

Research Article



Detection of UPI Mule Accounts Using Machine Learning and Streamlit-Based Predictive Analytics

Ishita D. Sonawane,* Priyanshi R. Suyal, Shradha Chavan* and Rudresh Shirwaikar*

School of Computer Science and Information Technology, Symbiosis Skills and Professional University, Pune, Maharashtra, 412101, India

*Corresponding Author(s)

Ishita D. Sonawane

✉ ishita.sonawane@gmail.com

Shradha Chavan

✉ shradha.chavan@sspu.ac.in

Rudresh Shirwaikar

✉ rudresh.shirwaikar@sspu.ac.in

Received: 15 December 2026

Revised: 02 March 2026

Accepted: 09 March 2026

Published Online: 10 March 2026

Citation

I. D. Sonawane, P. R. Suyal, S. Chavan, R. Shirwaikar, Detection of UPI mule accounts using machine learning and Streamlit-based predictive analytics, *Journal of Information and Communications Technology: Algorithms, Systems and Applications*, 2026, 2(1), 26301, <https://doi.org/10.64189/ict.26301>.

Open Access

This article is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/), which permits the non-commercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as appropriate credit to the original author(s) and the source is given by providing a link to the Creative Commons License and changes need to be indicated if there are any. The images or other third-party material in this article are included in the article's Creative Commons License, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons License and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this License, visit: <https://creativecommons.org/licenses/by-nc/4.0/>

Abstract

By facilitating quick and easy money transfers, the Unified Payments Interface (UPI), in particular, and the rapid growth of digital payment systems in India have drastically changed the financial landscape. However, the growing use of UPI has also resulted in an increase in fraudulent activity, particularly through mule accounts, which are used to route or launder money while hiding the identity of scammers. This study suggests a machine learning-based framework for identifying UPI mule accounts using transactional data to solve this problem. The suggested approach uses the Light Gradient Boosting Machine (LightGBM) algorithm to determine whether a transaction is authentic or fraudulent. Transaction-related characteristics such as transaction amount, time, user demographics, and fraud risk labels are included in the dataset, which was taken from a publicly accessible Kaggle repository. Methods for preprocessing data, such as label encoding, feature scaling, and Synthetic Minority Oversampling Technique (SMOTE), were used to address class imbalance and enhance model performance. With a Receiver Operating Characteristic–Area Under the Curve (ROC-AUC) score of 0.9962 and an accuracy of 96.55%, the trained LightGBM model proved to have a strong ability to distinguish between fraudulent and legitimate transactions. Additionally, a web application based on Streamlit was created to facilitate interactive model demonstration and real-time fraud risk prediction. The suggested framework offers a scalable and effective way to improve the UPI ecosystem's fraud monitoring systems.

Keywords: UPI; Mule account; Fraud detection; Machine learning; LightGBM; Streamlit; ROC-AUC.

1. Introduction

With India emerging as one of the top adopters of real-time digital payment infrastructures, the swift digitization of financial services has drastically changed the global payment ecosystem.^[1] The shift to a cashless economy has been sped up by the growing dependence on peer-to-peer transfers, merchant payments, and mobile-based transactions.^[2] The Unified Payments Interface (UPI), one of these systems, has emerged as a key component of India's digital payment infrastructure, facilitating quick, easy, and interoperable money transfers between banking platforms.^[3,4] Transaction volumes have increased exponentially in recent years because of their simplicity of use, round-the-clock accessibility, and compatibility with mobile applications.^[5] However, the UPI ecosystem's rapid growth has resulted in a corresponding increase in online financial fraud. The abuse of mule accounts is among the most serious risks in this ecosystem.^[6] Mule accounts are bank accounts that fraudsters use to conceal their true identities while receiving, transferring, or withdrawing illegal funds. These accounts might be owned by people who participate knowingly or unknowingly in exchange for cash rewards.^[7] Fraud detection is made more difficult by the layered transfer of funds through mule accounts, particularly when fraudulent transactions resemble patterns of legitimate behavior.^[8]

