industry [3]. Given that ML is becoming a popular method for banks to improve their decision-making evolution processes, creators can expect to see improvements in the accuracy of predictions and the rate of FI [10]. Humans are still making most of the decisions regarding bank lending based on their limited financial indicators, such as income type, job category, credit score, and collateral [1].

2020 *Mathematics Subject Classification:* 68T05, 91G40.
Submitted May 02, 2026. Published June 05, 2026.

There is a lot of value in all these evaluations, but they can all vary considerably, be subjective and take time, as noted by [13]. The traditional method of assessing creditworthiness often can limit or marginalize people who do not have a written financial record, such as low-income workers, those living in rural areas, or the informal economy [25]. The continued presence of a digital divide between financially privileged consumers and underserved consumers has arisen as a result of the limitations of current systems to provide credit access to low-income consumers [22]. Therefore, systematising and improving the loan approval process via ML is not only a technological improvement it is also an imperative plan that promotes the democratization of access to credit and provides meaningful financial inclusion [7].

Machine Learning (ML) acquires a broad range of databases [16] and applies them to create actionable perceptions resulting from the predicted repayment behaviour of debtors and the consequent result of the loan [8]. By analysing historic loan results and debtors' features, as well as other aspects [15] ML can originate predictive modelling with a diversity of constraints, inclusive of loan amount pursued, applicant's income, credit history and demographic data; hence, it can find intricate and non-linear contacts that are certainly external the realm of social assessors [18]. The ability to leverage ML for assessing loans is a significant benefit for banks since it can help reduce default risk and operational costs, while improving the transparency in the decision-making [4]. The rapid expansion of financial systems in emerging economies has generated a pronounced demand for inclusive and efficient allocation of funds through diverse mechanisms [23]. Financial institutions find themselves grappling with these competing goals: making money and remaining committed to inclusive lending [18]. Conventional credit risk evaluation tools have historically preferred conservative tendencies, causing an increase in loan rejection rates against borrowers with unorthodox credit records [28]. To be more specific, machine learning algorithm-inspired tools offer credit institutions an opportunity to incorporate more variables into credit risk determination: transaction behaviour, habits, employment stability, among others. As a result, they gain better insights into credit risk instead of solely relying on credit records and creditworthiness [17]. Some credit institutions use ML algorithms because they can identify people who would honour credit constraints, even if they have not yet shown high creditworthiness qualitatively [26]. Overall, ML algorithm-inspired credit risk tools allow credit institutions to rely on technology as an instrument of inclusion within society and on an economy level. Among various ML algorithm-inspired tools for credit risk, LR, NB, SVM, KNN, and DT algorithms, as well as RF, GBM, XGB, LightGBM, are recognized and highly successful at loan approval and credit defaults prediction [24]. These models vary in both the way they learn from data and how transparent they are, which gives practitioners the freedom to pick a single method or mix several to hit the sweet spot of predictive performance [12]. Integrating these algorithms into a comprehensive predictive framework can substantially improve loan decision accuracy and mitigate human bias [21].Traditional banking institutions still struggle with conventional decision-making processes, even though fintech companies have come up with sophisticated data-driven tools [2]. Manual screening remains resource-intensive and lacks consistency across evaluators, often leading to long-time approvals and a poorly performing loan approval portfolio [20]. Furthermore, an overreliance on human discretion exposes the loan approval process to psychological biases, including favouritism and risk aversion [9].

Machine Learning algorithms play a key role in automating and standardising credit risk assessment. This is required to build an efficient and just lending ecosystem [27]. Implementing ML-driven loan approval systems directly promotes financial inclusion by making credit more accessibility of credit to underrepresented populations [19]. Properly designed ML systems will identify qualified lenders that traditional scoring methods have missed [29]. In addition, the combination of interpretability with ML models ensures that decisions made by the algorithms can be held accountable and remain transparent, thus allowing for the enforcement of ethical standards and compliance requirements in the financial services sector [11]. In developing an all-encompassing machine learning model for Financial Inclusion, through the provision of loan approvals to banks, this represents a harmonious relationship between both technological innovations and social responsibility [15]. While there has been much investigation into using Machine Learning techniques for loan approvals, the majority of this research has looked only at the performance of different models compared to one another and not at the fairness and explainability components that are needed when making credit decisions responsibly and inclusively. Previous studies do not integrate bias mitigation/evaluation with interpretability of machine learning (ML) models, thus failing to demonstrate how algorithmic decisions can promote financial inclusion of underrepresented borrowing groups. Further, there is a lack of evidence demonstrating how Explainable Artificial Intelligence (XAI) techniques can be implemented into operationalised credit risk decision-making processes in order to achieve a balance between accuracy, fairness and transparency.

To bridge the above research gaps, this study makes 1) a fairness-aware explainable ML pipeline for automated loan approval that integrates ffairness mtrics (demographic parity differences, disparate impact, equal opportunity difference and SHAP-based model interpretability within the model evaluation process 2) provides empirical evidence showing that the proposed framework can reduce fairness disparity while maintaining high predictive performance, demonstrated through nine ML algorithms compared under stratified cross-validation 3) a practical deployment oriented evaluation, including calibration, robustness checks and sensitivity analysis, to support real-world implementation in banking and financial inclusion context 4) presents Interpretable feature-level insights that highlight the most influential drivers affecting credit approval outcomes, enabling transparent and accountable decision making.

This study will develop and analyse machine learning models aimed at improving customer access to financial services by optimising loan approval processes. The research will investigate how the use of classical machine learning (ML) algorithms, as well as combinations or ensembles of these algorithms, can maximise prediction accuracy and fairness when assessing loan eligibility. ML algorithms, including Logistic Regression, Decision Trees, Random Forests, Support Vector Machines (SVM), K-Nearest Neighbours (KNN), Naïve Bayes classifiers, AdaBoost, and Gradient Boosting, will be used to identify trends or patterns within applicant data that may predict the applicant's ability to repay the loan or to default on it. This research aims to create a reliable, interpretable, and scalable model to assist banks with automating the evaluation of credit applications and promoting inclusive financial growth.

