International Journal of Multidisciplinary Research and Growth Evaluation.

A Multi-Agent Framework for Personalized Credit Recommendations Using Interpretable Machine Learning and Large Language Models

Sai Prashanth Pathi
Independent Researcher, USA

* Corresponding Author: **Sai Prashanth Pathi**

Article Info

ISSN (Online): 2582-7138
Impact Factor (RSIF): 7.98
Volume: 06
Issue: 04
July - August 2025
Received: 12-06-2025
Accepted: 10-07-2025
Published: 04-08-2025
Page No: 1406-1410

Abstract
We present a novel multi-agent framework that integrates interpretable machine learning (ML) models with large language models (LLMs) to deliver personalized credit risk recommendations. Using the Lending Club Loan dataset, we train an XGBoost based credit risk model, applying SHAP (SHapley Additive exPlanations), we extract the top risk factors for individual borrowers. Subsequently, these personalized explanations are provided to an LLM, which generates actionable recommendations to help customers improve their credit profiles. A structured multi-agent LLM review pipeline is introduced to assess each recommendation for tone, compliance, and legal soundness. Recommendations scoring above a set threshold in all review dimensions are only shared with customers, ensuring fairness, transparency, and regulatory adherence. The experimental results demonstrate the effectiveness of the framework in providing clear, actionable, and compliant recommendations.

DOI: https://doi.org/10.54660/.IJMRGE.2025.6.4.1406-1410

Keywords: Interpretable Machine Learning, SHAP Explanations, Credit Risk Modeling, Large Language Models (LLMs), Multi-Agent Systems, AI in Finance, Explainable AI

1. Introduction
Credit risk assessment plays a crucial role in the financial ecosystem, determining access to credit for millions of borrowers. While advanced ML models such as XGBoost deliver high predictive accuracy, their complexity often limits interpretability [3]. Explainability methods like SHAP provide transparent, individualized reasoning behind risk predictions but limit borrowers' ability to understand or improve their creditworthiness. To address this, we propose a novel framework that combines interpretable ML with LLM-based natural language generation to produce personalized, actionable credit recommendations. Our approach ensures regulatory compliance and fairness through a rigorous multi-agent AI review process before delivering these recommendations to customers.

2. Related Work
Explainable AI (XAI) has gained prominence in financial risk modeling, with methods like SHAP widely adopted to provide local interpretability for individual predictions [3, 4]. Additionally, LLMs are increasingly explored for generating human readable explanations and recommendations in various domains, including finance [5, 6].
However, limited work has explored integrating these technologies to offer personalized credit recommendations with built-in compliance checks. Our framework fills this gap by combining SHAP-based explainability, LLM driven generation, and multi-agent review to deliver compliant, customer-friendly recommendations.

3. Data and Preprocessing
We use the publicly available Lending Club Loan dataset [1], which includes anonymized borrower information, loan performance data, and credit bureau-derived features. To preserve privacy and ensure focus on creditworthiness, we:

- Remove any personally identifiable attributes.
- Use only bureau-derived features such as credit utilization.
- Delinquencies, and payment history
- Clean, impute missing values, and scale features as needed.
- The dataset is split into a training set (80%) and a testing set (20%), preserving class balance via stratification.

4. Methodology

4.1. Credit Risk Modeling

We train an XGBoost classifier to predict credit default (denoted by the 'default flag' variable). Hyperparameters are tuned for optimal performance using grid search. The model outputs a probability of default for each customer. A higher predicted probability indicates a higher likelihood of loan default.

4.2. Model Explainability via SHAP

SHAP values are computed to explain the contribution of each feature to individual predictions. We extract the top three risk factors per customer that most negatively impact their credit profile, providing a transparent basis for personalized recommendations.

4.3. Personalized Recommendation Generation

The top three SHAP-derived factors are provided as structured inputs to an LLM [7, 8]. In this study, we employed **Gemini 2.5 Flash**, a multi- modal LLM developed by Google, known for its high efficiency, multimodal reasoning capabilities, and large context window of up to 1 million tokens [9, 10]. Carefully designed prompts instruct the LLM to generate empathetic, actionable recommendations to remedy the affecting factors.

