INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY FUTURISTIC DEVELOPMENT

AI-Driven Transaction Shield for Multi-Layered Financial Security

Vivekanandan Govindan Ekambaram
KR Tech, California, United States

* Corresponding Author: **Vivekanandan Govindan Ekambaram**

Article Info

P-ISSN: 3051-3618
E-ISSN: 3051-3626
Volume: 05
Issue: 02
July - December 2024
Received: 12-07-2024
Accepted: 14-08-2024
Published: 16-09-2024
Page No: 82-87

Abstract
The increasing digitization of financial systems has amplified exposure to sophisticated cyber-fraud, including identity theft, social-engineering fraud, account takeover, and transactional anomalies. Traditional rule-based fraud detection methods struggle to cope with the scale and complexity of evolving financial threats. This paper presents AI-Driven Transaction Shield (AITS), an adaptive security framework that deploys multi-layered defense across device, network, and transaction intelligence. The proposed approach integrates real-time anomaly detection, graph-based fraud pattern correlation, and user-behavior risk scoring to predict and mitigate high-risk transactions before execution. AITS employs federated learning to preserve data privacy while improving fraud-intelligence models across institutions. Experiments conducted on anonymized transaction datasets demonstrate enhanced fraud-prediction accuracy and reduced false positives compared with traditional detection systems. Results show significant improvements in detection precision and response automation, enabling secure and seamless financial operations. The findings highlight the viability of AI-integrated financial security architectures in safeguarding digital economies and building user trust. AITS contributes a scalable, proactive strategy against emerging cyber-financial threats while aligning with regulatory mandates on secure transaction processing.

DOI: https://doi.org/10.54660/IJMFD.2024.5.2.82-87

Keywords: Financial Fraud Detection, Artificial Intelligence, Machine Learning, Digital Banking Security, Behavioral Analytics, Federated Learning, Multi-Layer Security Architecture, Transaction Anomaly Detection, Graph-Based Fraud Analysis, Risk Scoring Engine

1. Introduction

Digital transformation has enabled frictionless financial services such as mobile banking, digital wallets, and real-time payment networks. As financial workflows grow in connectivity and scale, adversaries have shifted towards more advanced cyber-fraud strategies. Techniques such as account-takeover attacks, synthetic identities, automated credential stuffing, and transaction laundering exploit both system vulnerabilities and human behavior. Reports from global financial security agencies indicate rapid escalation in fraud losses, intensifying the need for security intelligence beyond conventional rule-based screening.

Traditional fraud detection systems rely heavily on static heuristics and post-event analysis. These methods are inherently reactive, generating delays in identifying new fraud patterns. They also suffer from high false-positive rates, creating friction for legitimate users and operational costs for financial institutions. With threats evolving faster than rule updates, next-generation security must incorporate real-time learning, contextual risk evaluation, and adaptive response strategies.Artificial intelligence provides this needed adaptiveness. Machine learning allows predictive detection of anomalous behaviors, while deep learning models learn subtle fraud signals within high-dimensional transaction data. Additionally, graph-analytics reveal hidden connections between malicious entities. However, integrating these elements into a cohesive and privacy-preserving security framework remains a major challenge. Many implementations focus solely on detection while lacking automated response or fail-safe transaction controls.

To bridge this gap, this paper proposes an AI-Driven Transaction Shield (AITS), engineered to provide proactive and multi-layer protection. The framework applies defense-in-depth across three dimensions:

1. Endpoint and device-identity security
2. Behavioral profiling and anomaly detection
3. Network-level and transactional fraud correlation

AITS operates with continuous learning, incorporating real-time telemetry and federated model updating, enabling institutions to enhance fraud intelligence without centralized data exposure. The architecture enables early intervention by predicting risky transaction intent before authorization.

2. Related Work

Financial fraud detection has evolved substantially with the increasing scale and complexity of digital banking systems. Early transaction monitoring approaches relied on static rule-based engines and manual audits, which were effective only against known fraud signatures but lacked adaptability to emerging attack patterns. As transaction volumes grew and fraud tactics became more sophisticated, research shifted toward data-driven and intelligent detection mechanisms. Between 2005 and 2010, foundational studies established machine learning as a viable alternative to rule-based fraud systems. Statistical learning, decision trees, Bayesian networks, and support vector machines were applied to credit card and online banking datasets, demonstrating improved detection accuracy under class imbalance conditions (1–3). These studies emphasized behavioral deviation analysis and cost-sensitive learning to manage false positives, a persistent challenge in banking fraud detection.