Traditionally, fraud detection systems in banking have used rule-based or threshold-based methods.^[9] These systems flag suspicious transactions on the basis of set criteria, such as unusually high amounts, abnormal frequency, or geographic inconsistencies. While rule-based systems are easy to understand and implement, they have limited flexibility. Fraudsters keep changing their tactics, which makes static rules less effective over time.^[10] As a result, these systems often have high false positive rates and may miss complex mule account activities that look like real transactions.^[11] This finding shows that traditional fraud detection methods need smarter and more adaptable approaches. Machine Learning (ML) provides a data-intensive method for fraud detection, which involves the discovery of hidden patterns and correlations in large-scale transaction data. Supervised learning algorithms can be trained to differentiate between legitimate and fraudulent transactions on the basis of labeled data. Machine learning-based systems enhance the accuracy of fraud detection by modeling nonlinear relationships and adapting to new patterns of financial fraud. In the context of financial fraud detection, compared with traditional statistical methods, ensemble learning algorithms and gradient boosting algorithms have been found to perform better.^[12]

In the proposed research, a light gradient boosting machine (LightGBM)-based detection system is

developed for the detection of mule accounts in UPI transactions. LightGBM^[13] is chosen for its efficiency, scalability, and robustness on structured tabular data. The system uses the concept of gradient boosting to learn from mistakes in classification and reduce them while addressing large datasets and class imbalance issues. The proposed system combines preprocessing methods, feature extraction, and the synthetic minority over-sampling technique (SMOTE)^[14] to improve the ability of the system to detect imbalanced fraud data. The dataset employed in this research was sourced from a Kaggle repository^[15] and includes attributes such as the transaction hour, amount, category, state, user age, and fraud risk. These attributes make it possible to perform a behavior analysis of the transactions, including their temporal behavior, frequency, and value-based anomalies. The data were preprocessed in a manner that involved scaling, encoding, and balancing.

The main aim of this study is to develop a high-performance fraud detection model capable of distinguishing between legitimate transactions and mule account activities. The performance of the proposed model is evaluated using standard classification metrics such as accuracy, precision, recall, F1 score, confusion matrix, and receiver operating characteristic-area under the curve (ROC-AUC). The target ROC-AUC score is above 0.9, indicating strong discriminatory ability. To enhance practical applicability, the trained model is integrated into a Streamlit-based web application that enables real-time predictive analytics. Through this application, users can input transaction parameters and instantly receive predictions regarding the risk of fraud. The deployment of the machine learning model within the web application demonstrates its potential use in real-world financial monitoring systems.

2. Literature review

Few studies have investigated the application of machine learning for the detection of fraud in digital payments. Traditional machine learning models such as support vector machines (SVMs), decision trees, logistic regression, and random forests were used to detect anomalies in the data. Traditional machine learning models require labeled data to train the models to learn patterns of legitimate and fraudulent activities. Traditional machine learning models have shown moderate results, with an average accuracy and F1 scores of 80–90% on financial datasets. Traditional machine learning models perform poorly on highly imbalanced fraud datasets and tend to have higher false positive rates without the use of sophisticated resampling methods.^[16]

2.1 UPI-specific fraud detection studies

Research in the area of fraud detection in the UPI environment has recently become popular at a rapid pace

and has consequently given rise to fraud-related issues. Various studies have been conducted in the past few years to explore the use of ML-based solutions for UPI fraud detection on the basis of transactional parameters and evaluation metrics such as accuracy, precision, recall, F1 score, and ROC-AUC.

For instance, studies conducted using random forest and SVM classifiers on UPI transactions have shown the accuracy and recall of classification to be indicative of the efficiency of ML-based solutions over rule-based solutions.^[17,18] Various studies have conducted comparative analyses using algorithms such as logistic regression, support vector machine (SVM), and gradient boosting to evaluate their relative efficiency on the basis of performance metrics.^[19,20] These studies indicate that ensemble methods generally outperform individual classifiers in terms of the ROC-AUC and F1 score.

Another gradient boosting algorithm, CatBoost, has also been used successfully on UPI fraud datasets, and it has shown high AUC scores, which indicate a strong ability to distinguish between fraudulent and genuine transactions on categorical data.^[21] These recent studies, which are UPI focused, confirm once again that ML algorithms perform better than threshold rules do.

