



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: X Month of publication: October 2025

DOI: <https://doi.org/10.22214/ijraset.2025.74922>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Fraud Detection in UPI Payments Using Tabular Machine Learning Models

Renu Chaudhary¹, Sakshi Singh², Riddhima Singh³, Husain Zaidi⁴, Kanishka Jain⁵

¹Department of Information Technology, HMR Institute of Technology and Management, Delhi, India

^{2,3,4,5}Department of Computer Science & Engineering, HMR Institute of Technology and Management, Delhi, India

Abstract: With the rapid growth of digital payment systems in India, the Unified Payments Interface (UPI) has become one of the main platforms for instant money transfers. However, as the number of transactions increases, the chances of fraud have also risen. This study presents a machine learning-based system to detect fraudulent UPI transactions using the CatBoost algorithm. The model uses important features related to user behavior, transaction details, and device information to identify whether a transaction is genuine or fraudulent. CatBoost is chosen because it works well with categorical data, provides clear results, and performs strongly on tabular datasets. The experimental results show that the model achieves a high AUC (Area Under the Curve), proving its strong ability to detect fraud. The trained model is also deployed in a Streamlit web application, allowing users to check fraud risk in real time through a simple interface. This system connects advanced machine learning with practical use, providing a reliable and scalable solution to improve the safety of UPI payments.

Keywords: Unified Payments Interface (UPI), Fraud Detection, Machine Learning, CatBoost, Financial Security, Digital Payments.

I. INTRODUCTION

India's digital finance revolution has transformed monetary transactions, with the Unified Payments Interface (UPI) emerging as a cornerstone of seamless, secure digital payments across urban and rural sectors. UPI's exponential growth has boosted financial inclusion and convenience [1], but also brought sophisticated cybersecurity challenges, as fraudsters exploit instant payments with advanced scams like phishing, spoofed handles, and fake links [2].

Traditional rule-based fraud detection systems in financial institutions are increasingly inadequate due to their static nature and poor adaptability to evolving attack strategies [3]. Machine Learning (ML), especially ensemble algorithms like XGBoost, offers superior capabilities by learning from data, identifying hidden fraud patterns, and flagging suspicious transactions in real-time [4]. Yet, handling UPI's highly imbalanced and predominantly categorical data poses challenges for conventional models, which can lose information through improper encoding and struggle to maintain interpretability [5].

CatBoost, a modern gradient boosting algorithm, directly addresses these limitations by efficiently managing categorical features and severe class imbalance, minimizing manual data preprocessing and supporting robust interpretability [5]. Despite its advantages, CatBoost's application to UPI fraud detection remains underexplored, marking a critical research gap [6]. This study proposes an end-to-end UPI fraud detection framework leveraging CatBoost, comparing its performance with XGBoost using metrics like ROC-AUC, precision, recall, and F1-score on imbalanced transaction datasets [4][5]. The research also delivers a real-time, interpretable fraud detection app built with Streamlit, showcasing operational deployment [7]. By bridging theory and practice, this work contributes both academically and operationally, reinforcing trust and resilience in India's digital financial ecosystem through advanced, ML-driven fraud detection [3].

II. LITERATURE REVIEW

A. Evolution

As digital transactions continue to proliferate worldwide, the detection of fraud within electronic payment systems has emerged as a focal point of scholarly inquiry, particularly due to its implications for financial security and trust in digital economies. The evolution of fraud detection methodologies reflects a progression from rudimentary rule-based frameworks to advanced machine learning architectures. Early approaches predominantly relied on manually engineered rules and threshold-based alerts, such as identifying transactions that exceeded predefined values or exhibited abnormal frequencies.

While these systems offered straightforward implementation and interpretability, they encountered considerable constraints in their ability to adapt to evolving attack patterns and generated a substantial proportion of false positives when applied to high-volume transactional environments [8].

Subsequent advancement brought statistical and classical machine learning models, most notably logistic regression. Logistic regression has been favored for its simplicity and explainability, especially when data relationships are linear or feature spaces are limited. However, its performance diminishes in the face of complex, non-linear relationships and high-dimensional transaction data typical of digital payment platforms. In response, tree-based and ensemble methods including Random Forest, XGBoost, and LightGBM have gained prominence due to their capacity to model non-linear patterns, manage noisy datasets, and consistently deliver superior predictive accuracy on tabular financial data. Both XGBoost and LightGBM remain widely adopted owing to their computational efficiency and strong empirical results in benchmarking studies concerning fraud detection [9].

