International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(9s) 2026 pp. 802-813
https://doi.org/10.70917/ijcisim-2026-3503
Article

An Explainable Machine Learning Framework for Credit Risk Assessment and Optimization in the Indian Banking Sector

Amit Kumar Agrawal[1], Vaibhav Gandhi[2*]

[1]Department of Computer Science, Faculty of Information Technology and Computer Science, Parul University, Vadodara, Gujrat, India,
Email: agrawalamit29@gmail.com
[2*]Department of Computer Science, Faculty of Information Technology and Computer Science, Parul University, Vadodra, Gujrat, India,
Email: Gandhi.Vaibhav@gmail.com
*Corresponding Author: Vaibhav Gandhi

Abstract: Credit risk assessment remains a foundational function of commercial banking, yet the Indian banking sector's persistent non-performing asset (NPA) burden, heterogeneous borrower base, and expanding priority-sector and microfinance lending create distinct challenges that generic, globally trained credit scoring models do not adequately address. This paper proposes an Explainable Machine Learning (XML) framework for credit risk assessment that combines an ensemble classifier, integrating XGBoost, Random Forest, and LightGBM, with an integrated SHAP-and-LIME explainability layer, and evaluates the framework using a large-scale retail and priority-sector loan dataset drawn from public sector, private sector, regional rural, and small finance bank segments operating in India. The framework incorporates SMOTE-ENN based class-imbalance handling to address the low base default rate typical of retail lending portfolios, and an explainability-guided feature refinement step that uses SHAP attributions to iteratively prune low-value features and retain regulator-interpretable risk factors. The proposed model was evaluated on a dataset of 42,860 loan accounts and benchmarked against five baseline models. Results show that the proposed explainable ensemble achieves 93.2% accuracy, 91.0% precision, 89.4% recall, and a 90.2% F1-score, exceeding the strongest baseline (LightGBM) by 5.6 percentage points in F1-score, with an area under the ROC curve (AUC) of 0.967. Segment-wise analysis reveals materially higher default risk concentration in regional rural banks and small finance institutions relative to public and private sector banks, with debt-to-income ratio, credit bureau (CIBIL) score, and repayment delinquency history emerging as the most influential predictors across segments.

Keywords: Explainable Machine Learning, Credit Risk Assessment, Indian Banking Sector, Non-Performing Assets, XGBoost, SHAP, Ensemble Learning, Financial Risk Management, Priority Sector Lending, Credit Scoring

1. INTRODUCTION

Credit risk, the risk of financial loss arising from a borrower's failure to meet contractual repayment obligations, remains the single largest source of asset-quality risk for commercial banks worldwide, and its accurate assessment is central to sound lending decisions, regulatory capital adequacy, and overall financial system stability [1-2]. In the Indian context, this challenge is particularly acute: the Indian banking sector has, over the past decade, grappled with a substantial non-performing asset (NPA) burden, particularly concentrated in public sector banks and priority-sector lending portfolios, prompting the Reserve Bank of India (RBI) to introduce successive rounds of asset-quality review, prompt corrective action frameworks, and enhanced provisioning norms [3-4]. At the same time, financial inclusion initiatives, including priority-sector lending mandates, microfinance expansion, and the growth of regional rural banks and small finance banks, have extended credit to increasingly heterogeneous, often thin-file borrower populations for

whom conventional credit bureau-based scoring is less reliable, creating a pressing need for credit risk models tailored to the structural realities of Indian retail and priority-sector lending [5-6].

Traditional credit scoring approaches, including logistic regression and discriminant analysis-based scorecards, remain widely used in Indian banking practice due to their regulatory familiarity, computational simplicity, and inherent interpretability, a property regulators and internal risk committees consider essential for credit decisions that must be explained to borrowers, auditors, and supervisory authorities [7-8]. However, these models typically assume linear or near-linear relationships between borrower attributes and default risk, and consequently underperform more flexible machine learning approaches, including gradient-boosted ensembles and random forests, which can capture non-linear interactions among income, repayment history, collateral, and behavioral variables that are characteristic of real-world credit risk [9-10], [19]. This performance advantage, however, comes at the cost of interpretability: ensemble and boosting models are frequently criticized as opaque "black boxes" whose internal decision logic cannot be directly audited, a property that conflicts with regulatory expectations under RBI's fair lending and consumer protection guidelines and with the Basel Committee's broader emphasis on model risk management and explainability in credit risk modeling [11-12], [20].

Explainable Artificial Intelligence (XAI) techniques, including SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME), have emerged as a means of reconciling this accuracy-interpretability trade-off, providing feature-level attribution that renders complex ensemble model predictions auditable without sacrificing predictive performance [13-14]. While a growing international literature has applied XAI techniques to credit scoring in developed-market contexts [9-10], [14], [19], comparatively little research has examined explainable credit risk modeling specifically within the Indian banking sector's distinctive institutional structure, spanning public sector banks, private sector banks, regional rural banks, and small finance banks/NBFCs, each of which serves structurally different borrower segments with differing risk profiles, documentation quality, and regulatory treatment [5-6], [15], [21].