2.2 Advanced approaches and hybrid methods

In addition to traditional ML, advanced methods that combine deep learning, transformers, and federated learning have been suggested to improve detection results even further. A more recent method that combined causal inference, transformers, and federated

learning reported a precision and recall of 80–90%, which outperformed traditional baselines on UPI datasets.^[22] Other systematic literature reviews that have examined digital payment fraud detection in general suggest that neural networks such as Long Short-Term Memory (LSTM) and convolutional neural networks (CNNs) can detect fraud with accuracies above 99% when used on sequential transaction data.^[23] These hybrid and deep learning methods overcome the shortcomings of traditional ML by modeling the temporal dynamics of transactions and the nonlinear interactions of features. Nevertheless, these methods may require additional computational resources and more labeled data to prevent overfitting.

2.3 Comparative analysis of performance metrics

This comparison indicates that while traditional models provide interpretable baselines, the ensemble and boosting methods generally achieve better performance metrics (higher ROC-AUC and F1 score) on imbalanced fraud datasets. Hybrid and neural approaches show potential for further improvement but pose practical challenges for deployment in real-time systems such as UPI. Table 1 provides a comparative perspective on different fraud detection models.

3. Methodology

3.1 Overview

The system proposed here for UPI Mule Account Detection follows a structured pipeline of data preprocessing, exploratory data analysis, feature

Table 1: A comparative perspective on fraud detection models highlights several trade-offs.

Model Type	Typical Metrics	Strengths	Limitations
Logistic Regression/SVM	Acc ~80–88%, F1 ~70–85%	Interpretable, low complexity	Struggles with nonlinearity
Decision Trees	Acc ~85–90%, Precision ~80–88%	Handles categorical data	Risk of overfitting
Random Forest	Acc ~88–92%, ROC-AUC ~0.91	Good generalization	Higher computation
Gradient Boosting (XGBoost/CatBoost)	Acc ~90–96%, ROC-AUC ~0.93–0.99	Strong performance on tabular data	Requires hyperparameter tuning
Transformer/Deep Learning	Acc ~90–99%, ROC-AUC ~0.94+	Captures complex patterns	Data & compute intensive




Fig. 1: Block diagram of the proposed UPI mule account detection architecture.

engineering, model training on LightGBM, and deployment using Streamlit. Each step ensures that the fraud detection model has maximum accuracy, interpretability, and applicability in a real-world scenario. The workflow is inspired by contemporary research that integrates machine learning into fintech systems for identifying suspicious digital payment patterns. The end-to-end pipeline of the proposed UPI mule account detection architecture is shown in Fig. 1.

3.1.1 Dataset description

Dataset Source: The dataset utilized in this research was obtained from a publicly accessible Kaggle notebook by Udaykumar Dhokia (2025), titled “UPI Fraud Detection” (<https://www.kaggle.com/code/udaykumardhokia/upi-fraud-detection/notebook>). The dataset comprises 2,666 UPI transactions with 11 input features and one target variable (fraud_risk).^[12]

Dataset attributes

- Transaction identifiers: Transaction ID, UPI number (anonymized)
- Temporal features: trans_hour, trans_day, trans_month, trans_year
- Categorical features: category (e.g., retail, utility, peer-to-peer), state, zip code
- Numerical features: age and transaction amount (in INR)
- Target variable: fraud_risk (0: legitimate, 1: fraudulent)

Class distribution: The dataset exhibits class imbalance:

- Legitimate transactions (class 0): 2,074 (77.8%)
- Fraudulent transactions (class 1): 592 (22.2%)

3.2 Data preprocessing

The dataset has several attributes, including Id, trans_hour, trans_day, trans_month, trans_year, category, upi_number, age, trans_amount, state, zip, and fraud_risk. The data are then cleaned and transformed to maintain consistency and quality before model training. Missing values were treated using median imputation for numeric fields and mode imputation for categorical variables. Categorical features such as state, category, and upi_number were encoded into Label Encoding to make them machine learning algorithm friendly. Maintaining ordinal relationships by this approach keeps the computational efficiency intact. Continuous features such as transaction amount, user age, and transaction frequency are scaled using StandardScaler, and their feature distributions are normalized for better convergence and stability of the model at training itself. The count of legitimate transactions (class 0: 2,074) and fraudulent transactions (class 1: 592) from the dataset of 2,666 UPI transactions are shown in Fig. 2. The class distribution of 77.8% legitimate and 22.2% fraudulent confirms the class imbalance that required SMOTE application.