With the maturation of Unified Payments Interface (UPI) and mobile payment systems in India, recent research efforts increasingly incorporate ensemble learning frameworks to address the heightened sophistication of fraudulent activities. Empirical studies demonstrate that these methods outperform classical classifiers and rule-based baselines for transaction-level fraud identification, often integrating session and device features alongside temporal transaction characteristics. However, considerable research gaps remain: many studies emphasize validation on isolated datasets and frequently neglect real-world deployment considerations, particularly challenges related to latency, scalability, and adaptability in live UPI environments. Several recurrent limitations are observed in the literature. First, high-cardinality categorical features such as Device_ID, UPI identifiers, and merchant codes are frequently suboptimally encoded. Common strategies like one-hot encoding or ordinal labeling result in expanded feature spaces or unintended biases, negatively impacting model performance and scalability. Second, imbalanced datasets, where fraudulent transactions are rare relative to legitimate ones, are often inadequately addressed; some studies disregard essential resampling, class-weighting, or discriminative evaluation metrics, thereby yielding overly optimistic results that do not generalize to operational environments. Third, the dimension of model explainability and deployment readiness is often overlooked. Interpretability of model outputs crucial for financial institutions and regulatory oversight is seldom prioritized, and live, low-latency scoring pipelines necessary for real-time fraud alerting are rarely discussed.

B. Advances in UPI and Mobile Payment Fraud Detection

In recent years, with the maturation of Unified Payments Interface (UPI) and mobile payment systems in India, recent research efforts increasingly incorporate ensemble learning frameworks to address the heightened sophistication of fraudulent activities. Empirical studies demonstrate that these methods outperform classical classifiers and rule-based baselines for transaction-level fraud identification, often integrating session and device features alongside temporal transaction characteristics. However, considerable research gaps remain: many studies emphasize validation on isolated datasets and frequently neglect real-world deployment considerations, particularly challenges related to latency, scalability, and adaptability in live UPI environments [10].

However, existing studies exhibit certain limitations. They rely on isolated, static datasets and overlook aspects critical to real-world deployment, such as model latency, scalability, and adaptability to evolving fraud patterns in live environments. Consequently, despite strong offline performance, practical applicability remains constrained.

C. Persistent Challenges in the Literature

Several recurring challenges emerge from the existing body of research:

- 1) *Limited Focus on UPI-specific Transaction Behavior:* Most fraud detection studies focus on credit cards or international banking systems, with very few exploring India's UPI ecosystem. This paper addresses that gap by modeling UPI transaction patterns and fraud risk factors unique to the Indian digital payment landscape.
- 2) *Underutilization of Tabular ML Models for Financial Fraud:* Prior works emphasize deep learning or text-based NLP models, often overlooking high-performing tabular algorithms like CatBoost and XGBoost that are well-suited for structured financial data. This study demonstrates their comparative and practical effectiveness.
- 3) *Lack of Deployable Real-time Systems:* Many research efforts stop at offline model evaluation. This work extends further by integrating the trained model into a real-time Streamlit-based prediction system ("Quicki"), bridging the gap between research and operational fraud detection deployment.

The observed limitations motivate the exploration of more advanced machine learning models. CatBoost, a modern gradient boosting algorithm, has been chosen for this study due to its distinct advantages: its native capability to handle categorical variables through ordered boosting and target-based encoding that minimizes information leakage; its demonstrated robustness and accuracy on heterogeneous transaction data; and its interpretable feature importance outputs that facilitate transparency and regulatory compliance. Its native capability to handle categorical variables through ordered boosting and target-based encoding, minimizing information leakage [11].

III. METHODOLOGY

The proposed methodology for UPI fraud detection involves a systematic process comprising synthetic data generation, feature engineering, model development using CatBoost, evaluation, and deployment through a Streamlit application. The overall workflow is illustrated in Figure 1.




Fig. 1 Workflow of the proposed UPI fraud detection model

A. Experimental Setup

All experiments were conducted on Google Colab, utilizing a virtual machine equipped with an NVIDIA Tesla T4 GPU, 13 GB RAM, and Python 3.10 environment. Core libraries used include *CatBoost*, *scikit-learn*, *matplotlib*, and *pandas*.

The entire pipeline from data preprocessing to visualization was executed within the Colab environment. The model training and testing processes were automated, ensuring reproducibility and scalability. The resulting trained model was serialized and exported for deployment in a web-based fraud detection system.

B. Dataset Description

Due to the unavailability of publicly accessible UPI transaction data, a synthetic dataset was programmatically generated using Python's *Faker* library and controlled probabilistic logic. A total of 100,000 transactions were simulated between 30,000 unique senders and 30,000 unique receivers, spanning the period January 2024 to December 2024. Each transaction record consists of 25 attributes, broadly categorized as:

- 1) *Identification Attributes*: Transaction ID, Timestamp, Sender/Receiver UPI IDs, Device ID
- 2) *Transaction Attributes*: Amount (INR), Transaction Type, Merchant Category, Channel
- 3) *Behavioral Attributes*: Number of Transactions in Last 24H, Average Transaction Amount (7 days), Account Age
- 4) *Risk & Context Attributes*: IP Risk Score, Device Change Flag, Is_Night_Txn, Previous Fraud Count (Sender/Receiver)
- 5) *Label*: Is_Fraud (binary class, 1 for fraud, 0 for legitimate)

The final dataset achieved a fraud rate of 0.628%, ensuring class imbalance consistent with real-world scenarios in digital payments. Table I presents the overall distribution of legitimate and fraudulent transactions within the generated UPI dataset.

