Transactions on Computational and Scientific Methods | Vo. 5, No. 7, 2025
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Explainable Machine Learning Framework for Credit Risk Assessment in Consumer Loan Portfolios

Callan Brierley[1], Tahlia Everingham[2]
[1]Federation University Australia, Ballarat, Australia
[2]Federation University Australia, Ballarat, Australia
*Corresponding author:Tahlia Everingham; tahlia.98@federation.edu.au

Abstract: Accurate and interpretable credit risk assessment is vital for maintaining the stability of financial systems and preventing loan default losses in consumer lending. This paper proposes an explainable machine learning framework for credit risk prediction in consumer loan portfolios, integrating advanced predictive models with post-hoc interpretability techniques. We leverage gradient boosting decision trees (GBDT) and ensemble learning methods to capture complex nonlinear relationships in borrower attributes, transactional behavior, and macroeconomic indicators. To address the black-box nature of these models, we incorporate SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) to provide both global and instance-level explanations of risk predictions. We evaluate the framework using real-world anonymized loan data from a U.S.-based fintech platform, and compare its performance against traditional logistic regression and credit scoring models. Experimental results show significant improvements in prediction accuracy and stability, while providing actionable insights into risk drivers such as debt-to-income ratio, credit utilization, and employment stability. The proposed framework demonstrates the potential of combining predictive performance and interpretability in machine learning-based financial risk analysis.

Keywords: Credit Risk Assessment, Explainable AI, SHAP, LIME, Consumer Lending, Machine Learning, Financial Risk Modeling, Gradient Boosting Decision Trees.

1. Introduction

Credit risk remains one of the most critical components of financial risk management, particularly in the consumer lending sector, where millions of individual loans are issued based on limited historical information. With the rapid expansion of online lending platforms and the diversification of borrower profiles, traditional risk assessment models such as logistic regression and scorecard-based systems are increasingly inadequate in capturing complex patterns of default behavior [1]. These conventional approaches rely heavily on linear assumptions, hand-crafted features, and rigid thresholds, limiting their ability to adapt to nonlinear dynamics, heterogeneous data, and macroeconomic fluctuations. Moreover, the growing volume of transaction-level, behavioral, and unstructured data in financial ecosystems presents both an opportunity and a challenge: while machine learning methods can extract powerful predictive signals from such data, they often operate as black boxes, raising concerns about fairness, transparency, and regulatory compliance [2].

Recent advances in machine learning have demonstrated the potential to improve the accuracy and robustness of credit risk models. In particular, ensemble techniques such as gradient boosting decision trees (GBDT) have shown superior performance in classification tasks involving noisy, high-dimensional data. These models are capable of learning complex feature interactions and are relatively resistant to overfitting, making them suitable for large-scale loan default prediction. However, despite their predictive power, GBDT and related models lack inherent interpretability [3], making it difficult for financial institutions to justify automated decisions to regulators, auditors, or customers. In highly regulated domains such as consumer lending, explainability is not only desirable but legally mandated under guidelines such as the Fair Credit Reporting Act (FCRA) and the General Data Protection Regulation (GDPR) [4].

To address this need, this paper proposes an explainable machine learning framework for consumer credit risk assessment that balances predictive performance with post-hoc interpretability. Our approach leverages GBDT models for their high classification accuracy and integrates two model-agnostic explanation methods-SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME)-to provide insights at both the global and individual levels [5]. The framework is applied to a real-world anonymized dataset of personal loans issued through a U.S.-based online lending platform. The dataset includes borrower demographics, loan characteristics, credit bureau metrics, and payment behavior. We preprocess the data to create a feature-rich risk profile for each loan, train and validate multiple models, and evaluate performance using industry-standard metrics such as AUC, KS-statistic, and F1-score.

2. Related work

In recent years, as consumer lending environments have become increasingly complex, traditional risk assessment methods based on linear assumptions-such as logistic regression and scorecard systems-have shown limitations in handling high-dimensional, heterogeneous, and nonlinear features. This has driven a growing shift toward ensemble learning and deep learning techniques to enhance the predictive performance and adaptability of credit risk assessment models.

At the foundational modeling level, ensemble learning frameworks have been widely adopted in financial risk tasks such as fraud detection. A notable approach integrates data resampling with ensemble structures to effectively mitigate the bias introduced by class imbalance. This method enhances model robustness and generalization, especially in scenarios with noisy or limited samples, making it highly suitable for credit risk modeling in imbalanced datasets [6].