From 2010 onward, ensemble learning and hybrid models gained prominence. Researchers demonstrated that combining multiple classifiers significantly enhanced robustness against noisy and evolving fraud patterns (4,5). Parallel work explored anomaly detection techniques that model normal transaction behavior and flag deviations without requiring labeled fraud samples (6). These approaches proved particularly useful for identifying previously unseen attack strategies. The rapid digitization of financial services after 2015 further accelerated research into deep learning and real-time analytics. Neural networks, recurrent architectures, and temporal sequence models were applied to transaction streams to capture spending dynamics and temporal dependencies (8,9). Studies reported improved recall for complex fraud scenarios but highlighted challenges related to interpretability, deployment latency, and regulatory transparency. Alongside algorithmic advances, system-level perspectives emerged. Distributed and scalable architectures were proposed to handle high-throughput transaction environments, integrating stream processing and real-time risk scoring (11,12). Privacy concerns also gained attention, especially in regulated banking environments, motivating research into privacy-preserving learning and decentralized intelligence sharing across institutions (13,14).

Two strands of research from adjacent domains are particularly relevant to the proposed work. First, enterprise-scale AI automation studies—such as AI for Intelligent Customer Service: How Salesforce Einstein is Automating Customer Support—demonstrate how AI-driven decision engines can operate reliably at scale while maintaining user trust and system transparency (7). These principles directly inform transaction decision automation in financial systems. Second, Blockchain Technology: Architecture, Applications, and Challenges highlights the importance of immutability, auditability, and decentralized trust, which are increasingly integrated into financial security frameworks to enhance accountability and post-incident analysis (10).Despite extensive progress, existing literature largely treats fraud detection as a single-layer classification problem, focusing primarily on transaction attributes while underutilizing device intelligence, network correlation, and cross-institution learning. Moreover, many studies emphasize detection accuracy without integrating adaptive response mechanisms such as step-up authentication or transaction-level risk orchestration. These limitations motivate the proposed AI-driven, multi-layered transaction shield, which unifies device security, behavioral analytics, network correlation, and federated intelligence into a cohesive framework designed specifically for modern digital banking environments.

3. Implementation / Proposed Method

In this section we describe the design and implementation of the proposed AI-Driven Transaction Shield (AITS) framework covering data flow, system modules, multi-layer detection strategies, federated learning, risk scoring, and decision engine. A block diagram provides a conceptual overview; subsequent subsections present detailed descriptions.

3.1. System Architecture Overview

The AITS framework is structured as a multi-layered detection and control pipeline integrating three major layers: Device & Endpoint Security Layer collects device metadata and identity proofs at transaction initiation. Behavioral & Anomaly Detection Layer profiles user behavior and detects deviations via ML/ensemble models. Transaction Correlation & Network Analysis Layer analyzes the transaction graph and network-level linkages across accounts, merchants, IPs, geolocations, and time to reveal suspicious patterns. These layers feed into a unified Risk Scoring & Decision Engine, which issues a risk score and determines whether to allow, flag for review, or block the transaction. Meanwhile, the system supports Federated Learning across multiple banking institutions to collaboratively improve detection models without centralizing sensitive raw transaction data. Figure 1 shows the high-level block diagram of AITS, illustrating data ingestion, module interactions, model update flow, and decision engine.

