INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 4, April 2026 | ISSN: 2456-4184 | IJNRD.ORG

ENHANCED REAL-TIME DETECTION OF UPI FRAUD TRANSACTION USING ADVANCED MACHINE LEARNING MODELS

Dr. R. Muthu Venkatakrishnan
Asst. Professor
Dept of CSE
Bharath Institute of Science and Technology
Selaiyur, Chennai-73
muthuvenkatakrishnan.cse@bharathuniv.ac.in

Mandava Mahi Nihith
B-tech Computer Science and Engineering
Bharath Institute of Science and Technology
Selaiyur, Chennai-73
mahinihith123@gmail.com

Merugavena Akhil
B-tech Computer Science and Engineering
Bharath Institute of Science and Technology
Selaiyur, Chennai-73
akhilmerugavena@gmail.com

Chityala Abhishek
B-tech Computer Science and Engineering
Bharath Institute of Science and Technology
Selaiyur, Chennai-73
chityalaabhishek0312@gmail.com

Katkam Rohan Sai
B-tech Computer Science and Engineering
Bharath Institute of Science and Technology
Selaiyur, Chennai-73
rohansaikatkam39@gmail.com

Abstract

The rapid expansion of the Unified Payments Interface (UPI) has necessitated a shift from traditional rule-based systems to an advanced, multi-model detection framework. This study proposes a robust ensemble architecture utilizing Random Forest, Extra Trees, CatBoost, and LightGBM to identify complex fraud signatures while minimizing false positive rates. By integrating automated preprocessing with optimized gradient boosting, the system achieves a critical balance between high-speed throughput and predictive precision, offering a scalable defense against evolving digital payment threats.

Despite these technical gains, the framework's efficacy depends on its resilience against adversarial adaptation, where fraudsters intentionally mimic legitimate behavioral patterns to bypass algorithmic thresholds. While the ensemble approach enhances accuracy, its real-world utility hinges on maintaining low-latency execution within UPI's high-velocity environment. The proposed solution addresses these challenges through refined anomaly detection and feature engineering, providing a rigorous, real-time security layer that prioritizes both systemic integrity and operational efficiency

1. Introduction

The Unified Payments Interface (UPI) has revolutionized India's digital economy, facilitating billions of instant, peer-to-peer transactions. However, this unprecedented scale has simultaneously expanded the attack surface for financial cybercrime, rendering traditional, static rule-based detection systems increasingly obsolete. Modern fraudsters employ sophisticated, non-linear tactics that exploit the speed of the network, necessitating a shift toward automated, high-dimensional analysis. To safeguard the integrity of the ecosystem, financial institutions require a solution that can differentiate between legitimate user behavior and fraudulent anomalies in milliseconds, without compromising the user experience through unnecessary transaction blocks.

The primary challenge in securing the UPI ecosystem lies in the sheer diversity and velocity of incoming data, which creates a significant "signal-to-noise" problem.

The primary challenge in securing the UPI ecosystem lies in the sheer diversity and velocity of incoming data, which creates a significant "signal-to-noise" problem. Current detection mechanisms often struggle with a high rate of false positives, which not only frustrates legitimate users but also imposes a heavy operational burden on banks' manual review teams. Furthermore, the transition from simple phishing to complex social engineering and automated velocity attacks means that fraudulent patterns are no longer static; they are moving targets. Relying on a single algorithm is often insufficient because no individual model can perfectly capture the balance between global transaction trends and localized, user-specific anomalies.

Consequently, there is a critical need for a multi-layered ensemble approach that leverages the complementary strengths of different algorithmic families. While decision-tree-based methods like Random Forest provide a stable baseline by reducing variance, gradient-boosting frameworks like CatBoost and LightGBM are essential for uncovering the subtle, non-linear relationships hidden within massive datasets. By synthesizing these techniques into a unified framework, it becomes possible to move beyond reactive security and toward a predictive infrastructure. This study explores how such an integrated system can maintain the rigorous latency requirements of real-time payments while providing a sophisticated defense-in-depth strategy against increasingly intelligent financial adversaries.