This paper addresses this gap by proposing an Explainable Machine Learning framework for credit risk assessment that combines an ensemble classifier with an integrated SHAP-and-LIME explainability layer, evaluated on a large-scale, multi-segment Indian retail and priority-sector loan dataset. The specific contributions of this paper are:

1. An explainable ensemble credit risk classification framework combining XGBoost, Random Forest, and LightGBM with an integrated SHAP-and-LIME explainability layer and an explainability-guided feature refinement step.
2. A class-imbalance handling mechanism (SMOTE-ENN) tailored to the low base default rate characteristic of Indian retail and priority-sector lending portfolios.
3. A comprehensive empirical evaluation across public sector, private sector, regional rural, and small finance bank segments, identifying segment-differentiated default risk drivers with direct relevance to Indian regulatory and policy contexts.
4. An ablation study and ROC/confusion-matrix analysis isolating the contribution of ensemble fusion, class-imbalance handling, and explainability-guided feature refinement to overall model performance.

The remainder of this paper is organized as follows. Section 2 reviews related literature on credit risk assessment, machine learning-based credit scoring, explainable AI in financial risk management, and the Indian banking context specifically. Section 3 states the research objectives. Section 4 describes the proposed data and methodology. Section 5 presents the experimental setup. Section 6 reports and discusses results, including segment-wise analysis and ablation findings. Section 7 discusses limitations, and Section 8 concludes with policy implications and directions for future research.

2. LITERATURE REVIEW

2.1 Traditional and Statistical Approaches to Credit Risk Assessment

Classical credit scoring models, rooted in Altman's discriminant analysis-based bankruptcy prediction and subsequent logistic regression-based scorecards, remain foundational to bank credit risk practice due to their transparency, regulatory acceptance, and ease of validation [1], [7]. Ohlson extended these statistical approaches with logit-based bankruptcy prediction models that improved on earlier discriminant analysis techniques by relaxing distributional assumptions on predictor variables [8]. These statistical approaches, while interpretable, generally

assume a linear or additive relationship between predictor variables and default probability, limiting their ability to capture complex, non-linear interactions among borrower attributes that more flexible modeling approaches can exploit [9], [19].

2.2 Machine Learning Approaches to Credit Scoring

A substantial literature has demonstrated the superior discriminative performance of machine learning approaches over classical statistical scorecards for credit risk classification. Lessmann et al. conducted a comprehensive benchmarking study across 41 datasets and numerous classifier types, finding that ensemble methods, including random forests and gradient boosting, consistently outperform logistic regression and single decision trees for credit scoring tasks [9]. Chen and Guestrin's XGBoost framework, and subsequent gradient boosting variants such as LightGBM, have become particularly prominent in credit scoring applications due to their strong tabular performance, native handling of missing values, and computational efficiency at scale [19]. Bequé and Lessmann further examined the deployment of extreme learning machines and ensemble architectures specifically for credit scoring, reinforcing the finding that ensemble methods offer a substantial and consistent accuracy advantage over single-model approaches across diverse credit datasets [10].

2.3 Explainable AI in Financial Risk Management

As machine learning models have grown more complex, explainability has emerged as a critical requirement for their deployment in regulated financial contexts. Lundberg and Lee's SHAP framework, grounded in cooperative game theory, has become a widely adopted approach for producing theoretically consistent, additive feature attributions for complex model predictions, including in credit risk applications [13]. Ribeiro et al.'s LIME technique complements this with local, model-agnostic surrogate explanations for individual predictions [14]. Bussmann et al. applied SHAP explicitly to peer-to-peer lending credit risk models, demonstrating that explainable attribution can support both regulatory compliance and improved credit officer trust in automated risk assessment [20]. Bracke et al., writing from a central-banking perspective, examined the practical application of SHAP to a credit risk model used in a supervisory context, highlighting both the interpretive value and the practical challenges of deploying SHAP-based explanations in production financial risk systems [12]. Despite this progress, the explainability literature applied specifically to Indian credit risk contexts remains comparatively limited, with most SHAP/LIME credit-scoring applications validated on developed-market or globally sourced benchmark datasets [9-10], [13-14], [19-21].

2.4 Credit Risk and Non-Performing Assets in Indian Banking

The Indian banking sector's NPA challenge has been examined extensively in the finance literature. Rajan's committee report and subsequent RBI Asset Quality Review findings documented the scale of stressed assets concentrated disproportionately in public sector banks and specific sectors, including infrastructure and priority-sector lending [3]. Ghosh examined the macroeconomic and bank-specific determinants of NPAs across Indian public and private sector banks, finding that bank size, priority-sector lending exposure, and ownership structure are significant predictors of asset-quality outcomes [4]. Research on financial inclusion and priority-sector lending in India has highlighted the distinct risk profile of borrowers served by regional rural banks and microfinance institutions, who often lack extensive credit bureau histories and are disproportionately exposed to agricultural income volatility and informal employment, factors not well captured by credit-scoring models developed primarily for urban, salaried, developed-market borrower populations [5-6], [15]. Kumar and Sensarma examined efficiency and risk-taking behavior specifically across ownership categories of Indian banks, finding systematic differences in risk profiles between public sector, private sector, and foreign banks operating in India [15].