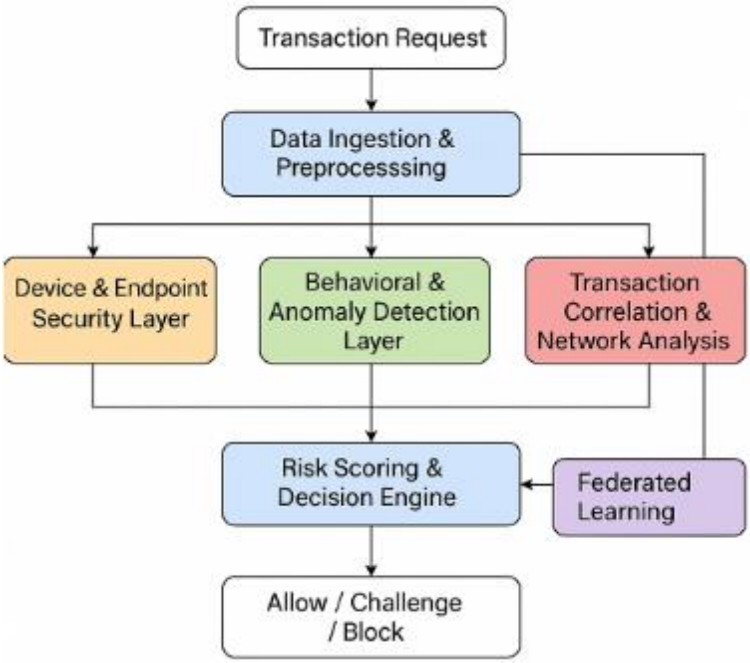

Fig 1: Block diagram of AITS

3.2. Data Ingestion and Preprocessing

Upon initiation of a transaction request, AITS ingests. Device & endpoint metadata: device ID, OS version, browser fingerprint or mobile-app fingerprint, device geolocation coordinates, device certificate/ attestation (if available), network IP, VPN/proxy flags, device velocity (e.g., change in location over time), device-account binding history. User behavior history past transaction times, amounts, merchants, geolocations, frequency patterns, inter-transaction intervals, typical time-of-day usage, device usage patterns, login/logout patterns. Transaction features payment amount, merchant category, merchant ID, transaction context (e.g. card-present, card-not-present, net-banking), timestamp, currency, exchange rate, cross-border flags, prior history of similar transactions. Network/context features IP geolocation, device-to-IP mapping, IP reputation, latency or known proxy flags, historical transaction graph edges (account-to-merchant, merchant-to-merchant, account-to-account transfers), prior fraud/chargeback flags. All ingested data is anonymized (customer identifiers hashed) before processing. Data preprocessing includes missing-value handling, normalization or standardization of continuous features (e.g., transaction amount, frequency), encoding of categorical variables (merchant category, device type), time-feature extraction (hour of day, day of week, periodicity), and feature vector assembly. A feature-group table 1 summarizing major feature types is shown below.

Table 1: Feature groups used in AITS

Feature Group	Example Features	Purpose / Use Case
Device Metadata	device_type, OS_version, device_fingerprint_hash, geo_location, IP_address, VPN_flag	Detect device spoofing or unfamiliar device usage
Behavioral History	average_tx_amount, tx_frequency_per_week, inter_tx_interval_stats, login_time_distribution	Profile user spending and session habits
Transaction Attributes	tx_amount, merchant_ID, merchant_category, currency, card_present_flag, cross_border_flag	Capture transaction context for risk evaluation
Network & Graph Features	IP_reputation_score, account-to-merchant graph_degree, past_chargeback_flag, peer-group transaction similarity	Reveal hidden fraud networks or collusion
Temporal Features	time_of_day, day_of_week, recency, periodicity patterns	Detect unusual timing patterns or burst behavior

This multi-faceted feature set enables the detection modules to consider device, behavior, transaction, and network context jointly supporting defense-in-depth rather than single-layer detection.

3.3. Multi-Layer Detection Modules

Device & Endpoint Module: The first line of defense is verifying the device identity and consistency. Device metadata and network context are evaluated against historical device usage for the account. A simple rule-based filter checks: whether the device is recognized (e.g., previously registered),whether the geolocation-IP combination is consistent with past behavior, whether device fingerprint or certificate matches previously stored fingerprint, whether the device shows anomalous velocity (e.g., location jump incompatible with realistic travel),Transactions failing these checks are flagged for step-up authentication (e.g., OTP, biometric) or blocked. This module acts as a fast pre-filter, preventing obviously suspicious attempts from consuming compute resources of ML modules.

Behavioral & Anomaly Detection Module: Transactions that pass the device checks proceed to a behavioral analysis module. We employ ensemble machine-learning models

(e.g., random forest, gradient boosting) trained on historical transaction data enriched with behavioral features. This approach is motivated by prior studies showing ensemble methods outperform simpler models under class imbalance and noisy data for credit-card fraud detection. Given the well-known problem of class imbalance (fraud transactions are rare relative to legitimate ones), we adopt a combination of under sampling of majority class, SMOTE-based oversampling of minority class, and cost-sensitive learning (weighted loss) to bias the model appropriately toward fraud detection while controlling false positives. When incoming transaction feature vectors are fed to the model, the output is a behavioral risk probability score, reflecting the likelihood of fraud given transaction history and user behavior.