3.3 Feature engineering

To extract better insights, feature engineering was performed on the raw transaction data. Temporal features such as trans_hour, trans_day, and trans_month were explored to capture user spending behavior at different times. Therefore, aggregated metrics, including the average transaction amount per user, transaction frequency, and high-value transaction flags, were created to enhance model discriminability. These features have been included because evidence suggests that mule accounts often depict irregular timings and inconsistency




Fig. 2: Distribution of fraudulent vs non-fraudulent transactions.




Fig. 3: Transaction amount distribution by fraud risk.

in spending patterns compared with legitimate user patterns.^[21] Furthermore, duplicate or sequentially increasing UPI IDs were also analyzed to spot automated activity within the system.

The distribution of the transaction amounts for the fraudulent and legitimate classes is shown in Fig. 3. Box plot comparing transaction amounts for legitimate (class 0) and fraudulent (class 1) transactions. The y-axis shows transaction amounts ranging from 0 to more than 3,500 INR. Compared with legitimate transactions, fraudulent transactions exhibit higher median values and greater variability, confirming the discriminative value of amount-based features.

The correlation between the engineered and original features is represented in Fig. 4. Correlation matrix showing relationships between all features in the dataset, including Id, trans_hour, trans_day, trans_month, trans_year, category, upi_number, age, trans_amount, state, zip, and fraud_risk. The color bar ranges from -0.6 (negative correlation) to 1.0 (positive correlation). Features with stronger correlations to fraud_risk are more valuable for classification.

3.4 Handling imbalanced data

In most fraud datasets, including financial datasets, there are far fewer fraudulent transactions than legitimate transactions. Owing to this imbalance, the synthetic minority oversampling technique was utilized to generate synthetic samples for the minority class of interest (fraud).^[23] Financial fraud datasets typically exhibit severe class imbalance, with fraudulent transactions substantially outnumbered by legitimate ones. To address this, the synthetic minority oversampling technique (SMOTE) was applied to

generate synthetic samples for the minority (fraudulent) class.^[16]

SMOTE Implementation:

- Before SMOTE: 2,133 training samples (1,659 legitimate, 474 fraudulent)
- SMOTE parameters: k-neighbors = 5, sampling strategy = 1.0 (balance classes)
- After SMOTE: 3,318 training samples (1,659 legitimate, 1,659 fraudulent)

3.5 Model selection and training

Because of its high efficiency, ability to handle large datasets quickly, and good performance on tabular data, model training was performed using the light gradient boosting machine algorithm.^[22] LightGBM works according to the principle of gradient boosting: Trees are constructed consecutively, so every new tree corrects the mistakes of the previous ones. The model was optimized for the ROC-AUC metric during training because it provides a robust measure of the performance of classification models when the dataset is imbalanced. To set the optimal hyperparameters, cross-validation was applied by tuning the learning rate, number of leaves, maximum depth, and feature fraction. The model achieved an ROC-AUC score of more than 0.9, indicating its efficiency in classifying mule and genuine transactions. This performance metric shows a balance between sensitivity and specificity, hence making the system reliable for real-world fraud detection in the UPI ecosystem.

3.6 Model evaluation

The performance of the trained LightGBM model was assessed using common classification performance




Fig. 4: Feature correlation heatmap.

metrics such as the confusion matrix, accuracy, precision, recall, F1 score, and receiver operating characteristic-area under the curve (ROC-AUC). The abovementioned metrics are used to assess the performance of classification models, especially in fraud detection tasks where class imbalance is a common issue.