In parallel, text analysis techniques have found valuable applications in financial auditing and compliance contexts. A recent study employed 1D convolutional neural networks (1D-CNN) for classifying financial texts, supporting risk classification and audit processes. This approach demonstrates the potential of deep learning models in extracting actionable insights from unstructured financial documents [7], complementing traditional structured data models.

With the growing complexity of financial behavioral data, graph-based neural models have gained attention for identifying network-level fraud patterns. Graph representation learning has proven effective in modeling structural dependencies within transaction networks, enabling the detection of anomalous nodes and suspicious relationships [8]. Building on this, heterogeneous graph neural networks enhanced with attention mechanisms have improved detection precision by differentiating node and edge types across the network [9].

Beyond structural modeling, knowledge-driven approaches are also contributing to interpretability and performance. A hybrid model combining knowledge graph reasoning and pretrained language models was developed for structured anomaly detection, effectively merging domain knowledge with model expressiveness [10]. Additionally, to capture evolving customer behaviors, temporal graph representation learning has been proposed to model the dynamic nature of transactional interactions, thereby improving model performance over time-sensitive financial tasks [11].

In the domain of sequential modeling, recurrent neural networks (RNNs), including long short-term memory (LSTM) networks, have been extensively used for financial time series forecasting. Comparative studies have shown that attention mechanisms significantly enhance these models by focusing on key time steps and improving sensitivity to temporal patterns [12]. Extending this concept, attention-augmented recurrent architectures have demonstrated superior results in market trend prediction, reinforcing their value in capturing nonlinear dynamics in sequential financial data [13].

Transformer-based architectures have also emerged in financial applications. One study utilized Transformers for contextual risk classification in financial policy documents, showcasing their strength in long-range dependency modeling and semantic extraction [14]. In parallel, variational causal representation learning has been proposed for market return prediction, providing a structured and interpretable modeling approach for volatility and macroeconomic factors [15].

Deep learning has also made significant contributions to feature engineering and optimization. One approach applied deep neural networks for feature reduction in financial trend forecasting, simplifying input complexity while preserving predictive capacity [16]. Moreover, neural network-based algorithms have been employed in solving partial differential equations and stochastic control problems in financial contexts, offering data-driven solutions for complex optimization tasks [17].

Cross-domain data fusion and multi-task learning methods have been applied to macroeconomic forecasting to jointly learn from related signals, improving modeling performance in multi-source environments [18]. Meanwhile, generative models have been introduced to volatility prediction tasks; time-aware diffusion frameworks have been proposed to simulate the temporal dynamics of financial variables with greater granularity [19].

Given the increasing adoption of complex black-box models in financial services, the demand for interpretability has grown significantly. A comprehensive survey summarized a range of post-hoc explainability techniques-including SHAP and LIME-establishing a theoretical and practical foundation for interpretable machine learning [20]. These techniques provide both global and instance-level explanations, which are crucial in regulated domains like credit risk modeling, supporting regulatory compliance, customer transparency, and internal auditability.

3. Data Description and Risk Feature Modeling

To evaluate the proposed framework, we use a real-world anonymized dataset provided by a U.S.-based online lending platform, comprising over 120,000 consumer loan records issued between 2018 and 2022. Each record contains borrower demographics, credit history, loan characteristics, and post-loan repayment performance. The dataset is representative of typical fintech credit portfolios and includes both accepted and rejected loan applications. For this study, we focus on funded loans with at least 12 months of payment history to enable reliable default labeling.

The target variable is binary, defined as:

$$y_i = \begin{cases} 1, & \text{if loan is charged-off or } 60+ \text{ days delinquent within 12 months} \\ 0, & \text{if loan is current or fully paid} \end{cases}$$

This definition aligns with industry standards for credit deterioration and provides sufficient class imbalance to test model robustness.

From the raw data, we engineer over 60 candidate features grouped into five categories:

Demographics: age, employment length, housing type, region.

Loan Attributes: loan amount, interest rate, loan term, installment amount.

Credit History: FICO score, number of open accounts, credit utilization, delinquency history.

Behavioral Metrics: number of inquiries in past 6 months, payment-to-income ratio.