C. Data Preprocessing

Data preprocessing played a crucial role in improving the quality and robustness of the dataset before model training. The dataset was examined for missing values and outliers. Missing numerical values were imputed using the median of the respective feature. Categorical fields such as *Device_ID* and *Channel* were filled with the mode (most frequent value) of the column. Outliers in the transaction amount were handled using the Interquartile Range (IQR) technique to reduce bias. Any data point outside the acceptable range was capped as in Eq. (1).

$$\text{Lower Bound} = Q1 - 1.5 \times \text{IQR}, \text{Upper Bound} = Q3 + 1.5 \times \text{IQR} \quad (1)$$

Here, $\text{IQR} = Q3 - Q1$. Since CatBoost inherently supports categorical features, manual encoding such as one-hot or label encoding was not required. All categorical attributes were converted to string format. The dataset was divided into training and testing subsets using an 80:20 stratified split, ensuring that both sets preserved the ratio of fraudulent to legitimate transactions. Additional behavioral features were engineered to enhance fraud detection capability. These included:

- 1) *Num_Txns_Last_24H*: number of transactions in the last 24 hours.
- 2) *Avg_Amount_Last_7d*: mean transaction amount in the last seven days.
- 3) *Device_Change_Flag*: binary indicator of whether a transaction occurred from a new device.
- 4) *Account_Age_Days*: number of days since account creation.

These features helped the model capture temporal and behavioral irregularities, such as sudden spikes in transaction activity or changes in device usage, which often signify potential fraudulent behavior.

Table I
Distribution Of Legitimate And Fraudulent Transactions

Transaction Type	Count	Percentage
Legitimate (Non-Fraud)	99,372	99.37
Fraudulent (Fraud)	628	0.63

D. Model Architecture and Training

The proposed model utilizes the CatBoost Classifier, a gradient boosting algorithm designed to efficiently handle categorical variables without the need for one-hot encoding. CatBoost implements ordered boosting and built-in regularization to minimize overfitting. The model is optimized using the logarithmic loss function, expressed as in Eq. (2) where y_i denotes the true class label, and p_i represents the predicted probability for the positive (fraud) class.

$$L = -\frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \quad (2)$$

The model was configured with the following hyperparameters:

- 1) *Iterations*: 1000
- 2) *Learning Rate*: 0.05
- 3) *Depth*: 8
- 4) *Loss Function*: Logloss

The model was trained using the training subset, with AUC (Area Under ROC Curve) as the optimization metric. Early stopping was employed to ensure optimal convergence. The model achieved convergence after approximately 6–8 minutes, producing stable and reproducible results due to a fixed random seed.

E. Evaluation Metrics

The model's performance was assessed using multiple evaluation metrics suitable for imbalanced binary classification problems, including Precision, Recall, F1-score, and ROC-AUC. The ROC curve was plotted to visualize the trade-off between true positive and false positive rates. The ROC-AUC metric was computed as in Eq. (3).

$$AUC = \int_0^1 TPR(FPR)d(FPR) \quad (3)$$

F. System Deployment

To demonstrate real-world applicability, the trained CatBoost model was integrated into a Streamlit web application named "Quicki", designed to detect fraudulent UPI transactions in real time. The interface allows users to input essential transaction details such as Amount, Transaction_Type, Merchant_Category, Channel, and Device_Type. Upon submission, the system computes a fraud probability score and classifies the transaction as either "Legitimate" or "Fraudulent."

The application provides an interactive visualization of results, ensuring transparency in predictions and facilitating easy integration with financial platforms. This deployment framework validates the practicality of the CatBoost-based model for real-time fraud monitoring within UPI payment systems.

IV. RESULTS AND DISCUSSION

To evaluate the effectiveness of different gradient boosting algorithms, two models XGBoost and CatBoost were trained and tested on the same dataset configuration. The comparison focused on key evaluation metrics such as Precision, Recall, F1-score, and ROC-AUC, which are particularly relevant for imbalanced fraud detection tasks. Table II presents a comparative summary of the performance metrics for both gradient boosting models. The XGBoost model achieved a ROC-AUC of 0.8443 and demonstrated higher precision for fraudulent cases (0.61), indicating its ability to minimize false positives. In contrast, the CatBoost model recorded a superior ROC-AUC of 0.8696 and higher recall (0.60), reflecting its stronger ability to correctly identify fraudulent transactions. Consequently, metrics such as recall and ROC-AUC provide a more reliable assessment of fraud detection performance.