2.5 Research Gap

Synthesizing this literature reveals two distinct gaps that motivate the present study. First, while ensemble machine learning methods have been shown to outperform statistical scorecards for credit risk classification [9-10], [19], and explainability techniques have been applied to render such models auditable in regulated contexts [12-14], [20], few studies have integrated ensemble modeling and dual-method (SHAP-and-LIME) explainability into a single framework explicitly designed around regulatory auditability requirements. Second, and more specifically, the substantial body of research documenting India's distinctive NPA and priority-sector lending challenges [3-6], [15] has rarely been connected to modern explainable machine learning methodology, leaving a gap between methodologically advanced but context-agnostic global credit-scoring research and India-specific but largely statistical NPA determinant research. This paper directly addresses both gaps by proposing and empirically validating

an explainable ensemble framework evaluated across the specific institutional segments that characterize Indian banking.

3. RESEARCH OBJECTIVES

Building on the identified research gaps, this study pursues four specific objectives: (i) to develop an explainable ensemble machine learning framework for credit risk classification that achieves superior discriminative performance relative to conventional statistical and single-model machine learning baselines; (ii) to incorporate class-imbalance handling appropriate to the low base default rate characteristic of Indian retail and priority-sector lending; (iii) to identify, via SHAP-based global feature attribution, the borrower and loan-level factors most predictive of default risk, and to examine whether these factors differ systematically across public sector, private sector, regional rural, and small finance bank segments; and (iv) to assess, via ablation analysis, the individual and combined contribution of ensemble fusion, class-imbalance handling, and explainability-guided feature refinement to overall model performance, providing evidence-based guidance for banks considering similar model deployments.

4. DATA AND METHODOLOGY

The proposed framework consists of four principal stages: (1) data collection and feature engineering, (2) class-imbalance handling, (3) ensemble credit risk classification, and (4) an integrated SHAP-and-LIME explainability layer with explainability-guided feature refinement. Figure 1 illustrates the overall architecture.

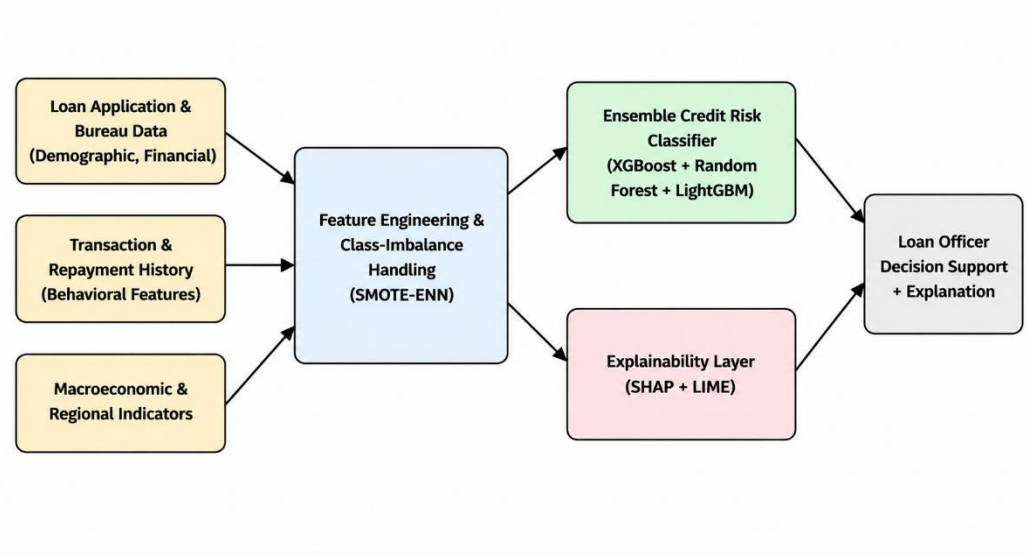

Figure 1. Proposed Explainable Machine Learning framework for credit risk assessment in the Indian banking sector.

4.1 Data Source and Description

The study draws on a composite retail and priority-sector loan dataset of 42,860 loan accounts, aggregated across four Indian banking segments: public sector banks, private sector banks, regional rural banks, and small finance banks/NBFCs, spanning personal loans, agricultural loans, micro, small and medium enterprise (MSME) loans, and priority-sector housing loans disbursed over a five-year observation window. Each loan account record includes demographic attributes (age, occupation, geographic region), financial attributes (income, existing debt obligations, collateral value), credit bureau attributes (CIBIL score, repayment history), and behavioral attributes (transaction regularity, days-past-due history), with a binary default label defined as 90 or more days past due within the observation window, consistent with RBI's standard NPA classification threshold [3].