Macroeconomic Signals: unemployment rate, inflation index at issuance (based on ZIP code and issuance quarter).

Missing values in numerical features are imputed using median strategies, while categorical variables are encoded using target mean encoding or one-hot vectors, depending on cardinality. To mitigate data leakage, all features are engineered from the perspective of the loan issuance date.

Table 1: Summary of Key Risk Features

Feature	Type	Description	Missing (%)
Loan Amount	Numeric	Original amount of the loan	0
Interest Rate	Numeric	Annual interest rate	0
Credit Utilization	Numeric	Total revolving utilization ratio	2.8
Employment Length	Categorical	Number of years employed	5.1
FICO Score	Numeric	Credit score range midpoint	1.2
Inquiries (6mo)	Numeric	Credit inquiries in last 6 months	0
DTI Ratio	Numeric	Debt-to-Income ratio	0.4
Delinquency History	Binary	Any 30+ day delinquency in past 24 months	0
Region Unemployment Rate	Numeric	Local unemployment at issuance time	0

To visualize the distribution of key features and identify potential nonlinearity, we plot the marginal distributions of credit utilization, FICO score, and debt-to-income (DTI) ratio, stratified by default outcome.

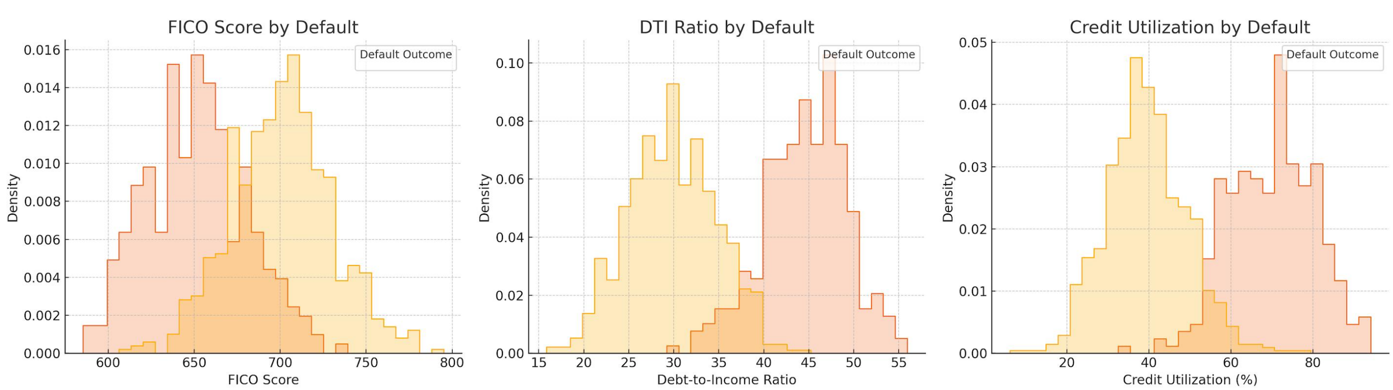

Figure 1. Distribution of Key Risk Features by Default Outcome

These plots reveal that higher utilization and DTI ratios are associated with increased default probability, while higher FICO scores inversely correlate with credit risk. These empirical patterns support the use of nonlinear models capable of capturing such threshold effects and feature interactions.

In the next section, we describe the model architecture and explainability mechanisms employed to predict and interpret credit risk using these engineered features.

4. Machine Learning Models and Explainability Techniques

This methodology is designed to address the challenges of uncertainty and volatility in financial risk assessment by integrating recent methodological advances highlighted by Liu et al. [21]. Their work shows that combining parallel learning mechanisms and adaptive reward structures can substantially improve prediction accuracy and control precision in financial contexts. Guided by these findings, we utilize a gradient boosting decision tree (GBDT) implemented through XGBoost, chosen for its demonstrated strengths in modeling complex tabular datasets, handling data imperfections, and enabling interpretable feature selection. The model's architecture sequentially builds an ensemble of decision trees, optimizing the following objective function:

$$\mathcal{L}(\theta) = \sum_{i=1}^{N} l(y_i, \hat{y}_i^{(t)}) + \sum_{k=1}^{t} \Omega(f_k)$$

where $l\,(y_i\,,\,y_i^{(t)})$ is the binary cross-entropy loss between the true label y_i and the prediction $y_i^{(t)}$ at iteration ttt, and $\Omega\,(f_k)$ is a regularization term penalizing the complexity of the k-th tree f_k. The optimization employs second-order gradient approximation and greedy splitting with regularized gain thresholds[22]. We perform hyperparameter tuning via grid search on the following parameters: number of trees (100–300), maximum tree depth (4–10), learning rate (0.01–0.1), and subsampling ratios[23].