3.6.1 Confusion matrix

The confusion matrix represents the distribution of correctly and incorrectly classified instances and is defined as follows:

	Predicted Legitimate (0)	Predicted Fraud (1)
Actual Legitimate (0)	TN	FP
Actual Fraud (1)	FN	TP

where,

TP (True Positive): Fraud transactions correctly classified

TN (True Negative): Legitimate transactions correctly classified

FP (False Positive): Legitimate transactions incorrectly classified as fraud

FN (False Negative): Fraud transactions incorrectly classified as legitimate

On the basis of the model predictions, the confusion matrix obtained from the test dataset is shown in Fig. 5. Visualization of classification results on the test data (533 transactions). The matrix shows 398 true negatives, 114 true positives, 17 false positives, and 4 false negatives,

yielding 96.06% accuracy. The model demonstrates strong performance with minimal misclassifications.




Fig. 5: Confusion matrix.

3.6.2 Performance evaluation metric

The evaluation metrics are computed as follows:

Accuracy: Measures overall correctness of classification

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

$$Accuracy = \frac{512}{533} = 96.06\%$$

Precision (Fraud Class): Proportion of predicted fraud that is actually fraudulent

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

$$Precision = \frac{114}{114 + 17} = \frac{114}{131} = 0.87$$

Recall (Sensitivity) (Fraud Class): Proportion of actual fraud correctly identified

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

$$Recall = \frac{114}{114 + 4} = \frac{114}{118} = 0.966$$

F1-Score (Fraud Class): Harmonic mean of precision and recall

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

$$F1-Score = 2 \times \frac{0.87 \times 0.966}{0.87 + 0.966} = 2 \times \frac{0.8404}{1.836} = 0.916$$

3.6.3 Computed results

Using the above formulas, the LightGBM model achieved the following performance:

- Accuracy: 96.55%
- Precision (Fraud Class): 0.98
- Recall (Fraud Class): 0.96
- F1-Score (Fraud Class): 0.97
- ROC-AUC Score: 0.9962

The Receiver Operating Characteristic (ROC) curve corresponding to the trained model is illustrated in Fig. 6. ROC curve of the true positive rate against the false positive rate across different classification thresholds. The curve approaches the top-left corner, with an area under the curve (AUC) of 0.9962, demonstrating excellent discriminative ability between fraudulent and legitimate transactions.




Fig. 6: ROC curve.

3.7 Model deployment

The final trained model was serialized into .pkl files, along with preprocessing artifacts such as label encoders and

(2) scalers using Joblib. These were integrated into a Streamlit-based web application that enabled users to interactively test the system with new transaction data.

The app allows users to input transaction details and provides instantaneous predictions on whether a transaction or account could be fraudulent. Streamlit was selected because it is easily integrated with Python and is well positioned for rapid prototyping of machine learning interfaces. This deployment enhances model accessibility and showcases practical implementation during demonstrations and evaluations.

3.8 Summary

This methodology provides an end-to-end pipeline, from data collection and preparation to model deployment, with a focus on practical fraud detection in UPI transactions. The use of SMOTE for balancing, LightGBM for model training, and Streamlit for deployment ensures both technical robustness and demonstration feasibility, following modern standards in research on fintech fraud detection.^[24]

4. Results and analysis

4.1 Performance metrics

The performance of the trained LightGBM model was evaluated on various performance indicators—accuracy, precision, recall, F1 score, and receiver operating characteristic (ROC)-AUC, which help determine the performance of the model in detecting fraudulent UPI transactions. In the case of the overall accuracy for both fraudulent and legitimate transactions, a high rate of 96.55% was achieved. Additionally, the ROC-AUC score is very high (0.9962), demonstrating the great ability of the model to discriminate between the two classes. The classification report further highlights the classwise performance. For class 0 or legitimate transactions, the precision, recall, and F1 score were 0.94, 0.97, and 0.96, respectively, whereas for class 1 or fraudulent transactions, these values were 0.98, 0.96, and 0.97, respectively. A macro average F1 score of 0.96 and a weighted average F1 score of 0.97 confirm that the performance of the model is consistent under both balanced and imbalanced data conditions.