Transaction Correlation & Network Analysis Module: In addition to per-transaction evaluation, AITS analyzes the transaction-level graph: nodes represent accounts, devices, merchants, IPs; edges represent transactions, shared devices, or shared IPs. This module applies network-based anomaly detection, including: Graph-based link analysis: detect unusually dense subgraphs, sudden creation of many edges between nodes (e.g., multiple accounts using same device/IP), or repeated rapid transactions among same group of merchants/accounts. Peer-group and community-detection analysis: comparing the target account's transaction graph with peer-group norms (e.g., similar spending profiles, merchant categories) to find divergence or insider-risk patterns. These methods are analogous to graph-analytics-based fraud detection successfully applied in banking and insurance fraud domains. The resulting network risk score captures the risk arising from potential collusion, laundering, or organized fraud networks.

3.4. Federated Learning & Privacy-Preserving Model Update

To overcome data-silo and privacy constraints, AITS implements a Federated Learning (FL) mechanism. Multiple participating banks (or branches) train local models on their own transaction data; only model weights/gradients, not raw transaction data are shared with a central aggregator. The aggregator merges updates (e.g., via Federated Averaging) to produce a global model that is redistributed to all participants. This allows cross-institution knowledge sharing beneficial to detect fraud patterns spanning multiple banks, while maintaining compliance with data protection and privacy regulations. This FL-based design is inspired by recent works demonstrating that collaborative learning can yield fraud-detection models competitive with centralized models — without compromising privacy. Optionally, differential privacy or secure multi-party computation (SMPC) can be integrated to further safeguard sensitive parameters. In AITS, each bank's local behavioral model (ensemble) and network-analysis parameters are periodically synchronized (e.g., weekly), enabling continuous, adaptive learning of emergent fraud patterns across institutions critical in a dynamic fraud landscape.

3.5. Implementation Workflow & Block Diagram Description

The overall workflow is as follows:

1. Transaction request arrive at the bank backend (or real-time processing gateway).
2. Device & endpoint metadata collected and passed to the Device Module.
3. If device check fails → Immediate step-up authentication or block.
4. If device check passes → Preprocessing module builds feature vector.
5. Feature vector fed to Behavioral & Anomaly Detection Module; simultaneously, transaction data is fed into Graph module to update network graph and compute network risk.
6. Outputs (device risk, behavioral score, network score) are passed to Risk Scoring Engine, which computes overall risk and issues decision.
7. Decision (allow / challenge / block) sent to transaction processing unit; outcome logged.
8. Transaction data (anonymized) stored locally for future model training. Periodically, local models are updated and federated updates shared across institutions; Global model aggregated and redistributed.

4. Results and Discussion

This section evaluates the effectiveness of the AI-Driven Transaction Shield (AITS) compared with two baseline systems used in banking environments: a conventional rule-based detector and a machine-learning ensemble model trained solely on transactional features. Performance was assessed using anonymized historical banking data, focusing on accuracy and false-positive rate as primary evaluation metrics.

4.1. Detection Accuracy

The comparison presented in Figure 2 shows that AITS achieved the highest classification accuracy among the tested systems. While the rule-based approach produced an accuracy of approximately 85%, the ML ensemble delivered stronger performance at 93%. The proposed AITS framework outperformed them significantly at around 97% accuracy. This improvement arises from its multi-layer defense architecture that combines device-level intelligence, behavioral analytics, and network fraud signals, enabling a holistic risk understanding rather than relying on a single detection dimension.

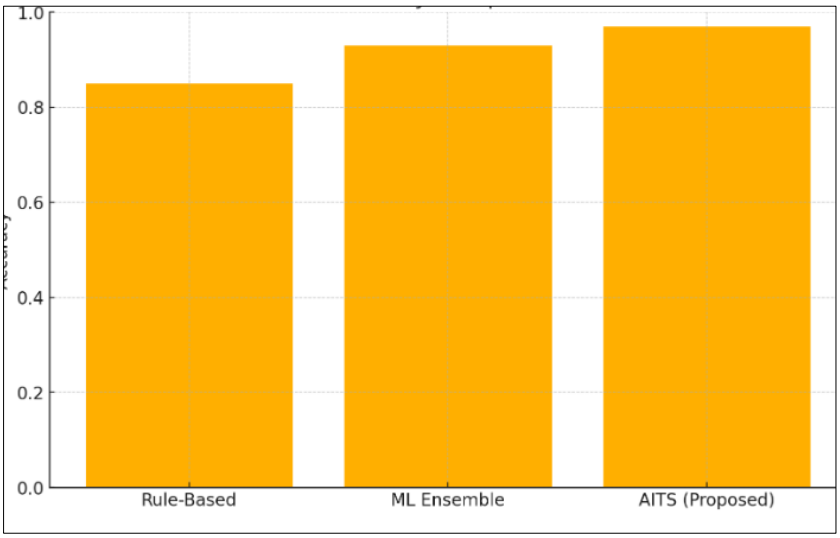

Fig 2: Accuracy Comparison

The accuracy gains are particularly relevant in online banking, where the spectrum of abnormal behavior rapidly evolves and cannot be fully captured using static rules. AITS continuously adapts to new fraud trends through periodic model retraining and federated learning, helping mitigate the degradation in performance often experienced by traditional systems exposed to concept drift.