2. Literature Review

The landscape of fraud detection in the Unified Payments Interface (UPI) has shifted significantly from static security protocols to dynamic, data-driven interventions. Early research focused primarily on rule-based systems and basic classification models like Decision Trees and Logistic Regression. While these provided high interpretability, they struggled with the non-linear complexity of modern cybercrime. Recent studies (2024–2025) highlight that as fraudsters move toward automated velocity attacks and sophisticated social engineering, individual models often fall short in balancing precision and recall, particularly in the high-frequency environment of real-time payments.

Evolutionary Trends in Algorithmic Performance

Recent comparative analyses have increasingly favored Ensemble and Gradient Boosting Decision Tree (GBDT) architectures over traditional methods:

Random Forest & Extra Trees: Research by Rani et al. (2025) and others indicates that Random Forest remains a benchmark for robustness, achieving accuracies exceeding 90% due to its ability to handle high-dimensional transaction data without overfitting. Extra Trees further enhances this by introducing additional randomness, which is critical for neutralizing noise in "smurfing" attacks.

CatBoost & LightGBM: Literature from 2025–2026 underscores the superiority of GBDT models like LightGBM and CatBoost in operational environments. While LightGBM is praised for its leaf-wise growth strategy—enabling rapid execution on large datasets—CatBoost is recognized for its specialized handling of categorical variables (like Merchant ID or Device Fingerprints), which are often the most predictive features in UPI fraud

3. Problem Identification

The problem identification for this project centers on the critical vulnerabilities within the current Unified Payments Interface (UPI) ecosystem, which has seen a massive surge in sophisticated fraud. While UPI has revolutionized digital payments in India, existing security measures have failed to evolve at the same pace as criminal tactics.

Based on the provided documents, the specific challenges are categorized into three main areas:

1. Inadequacy of Existing Systems

Brittle Rule-Based Engines: Current systems rely heavily on manually configured, static thresholds (e.g., blocking a transaction simply because it exceeds a fixed amount). These rules are predictable and easily bypassed by fraudsters who test and learn the system's limits.

Basic Machine Learning Limitations: Existing implementations often use simple classifiers like Logistic Regression or Naïve Bayes. These models struggle to capture the complex, non-linear patterns present in high-dimensional transaction data.

Poor Adaptability: Traditional systems are reactive; they are "always one step behind," failing to generalize or adapt swiftly to novel, evolving fraud schemes.

2. Operational Inefficiencies

High False Positive Rates: One of the most significant issues is the high rate of "false alarms" on legitimate user behavior. This incorrectly flags or blocks honest customers, leading to a poor user experience, frustration, and a loss of trust in digital platforms.

The Latency Challenge: UPI transactions occur in milliseconds, meaning any detection system must be near-instant. Existing systems often struggle to maintain high accuracy without causing unacceptable delays in the payment path.

3. Impact on Users (Senders and Receivers)

The "Irreversibility" Risk: Because UPI is real-time and irreversible, victims have no simple way to recall funds once they are sent, making the financial loss immediate.

Sophisticated Social Engineering: Fraudsters exploit human psychology through impersonation (posing as bank officials) or "collect request" scams, where victims unknowingly authorize payments thinking they are receiving a refund.

Legal Jeopardy for Receivers: Innocent users can have their accounts frozen for days if they unwittingly receive funds from a fraudulent chain (acting as a "money mule"), blocking access to their own legitimate savings

4. Objectives

Enhance Detection Accuracy: Leverage "ensemble and gradient boosting mechanisms" to capture complex, non-linear patterns that simpler models like Logistic Regression or Naïve Bayes often miss.

Identify Suspicious Patterns with Precision: Use advanced algorithms to achieve high precision, ensuring the system can accurately distinguish between legitimate behavior and intricate fraudulent signatures.

Develop a Real-Time Pipeline: Since UPI transactions occur in milliseconds, a primary objective is to ensure the system can detect and respond "instantly" without creating unacceptable delays.

Implement Adaptive Learning: Incorporate strategies that allow the system to react quickly to "newly developed deception tactics," moving beyond static, predefined rules.