This study will answer three major research questions: 1) Which ML algorithms yield the best predictive performance for loan approval? 2) Do ensemble-based models provide a statistically significant improvement in predictive performance over classical-based models? 3) How can improved prediction accuracy lead to the likelihood of responsible and sustainable lending practices? The major contributions of this study are a detailed side-by-side comparison of nine ML techniques using cross-validation, a repeatable end-to-end data analysis workflow, and a discussion of the social and ethical responsibilities of increasing access to financial services through the use of technology.

2. Methodology

Data Protocol

:

The purpose of this research is to create an evidence-based framework for credit assessments to improve accuracy, transparency, and fairness. The framework consists of several stages that include the exploration and the exploratory data analysis (EDA) process, identifying the EDA through missing values, outliers, and patterns; the pre-processing of data on the EDA so that the predictions are made based on reliable pre-processed data; the Transformation and Selection of Variables (Feature Engineering) uses, through fair and reliable data sources, calls for transforming and selecting significant variables; the Training and Testing of Machine learning (ML) algorithms (Trained MLs) or Ensemble Algorithms, based on performance measures of accuracy, precision, recall, F1-score and Finally Model Validation and Testing to measure the robustness and fairness of the model to be able to generalise to various socio-economic segments. This framework allows not only for increased predictive accuracy of loan approvals but also promotes financial inclusion by creating equitable and data-driven credit decisions.

Unlike conventional model comparison studies, this study introduces a novel end-to-end pipeline integrating performance, fairness and explainability in a unified evaluation framework. The pipeline includes preprocessing, imbalance handling, cross-validation model training, fairness evaluation using group fairness metrics and SHAP explainability to analyse feature-level division drivers. This enables systematic analysis of the accuracy fairness trade-off in automated lending applications.

The proposed work for Bank Loan Approval Prediction

:

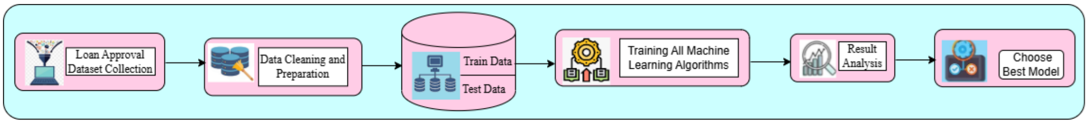

Figure 1: Block diagram of the proposed work

Figure 1 depicts a whole end-to-end machine learning process for loan approval prediction, beginning with the collection of the loan approval dataset and progressing through data cleaning and preparation to ensure data quality. The processed data is then split into training and test sets, enabling the construction and evaluation of a range of machine learning techniques, including Logistic Regression, Decision Trees, Random Forests, and sophisticated boosting models. After training, the models are subjected to extensive result analysis utilising measures such as accuracy, precision, recall, F1-score, confusion matrix, and ROC/AUC to evaluate their prediction ability. After a thorough evaluation, the best-performing model is chosen to provide dependable and accurate loan approval predictions.

Data collection for this project includes a Kaggle dataset on loan prediction https://www.kaggle.com/datasets/architsharma01/loan-approval-prediction-dataset. The dataset contains a number of important borrower attributes that affect decisions about loan approval. While family size indicates the number of dependents that may affect disposable income, education reflects the applicant's academic background and potential financial awareness. Employment status, whether self-employed or employed by a company, helps determine income stability, and the applicant's ability to repay is directly measured by their annual income. The borrower's financial commitment and repayment period are specified by the loan request amount and loan term. An unbiased assessment of prior credit behaviour and general creditworthiness is offered by the government credit score. Real estate, business, luxury, and bank assets also demonstrate the applicant's stability, strength, and capacity to support repayment or offer collateral, all of which are important factors in determining loan approval eligibility.

Selection of the Best Model Processing

:

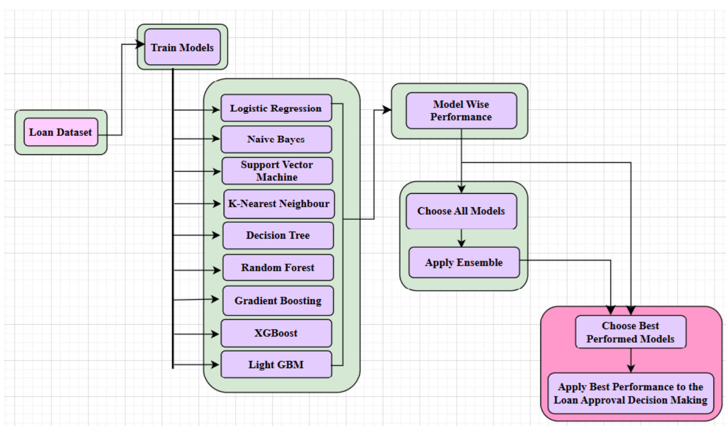

Figure 2: Process of choosing the Best Model

Figure 2 mentioned below, demonstrates using a loan dataset, the system trains several classification models, including logistic regression, decision trees, random forests, SVM, and ensemble techniques. The system then evaluates all models, applies ensemble methods with all models or with the top three performers, and finally chooses the configuration that offers the best predictive performance for deployment in a user interface. This process ensures accurate and effective loan approval decisions through rigorous model selection and combination strategies, as depicted in the diagram.