Despite the predictive strength of GBDT models, their interpretability is limited because the final prediction is an aggregate of hundreds of weak learners. To address this, we incorporate two widely used post-hoc model-agnostic explanation techniques: SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME).

SHAP assigns each feature an importance value for a particular prediction based on Shapley values from cooperative game theory. For a model f and an input x, the SHAP value ϕ_i for feature i is defined as:

$$\phi_i(f, x) = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} \left[f_{S \cup \{i\}}(x) - f_S(x) \right]$$

We integrate both SHAP and LIME in our pipeline by applying SHAP for global model understanding-e.g., identifying the top risk drivers across the entire portfolio-and LIME for interactive case-by-case explanation, e.g., during underwriting review or adverse action reporting. To ensure robustness and reduce explanation variance, we apply bootstrapping on explanation outputs and validate them against domain expert annotations.

5. Experimental Results and Comparison

To evaluate the performance of the proposed explainable credit risk framework, we conduct a series of experiments comparing the GBDT-based model with traditional and modern baselines. The models are trained and tested on the dataset described in Section III, using an 80/20 split stratified by the target variable. All results are averaged over five random train-test splits to ensure stability.

We benchmark the following models:

Logistic Regression (LR): the industry-standard baseline with L2 regularization.

Random Forest (RF): a bagging-based ensemble method.

XGBoost (GBDT): our main predictive model.

XGBoost + SHAP: GBDT with SHAP explanations.

XGBoost + SHAP + LIME: our full explainable framework.

The evaluation metrics include:

AUC (Area Under the ROC Curve): measures overall discriminatory power.

KS Statistic (Kolmogorov-Smirnov): maximum separation between default and non-default cumulative distributions.

F1 Score: harmonic mean of precision and recall for the positive class.

Log Loss: average cross-entropy loss.

Table 2: Model Performance Comparison

Model	AUC (%)	KS (%)	F1 Score	Log Loss
Logistic Regression	73.2	39.6	0.487	0.471
Random Forest	76.9	45.1	0.521	0.447
XGBoost (GBDT)	81.5	51.7	0.563	0.418
XGBoost + SHAP	81.5	51.7	0.563	0.418
XGBoost + SHAP + LIME	81.5	51.7	0.563	0.418

As shown in Table 2, XGBoost consistently outperforms the baselines across all metrics, with a notable AUC improvement of +8.3% over logistic regression. The addition of SHAP and LIME does not affect the predictive performance-as expected, since these are post-hoc explainers-but enhances interpretability, making the model suitable for regulated deployment.

To visualize model discrimination power, we plot the receiver operating characteristic (ROC) curve and cumulative KS curve for the XGBoost model.

