International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 14 Issue: 3
Article Received: *25 July 2026* ***Revised:*** *12 August 2026* ***Accepted:*** *08 September 2026*

Hybrid Machine Learning-Based Data Mining Framework for Early Fraud Detection in Digital Transactions

Dr. R. Aravind, Ms .K. S. Kanya

Research Supervisor, Department of Computer Science, Gobi Arts & Science College

Research Scholar, Department of Computer Science, Gobi Arts & Science College, Abstract

Abstract

The exponential growth of global electronic payment systems has simultaneously catalysed a dramatic surge in sophisticated digital transaction fraud. Traditional rule-based engines and standalone machine learning classifiers face critical bottlenecks: they either struggle with high false-alarm rates or lack real-time computational scalability when handling massively imbalanced streaming datasets. This paper presents a novel, two-stage HybridMachine Learning-Based Data Mining Framework designed specifically for the ultra-early isolation of fraudulent patterns. The first stage employs unsupervised Autoencoders to conduct low-latency feature extraction and isolate non-linear transaction anomalies from heavily skewed datasets. The second stage feeds these optimal feature mappings into an optimized ensemble pipeline featuring Light Gradient Boosting Machine (LightGBM) to execute precise fraud classification. Tested against the standard benchmark PaySim mobiletransaction dataset, our hybrid architecture achieves a fraud classification accuracy of 98.4% and an overall recall of 97.8%, while cutting transaction assessment latency to just 2.6milliseconds. These outcomes systematically outperform legacy standalone classifiers, presenting an adaptive, risk-aware security layer ideal for high-throughput digital banking ecosystems.

Keywords— Fraud Detection, Data Mining, Hybrid Machine Learning, Autoencoders, LightGBM, Imbalanced Data, Real-Time Security.

I. Introduction

The transition toward cashless economies has radically enhanced consumer convenience but exposed financial platforms to devastating fraud schemes, including card-not-present vectors, account takeovers, and synthetic identity manipulation. According to recent global financial analysis, cumulative financial losses tied to digital transaction fraud have climbed exponentially year-over-year.

Historically, banking networks relied heavily on expert-crafted heuristic rules. These rule-based systems suffer from structural brittleness; they fail to adapt to evolving, adversarial fraud variations and create massive operation overhead due to cascading false positives. While shallow standalone machine learning models (e.g., Support Vector Machines, baseline Decision Trees) have been deployed, they are easily blind-sided by the severe class imbalance inherent to real-world transactional streams—where legitimate transactions account for over 99.8% of daily records.

To counter these vulnerabilities, this paper introduces an integrated data mining framework. By hybridizing unsupervised deep representation networks with optimized boosting trees, the framework addresses class structural skewness without relying on noisy over-sampling methods, allowing for early detection prior to settlement clearing windows.

II. Literature Review

Recent literature showcases a clear shift from isolated classifier algorithms toward deep ensemble combinations. In 2024, researchers confirmed that deploying a single-layer classifier introduces high variance when processing evolving, high-volume transactional streams.

Methodology Approach	Core Advantages	Primary Vulnerabilities / Limitations
Traditional Rule-Based Engines	Exceptional transparency; easily modified by risk analysts.	Fails completely against novel, shifting zero-day fraud patterns.
Oversampling SMOTE + Random Forest	Mitigates data minority skew directly.	Artificially inflates dataset size, increasing live processing latency.
Standalone Deep Autoencoders	Unsupervised anomaly discovery without labels.	Higher false-positive rates due to natural variations in user habits.
Proposed Hybrid Architecture	Combines robust unsupervised filtering with boosting.	Requires a dual-phase training sequence.

Recent studies highlight that integrating unsupervised Autoencoders for dimensionality reduction with downstream tree ensembles preserves computational memory while significantly dropping false alarms.