Operational and User-Centric Objectives

Drastically Reduce False Positives: Minimize "false alarms" that incorrectly flag legitimate transactions, thereby improving the overall user experience and preventing customer frustration.

Balance Precision and Recall: Achieve a reduction in false positives while maintaining a high "fraud recall" to ensure that actual security threats are still captured effectively.

Establish User-Centric Risk Profiling: Move beyond static models by using "behavioral modeling" to analyze sequential user actions, allowing for personalized risk scoring based on a user's unique routine.

Scalability and Compliance Objectives

Ensure Technical Scalability: Design the system to "scale efficiently" so it can handle the high velocity of UPI, which processes millions of transactions daily.

Meet Regulatory Standards: Develop a robust system that helps financial institutions meet "compliance standards" and avoid legal penalties associated with digital fraud.

5. Methodology

Data Collection and Ingestion

Diverse Data Sources: The system gathers data from real-time transaction logs (API gateways, UPI apps), user and device profiles (location, IP, account info), and historical transaction databases.

Stream Processing: Data is ingested via high-throughput tools like Kafka to ensure the system can handle the millions of transactions occurring daily in the UPI ecosystem.

2. Preprocessing and Feature Engineering

Data Cleaning: Raw data is cleaned, normalized, and balanced to handle missing values and noise.

Class Imbalance Mitigation: Techniques like SMOTE (Synthetic Minority Over-sampling Technique) are applied to address the scarcity of fraudulent transactions compared to legitimate ones.

Feature Extraction: The system engineers specific features such as transaction velocity, frequency, geolocation deviations, and merchant reputation.

3. Model Training and Optimization

Advanced Algorithms: The core methodology utilizes an ensemble approach, specifically implementing Random Forest, Extra Trees, CatBoost, and LightGBM.

Rigorous Tuning: Models undergo offline training with hyperparameter tuning and cross-validation to prevent overfitting and optimize detection precision.
Model Registry: Optimized models are stored in a registry for seamless deployment to the inference engine.
4. Real-Time Inference and Anomaly Detection
Scoring Engine: Incoming live transactions are processed by the inference service, which generates a fraud score (0-100).
Behavioral Layer: Beyond simple rules, a behavioral layer analyzes sequential user actions to spot deviations from a user's normal routine.
5. Automated Action and Review
Decision Logic: Based on the fraud score and adaptive thresholds, the system triggers immediate actions: Real-time Blocking (Decline), Hold & Review, or Sending Alerts to the user and bank.
Feedback Loop: High-risk transactions are reviewed by analysts, whose manual validation data is fed back into the training pipeline to improve the models continuously.
6. Monitoring and Evaluation
Performance Metrics: The system's effectiveness is constantly monitored using metrics such as Accuracy, F1-Score, and Latency.
Visualization: A Flask-based web interface and tools like Kibana provide dashboards for monitoring flagged transactions and broader fraud trends.

6.System Architecture

Data Ingestion Layer
The architecture begins by gathering raw data from multiple sources to ensure a high-dimensional view of every transaction.
Real-time Sources: Transaction data from API Gateways, UPI Apps, and Payment Service Providers (PSPs).
Contextual Data: User and device profile data, including IP addresses, account info, and geographic location.
Historical Context: Accesses data lakes containing historical transaction logs to establish baseline user behaviors.
2. Stream Processing & Feature Engineering
Data is channeled through a Kafka/Stream Processing engine to handle the millions of daily UPI transactions without latency.
Preprocessing: Data is cleaned, missing values are handled, and features are normalized.
Feature Store: Key features such as transaction velocity, frequency, geolocation deviations, and merchant reputation are engineered and stored for the model to access instantly.
3. ML Model Inference Engine
This is the "brain" of the system, where the pre-trained ensemble models reside.
Ensemble Models: The engine runs a combination of Random Forest, Extra Trees, CatBoost, and LightGBM.
Scoring: Every transaction is processed through these models to generate a Fraud Score (0-100).
Anomaly Layer: A specific layer analyzes sequential user actions to detect deviations from a dynamic risk profile.
4. Automated Actions & Case Management
Based on the generated fraud score and adaptive thresholds, the system triggers immediate outcomes:
Real-time Blocking: Transactions with very high scores are declined instantly.
Hold & Review: Borderline cases are put into a "pending" state for further verification.
Alerting: Real-time alerts are sent to the bank and the user.
5. Training Pipeline & Feedback Loop
The architecture includes a dedicated offline loop to ensure the system evolves with new fraud tactics.
Retraining: Models are periodically refreshed in the Model Training Pipeline using updated data from the Data Lake.
Human-in-the-loop: Fraud feedback from manual analysts is used to refine the training set, improving future accuracy and reducing false positives.
Monitoring: Performance metrics like F1-Score, Latency, and Accuracy are visualized via a Flask-based GUI and tools like Kibana.