4.4. Multi-Agent Review Pipeline

To ensure recommendations meet high standards of quality and compliance, we implement a multi-agent LLM-based review system consisting of

1. **Tone Agent:** Evaluates the recommendation's clarity, empathy, and readability.
2. **Compliance Agent:** Checks adherence to financial regulations and fair lending laws.
3. **Legal Agent:** Verifies legal soundness and regulatory alignment.
4. **Approval Agent:** Approves the recommendation only if it scores above a set threshold of 90% in all the tone, compliance and legal checks

5. Experimental Results

5.1. Model Performance

To benchmark the predictive performance of our model, we compared the XGBoost classifier against two common baseline models—Logistic Regression and Random Forest.

Figure 1 shows the ROC (Receiver Operating Characteristic) curves for all three models, with the XGBoost model achieving the highest AUC of 0.6952, outperforming both Logistic Regression (AUC = 0.6560) and Random Forest (AUC = 0.6747). This demonstrates the superior discriminatory capability of the XGBoost model in predicting credit default risk, even when constrained to bureau-derived features.

The ROC curve also highlights that XGBoost consistently achieves a higher True Positive Rate across most False Positive Rate thresholds, making it a preferable choice for this application.

The lift graph in Figure 2 shows that the model has good ranking power and strong discriminatory power that differentiates bad and good.

5.2. Explainability and Feature Insights

SHAP summary plots help you understand which features are most influential in model's predictions, offering transparency in otherwise complex models. The SHAP summary plot in Figure 3 provides a clearer explanation of the model's decisions, facilitating communication with stakeholders, regulators, and end-users, which is crucial in sensitive applications like finance or healthcare.

SHAP analysis consistently identified the following key risk factors:

1. High credit utilization
2. Number of recently opened trades
3. History of Delinquencies
4. Total balances owned

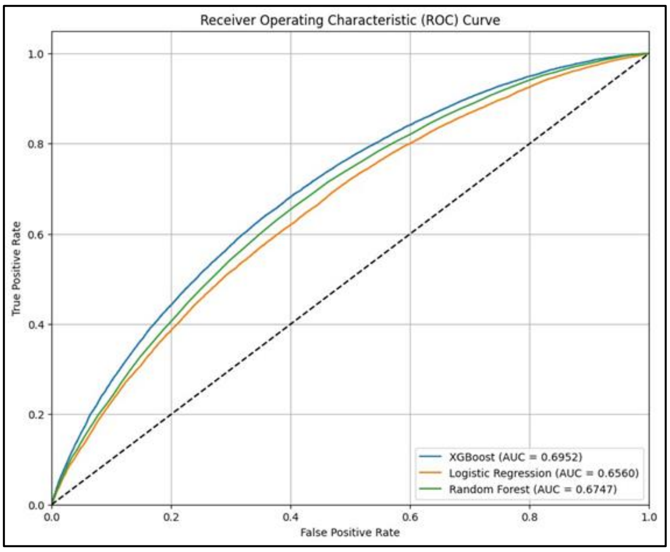

Fig 1: Comparison of ROC curves for XGBoost, Logistic Regression, and Random Forest models. XGBoost demonstrates superior performance with an AUC of 0.6952, indicating better discriminatory power in credit risk classification

5.3. Recommendation Generation and Review Pipeline

Personalized recommendations were generated by the LLM based on the top SHAP-identified risk factors for each customer. To ensure high standards of quality, clarity, and compliance, we employed a multi-agent review pipeline comprising four independent agents:

- **Tone Agent:** Assesses the recommendation's empathy, clarity, and overall customer-friendliness.
- **Compliance Agent:** Evaluates adherence to financial regulations, fair lending laws, and credit reporting guidelines.