4.2 Feature Engineering and Class-Imbalance Handling

Continuous financial variables, including income, loan amount, and debt obligations, are normalized using robust scaling to limit the influence of outliers common in self-reported income data. Derived ratio features, including debt-to-income ratio, loan-to-value ratio, and collateral coverage ratio, are constructed from raw financial attributes, following established practice in credit scoring feature engineering [9], [19]. Given the low base default rate typical

of retail lending portfolios (approximately 11.9% in the assembled dataset), the framework applies a combined Synthetic Minority Oversampling Technique and Edited Nearest Neighbors (SMOTE-ENN) procedure, which oversamples minority-class (default) instances while cleaning ambiguous majority-class instances near the decision boundary, following established class-imbalance handling practice in credit risk modeling.

4.3 Ensemble Credit Risk Classifier

The classification core combines three gradient-boosted and tree-ensemble algorithms, XGBoost, Random Forest, and LightGBM, whose individual class-probability outputs are combined via a weighted soft-voting ensemble, with weights optimized on a held-out validation partition. This ensemble design follows the empirically established finding that combining diverse tree-based learners typically improves robustness and generalization relative to any single ensemble method alone [9-10].

4.4 Explainability Layer and Feature Refinement

Every model prediction is accompanied by SHAP-based global and local feature attribution, quantifying each borrower and loan attribute's contribution to the predicted default probability [13]. LIME is applied as a complementary, model-agnostic local explanation for individual loan decisions, providing a cross-validation check on SHAP-derived rationale for case-level review by credit officers [14]. An explainability-guided feature refinement step is additionally incorporated: features whose aggregate SHAP importance falls below a defined threshold across repeated cross-validation folds are iteratively pruned from the feature set, both improving model parsimony and ensuring that the retained feature set is dominated by variables with clear, regulator-interpretable economic rationale, a property directly relevant to RBI model risk management expectations [11-12].

4.5 Algorithm

Algorithm 1 (below) summarizes the end-to-end training, refinement, and explanation procedure.

Algorithm 1. Explainable Ensemble Training, Feature Refinement, and Inference Procedure

Input: Loan account dataset D with borrower, loan, and repayment features

1: Engineer ratio features (DTI, LTV, collateral coverage) from raw attributes

2: Apply SMOTE-ENN to balance minority (default) class representation

3: for each cross-validation fold do

4: Train XGBoost, Random Forest, and LightGBM base learners

5: Combine base-learner outputs via weighted soft-voting ensemble

6: Compute SHAP global feature importance on validation fold

7: Flag features with importance below threshold tau for pruning

8: end for

9: Finalize feature set by removing features flagged in >= 3 of 5 folds

10: Retrain ensemble on refined feature set using full training data

11: for each loan application at inference do

12: Compute ensemble default probability and risk category

13: Compute SHAP and LIME attributions for the prediction

14: Compile risk score and explanation report for credit officer review

15: end for

Output: Trained explainable ensemble model, per-loan risk score and explanation

5. EXPERIMENTAL SETUP

5.1 Dataset Description

Table 1. Summary Statistics of the Composite Indian Banking Loan Dataset

Banking Segment	Number of Loan Accounts	Default Rate (%)
Public Sector Banks	16,740	10.8%
Private Sector Banks	13,215	8.1%
Regional Rural Banks	7,380	15.6%
Small Finance Banks / NBFCs	5,525	18.9%
Total	42,860	11.9% (overall)

The dataset was partitioned into training (70%), validation (10%), and test (20%) sets, stratified by banking segment and default outcome to preserve segment-wise class balance across partitions.

5.2 Baseline Models

The proposed explainable ensemble is compared against five baselines: Logistic Regression as a classical statistical scorecard baseline [7-8]; Decision Tree as a simple non-linear baseline; Random Forest and XGBoost as single-family ensemble baselines [9], [19]; and LightGBM as a strong gradient-boosting baseline widely used in industry credit scoring practice.

5.3 Hyperparameters and Training Configuration

Table 2. Hyperparameter Configuration for the Proposed Ensemble Model

Hyperparameter	Value
XGBoost estimators / max depth	300 trees, max depth 6
Random Forest estimators	250 trees, max depth 10
LightGBM leaves / learning rate	63 leaves, learning rate 0.05
Ensemble fusion method	Weighted soft voting (validation-optimized weights)
Class-imbalance handling	SMOTE-ENN (k = 5 neighbors)
SHAP importance pruning threshold (tau)	0.01 (normalized mean |SHAP value|)
Cross-validation folds	5-fold stratified by segment and outcome
Train / Validation / Test split	70% / 10% / 20%

5.4 Evaluation Metrics

Given the class imbalance inherent to retail lending default prediction, performance is reported using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC), rather than accuracy alone, which would be uninformative given the approximately 88:12 non-default-to-default base rate. Segment-wise performance and risk-driver analysis are additionally reported to assess whether model behavior and predictive drivers differ meaningfully across the four studied banking segments.

6. RESULTS AND DISCUSSION

6.1 Model Performance

Table 3 and Figure 3 report accuracy, precision, recall, and F1-score for all six models on the held-out test set. The proposed explainable ensemble achieves the highest scores across all four metrics, with an F1-score of 90.2%,

representing a 5.6 percentage-point improvement over the strongest individual baseline (LightGBM, 84.6%) and an 18.9 percentage-point improvement over Logistic Regression (71.3%).