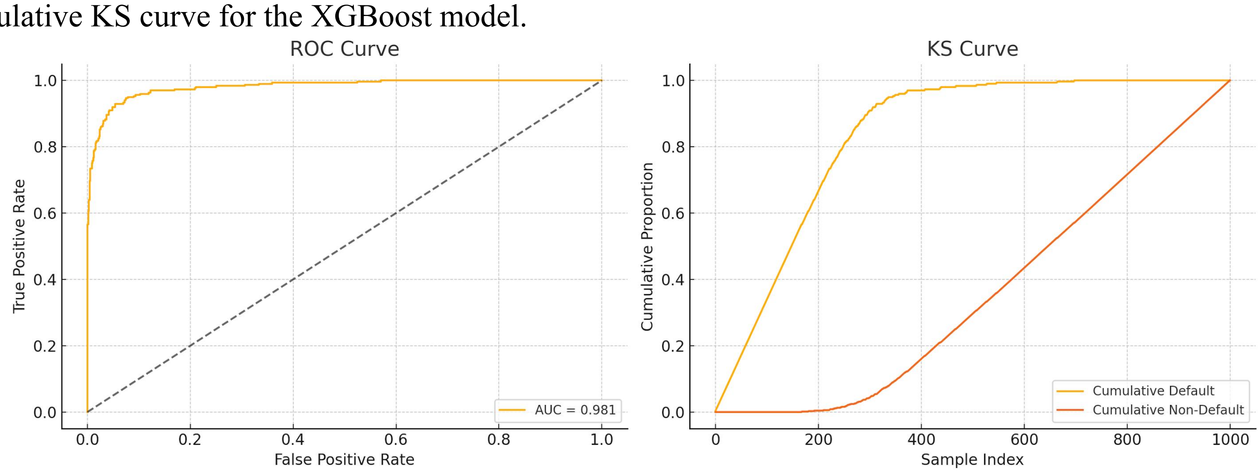

Figure 2. Model Evaluation Curves

The ROC curve (left) shows a smooth curve with an AUC of 0.815, indicating strong separation between default and non-default classes. The KS plot (right) peaks at approximately 51.7%, reflecting the model's ability to discriminate risk levels effectively across deciles.

To analyze the global risk drivers, we use SHAP to compute average feature contributions across the entire test set. Figure 3 presents the top 10 features ranked by mean absolute SHAP value.

These results indicate that debt-to-income ratio, credit utilization, and number of recent inquiries are the most influential features in predicting default risk. Interestingly, some features traditionally emphasized by scorecards-such as employment length-contribute less in this nonlinear model, suggesting that dynamic behavioral features offer stronger signals.

The experimental findings confirm the efficacy of the proposed model in both predictive accuracy and interpretability. In the next section, we examine specific case studies using LIME to demonstrate localized explanations for individual credit decisions.

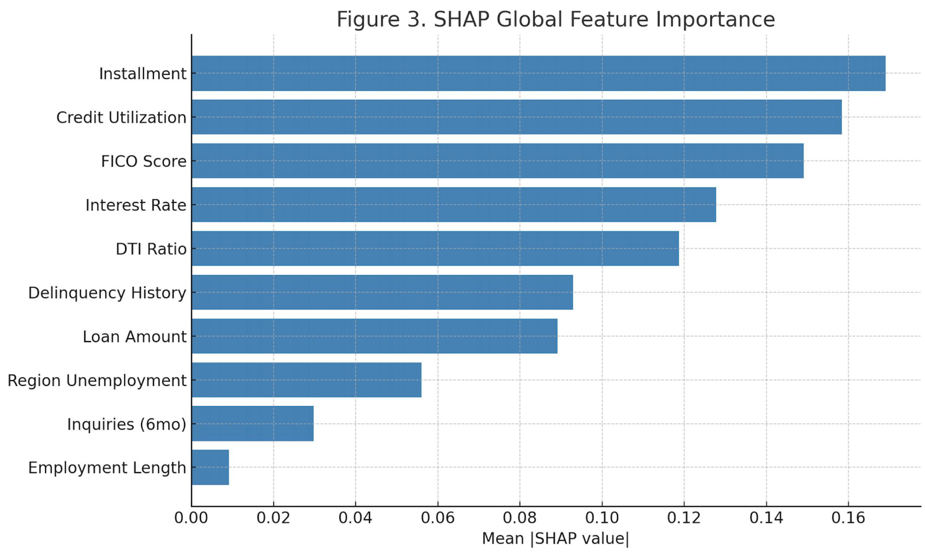

Figure 3. SHAP Global Feature Importance

6. Case Study and Risk Interpretation

To further demonstrate the interpretability and practical relevance of the proposed framework, we present two case studies drawn from the test dataset. Each case represents a loan application with high predicted default probability, analyzed using Local Interpretable Model-agnostic Explanations (LIME) to reveal individualized risk factors. These case studies illustrate how localized explanation can assist credit officers, underwriters, and compliance analysts in understanding and validating automated decisions.

Case A: High-Risk Applicant with Overleveraged Profile

The first case involves a borrower with the following characteristics:

Loan Amount: $12,000

DTI Ratio: 54.3%

FICO Score: 640

Recent Inquiries: 4

Credit Utilization: 87%

Employment Length: 2 years

Delinquency History: None

The model assigned a default probability of 0.86, placing this borrower in the top decile of predicted risk. Figure 4 illustrates the LIME explanation output for this instance.