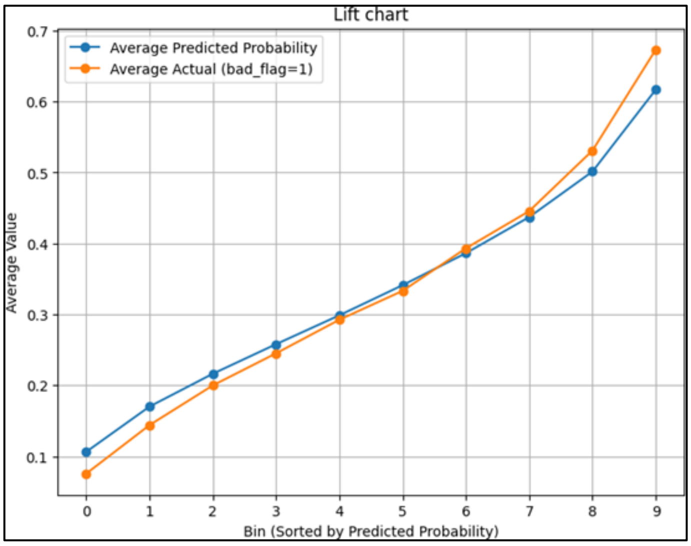

Fig 2: Lift chart illustrating the strong ranking capability of the XGBoost credit risk model, highlighting separation between good and bad credit profiles

- **Legal Agent:** Reviews legal defensibility and checks for potential legal risks.
- **Approval Agent:** Aggregates the scores from the three review agents and makes the final approval decision.

Each of the first three agents independently scores the recommendation on a scale from 0 to 10, acting as an autonomous reviewer focused on its respective evaluation criterion. The Approval Agent collects these scores and approves the recommendation only if the average score across all agents exceeds a threshold of 90% (equivalent to a minimum average score of 9.0). Through this rigorous process, approximately 15% of the generated recommendations were filtered out, primarily due to issues related to tone, clarity, or regulatory concerns. This strict multi-agent filtering ensured that only actionable, empathetic, and legally sound recommendations were delivered to customers.

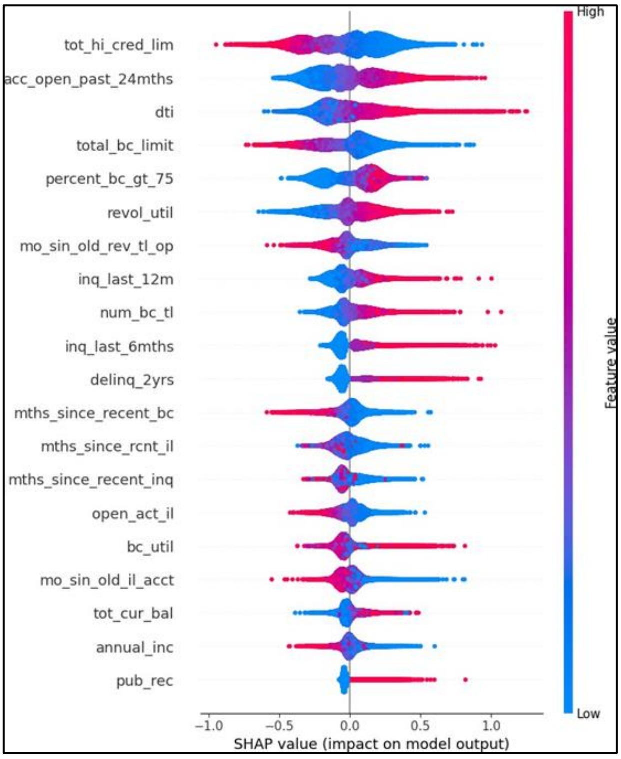

Fig 3: SHAP summary plot

5.3.1. Example Recommendation Output

To illustrate the nature and tone of the personalized recommendations generated by our framework, we present below a sample output produced for a customer whose top SHAP risk factors included:

1. Revolving line utilization rate, or the amount of credit the borrower is using relative to all available revolving credit.
2. A ratio calculated using the borrower's total monthly debt payments on the total debt obligations, excluding mortgage, divided by the borrower's self-reported monthly income.
3. Percentage of all bankcard accounts greater than 75% of limit.