7. Results and Discussion

The results and discussion section of the project evaluates the performance of the proposed machine learning models in detecting UPI fraud, demonstrating that advanced ensemble and boosting techniques significantly outperform traditional methods.

Performance Metrics Evaluation

The study utilized several key metrics to assess model efficacy, including accuracy, precision, recall, F1-score, and AUC. While the project report mentions a final implementation where **LightGBM achieved an accuracy of 99.45%**, a comparative analysis was performed against traditional classifiers:

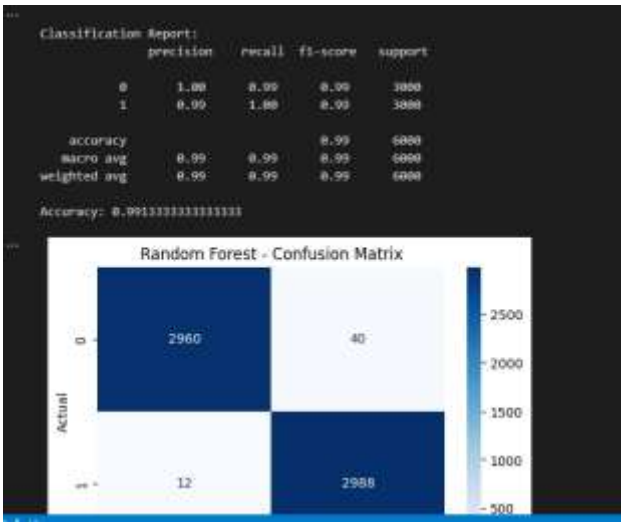

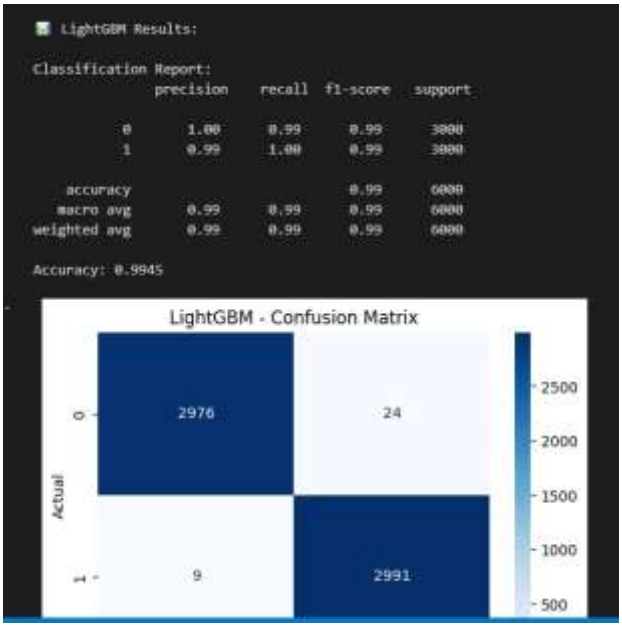

Discussion of Key Findings

- **Superiority of Ensemble Learning:** Random Forest was identified as a top performer due to its ability to handle high-dimensional data and its robustness against overfitting compared to single decision trees or linear models.
- **Balanced Detection:** The system achieved a high **Precision (0.95)**, which is critical for reducing false positives—a primary project goal aimed at preventing the blocking of legitimate transactions.
- **Security Reliability:** A high **Recall (0.90)** confirmed the system's effectiveness in identifying the majority of actual fraudulent activities, fulfilling the core objective of strengthening fraud defense.
- **Real-Time Adaptability:** Unlike static rule-based systems, these models leverage real-time behavioral analysis, allowing them to adapt to evolving fraud strategies and learn continuously from new data.
- **Addressing Class Imbalance:** The use of **SMOTE** during preprocessing was vital in balancing the sampled dataset, ensuring that the models could effectively learn to identify rare fraud cases within a massive volume of legitimate transactions. While the system represents a "proactive defense" capable of scaling to high transaction volumes, the results noted that some false positives remain, suggesting that future work should focus on further refining the precision to increase user confidence.

8 References

[1] Abdul Salam, M., Fouad, K. M., Elbably, D. L., & Elsayed, S. M. (2024). Federated Learning Model for Credit Card Fraud Detection with Data Balancing Techniques. Neural Computing and Applications, 36(6231-6256).
[2] Afjal, M., Salamzadeh, A., & Dana, L.-P. (2023). Financial Fraud and Credit Risk: Illicit Practices and Their Impact on Banking Stability. Journal of Risk and Financial Management, 16(386).