4.2 Analysis

The high accuracy and ROC-AUC scores indicate that the LightGBM algorithm managed to capture nonlinear relationships within the dataset and was thus able to distinguish subtle patterns of genuine users from mule account transactions. The slightly higher precision in the fraudulent class of 0.98 means that the model will be reliable in terms of not raising too many false positives—a prime necessity in a financial fraud detection system. For fraudulent cases, a recall value of 0.96 ensures that the majority of the actual fraud cases are picked up by the

system with few false negatives. Furthermore, from the results, it seems that the boosting approach in LightGBM handled the feature correlations and decision boundary complexities in UPI transaction data quite effectively. In comparison with standard classifiers such as logistic regression or decision trees, the ensemble-based approach of LightGBM helps reduce overfitting issues, hence increasing the generalization performance. The high value of the ROC-AUC confirms this further; even at various thresholds, the classifier shows great predictive separation between the classes.

The relative importance of the input variables in the classification decision is shown in Fig. 7. Ranking of features by their contribution to the LightGBM model's decisions. Transaction amount, temporal features (trans_hour, trans_day), and user age have the greatest importance, validating the value of behavioral feature engineering for detecting mule account activities.

Table 2: Comparative performance evaluation of logistic regression, random forest, and the proposed LightGBM model based on accuracy and ROC-AUC metrics.

Model	Accuracy	ROC-AUC
Logistic Regression	91.80%	0.94
Random Forest	94.20%	0.97
LightGBM (Proposed)	96.55%	0.9962

4.3 Interpretation

From a practical standpoint, the obtained metrics imply that this model is ready for real-world deployment in financial systems, especially for early fraud detection or identification of mule accounts. The close values of

precision and recall across both classes indicate a more balanced performance with minimal bias to any particular type of transaction.

The close-to-perfect ROC-AUC score of 0.9962 highlights the robustness of the LightGBM model because it can effectively prioritize suspicious transactions for manual review with minimal false alarms. Such accuracy during live deployment scenarios ensures appropriate allocations toward investigation resources for fast fraud mitigation. Integration of the model into the Streamlit-based web application provides visualizations of real-time predictions and batch analytics, making the results more interpretable for end users and decision makers. Its architecture is lightweight yet powerful, guaranteeing scalability across digital payment infrastructures—just what is needed for production-level fraud detection pipelines.

4.4 Baseline model comparison

To validate the effectiveness of the proposed LightGBM model, additional experiments were conducted using baseline classifiers, including logistic regression and random forest. All the models were trained using identical preprocessing procedures and evaluated using accuracy and ROC-AUC metrics on the test dataset.

The results indicate that LightGBM outperforms traditional classifiers in terms of both overall accuracy and discriminative ability. The superior ROC-AUC score demonstrates an enhanced ability to distinguish fraudulent transactions from legitimate transactions, confirming the suitability of gradient boosting for UPI mule account detection.




Fig. 7: Feature importance bar chart.

5. Future work

Although the model yielded very promising results, several avenues for future improvement still remain. First, the inclusion of more real-world large-scale transactional data in the dataset would help the model learn even richer behavioral patterns and improve generalizability across a wide range of banking environments. The incorporation of temporal and geospatial features, such as time windows between transactions and device fingerprints, might further increase the detection accuracy. Coupled with the integration of deep learning architectures, such as LSTM or transformer-based models, this approach may enable the tracking of sequential behavior and thus capture changing fraud patterns more effectively. Future iterations may also explore federated learning approaches that allow privacy-preserving collaboration among financial entities without centralized data sharing. Finally, turning the present system into a full-fledged, production-ready fraud prevention platform by enhancing the Streamlit application with real-time anomaly visualization dashboards and automated alert mechanisms will complete the vision. These improvements contribute to the broader goal of building intelligent, adaptive, and explainable AI solutions for the digital payment ecosystem.

6. Conclusion

The proposed research presents a robust and efficient machine learning-based framework to detect mule accounts in UPI transactions by employing the LightGBM algorithm. The model demonstrated an impressive 96.55% accuracy, with an ROC-AUC score of 0.9962, demonstrating a strong fraud detection capability with minimal misclassification errors. Extensive evaluation revealed that the model showed balanced performance concerning both the legitimate and fraudulent classes, ensuring reliability and fairness in financial transaction monitoring. Wrapping this pretrained model in an interactive application using Streamlit increases its practical value by offering a user-friendly interface for conducting real-time and batch-level fraud detection. This deployment highlights how data-driven methods can be effectively applied to real financial infrastructures to reduce the occurrence of fraud, enhance transaction security, and increase confidence in digital payment ecosystems. In this approach, the proposed solution effectively extracted complex transactional relationships with powerful gradient boosting in LightGBM and achieved strong generalization performance. Its near-perfect ROC-AUC score is an indication of the model's discriminative strength, making it a very reliable decision-support system for any financial institution in combat against UPI-based money mule operations.