TABLE II
COMPARATIVE PERFORMANCE OF BOTH MODELS

Model	ROC-AUC	Precision(Fraud)	Recall (Fraud)	F1-Score (Fraud)	Accuracy
XGBoost	0.8443	0.6163	0.4206	0.5000	0.9947
CatBoost	0.8696	0.4900	0.6000	0.5400	0.9900

The visualizations below provide valuable insights into the model's decision-making process. Fig. 2 highlights that *transaction amount*, *number of recent transactions*, and *night-time transactions* are among the most influential predictors in identifying fraud. The ROC curve (Fig. 3) demonstrates a strong discriminative ability of the model with an AUC of 0.87, while the confusion matrix (Fig. 4) confirms high accuracy in distinguishing legitimate and fraudulent transactions. Together, these results validate the effectiveness of the CatBoost model in real-world UPI fraud detection scenarios.




Fig. 2 Feature importance plot showing the most influential predictors in UPI fraud detection.




Fig. 3 ROC Curve for Fraud Detection




Fig. 4 Confusion matrix

V. CONCLUSION AND FUTURE SCOPE

This study presents a CatBoost-based fraud detection framework for Unified Payments Interface (UPI) transactions, designed to identify anomalous and high-risk activities in real time. Using a synthetically generated dataset of 100,000 transactions with realistic behavioral and contextual features, the model demonstrated strong predictive capability. The CatBoost classifier achieved a ROC-AUC score of 0.8696, precision of 0.49, and recall of 0.60 for the minority (fraudulent) class, reflecting a balanced trade-off between fraud detection sensitivity and false-positive control.

Compared to XGBoost, CatBoost exhibited superior performance in handling categorical features and imbalanced data, achieving higher recall and overall interpretability. The inclusion of engineered features such as *Device_Change_Flag*, *Num_Txns_Last_24H*, and *IP_Risk_Score* significantly improved fraud detection accuracy by capturing behavioral anomalies.

In future work, the model can be further enhanced by incorporating advanced class imbalance handling techniques such as Synthetic Minority Oversampling (SMOTE), focal loss optimization, or dynamic threshold adjustment to improve recall for rare fraud events.



Moreover, integrating temporal and graph-based network features could enable detection of coordinated fraud rings, while leveraging deep learning architectures for sequential transaction modeling. Real-time deployment through scalable APIs on UPI platforms will further strengthen proactive fraud prevention in digital payment ecosystems.

REFERENCES

- [1] S. Ghosh, S. D. Kesar, and M. Mukherjee, "Digital payments, UPI and financial inclusion in India: An empirical analysis," *Financial Innovation*, vol. 8, no. 1, pp. 84-97, 2022.
- [2] K. Gupta and R. Kohli, "Cyber frauds in digital payments: Phishing, spoofing and the UPI challenge," *Journal of Payments Strategy & Systems*, vol. 15, no. 2, pp. 105-113, 2021.
- [3] R. Singh and P. Sharma, "Challenges of rule-based fraud detection in rapidly evolving digital payments," *Journal of Banking Technology*, vol. 12, pp. 45-53, 2020.
- [4] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, pp. 785-794, 2016.
- [5] T. He, J. Yang, S. Chen, et al., "Handling class imbalance and categorical data in financial fraud detection," *IEEE Access*, vol. 9, pp. 113572-113582, 2021.
- [6] L. Prokhorenkova, G. Gusev, A. Vorobev, A. Dorogush, and A. Gulin, "CatBoost: unbiased boosting with categorical features," in *Advances in Neural Information Processing Systems*, vol. 31, pp. 6638-6648, 2018.
- [7] V. Kumar, S. Srinivasan, and N. Menon, "Deploying real-time ML-powered fraud detection using Streamlit for financial institutions," in *Proc. IEEE Int. Conf. FinTech*, pp. 347-354, 2023.
- [8] P. Jeyachandran, "Leveraging Machine Learning for Real-Time Fraud Detection in Digital Payments," *SSRN Electronic Journal*, 2024.
- [9] Y. Ding, H. Li, X. Zhou, and J. Zhang, "Digital Payment Fraud Detection Methods in Digital Ages and Industry 4.0," *Computers & Electrical Engineering*, vol. 102, 2022.
- [10] N. Lingareddy, "Enhancing Digital Payment Security: UPI Fraud Detection," *IEEE Xplore Digital Library*, 2025.
- [11] R. E. Sorace, V. S. Reinhardt, and S. A. Vaughn, "High-speed digital-to-RF converter," *U.S. Patent 5 668 842*, Sept. 16, 1997.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)