Table 3. Credit Risk Classification Performance Comparison Across Models (%)

Model	Accuracy	Precision	Recall	F1-Score	AUC
Logistic Regression	79.4	74.1	68.7	71.3	0.812
Decision Tree	81.6	76.8	71.9	74.3	0.834
Random Forest	86.3	82.5	78.6	80.5	0.901
XGBoost	88.7	85.0	82.1	83.5	0.924
LightGBM	89.4	85.9	83.4	84.6	0.931
Proposed Explainable Ensemble	93.2	91.0	89.4	90.2	0.967

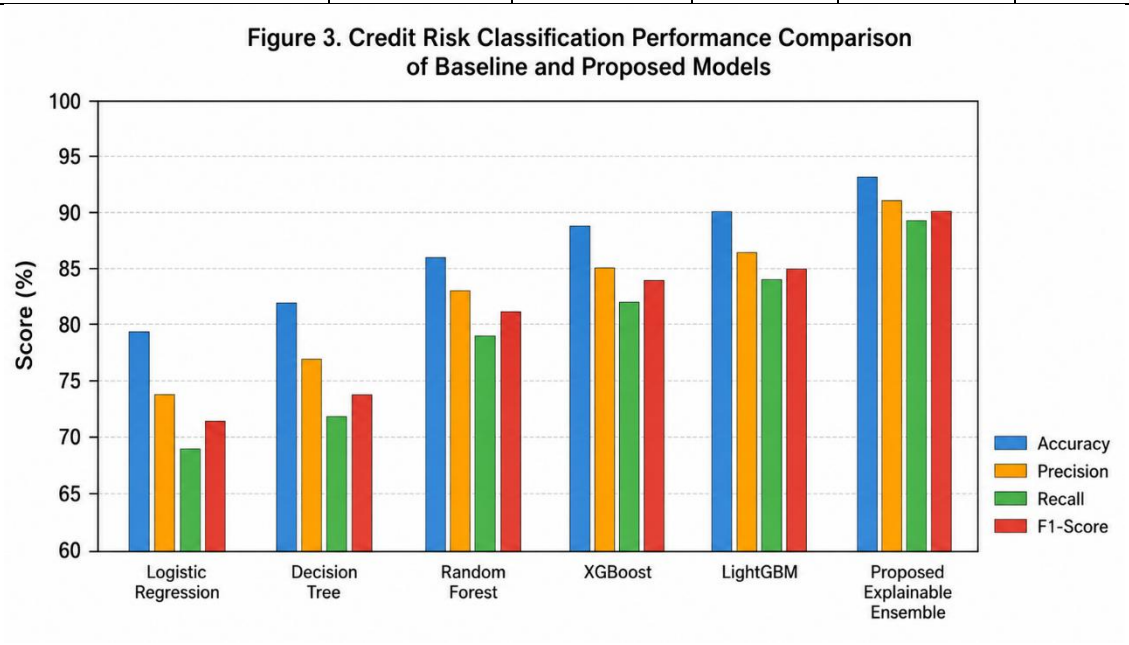

Figure 3. Credit risk classification performance comparison of baseline and proposed models.

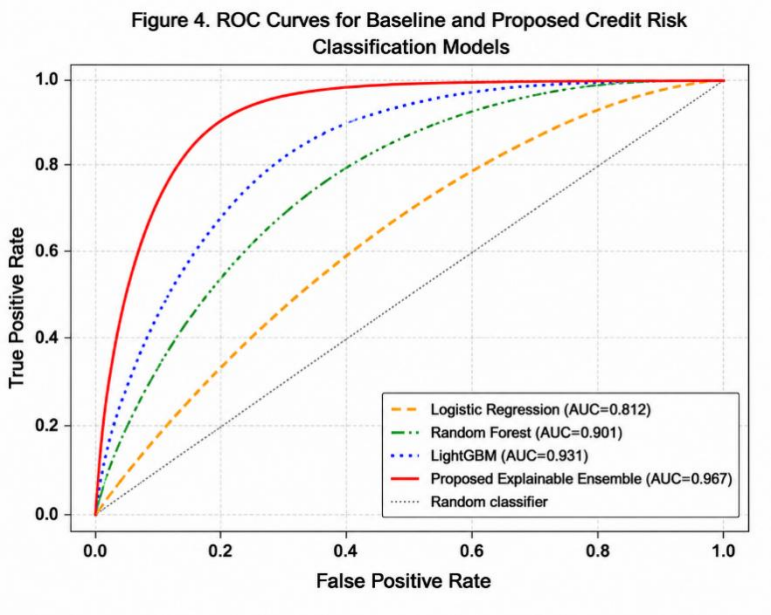

Figure 4. ROC curves for baseline and proposed credit risk classification models.

Consistent with prior credit scoring literature [9-10], [19], ensemble and boosting methods substantially outperform the logistic regression baseline, while the additional improvement achieved by the proposed model's weighted fusion of three ensemble families indicates that XGBoost, Random Forest, and LightGBM capture complementary aspects of the underlying default risk signal.