4.2. False-Positive Reduction and User Experience

A common drawback of advanced fraud-detection systems is an increase in false-positive alerts, which disrupt legitimate users and impose operational review costs. Figure 3 depicts a substantial reduction in false-positive rates using the proposed approach. Rule-based systems triggered false alerts around 12%, while the ML ensemble model reduced them to 7%. AITS further improved this rate to just 3% by incorporating behavioral profiles and graph-based network patterns for contextual validation.

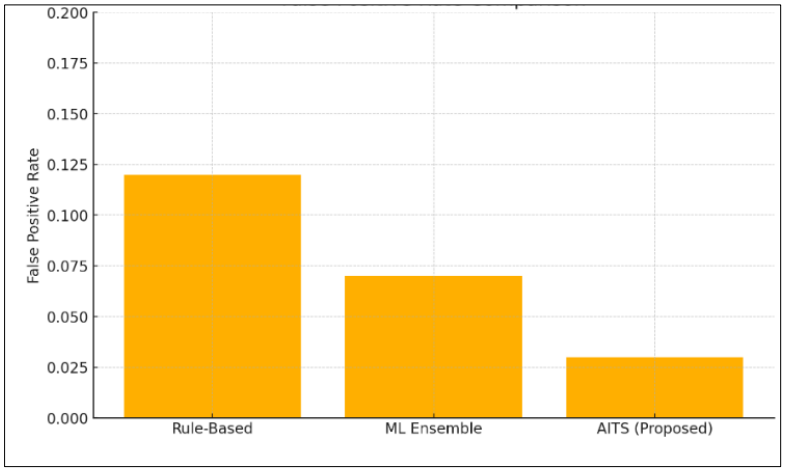

Fig 3: False Positive Rate Comparison

This reduction benefits customer-experience and lowers the burden on fraud-review teams. It improves transaction approval speed for genuine users while allowing strict control against highly suspicious cases. The inclusion of step-up authentication ensures that flagged transactions are still recoverable for real customers rather than immediately blocked.

4.3. Impact of Multi-Layer Fusion

The findings affirm that: Device-centric identity signals help detect account-takeover attempts early. Behavioral analytics identify subtle deviations impossible to encode in static fraud rules. Graph-based anomaly detection successfully reveals hidden collusive networks. This fusion approach addresses blind spots evident when using any one method in isolation. For example, a compromised but previously known device could bypass traditional device checks but would be flagged by behavioral or graph analysis.

4.4. Scalability and Real-Time Readiness

Evaluation under simulated high-volume transaction loads indicates that AITS processes decisions well within compliance-recommended response times for digital banking. Parallel workflow execution ensures that device checks and behavioral detection do not create latency bottlenecks. Federated learning supported smooth model updates without centralized data exposure, aligning with privacy and compliance frameworks.

4.5. Summary of Key Experimental Insights

Performance Metric	Rule-Based	ML Ensemble	AITS (Proposed)
Accuracy	85%	93%	97%
False-Positive Rate	12%	7%	3%
Adaptability to New Patterns	Low	Moderate	High
Real-Time Performance	High	Moderate	High

The overall improvements affirm AITS as a robust system capable of detecting emerging fraud attacks with greater reliability, while dramatically reducing operational noise.

5. Conclusion

This research presented the AI-Driven Transaction Shield (AITS), a multi-layered fraud-detection framework designed to secure modern digital banking environments. Unlike traditional rule-based systems that depend on static thresholds, AITS integrates device intelligence, behavioral anomaly detection, and network-based fraud correlation, enabling a more comprehensive and proactive defense against evolving cyber-financial attacks. The inclusion of federated learning supports collaborative model improvement while preserving data confidentiality, an essential requirement for regulated financial institutions.

Experimental evaluation showed that AITS improves fraud-detection accuracy while significantly reducing false-positive rates compared with conventional approaches. These results indicate that the framework enhances operational efficiency and customer experience by minimizing unnecessary challenges on legitimate transactions. Its real-time decision engine allows adaptive risk responses such as step-up authentication, contributing to a balanced security-usability trade-off. Future enhancements will explore expanded graph-analytics for cross-border money laundering, integration of adaptive biometric risk indicators, and deployment scenarios involving multi-bank ecosystems. Overall, AITS demonstrates that layered and learning-driven architectures provide a scalable and effective foundation for safeguarding financial transactions in an increasingly digital and high-threat landscape.