Data Exploration

:

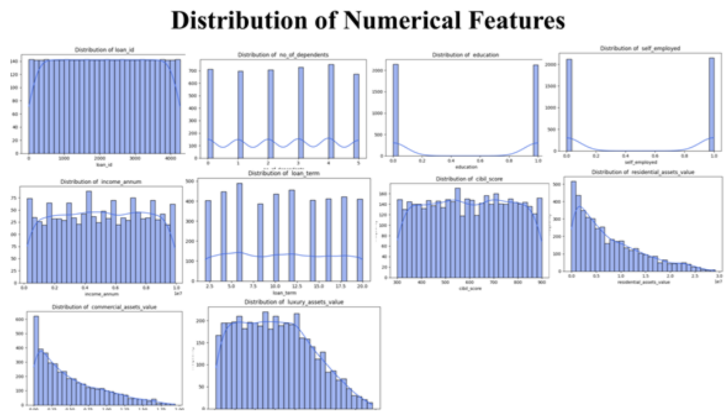

Figure 3: Numerical Feature Distribution for Loan approval

Figure 3 shows a visual representation of the 12 numerical features and their marginal distributions, which are presented in the form of histograms and trendlines that have been created using Kernel Density Estimates (KDE). Certain features, like Loan ID and CIBIL score, have distributions that are either approximately uniform (Loan ID) or normally distributed (CIBIL score). The other features, including Residential Assets Value and Commercial Assets Value, have distributions that exhibit a strong right skew, meaning they contain a significant number of observations with low values and then have long tails of observations at higher values. In addition to these distributions, Education and Self Employed have distributions that, while displayed on numerical axes, contain a large number of observations in the same two integer values (0 and 1). Therefore, they are likely ordinal or binary encoded variables that need to be treated differently when building models.

Distribution of Categorical Features

:

The box plots generated in figure 4 illustrate numerically relevant attributes across the two categories of the response variable 'Loan Status' (0 = Approved or Paid, 1 = Rejected or Default) so that these data can be compared visually. When comparing the boxes, there are major differences between the distributions for each category that would indicate they potentially contain information that could predict an outcome. For example, the 'Cibil Score' has the highest median and tightest range of values for those with loan status 0 compared to those with loan status 1. This indicates that the Cibil Score can likely differentiate between these two classifications very well. In contrast, 'Loan ID' and 'Number of Dependents' provide similar medians and ranges between categories, and they therefore may offer less of a predictive distinction. The various asset value attributes (e.g. Residential Assets, Commercial Assets) produce boxes with overlapping ranges between categories; however, some have distinctly different medians, which may warrant additional statistical analyses to assess their validity as predictors of Loan Status.

Correlation Heatmap

:

Figure 5 displays the correlation matrix of the numerical features of the dataset. It allows for an examination of how numerically related variables are to one another through a graphical representation. Some of the most apparent linear correlations are the strong positive correlations seen among multiple financial characteristics. For example, the income per annum variable has a very high positive correlation with the loan amount variable (r=0.93) as well as the bank asset value variable (r=0.79), indicating that the greater an individual's level of income, the larger the loan amount they tend to apply for and

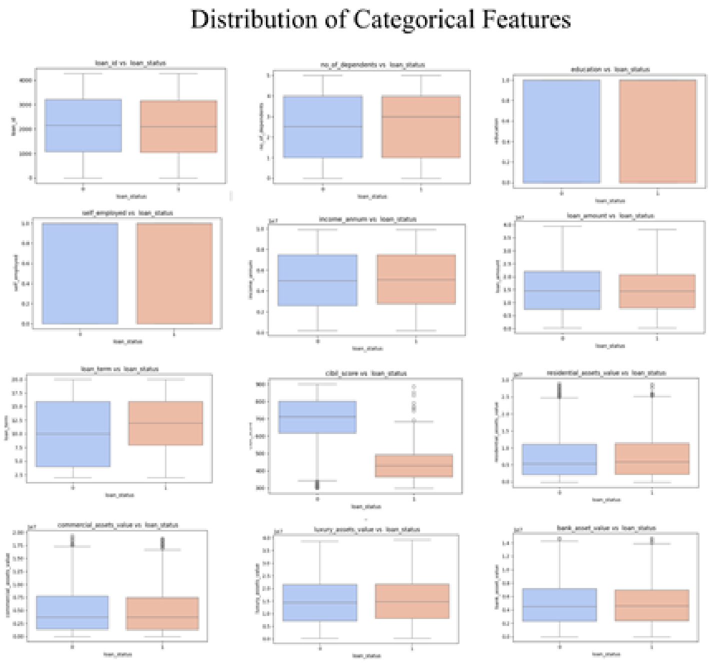

Figure 4: Categorical Features Distribution of Loan Approval

the more bank assets they tend to own. Likewise, the loan amount variable has very high positive correlations with both the bank asset value variable (r=0.79) and the residential assets value variable (r=0.64), suggesting that ownership of assets can influence how much borrowing one can do. Conversely, the credit score variable has a very high negative correlation with the loan status variable (r=-0.77), indicating that individuals who have a high level of credit are more likely to be successful in paying back their loans than individuals who have lower credit scores or are more likely to default. There are relatively weak correlations between the education/self-employed demographic features and the financial attributes included in this dataset. Therefore, it can be inferred that education/self-employment has a lesser impact when predicting the financial behaviours of the individuals encompassed in this dataset. In conclusion, the heatmap demonstrates that while many financial attributes are related to one another, the CIBIL score is particularly vital to predicting the repayment performance of loans within this dataset.

3. Machine Learning Models

Nine Machine Learning models for predicting loan approvals are compared, and as shown in Table 1, tree-based ensemble models have higher accuracy levels compared to traditional models. Classical models such as Logistic Regression, Naïve Bayes, and SVM showed accuracy levels of 0.91 to 0.94, while the accuracy level of KN was 0.88, which is the lowest of all the models. Using decision tree techniques yields a noticeable improvement in accuracy, starting with the Decision Tree at 0.97, while the accuracy levels of Random Forest, Gradient Boosting, XGBoost and Light GBM were the highest with all having accuracy levels of 0.98. The best overall performer was LightGBM with an accuracy of 0.98 and an F1 score of 0.985. This indicates that LightGBM has the best ability to model the complexities of the data. The top ensemble models have very high levels of precision and recall, which indicates that they accurately identify the most applicants that are eligible and accurately minimise the number of erroneous approvals. In conclusion, the top performers among the Gradient Boosting ensemble models (LightGBM, XGBoost) are the most effective and consistent performing models for use in automated loan approval prediction.