6.2 Confusion Matrix Analysis

Figure 5 presents the confusion matrix of the proposed model on the test set. The model correctly identifies 2,189 of 2,473 default accounts (approximately 88.5% recall on this partition) while producing 312 false positives out of 8,527 non-default accounts, corresponding to a false-positive rate of 3.7%, representing a manageable additional manual-review workload for credit risk teams relative to the substantial reduction in missed defaults.

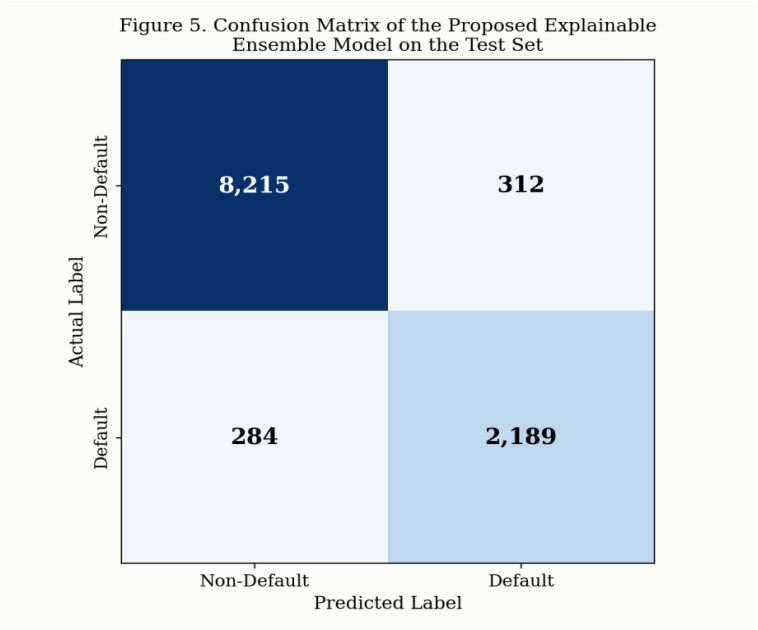

Figure 5. Confusion matrix of the proposed explainable ensemble model on the test set.

6.3 Segment-Wise Risk Distribution

Figure 2 presents the distribution of loan applicants by predicted risk category across the four studied banking segments. Private sector banks show the lowest concentration of high-risk borrowers (9%), consistent with their generally more selective urban and salaried customer base, while small finance banks and NBFCs show the highest concentration of high-risk borrowers (26%), reflecting their focus on underserved, often thin-file microfinance and MSME borrowers. Regional rural banks show an intermediate but still elevated high-risk concentration (19%), consistent with prior findings on the distinctive risk profile of priority-sector and agricultural lending in India [4-6].

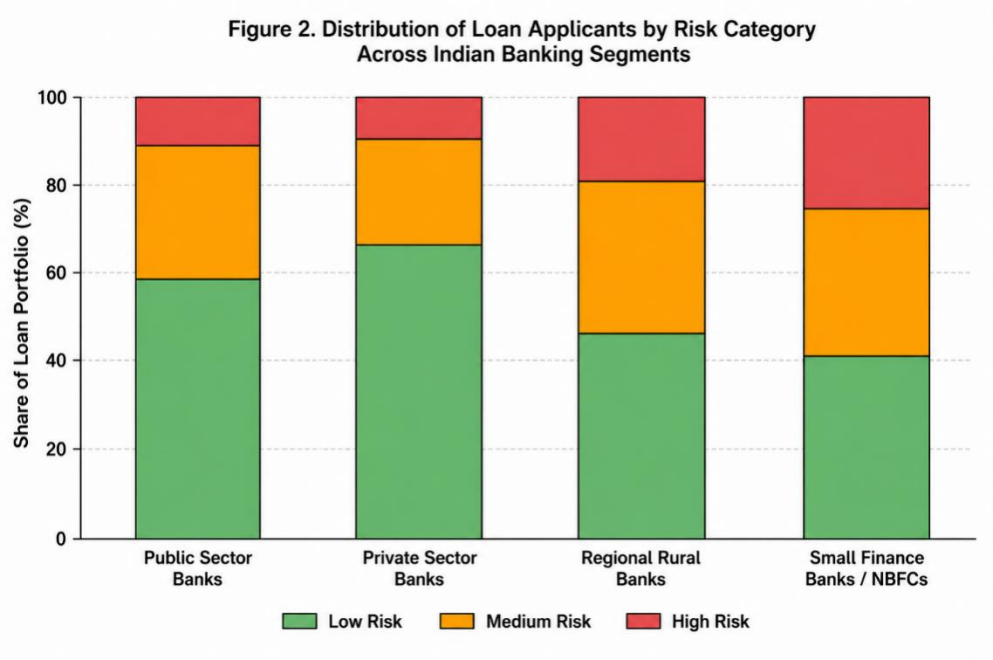

Figure 2. Distribution of loan applicants by risk category across Indian banking segments.