CRedit Author Contribution Statement

Ishita Sonawane: Conceptualization; Methodology; Formal analysis; Project administration; Writing-original draft, Writing – review & editing. **Shradha Chavan:** Methodology, Supervision, Writing - review & editing. **Rudresh Shirwaikar:** Investigation; Validation; Visualization. **Priyanshi Suyal:** Project administration; Writing-original draft, Writing-review & editing. All authors have read and agreed to the published version of the manuscript.

Acknowledgement

The authors would like to extend their gratitude to Symbiosis Skills and Professional University (SSPU) for providing all the necessary support, academic guidance, and institutional resources, which helped immensely in the completion of this research work. The commitment of the university to innovation, research, and academic excellence played a very critical role in shaping the direction and quality of this study. We value the access to scholarly resources, research facilities, and administrative assistance available during the course of this work. The encouragement and support from SSPU have been of immense help for undertaking effective analysis, critical thinking, and a methodological approach toward the fulfillment of the research objectives.

Funding Declaration

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data Availability Statement

The dataset used in this study is publicly available on Kaggle and can be accessed at the following link: <https://www.kaggle.com/code/udaykumardhokia/upi-fraud-detection/notebook>.

Conflict of Interest

There are no conflicts of interest.

Artificial Intelligence (AI) Use Disclosure

The authors confirm that there was no use of artificial intelligence (AI)-assisted technology for assisting in the writing or editing of the manuscript and that no images were manipulated using AI.

Supporting Information

Not applicable.

References

- [1] M. U. Farukh, M. Taqi, K. R. Vemavarapu, S. M. Fadel, N. A. Khan, Fintech innovations and the transformation of rural financial ecosystems in