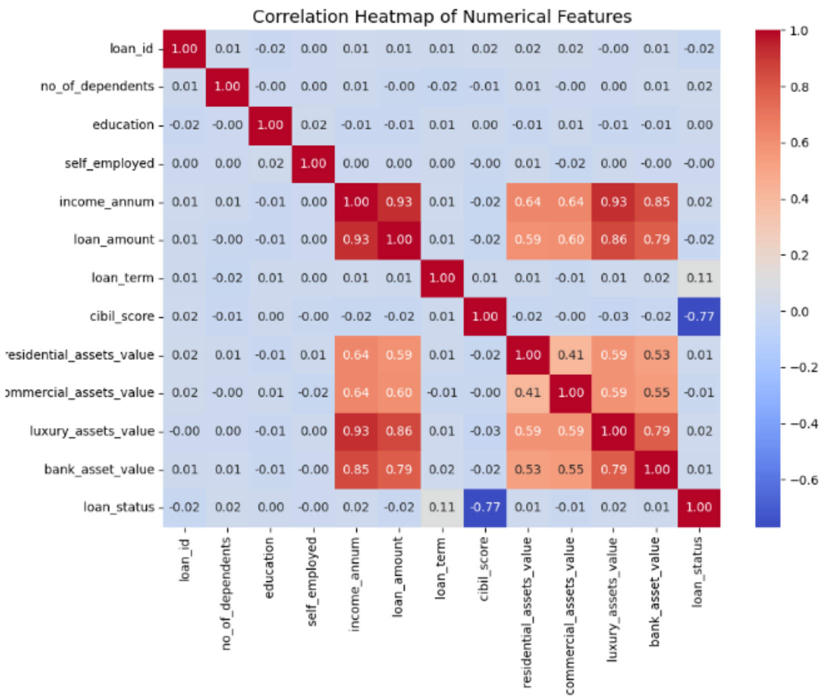

Figure 5: Correlation Heat Map based on Number

Accuracy Evaluation for all models

:

The bar chart, titled "Model Performance Evaluation," as shown in Fig. 6 presents a comparative assessment of various machine learning algorithms across four main metrics-Accuracy, Precision, Recall, and F1 Score-on the problem of bank loan approval prediction within the context of financial inclusion. The findings reveal that the ensemble-based classifiers consistently outperform traditional models in all evaluation metrics, including Random Forest, Gradient Boosting, XGBoost, and LightGBM, with an approximate value of 0.98 for each. This uniformity in performance reflects their capability to balance precision and recall effectively, thereby minimizing misclassification and improving the reliability of decisions related to financial risk assessment. The Decision Tree model follows next with 0.97 for each metric, which justifies its position as a base learner for most ensemble methods. However, conventional models such as SVM, Naive Bayes, and Logistic Regression register moderate accuracies ranging between 0.91 and 0.94, while the KNN algorithm exhibits the poorest performance at 0.88, indicating inadequacy with non-linear, high-dimensional data. Overall, the analysis substantiates the superior generalisation and robustness of tree-based ensemble algorithms, with particular emphasis on Gradient Boosting, XGBoost, and LightGBM, highlighting their suitability for accurate and unbiased data-driven modelling of loan approval for financial inclusion.

Confusion Matrices for all Models

:

Model 1: Logistic Regression (LR)

:

Logistic regression remains one of the fundamental models for binary classification tasks, including loan approval. By applying a logistic function to the estimated probability of loan acceptance from input variables, it produces outputs in the unit interval [0, 1]. Interpretability of the model allows researchers to estimate the marginal contribution of each variable, such as income or credit history, to the decision

Table 1: Model Performance Summary

Model Name	Accuracy	Precision	Recall	F1 Score
Logistic Regression	0.91	0.907	0.907	0.907
Naive Bayes	0.93	0.935	0.934	0.935
SVM	0.94	0.937	0.936	0.936
KNN	0.88	0.882	0.881	0.881
Decision Tree	0.97	0.974	0.974	0.974
Random Forest	0.98	0.979	0.979	0.979
Gradient Boosting	0.98	0.976	0.975	0.975
XGBoost	0.98	0.979	0.979	0.979
LightGBM	0.98	0.985	0.985	0.985

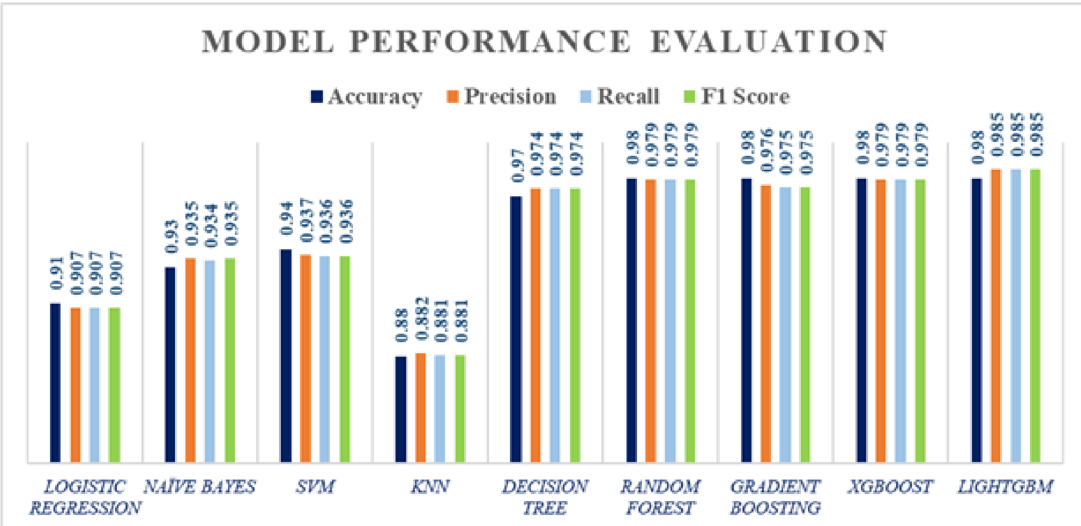

Figure 6: All Model Performance Summary

on a loan. Despite its simplicity, logistic regression assumes a linear relationship between predictors and log odds, which limits its potential in nonlinear, complex contexts. Yet, transparency and ease of implementation make it a necessary baseline in finance modelling and a typical benchmark for evaluating more advanced models.