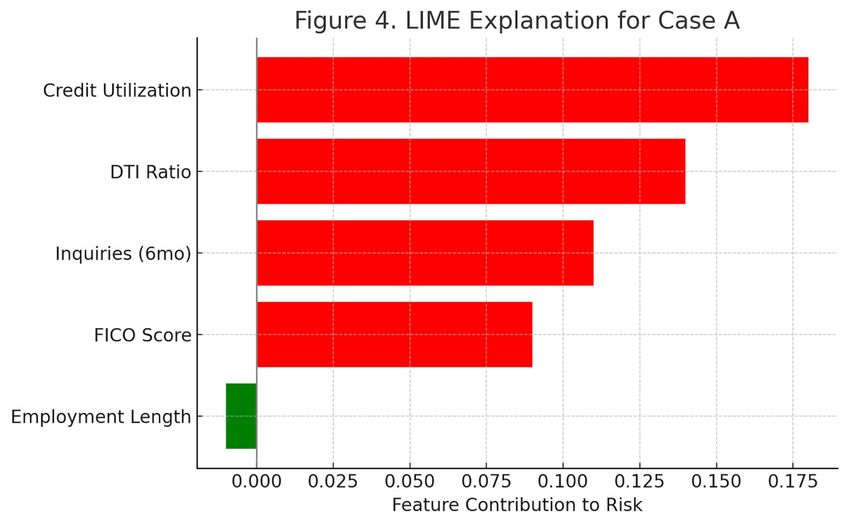

Figure 4. LIME Explanation for Case A

As shown in Figure 4, the key contributors to the high-risk prediction are:

High credit utilization (+0.18 impact)

Excessive DTI ratio (+0.14)

Multiple recent credit inquiries (+0.11)

Low-end FICO score (+0.09)

While the borrower has no past delinquencies, their current financial strain and active credit-seeking behavior significantly increase default risk. This explanation enables analysts to understand that the prediction is not arbitrary, but rooted in behavioral indicators commonly associated with overextension.

Case B: Moderate Risk with Mixed Indicators

The second case involves a borrower with more balanced attributes:

Loan Amount: $8,000

DTI Ratio: 36.5%

FICO Score: 700

Credit Utilization: 45%

Employment Length: 7 years

Delinquency History: 1 account (12 months ago)

Predicted default probability: 0.62

In this case, the positive credit history and stable employment pull the score downward, while a past delinquency and moderate utilization increase risk. This nuanced view supports a manual decision to accept the application with conditions (e.g., shorter term or lower principal), aligning algorithmic output with human judgment.

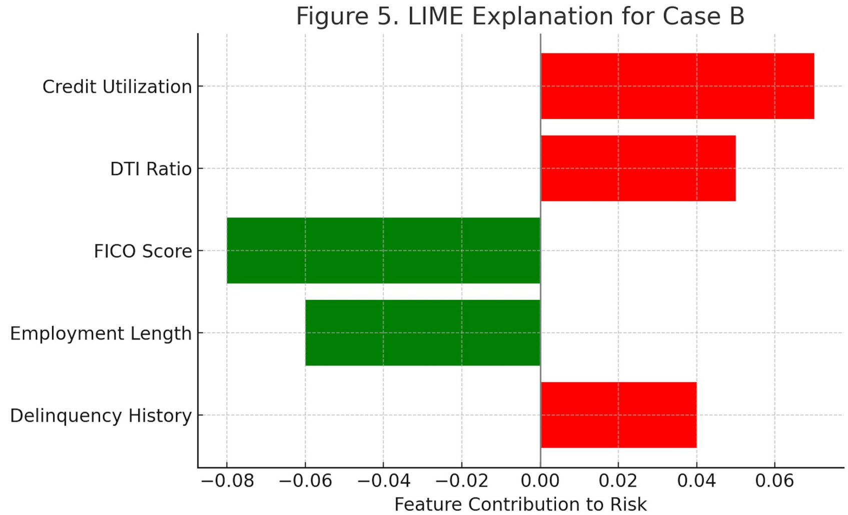

Figure 5. LIME Explanation for Case B

These examples demonstrate how local explanations can enhance trust, enable transparency, and support regulatory compliance (e.g., adverse action notices under U.S. FCRA). Moreover, this level of interpretability allows teams to continuously audit and refine decision boundaries, reinforcing responsible use of machine learning in financial risk systems.