Sample Personalized Credit Recommendation
Subject: Supporting Your Credit Journey: Personalized Tips from Your Bank Counselor Dear [Customer Name], I hope this email finds you well. As your financial partner, we're always looking for ways to support your journey towards financial well-being. We truly value you as a customer and appreciate the positive steps you're already taking in managing your finances. To help you continue building an even stronger credit foundation, we've identified a few areas where strengthening your credit practices could further enhance your overall financial profile. These are simply opportunities to help you thrive. **Personalized, Actionable Steps** 1. **Managing Your Overall Credit Card Balances:** Aim for balances under 30% of your total available credit limit across all revolving accounts. 2. **Reviewing Monthly Debt Payments:** Consider paying down non-mortgage debts, particularly those with high interest rates, to improve your debt-to-income ratio. 3. **Focusing on Individual Credit Card Utilization:** Prioritize reducing balances on credit cards where usage exceeds 75% of the available limit. *Remember, these are general guidelines. Implementing even one of these steps can contribute positively to your credit health over time. If you have any questions or would like to discuss further, please reach out to our Credit Counseling Team.* Sincerely, Your Credit Counseling Team

This example illustrates how complex risk factors identified by SHAP are translated into accessible, empathetic, and actionable recommendations for customers.

5.3.2. Mitigating LLM Risks

LLMs have a known risk of producing hallucinations or generating incorrect recommendations. To mitigate this, the Approval Agent's threshold has been deliberately set high to ensure that only thoroughly validated recommendations are delivered to customers.

6. Discussion

Our results demonstrate the feasibility and effectiveness of integrating interpretable ML with LLM-based recommendation systems in finance. Key strengths include:

- High interpretability and transparency via SHAP explanations.
- Actionable, personalized recommendations through LLM generation.
- Strong safeguards via multi-agent review for fairness, tone, legal and compliance.
- Future Work:
- Outputs could be further evaluated by customers or tested in real-world pilots.
- Conducting real-world user studies to measure customer impact.
- Testing on diverse datasets and industries beyond banking and lending.
- Extending the system for multilingual recommendations.

7. Conclusion

This study introduces a multi-agent AI framework that bridges ML explainability and personalized credit recommendations. By coupling SHAP-based risk factor analysis with LLM-driven recommendations, and rigorously reviewing outputs through a multi-agent system, we provide a robust, scalable, and compliant solution for improving credit decision transparency and fairness.

8. References

1. Lending Club. Lending Club loan data [Internet]. Kaggle; c2025 [cited 2025 Sep 4]. Available from: https://www.kaggle.com/datasets/adarshsng/lending-club-loan-data-csv?resource=download&select=LCDataDictionary.xlsx
2. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems (NeurIPS 2017) [Internet]; c2017 [cited 2025 Sep 4]. Available from: https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
3. Rongali SK. Enhancing machine learning models: addressing challenges and future directions. World Journal of Advanced Research and Reviews. 2025;25(1):190. doi:10.30574/wjarr.2025.25.1.0190
4. Lin Q, et al. Predicting the risk of heart failure after acute myocardial infarction using an interpretable machine learning model. Frontiers in Cardiovascular Medicine. 2025;12:1444323. doi:10.3389/fcvm.2025.1444323
5. Said A. On explaining recommendations with large language models: a review. Frontiers in Big Data. 2025;8:1-9.
6. Pineiro-Martín A, et al. Context is king: large language models' interpretability in divergent knowledge scenarios. Applied Sciences. 2025;15(4):1234-49.
7. Nandanwar H, Katarya R. Securing Industry 5.0: an explainable deep learning model for intrusion detection in cyber-physical systems. Computers and Electrical Engineering. 2025;110161. doi:10.1016/j.compeleceng.2025.110161
8. Shi X, Zhang Y, Yu M, Zhang L. Deep learning for

enhanced risk management: a novel approach to analyzing financial reports. PeerJ Computer Science. 2025;2661. doi:10.7717/peerj-cs.2661

9. Google Cloud, DeepMind. Gemini 2.5 Flash: Google's multimodal, "thinking" large language model [Internet]. Google Cloud; 2025 [cited 2025 Sep 4]. Available from: https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-5-flash
10. Doshi T. We're expanding our Gemini 2.5 family of models [Internet]. Google Blog; 2025 Jun [cited 2025 Sep 4]. Available from: https://blog.google/products/gemini/gemini-2-5-model-family-expands/