Figure 7(a) shows the confusion matrix obtained when assessing the performance of the Logistic Regression model for binary classification. The matrix suggests that the model has correctly classified 500 instances of class 0 (true negatives) and 275 instances of class 1 (true positives), while misclassifying 36 instances of class 0 to class 1 (false positives) and 43 instances of class 1 to class 0 (false negatives). A trend from these results is that there was strong overall predictive accuracy in the two classes, with slightly higher accuracy in identifying negative cases than identifying positive ones. The low counts of false positives and false negatives suggest that the model will keep a balanced trade-off between sensitivity and specificity. In an academic setting, these results are evidence of a well-calibrated Logistic Regression model that should be capable of distinguishing effectively between the two classes, which is of particular importance when the tasks at hand are, for example, credit risk assessment or loan approval prediction, both of which bear meaningful implications with each type of misclassification.

Model: 2 Naïve Bayes (NB)

:

The Naïve Bayes classifier applies Bayes' theorem by assuming independence of features, which reduces

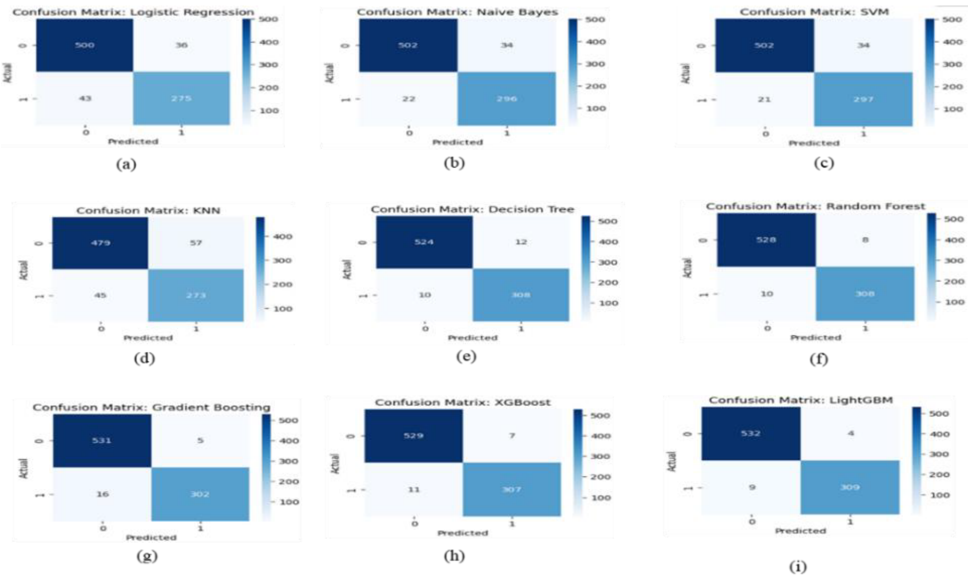

Figure 7: Confusion Matrices

computational load but still produces probabilistic output. Although financial attributes in real-life scenarios are often correlated, Naïve Bayes performs well with moderately sized data and offers a fast and interpretable solution for preliminary screening. It is particularly useful in developing wide trends about applicants and hence constitutes a useful component in multi-model ensemble models.

Figure 7(b) Confusion Matrix for Naïve Bayes, as shown above, depicts the performance of the model with respect to binary outcome prediction. The model correctly predicted 502 class 0 instances as true negatives and 296 class 1 instances as true positives, while misclassifying 34 class 0 instances as class 1 (false positives) and 22 class 1 instances as class 0 (false negatives). This outcome indicates that the Naïve Bayes model can yield high predictive accuracy and prove to be very efficient at identifying both positive and negative classes correctly. Given the total numbers involved, coupled with an even lower magnitude of misclassifications, especially false negatives, it should be noted that a model like this would capture the probabilistic relation lying beneath the data quite well. Measured against typical benchmarks, the models have performed well in establishing a correct balance between sensitivity and specificity, enjoying a slight edge in the proper identification of the negative class. Essentially, the confusion matrix shows the robustness and efficiency of the Naïve Bayes classifier when dealing with this dataset, making it very reliable in any application based on categorical or probabilistic decision-making processes, such as credit risk evaluation or customer classification.

Model: 3 Support Vector Machine (SVM)

:

Support Vector Machine is a robust and non-parametric technique that develops the best separating hyperplanes within the multi-dimensional space to classify approved and disqualified applicants. Its structure, based on a kernel, allows it to identify nonlinear boundaries. This characteristic proves very useful in credit risk analysis because credit approval or rejection can be determined with a high level of accuracy using a smaller sample size. However, with larger datasets, there are computational complexities

The confusion matrix depicted in Fig. 7(c) illustrates an SVM classification result on a binary classification problem. It shows that out of the total number of true negatives, 502 were predicted as belonging to class 0, with 34 predicted as belonging to class 1. Also, out of the total number of true positives, 297 were predicted as belonging to class 1, with 21 predicted as belonging to class 0. From a research perspective, confusion matrices are very informative for diagnostics with regard to SVM performance. The fact that there are fairly low numbers for false positives and false negatives, with 34 and 21, respectively, would indicate a good trade-off for sensitivity and specificity. It would appear there are high levels of sensitivity and specificity, given that there are high numbers of true negatives and positives, reflecting on the success of the algorithm in suppressing type I and type II errors.