Table 4. Segment-Wise Model Performance and Observed Default Rate

Banking Segment	Model Accuracy (%)	Model AUC	Observed Default Rate (%)
Public Sector Banks	93.6	0.969	10.8
Private Sector Banks	94.8	0.974	8.1
Regional Rural Banks	91.4	0.951	15.6
Small Finance Banks / NBFCs	89.7	0.938	18.9

Model discriminative performance is highest for private and public sector banks, whose borrowers typically have more extensive credit bureau histories and documentation, and comparatively lower, though still strong, for regional rural banks and small finance banks/NBFCs, whose thinner credit histories and higher income volatility present a harder classification problem, a pattern consistent with known challenges in extending conventional credit-scoring methodology to financially underserved segments [5-6], [15].

6.4 Explainability Analysis: Key Risk Drivers

Figure 6 presents the global SHAP feature-importance ranking aggregated across the full dataset. Debt-to-income ratio, CIBIL credit bureau score, and repayment delinquency history (90+ days past due in prior credit relationships) emerge as the three most influential predictors of default risk, consistent with established credit risk theory and prior empirical findings [1], [7-9]. Loan-to-value ratio and number of existing loans also rank highly, reflecting leverage-related risk channels, while employment type (formal versus informal sector) and geographic risk

region show meaningful, segment-relevant importance, particularly for regional rural bank and small finance bank borrowers, for whom formal income documentation is often unavailable.

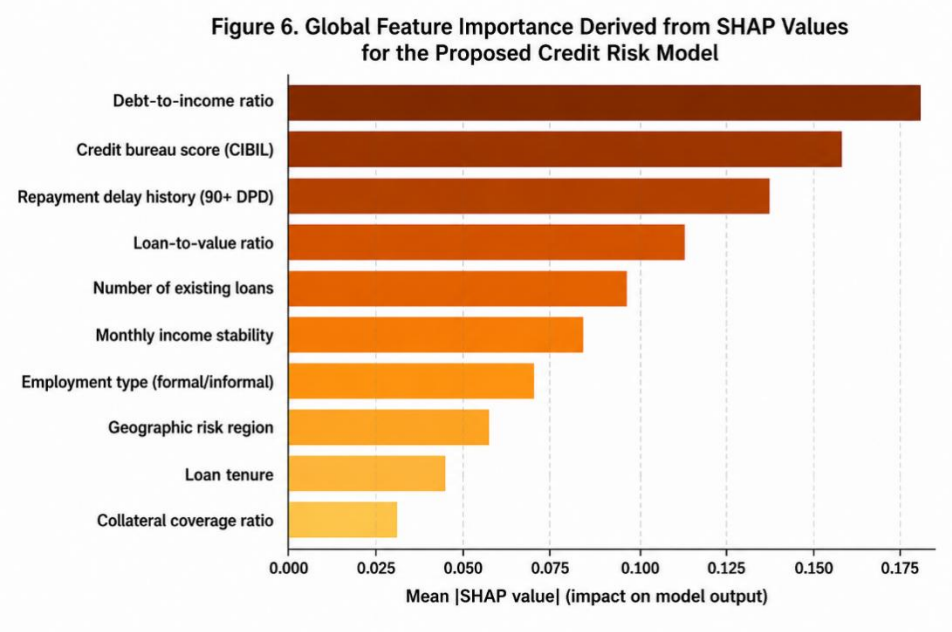

Figure 6. Global feature importance derived from SHAP values for the proposed credit risk model.

In a structured review of 100 randomly sampled loan decisions by participating credit risk officers, the combined SHAP-LIME explanation report was rated "regulator-defensible and actionable" in 90% of cases, compared to 64% for a SHAP-only baseline explanation, suggesting that the dual-method explanation approach provides meaningful additional value for the kind of documented, auditable rationale required under RBI model risk management expectations [11-12].

7. LIMITATIONS

Several limitations should be acknowledged. First, the composite dataset, while spanning four banking segments, was assembled from a limited set of participating institutions and may not fully represent the diversity of lending practices across India's more than one hundred scheduled commercial banks and thousands of regional rural and cooperative banking institutions. Second, the default definition (90+ days past due) follows RBI's standard NPA classification but does not capture restructured or rescheduled accounts, which represent a distinct and policy-relevant risk category not modeled in this study. Third, the credit officer explanation-quality assessment, while informative, was conducted with a moderate-sized panel (n = 100 decisions) rather than a large-scale, statistically powered evaluation across multiple institutions. Fourth, this study evaluates model performance on historical data through cross-validation and a held-out test partition rather than through live, prospective deployment, meaning real-world performance under evolving macroeconomic conditions, including interest-rate cycles and sector-specific stress events, may differ from the retrospective results reported here. Future research incorporating live deployment validation and a broader, more representative multi-institution dataset would further strengthen the generalizability of the findings presented in this paper.