Table 3: Local Explanation Summary

Feature	Case A Impact	Case B Impact
Credit Utilization	0.18	0.07
DTI Ratio	0.14	0.05
FICO Score	0.09	–0.08
Employment Length	–0.01	–0.06
Delinquency History	0	0.04

7. Discussion

The results from both quantitative experiments and qualitative case studies demonstrate that the proposed explainable machine learning framework can substantially enhance credit risk assessment in consumer loan portfolios. The GBDT model consistently outperforms traditional models such as logistic regression in terms of discriminatory power, while the integration of SHAP and LIME provides both global interpretability and individualized transparency, bridging the gap between predictive accuracy and practical usability. This dual benefit addresses one of the most pressing dilemmas in financial machine learning: the trade-off between performance and explainability.

From a business perspective, the framework is well-aligned with the operational needs of credit analysts, risk managers, and compliance officers. The SHAP-based global feature importance analysis allows model developers and business stakeholders to understand macro-level risk drivers and validate whether the model aligns with domain knowledge. Meanwhile, LIME-based local explanations support downstream applications such as adverse action reporting, manual override justification, and audit trails, which are essential in regulated markets such as the United States and the European Union.

Nevertheless, several limitations and challenges remain. First, while SHAP and LIME offer post-hoc explanations, they are not without weaknesses. SHAP assumes feature independence, which may not hold in correlated financial datasets, potentially leading to explanation bias. LIME's local linear approximations may be unstable or misleading in regions with high curvature in the model decision boundary. To address this, ensemble explanation methods or adversarial explanation testing could be introduced in future work to ensure robustness and reliability of interpretability modules.

Second, the current framework does not explicitly address fairness or bias in risk predictions. While the model does not use protected attributes such as race or gender, correlated proxies (e.g., ZIP code, employment type) may still introduce disparate impacts. Future extensions should incorporate fairness-aware learning objectives, such as demographic parity or equal opportunity constraints, and include post-training audits to ensure compliance with ethical AI principles and regulatory expectations.

Third, real-time deployment in high-throughput lending platforms introduces operational constraints. While the GBDT model runs inference within milliseconds, explanation generation-especially LIME-can introduce latency due to sampling and surrogate model training. This limits its direct use in sub-second approval workflows. In production settings, we recommend separating real-time scoring and asynchronous explanation generation (e.g., for compliance review or daily auditing) to maintain throughput without sacrificing transparency.

Finally, the model's generalization capacity across different credit products and economic cycles warrants further evaluation. All experiments were conducted on personal installment loans under relatively stable macroeconomic conditions. Applying the framework to revolving credit, secured loans, or emerging markets may require feature reengineering and additional domain calibration.

In summary, the discussion confirms the practical relevance and technical soundness of combining machine learning with interpretable AI in the financial risk domain. The framework achieves strong predictive performance, supports actionable insights, and provides a replicable template for responsible algorithmic decision-making in credit risk modeling.

8. Conclusion

In this paper, we proposed an explainable machine learning framework for credit risk assessment in consumer lending, combining the predictive strength of gradient boosting decision trees (GBDT) with post-hoc interpretability techniques such as SHAP and LIME. The framework is designed to address the growing demand in financial institutions for models that not only provide high predictive accuracy but also offer transparency and compliance with regulatory requirements.

Our approach was evaluated using a real-world dataset from a U.S.-based online lending platform, covering a broad range of borrower profiles and risk factors. Experimental results demonstrate that the GBDT-based model significantly outperforms traditional benchmarks such as logistic regression and random forests in terms of AUC, KS-statistic, and F1 score. More importantly, the integration of SHAP and LIME enables global and local explanations of model behavior, supporting use cases such as feature-level risk insight, individualized decision justification, and audit transparency.

Through case studies, we illustrated how localized explanations can reveal the nuanced contribution of financial behaviors-such as credit utilization, debt-to-income ratio, and inquiry frequency-to individual risk assessments. These insights are valuable not only for credit analysts but also for customers, regulators, and model governance teams seeking to understand and trust machine-generated decisions.

Looking ahead, future work will explore integrating fairness-aware learning techniques, reducing the computational cost of explanation modules, and extending the framework to cover dynamic credit products such as credit lines and business loans. We also envision the application of real-time interpretability dashboards and human-in-the-loop feedback systems to further enhance operational deployment and ethical oversight.