Model: 4 k-Nearest Neighbors (KNN)

:

K-Nearest Neighbors is another instance-based classifier that categorises new applicants by considering their feature profile similarity with the already labelled examples present in the data. The assumption made is pretty simple: applicants whose financial characteristics resemble each other tend to show similar behaviours in terms of repayment. The performance of KNN depends on factors such as how the data is scaled, what distance metric is used, and what value of k is chosen. It gives intuitive interpretability but does not scale as well to very large financial datasets.

Figure 7(d) shows the confusion matrix for KNN; it gives a complete picture of how the algorithm performs on a binary decision task. It shows that 479 instances were correctly identified as true negatives and 273 as true positives, while on the other side, it mislabeled 57 actual class 0 cases as class 1 and 45 actual class 1 cases as class 0. These numbers are indicative, from an academic point of view, of the fact that KNN is able to separate the two classes quite reasonably, even though its discriminative power is not as strong as that of more robust algorithms. These higher numbers of false positives and false negatives, compared to true positives and true negatives, respectively, point toward limitations with regard to sensitivity and specificity, signalling a moderate trade-off between Type I and Type II errors. More concretely, these misclassifications will visibly affect the true negative rate and the true positive rate and, therefore, the specificity and sensitivity of this algorithm, respectively, showing that there is definitely a need for optimisation of hyperparameters or feature engineering for the KNN model in this context. This would, however, be essential for making any recommendations on improvements and validation of the general classification chain.

Model: 5 Decision Tree (DT)

:

Decision Trees (DTs) are a type of algorithm used for classification problems. Decision Trees are able to partition data hierarchy based on the feature thresholds, which allows for the creation of an interpretable set of rule-based decision paths. The internal nodes of the tree represent the decisions an algorithm makes to assess the input (the features) and determine a result, while the terminal nodes represent the various outcomes or predictions, i.e loan approved or loan denied. While Decision Trees provide visual transparency and ease of interpretation, they can also suffer from issues like over-fitting, particularly when working with datasets with high levels of noise or class imbalance. A common way to mitigate these issues is through Pruning or using Ensemble-based approaches such as Random Forest or Extra Trees.

The confusion matrix for model 5 Fig. 7(e) provides insight into the efficacy of a Decision Tree classifier when applied to a binary classification scenario. As indicated by the values presented in the matrix, 524 true-negative events were identified correctly by the Decision Tree, 308 true-positive events were identified correctly; however, the model misclassified only 12 actual class 0 cases as class 1 and 10 actual class 1 cases as class 0. The extremely low number of misclassifications demonstrates that the Decision Tree achieved both a high sensitivity and specificity on this particular dataset and thus, from an academic point of view, suggests that the model has excellent discrimination power and is very effective in reducing Type I and Type II errors related to this classification task. In addition, the relatively small differences seen between the diagonal entries and off-diagonal entries of the confusion matrix demonstrate that the model performs well at correctly dividing the feature space and generalising to unknown data points. Therefore, the Decision Tree appears to be a highly dependable classifier for the dataset under consideration and should serve as a reference for future research and model creation.

Model: 6 Random Forest (RF)

:

A Random Forest is an ensemble model that creates several decision trees, each of which has been trained using random samples of the data and features from the data. The final output of all trees is used to decide the predicted output for a new input through majority vote. Using the RAM allows for the reduction of variance and provides greater stability to the model. Additionally, the Random Forest can capture the complex interaction effects of many financial variables as they relate to the level of creditworthiness. Also, the Random Forest generates importance values for features, allowing the analyst of a financial institution the ability to determine which financial variables are most responsible for a customer's creditworthiness. Because it can balance between interpretability and accuracy, the Random Forest is very popular among banks and other types of financial institutions.

Figure 7 (f) presents the confusion matrix for the Random Forest classifier, detailing the model's performance when used as a binary classifier. In the confusion matrix, there are 528 true negatives, 308 true positives, 8 false positives, and 10 false negatives reported. The classifiers achieved an overall accuracy of 97.9%, a precision for the positive class of 97.4%, and a recall on the positive class of 96.9%, providing a strong indication that the Random Forest can differentiate accurately between the two classes, reducing both type I and type II error rates. These metrics confirm that the Random Forest can provide a high degree of reliability and suitability for those fields that require accurate and reliable classification results.

Model: 7 Gradient Boosting

:

Gradient Boosting: Gradient Boosting extends the technique of boosting and refines residuals from previous models via gradient descent. It works well as it adapts a learner on the gradient of a loss function. GB models perform well with financial datasets with complex interactions among variables but require proper parameter tuning to avoid overfitting. Due to adaptability and accuracy, these models are preferred for structured finance problems, such as the prediction of loan defaults and approvals.

Figure 7(g) below shows the confusion matrix for the Gradient Boosting classifier, indicating significant predictive power for a binary classification problem. From the matrix, there were 531 correct assignments of True Negatives and 302 correct assignments for True Positive instances, with a significantly low 5 instances for False Positives and 16 instances for False Negatives. These observations demonstrate that the classifier manages to make distinctions between these two classes with excellent accuracy and very low levels of misclassifications for negatives as well as positive instances. From a statistical perspective, it can be seen that precision, recall, and accuracy statistics from this confusion matrix demonstrate the credibility and reliability of this classifier. The extremely low rate of false positives reveals that it is significantly effective at preventing wrong assignments within positive classifications. Misclassifications are extremely costly within specific realms. The low rate on false negatives, on the other hand, implies that it performs remarkably well at capturing most instances within positive classifications, thus preventing potentially costly misses within these classifications.

Model: 8 XG Boost

:

XGBoost is an exceptionally potent and efficient algorithm that can be employed for classification or regression problems. The methodology of creating an Ensemble of Decision Trees is sequential. Every newly created tree tries to correct the errors of the trees before it. Traditional methods of boosting do not provide for regularisation. XGBoost incorporates a variety of regularisation techniques to limit the potential for "over-fitting" and to improve generalisation. The algorithm also optimises for speed and memory through a variety of techniques, including parallel processing, tree pruning, and handling sparse data. As a result of its superior predictive success and ability to scale with the need for large amounts of data, it remains one of the most favoured choices of Data Scientists competing in Applied Data Science and in Practical Data Science tasks such as Credit Scoring, Fraud Detection, and Recommendation Systems.