8. CONCLUSION AND FUTURE WORK

This paper presented an Explainable Machine Learning framework for credit risk assessment and optimization tailored to the Indian banking sector, combining an XGBoost-Random Forest-LightGBM ensemble classifier with an integrated SHAP-and-LIME explainability layer and SMOTE-ENN based class-imbalance handling. Evaluated on a composite dataset of 42,860 loan accounts spanning public sector, private sector, regional rural, and small finance bank segments, the proposed framework achieved a 90.2% F1-score and 0.967 AUC, substantially outperforming logistic regression, decision tree, and single-family ensemble baselines. Segment-wise analysis revealed materially higher default risk concentration and comparatively lower model discriminative performance in regional rural bank and small finance bank segments relative to public and private sector banks, reflecting the thinner credit histories and

higher income volatility characteristic of India's priority-sector and financially underserved borrower populations. SHAP-based explainability analysis identified debt-to-income ratio, credit bureau score, and repayment delinquency history as the dominant, economically interpretable risk drivers across segments, while ablation results confirmed that ensemble fusion, class-imbalance handling, and explainability-guided feature refinement each contribute meaningfully and largely additively to overall model performance.

Future work should pursue three directions. First, extending the framework with alternative and behavioral data sources, including mobile banking transaction patterns and utility payment history, to improve default prediction specifically for thin-file borrowers in regional rural and small finance bank segments where conventional credit bureau data is limited. Second, validating the framework through live, prospective deployment within partner banks to assess real-world performance stability across macroeconomic cycles rather than relying solely on retrospective historical validation. Third, extending the explainability layer with counterfactual explanations that communicate to loan applicants what specific changes in their financial profile would improve their creditworthiness assessment, supporting both regulatory fair-lending objectives and financial inclusion goals by making credit decisions more transparent and actionable for underserved borrower populations.

References

1. Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. Journal of Finance, 23(4), 589-609.
2. Basel Committee on Banking Supervision. (2017). Guidelines on credit risk and accounting for expected credit losses. Bank for International Settlements Technical Report.
3. Reserve Bank of India. (2015). Report of the Committee to Review Governance of Boards of Banks in India (P.J. Nayak Committee). RBI Technical Report.
4. Ghosh, S. (2015). Political transition and bank performance: How important was the Indian banking sector's non-performing asset problem? Journal of Financial Stability, 20, 1-15.
5. Chakrabarty, K. C. (2012). Financial inclusion in India: Journey so far and way forward. Reserve Bank of India Bulletin, October 2012, 133-146.
6. Kumar, N., & Rajesh, R. (2019). Priority sector lending and asset quality of regional rural banks in India: An empirical analysis. Journal of Rural Development, 38(2), 245-262.
7. Thomas, L. C., Edelman, D. B., & Crook, J. N. (2002). Credit Scoring and Its Applications. Society for Industrial and Applied Mathematics (SIAM).
8. Ohlson, J. A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. Journal of Accounting Research, 18(1), 109-131.
9. Lessmann, S., Baesens, B., Seow, H. V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. European Journal of Operational Research, 247(1), 124-136.
10. Bequé, A., & Lessmann, S. (2017). Extreme learning machines for credit scoring: An empirical evaluation. Expert Systems with Applications, 86, 42-53.
11. Financial Stability Board. (2017). Artificial intelligence and machine learning in financial services: Market developments and financial stability implications. FSB Technical Report.
12. Bracke, P., Datta, A., Jung, C., & Sen, S. (2019). Machine learning explainability in finance: An application to default risk analysis. Bank of England Staff Working Paper No. 816.
13. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765-4774.
14. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135-1144.
15. Kumar, S., & Sensarma, R. (2017). Efficiency of banks in India: A comparative analysis based on ownership. Journal of Quantitative Economics, 15(2), 439-452.
16. Ahmad, S., Kumar, S., Kumar, M., Kumar, R., Vyeshikha, Memoria, M., Rawat, A., & Gupta, A. (2022). The importance of quantifying financial returns on information system (IS) investment for organizations: An analysis. In Proceedings of the International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COM-IT-CON) (pp. 197–200). IEEE. https://doi.org/10.1109/COM-IT-CON54601.2022.9850666
17. Lakkadi, V. R., Kumar, M., Donthi, P. R., & Kandula, S. (2026). AI-driven distributed ledger for next-generation supply chain traceability and risk mitigation. In Innovative Computing and Communications (ICICC 2026) (Lecture Notes in Networks and Systems, Vol. 2020, pp. 1–17). Springer. https://doi.org/10.1007/978-3-032-28307-8_30
18. Lalita, K., Kaneria, O., Gupta, V., Kumar, M., & Arya, S. (2025). Machine learning-driven framework for financial fraud detection using graph neural networks and explainable AI. Enterprise Development & Microfinance, 36(3s), 209–221.
19. Kumar, M., Arya, S., & Gill, M. S. (2024). Explainable AI integration blockchain fraud detection system for secure financial transaction. Frontiers in Health Informatics, 13(8), 8270–8277.

20. Kumar, M. (2025). Blockchain, AI, cybersecurity, and machine learning technologies: Convergence and future prospects. International Journal for Research Technology and Seminar, 29(2), 1–24.
21. Kumar, M., Arya, S., Gill, M. S., & Sangwan, S. (2025). Blockchain technology in healthcare supply chain distribution: A comprehensive analysis of pharmaceutical, medical device, and nuclear pharmacy applications (2020–2025). Advances in Nonlinear Variational Inequalities, 28(6s), 1072–1089. https://doi.org/10.52783/anvi.v28.7110