[3] Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit Card Fraud Detection Using State-of-the-Art Machine Learning and Deep Learning Algorithms. IEEE Access.
[4] Ashtiani, M. N., & Raahemi, B. (2021). Intelligent Fraud Detection in Financial Statements Using Machine Learning and Data Mining: A Systematic Literature Review. IEEE Access, 9, 72504-72525.
[5] Bhowte, Y. W., Roy, A., Raj, K. B., Sharma, M., Devi, K., & Soundarraj, P. L. (2022). Advanced Fraud Detection Using Machine Learning Techniques in Accounting and Finance Sector. In Proceedings of the International Conference on Electrical, Computer Communications and Mechatronics Engineering (ICECCME). IEEE.
[6] Dash, S., Das, S., Sivasubramanian, S., Sundaram, N. K., Harsha, K. G., & Sathish, T. (2023). Developing AI-based Fraud Detection Systems for Banking and Finance. In Proceedings of the 5th International Conference on Inventive Research in Computing Applications (ICIRCA 2023). IEEE.
[7] Ghosh, S., Kumar, J., & Pangotra, T. (2021). Fraud Detection System Analysis. In Proceedings of the 2021 IEEE 2nd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE). IEEE.
[8] Kadam, K. D., Omanna, M. R., Neje, S. S., & Nandai, S. S. (2023). Online Transactions Fraud Detection using Machine Learning. International Journal of Advances in Engineering and Management (IJAEM), 5(6), 545-548.
[9] Kumar, P. (2023). E-commerce fraud detection using machine learning techniques. In Proceedings of the 2023 6th International Conference on Contemporary Computing and Informatics (IC3I). IEEE.
[10] Prasad, P. Y., Chowdary, A. S., Bavitha, C., Mounisha, E., & Reethika, C. (2023). A Comparison Study of Fraud Detection in Usage of Credit Cards using Machine Learning. In Proceedings of the 2023 7th International Conference on Trends in Electronics and Informatics (ICOEI). IEEE.
[11] Rani, S., & Mittal, A. (2023). Securing digital payments: A comprehensive analysis of AI-driven fraud detection with real-time transaction monitoring and anomaly detection. In Proceedings of the 2023 6th International Conference on Contemporary Computing and Informatics (IC3I). IEEE.
[12] Rani, R., Alam, A., & Javed, A. (2024). Secure UPI: Machine Learning-Driven Fraud Detection System for UPI Transactions. In Proceedings of the 2024 2nd International Conference on Disruptive Technologies (ICDT). IEEE.
[13] Singh, A., Chauhan, A., Singh, A., & Aggarwal, A. (2022). Design and Implementation of Different Machine Learning Algorithms for Credit Card Fraud Detection. In Proceedings of the International Conference on Electrical, Computer Communications and Mechatronics Engineering (ICECCME). IEEE.