The XGBoost classification algorithm exhibits remarkable discriminatory capabilities, as evidenced by the confusion matrix seen in Figure 7 (h), with an accuracy rate of approximately 97.89%. The classification model performs very well against the actual classes, with 529 instances of true negatives and 307 instances of true positives. Notably, there are very low instances of misclassifications, with 7 instances of false positives, making the precision 0.9777. There are also very low instances of false negatives, at 11 instances, making it very supportive of a high-recall value of 0.9654. It thus shows remarkable capabilities as a model for identifying actual positives within the given data.

Model: 9 Light GBM

:

LightGBM, which stands for Light Gradient Boosting Machine, is an efficient and scalable ensemble learning algorithm based on gradient boosting decision trees. Its popularity arises from its good prediction accuracy and efficient computation. Unlike traditional boosting methods, LightGBM adopts Gradient-based One-Side Sampling and Exclusive Feature Bundling, two specially designed methods that reduce memory cost and accelerate training without compromising accuracy. Due to these capabilities of LightGBM, it can be very useful for handling large datasets and online prediction, such as financial inclusion, risk evaluation, and credit classification. By giving more weight to more informative data points and features, LightGBM tries its best to accelerate learning with guaranteed accuracy. Its applicability and strong performance have empowered it to be effectively used in several domains, including supply chain management, genomics, and malware classification.

The LightGBM classifier is depicted in Fig. 7(i), and it performs outstandingly on this classification task with an accuracy of approximately 98.48%. The confusion matrix shows its excellent discriminative capabilities with 532 True Negatives (TN) and 309 True Positives (TP). The rate of misclassifications is extremely low, with 4 False Positives (FP) and 9 False Negatives (FN), it achieves a Precision of approximately 0.9872 and a Sensitivity or Recall value of about 0.9717, signifying highly trustworthy positive predictions and an excellent collection of the positive class. Overall, these parameters make LightGBM a competent and superior-performing method on the classification task.

4. Results and Discussions

The comparison of approved and rejected loans depicted in the chart displays a clear majority of loans that were approved and clearly shows the skewed class distribution of loans approved versus loans rejected. More so, there is a total loan sample of approximately 4,200, with almost 2,600 of those loans being approved and roughly 1,600 of those loans being rejected as shown in Fig. 8. Thus, approximately 62% of the loan applications were approved, with 38% of the loan applications rejected. One reason the financial institution has such a large percentage of loan approvals may be a result of better applicant screening or more relaxed lending criteria, showing that it is likely that the bank is having success in attracting new customers. If one were to use this data in order to create models for loan approval prediction purposes, then the overwhelming collection of approved loans may cause a bias within the created loan prediction models unless either oversampling and/or undersampling techniques were utilised to create such models.

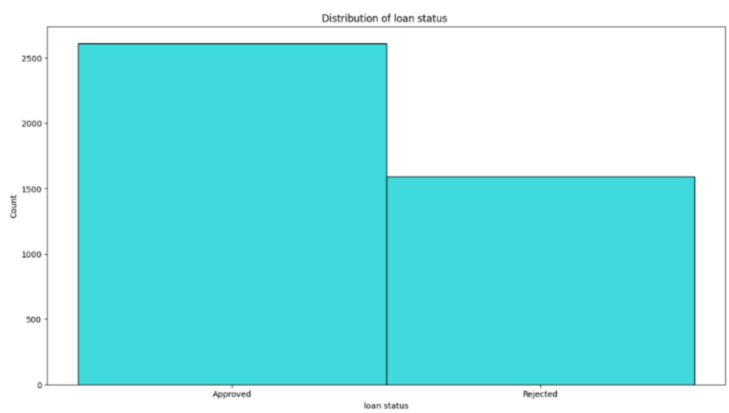

Figure 8: Approved and rejected Loans comparison

The results of the present study demonstrate how powerful machine learning could be in automating loan decisions in banks, especially when we use ensemble methods. Nine classification algorithms were put to the test: Logistic Regression, Naive Bayes, SVM, KNN, Decision Tree, Random Forest, Gradient Boosting, XGBoost, and LightGBM. These metrics were all computed for each model: accuracy, precision, recall, and F1-score. The ensemble methods shown here include Random Forest, Gradient Boosting, XGBoost, and LightGBM, which consistently outperform the classical models by reaching about 98%, proving strong across all metrics. For instance, Random Forest reached 0.98 accuracy, while its precision, recall, and F1-score were all over 0.97. LightGBM has the highest accuracy at 0.985. Their confusion matrices showed very few misclassifications, while false positives and false negatives were much fewer than the traditional algorithms. That is to say, this means a solid ability to correctly separate eligible from ineligible applicants by reducing both Type I and Type II errors. The Decision Tree performed equally well with 0.97, but was not that robust compared to the ensemble methods. On the other extreme, Logistic Regression, Naive Bayes, and SVM had a moderate accuracy with readings between 0.91 and 0.94, while KNN had the worst performance at an accuracy of 0.88, which is probably due to the struggles of this algorithm with non-linear and high-dimensional data. The superior performance of the ensemble models is related to their ability to capture complex interactions amongst financial indicators and reduce variance by combining multiple learners. On the whole, the findings point out that ensemble learning can substantially enhance the efficiency, reliability, and inclusiveness of automated loan approval systems while improving financial inclusion by expanding credit access to underserved groups.