- India, *FinTech*, 2026, **5**, 3, doi: 10.3390/fintech5010003.
- [2] B. Mukhopadhyay, Understanding cashless payments in India, *Financial Innovation*, 2016, **2**, 27, doi: 10.1186/s40854-016-0047-4
- [3] M. I. Haque, Abdul Azeez N. P., S. M. Jawed Akhtar, UPI and financial inclusion in rural India: A case study, *Development and Sustainability in Economics and Finance*, 2025, **6**, 100056, doi: 10.1016/j.dsef.2025.100056.
- [4] H. Dev, R. Gupta, S. Dharmavaram, D. Kumar, From cash to cashless: UPI's impact on spending behavior among Indian users and prototyping financially responsible interfaces, arXiv preprint, 2024, doi: 10.48550/arxiv.2401.09937.
- [5] National Payments Corporation of India, 2024, UPI Product Statistics 2023-24, NPCI Annual Report, <https://www.npci.org.in/statistics>.
- [6] M. Schidlow, Forced fraud: The financial exploitation of human trafficking victims, *Social Sciences*, 2025, **14**, 398.
- [7] A. Mirian, J. DeBlasio, S. Savage, G. M. Voelker, K. Thomas, Hack for hire: Exploring the emerging market for account hijacking, *Proceedings of The Worldwide Web Conference*, 2019, 1279-1289. doi: 10.1145/3308558.3313489.
- [8] M. I. Abdul Rani, S. N. Faiza Syed Mustapha Nazri, S. Zolkafli, A systematic literature review of money mule: its roles, recruitment and awareness, *Journal of Financial Crime*, 2024, **31**, 347-361, doi: 10.1108/JFC-10-2022-0243.
- [9] Y. Chen, C. Zhao, Y. Xu, C. Nie, Y. Zhang, Deep learning in financial fraud detection: innovations, challenges, and applications, *Data Science and Management*, 2025, In Press, doi: 10.1016/j.dsm.2025.08.002.
- [10] S. S. Baah, H.T. Adu-Twum, S.O. Adjei, G. Ampadu, A.O. Martins, B. Fonkem, Leveraging big data analytics to combat emerging financial fraud schemes in the USA: A literature review and practical implications, *World Journal of Advanced Research and Reviews*, 2024, **24**, 17-43, doi: 10.30574/wjarr.2024.24.1.2999.
- [11] H. Banirostan, T. Banirostan, M. M. Pedram, A. M. Rahmani, Analysis and evaluation of various fraud detection methods for electronic payment cards transactions in big data, *Journal of Signal Processing Systems*, 2024, **96**, 849-870, doi: 10.1007/s11265-025-01947-w.
- [12] A. Gandhar, K. Gupta, A. K. Pandey, D. Raj, Fraud detection using machine learning and deep learning, *SN Computer Science*, 2024, **5**, doi: 10.1007/s42979-024-02772-x.
- [13] M. Hajihosseini, A. Maghsoudi, R. Ghezelbash, A novel scheme for mapping of MVT-Type Pb-Zn Prospectivity: LightGBM, a highly efficient gradient boosting decision tree machine learning algorithm, *Natural Resources Research*, 2023, **32**, 2417-2438, doi: 10.1007/s11053-023-10249-6.
- [14] N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, SMOTE: Synthetic Minority Oversampling Technique, *Journal of Artificial Intelligence Research*, 2002, **16**, 321-357, doi: 10.1613/jair.953.
- [15] U. Dhokia, UPI Fraud Detection, 2025, [Kaggle Notebook], <https://www.kaggle.com/code/udaykumardhok/a/upi-fraud-detection/notebook>.
- [16] Q. Zheng, C. Yu, J. Cao, Y. Xu, Q. Xing, Y. Jin, Advanced payment security system: XGBoost, LightGBM and SMOTE integrated, arXiv preprint, 2024, doi: 10.48550/arxiv.2406.04658.
- [17] J. K. Afriyie, K. Tawiah, W. A. Pels, S. Addai-Henne, H. A. Dwamena, E. O. Owiredu, S. A. Ayeh, J. Eshun, A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions, *Decision Analytics Journal*, 2023, **6**, 100163, doi: 10.1016/j.dajour.2023.100163.
- [18] V. Kamra, V. Varshney, S. Yadav, A novice approach for UPI fraud detection by using machine learning techniques, 2025 IEEE 7th International Conference on Computing, Communication and Automation (ICCCA), Greater Noida, India, 2025, 1-6, doi: 10.1109/ICCCA66364.2025.11325722.
- [19] Z. Zhao, T. Bai, Financial fraud detection and prediction in listed companies using SMOTE and machine learning algorithms, *Entropy*, 2022, **24**, 1157, doi: 10.3390/e24081157.
- [20] Z. Liu, A comparative study of machine learning methods in financial fraud detection, Proceedings of the 2024 2nd International Conference on Finance, Trade and Business Management (FTBM 2024), doi: 10.2991/978-94-6463-546-1_44.
- [21] O. Koc, O. Ugur, A. S. Kestel, The impact of feature selection and transformation on machine learning methods in determining credit scoring, arXiv preprint, 2023, doi: 10.48550/arxiv.2303.05427.
- [22] A. Papasavva, S. Johnson, E. Lowther, S. Lundrigan, E. Mariconti, A. Markovska, N. Tuptuk, Application of AI-based models for online fraud detection and analysis, arXiv preprint, 2024, doi: 10.48550/arxiv.2409.19022.
- [23] G. Kou, Y. Lu, Fintech: A literature review of emerging financial technologies and applications, *Financial Innovation*, 2025, **11**, doi: 10.1186/s40854-024-00668-6.
- [24] L. W. Rizkallah, Enhancing the performance of gradient boosting trees on regression problems, *Journal of Big Data*, 2025, **12**, doi: 10.1186/s40537-025-01071-3.

Publisher Note: The views, statements, and data in all publications solely belong to the authors and contributors. GR Scholastic is not responsible for any injury resulting from the ideas, methods, or products mentioned. GR Scholastic remains neutral regarding jurisdictional claims in published maps and institutional affiliations.