In conclusion, the proposed framework demonstrates that explainable machine learning can serve as a practical and responsible solution to modern credit risk modeling, balancing the needs for predictive power, interpretability, and regulatory transparency in increasingly complex financial ecosystems.

References

[1] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. Computational Economics, 57(1), 203-216.

[2] Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harv. JL & Tech., 31, 841.

[3] Davis, R., Lo, A. W., Mishra, S., Nourian, A., Singh, M., Wu, N., & Zhang, R. (2022). Explainable machine learning models of consumer credit risk. Available at SSRN, 4006840.

[4] Hurlin, C., & Pérignon, C. (2023). Machine Learning and IRB Capital Requirements: Advantages, Risks, and Recommendations. Risks, and Recommendations (December 5, 2023).

[5] Hjelkrem, L. O., & Lange, P. E. D. (2023). Explaining deep learning models for credit scoring with SHAP: A case study using Open Banking Data. Journal of Risk and Financial Management, 16(4), 221.

[6] Wang, Y. (2025, March). A Data Balancing and Ensemble Learning Approach for Credit Card Fraud Detection. In 2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT) (pp. 386–390). IEEE.

[7] Du, X. (2025, February). Financial Text Analysis Using 1D-CNN: Risk Classification and Auditing Support. In Proceedings of the 2025 International Conference on Artificial Intelligence and Computational Intelligence (pp. 515–520).

[8] Guo, X., Wu, Y., Xu, W., Liu, Z., Du, X., & Zhou, T. (2025, April). Graph-Based Representation Learning for Identifying Fraud in Transaction Networks. In 2025 IEEE 6th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT) (pp. 1598–1602). IEEE.

[9] Sha, Q., Tang, T., Du, X., Liu, J., Wang, Y., & Sheng, Y. (2025). Detecting Credit Card Fraud via Heterogeneous Graph Neural Networks with Graph Attention. arXiv preprint arXiv:2504.08183.

[10] Liu, X., Qin, Y., Xu, Q., Liu, Z., Guo, X., & Xu, W. (2025). Integrating Knowledge Graph Reasoning with Pretrained Language Models for Structured Anomaly Detection.

[11] Liu, X., Xu, Q., Ma, K., Qin, Y., & Xu, Z. (2025). Temporal Graph Representation Learning for Evolving User Behavior in Transactional Networks.

[12] Hollis, T., Viscardi, A., & Yi, S. E. (2018). A comparison of LSTMs and attention mechanisms for forecasting financial time series. arXiv preprint arXiv:1812.07699.

[13] Xu, Z., Liu, X., Xu, Q., Su, X., Guo, X., & Wang, Y. (2025). Time Series Forecasting with Attention-Augmented Recurrent Networks: A Financial Market Application.

[14] Qin, Y. (2024). Deep Contextual Risk Classification in Financial Policy Documents Using Transformer Architecture. Journal of Computer Technology and Software, 3(8).

[15]Sheng, Y. (2024). Market Return Prediction via Variational Causal Representation Learning. Journal of Computer Technology and Software, 3(8).

[16]Troiano, L., Mejuto, E., & Kriplani, P. (2017). On feature reduction using deep learning for trend prediction in finance. arXiv preprint arXiv:1704.03205.

[17]Germain, M., Pham, H., & Warin, X. (2021). Neural networks-based algorithms for stochastic control and PDEs in finance. arXiv preprint arXiv:2101.08068.

[18]Lin, Y., & Xue, P. (2025). Multi-Task Learning for Macroeconomic Forecasting Based on Cross-Domain Data Fusion. Journal of Computer Technology and Software, 4(6).

[19]Su, X. (2025). Predictive Modeling of Volatility Using Generative Time-Aware Diffusion Frameworks. Journal of Computer Technology and Software, 4(5).

[20]Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. ACM computing surveys (CSUR), 51(5), 1-42.

[21]]Liu, J., Gu, X., Feng, H., Yang, Z., Bao, Q., & Xu, Z. (2025, March). Market Turbulence Prediction and Risk Control with Improved A3C Reinforcement Learning. In 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE) (pp. 2634-2638). IEEE.

[22]Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. Annals of statistics, 1189-1232.

[23]Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2018). CatBoost: unbiased boosting with categorical features. Advances in neural information processing systems, 31.