5. Conclusion

This research further supports the important role that machine learning, in particular the ensemble learning methods, plays in providing financial access for people through automating and optimising the way banks approve loans. A thorough investigation into several different types of classifiers found that the use of an ensemble classifier method [6] was much more accurate, reliable, and generalisable than traditional classifiers such as logistic regression, naive Bayes, and SVMs. Therefore, ensemble classifiers also allow financial institutions to decrease the number of incorrect classifications of loan applicants as well as reduce their operating expenses, allowing them to grant greater access to credit to underrepresented consumers, particularly those who have had limited formal banking relationships. The use of alternative data sources and the emphasis that most of the ensemble learning methods place on being able to interpret how the respective classifiers assign weight to various features contribute to a more equitable and transparent method of assessing a consumer's creditworthiness, creating a synergistic link between technological advancements and social responsibility.

This research demonstrates an innovative approach to equitable loan assessment through a fairness-aware, transparent, and fully explainable machine learning platform. The findings indicate that fairness can be achieved without sacrificing significant predictive power. In addition, by providing explanations of how machine learning models make decisions, this research contributes to a better understanding of machine learning technologies, which currently lacks transparency. Finally, all steps taken in the development process, including data preparation, strategies for validating models, the hyperparameter search space, the metrics used to evaluate models and quoting random state values, are clearly documented for other researchers wishing to replicate these methods.

Future research directions to further enhance and propel the benefits of ML to improve financial inclusion include the following: Increase the number of classes of data available by including additional groups of people within the socio-economic spectrum, which will allow for greater diverse models, thus providing equitable loan access to all people. Incorporate additional non-traditional data sources, including transaction data and behavioural attributes, to further expand upon and improve an inclusive model's accuracy and replicability. Continue to evolve machine learning models' interpretability through the use of AI-based explanation methods, which will facilitate greater understanding and trust from all parties involved, ensure compliance with regulatory obligations, and increase confidence in the use of machine learning. Finally, further investigate whether a company's existing banking systems can enable the optimum use of machine learning for operational or scalable efficiency through immediate data capture and automated processing. All of these advances in ML will enhance the reliability of automated lending decisions, while also advancing the mission to democratize access to banking services and support the growth and sustainability of more economically inclusive societies.

References

1. Abbas, S. K. (2025). *Lending by Algorithm: Fair or Flawed? An Information-Theoretic View of Credit Decision Pipelines*, 1–14.
2. Ahmed, A., Author, C., & Ahmed, T. (2025). *AI-Driven Innovations in Modern Banking: From Secure Digital Transactions to Risk Management, Compliance Frameworks, and AI-Based ATM Forecasting Systems*, 1145–1183.
3. Aniceto, M. C., Barboza, F., & Kimura, H. (2020). Machine learning predictivity applied to consumer creditworthiness. *Future Business Journal*, 6(1), 1–14.
4. Arun, K., Ishan, G., & Sanmeet, K. (2016). Loan Approval Prediction based on Machine Learning Approach. 18(3), 79–81.
5. Basnayake, D., Naranpanawa, A., Selvanathan, S., & Bandara, J. S. (2024). Financial inclusion through digitalisation and economic growth in Asia-Pacific countries. *International Review of Financial Analysis*, 96, 103596.
6. Dash, A., & Mohanta, G. (2024). Fostering financial inclusion for attaining sustainable goals. *Journal of Cleaner Production*, 449, 141731.
7. Ferrari, A. (2025). *Artificial Intelligence and Machine Learning in Finance: Addressing Complex Problems and ESG Applications.*
8. Gatete, O. (2025). Advancing Predictive Analytics. *IJLTEMAS*, XIV(IV), 169–189.
9. Gsenger, R., & Strle, T. (2021). Trust, automation bias and aversion. 19(4), 542–560.
10. Gupta, S., & Agarwal, S. (2023). Use of AI and ML in Banking & Finance.
11. Jie, W., et al. (2025). *A Comprehensive Review on Financial Explainable AI*, 1–49.
12. Linardatos, P., & Papastefanopoulos, V. (2021). *Explainable AI: A Review of Machine Learning Interpretability Methods.*
13. Maduakor, A. H. (2024). Equity in Financial Systems. 5(11), 3790–3803.
14. Murinde, V., Rizopoulos, E., & Zachariadis, M. (2022). FinTech revolution impact on banking. *International Review of Financial Analysis*, 81, 102103.
15. Nizam Uddin, M. K. U. et al. (2023). *Nazim Uddin.pdf.*
16. Nuka, T. F., & Ogunola, A. A. (2024). AI and ML for financial inclusion.
17. Nwaimo, C. S., et al. (2024). Predictive analytics for financial inclusion. 5(6), 1358–1373.
18. Pal, S., et al. (2025). Financial inclusion and economic growth. *World Development Sustainability*, 6, 100201.
19. Pandey, J. P., et al. Revolutionising Loan Approval using Machine Learning.
20. Pingili, R. (2025). AI-driven document processing.
21. Rômulo, J., et al. (2025). Towards Fair AI: Bias mitigation in credit decisions.
22. Sadok, H., et al. (2022). AI and bank credit analysis. *Cogent Economics & Finance*, 10(1).
23. Sargeant, H. (2023). Algorithmic decision-making in finance. *AI and Ethics*, 3(4), 1295–1311.
24. Sheth, V., et al. (2022). Comparative analysis of ML algorithms. *Procedia Computer Science*, 215, 422–431.
25. Tay, L., et al. (2022). Digital financial inclusion. *Heliyon*, 8(6), e09766.
26. Umeaduma, C. M., & I, A. A. (2025). AI-powered credit scoring models. 6(3), 6647–6661.
27. Wang, X. (2024). ML-driven lending decisions. 17(1), 1–19.
28. Yang, Y., et al. (2021). Bullwhip effect in supply chains. *International Journal of Production Economics*, 236, 108120.
29. Zhou, Y., et al. (2023). Two-stage credit scoring model. *International Review of Financial Analysis*, 89, 102755.

M. Aruna Prasanna Kumari,
VIT-AP School of Business,
VIT-AP University,
India.
E-mail address: aruna.23phd7175@vitap.ac.in

and

Shaiku Shahida Saheb,
VIT-AP School of Business,
VIT-AP University,
India.
E-mail address: shahid.sk@vitap.ac.in