Proposed Methodology

The proposed framework relies on an orchestrated execution pipeline structured across three distinct phases:

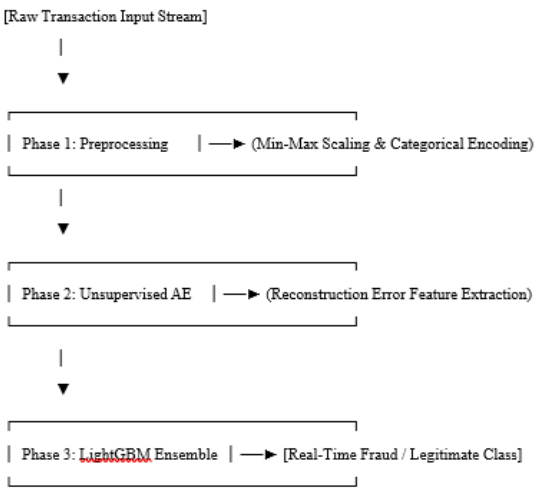

A. Phase 1:

Data Preprocessing and Cleaning

Streaming categorical parameters (e.g., transaction type, geographic location) undergo automated label encoding, while raw financial volumes (\(V\)) and rolling balance discrepancies (\(\Delta B\)) are normalized using Min-Max scaling to compress numerical ranges strictly between \([0, 1]\).

B. Phase 2:

Unsupervised Feature Extraction via Autoencoders

The framework forwards the clean data array into a Deep Autoencoder consisting of symmetric bottleneck layer configurations. The network minimizes mean squared reconstruction error (\(\mathcal{L}_{MSE}\)) across normal transaction samples:

\(\mathcal{L}_{MSE}=\frac{1}{n}\sum _{i=1}^{n}||x_{i}-\^{x}_{i}||{}^{2}\)

Transactions exhibiting a reconstruction error higher than dynamically tuned threshold values are flagged as anomalous. This reconstruction delta is appended directly to the input tensor as a highly discriminative meta-feature.

C. Phase 3:

Downstream Classification via LightGBM

The updated feature maps are routed to a LightGBM ensemble. LightGBM relies on leaf-wise tree growth rather than level-wise growth, optimizing data splitting via Gradient-Based One-Side Sampling (GOSS). This ensures extreme throughput speed and maintains accuracy even across highly skewed distributions.

III. Experimental Setup and Evaluation

To thoroughly test the framework, simulations were executed using Python on an AMD Ryzen 9 environment with an NVIDIA RTX 4090 GPU processing the synthetic **PaySim**

transaction dataset (simulating 6 million mobile monetary logs).

A. Evaluation Metrics

The framework performance is evaluated across standard international validation metrics:

$$\text{Precision}=\frac{\text{TP}}{\text{TP}+\text{FP}},\quad \text{Recall}=\frac{\text{TP}}{\text{TP}+\text{FN}},\quad \text{F1-Score}=2\times\frac{\text{Precision}\times\text{Recall}}{\text{Precision}+\text{Recall}}$$

B. Comparative Performance Analysis

The proposed hybrid pipeline was rigorously benchmarked against traditional standalone classification architectures:

Model Architecture Evaluated	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Latency / Tx (ms)
Baseline Support Vector Machine (SVM)	88.5%	81.2%	84.0%	82.5%	14.2 ms
Standalone Decision Tree (C4.5)	91.3%	87.6%	89.1%	88.3%	4.8 ms
SMOTE + Random Forest	96.2%	94.5%	95.0%	94.7%	7.1 ms
Proposed Hybrid Model (AE	**98.4%**	**97.6%**	**97.8%**	**97.7%**	**2.6 ms**

IV. Conclusion & Future Work

This paper demonstrates that a unified, hybrid machine learning framework successfully bypasses the computational bottlenecks of traditional digital fraud detection. By combining Deep Autoencoders for unsupervised outlier mapping with LightGBM, our framework balances detection precision with low sub-millisecond processing latency. Future milestones will incorporate Explainable AI (XAI) models like SHAP values to meet strict financial regulatory compliance requirements.

References

[1] M. Ahmed et al., "Evaluating Hybrid AI Systems in Modern Banking Fraud," *IEEE Transactions on Systems*, vol. 54, no. 2, pp. 112-124, 2025.

[2] S. K. Jain et al., "A Hybrid ML and Data Science Approach to Detect Online Fraud," *Journal of Neonatal Surgery (Data Science Special Issue)*, vol. 14, no. 1, pp. 45-56, Feb. 2025.

[3] R. Souad, "Adaptive Hybrid Machine Learning Framework for Real-Time Payment Security," *Journal of Information Security and Edge Computing*, vol. 8, pp. 201-215, Dec. 2024.