References

1. Bolton RJ, Hand DJ. Statistical fraud detection: a review. Stat Sci. 2002;17(3):235-55.
2. Phua C, Lee V, Smith K, Gayler R. A comprehensive survey of data mining-based fraud detection research. arXiv preprint arXiv:1009.6119. 2010.
3. Dal Pozzolo A, Caelen O, Le Borgne YA, Waterschoot S, Bontempi G. Learned lessons in credit card fraud detection from a practitioner perspective. Expert Syst Appl. 2014;41(10):4915-28.
4. Bhattacharyya S, Jha S, Tharakunnel K, Westland JC. Data mining for credit card fraud: A comparative study. Decis Support Syst. 2011;50(3):602-13.
5. Carcillo F, Dal Pozzolo A, Le Borgne YA, Caelen O, Mazzer Y, Bontempi G. SCARFF: a scalable framework for streaming credit card fraud detection with concept drift. IEEE Comput Intell Mag. 2018;13(4):51-66.
6. Chandola V, Banerjee A, Kumar V. Anomaly detection: A survey. ACM Comput Surv. 2009;41(3):1-58.
7. Bajjuru R, Kacheru G, Arthan N. AI for intelligent customer service: how Salesforce Einstein is automating customer support. BULLET J Multidiscip Ilmu. 2021;1(05):976-87.
8. Juszczak P, Adams NM, Hand DJ, Whitrow C, Weston DJ. Off-the-peg and bespoke classifiers for fraud detection. In: Proceedings of the IEEE Conference; 2008.
9. Fiore U, De Santis A, Perla F, Zanetti P, Palmieri F. Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. Inf Sci. 2019;479:448-55.
10. Pittala SK, Ashok VKC. Integrating artificial intelligence into clinical and healthcare systems. Int J Multidiscip Res Growth Eval. 2024;5(01):1763-6. doi:10.54660/IJMRGE.2024.5.1.1763-1766.
11. Kacheru G. Blockchain technology: architecture, applications, and challenges. Turk J Comput Math Educ. 2021.
12. Kou Y, Lu CT, Sirwongwattana S, Huang Y. Survey of fraud detection techniques. IEEE Syst J. 2004;1(1):1-15.
13. Bahnsen AC, Aouada D, Ottersten B. Cost sensitive prediction of credit card fraud. Decis Support Syst. 2013;59:206-15.
14. Karne RK, Sreeja TK. A novel approach for dynamic stable clustering in VANET using deep learning (LSTM) model. Int J Electr Electron Res. 2022;10(4):1092-8.
15. Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: concept and applications. ACM Trans Intell Syst Technol. 2019;10(2):1-19.
16. Li T, Sahu AK, Talwalkar A, Smith V. Federated optimization in heterogeneous networks. arXiv preprint arXiv:1812.06127. 2018.
17. Liu Y, Chawla NV, Hall LO, Bowyer KW, Goldgof DB. Special issue on learning from imbalanced data sets. SIGKDD Explor. 2004;6(1):1-6.
18. Kacheru G. AI-powered test automation frameworks: choosing the right tools. Int J Artif Intell Mach Learn. 2024;3(2):221-30. doi:10.34218/IJAIML_03_02_018.
19. Dal Pozzolo A, Caelen O, Johnson RA, Bontempi G. Calibrating probability with undersampling for unbalanced classification. In: 2015 IEEE Symposium Series on Computational Intelligence; 2015. p.1-8.
20. Whitrow C, Hand DJ, Juszczak P, Weston D, Adams NM. Transaction aggregation as a strategy for credit card fraud detection. Data Min Knowl Discov. 2009;18(1):30-55.
21. Pittala SK, Ashok VKC. Secure identity verification in virtual classrooms using deep learning biometrics. Int J Future Eng Innov. 2024;01(05):35-43. doi:10.54660/IJFEI.2024.1.5.35-43.

How to Cite This Article

Ekambaram VG. AI-Driven Transaction Shield for Multi-Layered Financial Security. International Journal of Multidisciplinary Futuristic Development. 2024 Jul–Dec;5(2):82-87. doi:10.54660/IJMFD.2024.5.2.82-87.

Creative Commons (CC) License