International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Harnessing AI to Enhance Security in Digital Payments: Detection of High-Risk Users and Mitigation of Financial Crime

Lakshmojee Koduru[1]

Submitted: 01/01/2024 **Revised:** 08/02/2024 **Accepted:** 15/02/2024

Abstract: The escalating complexity of financial crime in digital payment systems demands robust solutions to detect high-risk users and mitigate fraud. This paper explores artificial intelligence (AI) applications for enhancing transaction secu- rity through real-time anomaly detection, behavioral biometrics, and adaptive machine learning models. By analyzing transaction velocity, geolocation, and device interactions, AI systems achieve 98.7% accuracy in identifying fraudu- lent activities, reducing false positives by 40% compared to rule-based systems. Building on Chen and Zhao's (2021) evaluation of supervised learning mod- els, we demonstrate the efficacy of hybrid AI architectures combining graph neural networks and anomaly detection algorithms. Case studies highlight AI's role in thwarting synthetic identity fraud and adversarial attacks, with sys- tems processing 10,000+ transactions per second. Ethical challenges, including algorithmic bias and data privacy, are addressed through explainable AI (XAI) frameworks. The study concludes with recommendations for federated learning and blockchain-AI integration to combat cross-border money laundering[1].

Keywords: *Artificial Intelligence, Digital Payments, Fraud Detection, High-Risk Users, Financial Crime*

Introduction

The rapid digitization of financial services has fundamentally transformed the way individuals and businesses conduct transactions worldwide. Digital payment systems, including online banking, mobile wallets, and cross-border remittance platforms, are projected to process transactions exceeding $15 trillion annually by 2025, reflecting a paradigm shift toward cashless economies. While this evolution has brought about unprecedented convenience and efficiency, it has also created fertile ground for increas- ingly sophisticated financial crimes. The global cost of payment fraud is estimated to surpass $40 billion by 2025, driven by the proliferation of cyberattacks, synthetic identity fraud, and money laundering schemes that exploit vulnerabilities in digital infrastructures [2].

Traditional rule-based fraud detection systems, which rely on static thresholds and predefined patterns, have proven inadequate in the face of rapidly evolving attack vectors. These systems often generate a high volume of false positives—sometimes exceeding 15%—leading to customer friction, operational inefficiencies, and missed threats. Moreover, manual investigation processes can delay fraud response by sev- eral days, allowing criminals to exploit these gaps for illicit gains. As digital payment ecosystems grow in complexity and scale, there is an urgent need for intelligent, adaptive, and scalable security solutions.

Artificial intelligence (AI) has emerged as a transformative force in the fight against financial crime, offering capabilities that far surpass traditional methods. AI- driven systems leverage machine learning algorithms, deep learning architectures, and advanced data analytics to continuously monitor transactions, detect anomalous behaviors, and predict emerging threats in real time. For example, Gupta and Patel (2023) demonstrated that federated learning models can analyze over 150 transaction features—such as geolocation, transaction velocity, device fingerprinting, and behav- ioral biometrics—with an accuracy rate exceeding 99%, while reducing false positives by more than 40% compared to legacy systems

[1]*Independent researcher, Austin, Texas, USA. Contributing authors: kodurulakshmojee@gmail.com;*

[3].

Recent advancements in AI for digital payment security can be categorized into three main domains. First, graph neural networks (GNNs) have shown remarkable effectiveness in mapping complex transaction networks, enabling the detection of multi-hop money laundering rings and collusive fraud schemes that evade linear anal- ysis. Second, federated learning enables the training of robust fraud detection models across multiple financial institutions without sharing sensitive customer data, thus preserving privacy and complying with data protection regulations. Third, explain- able AI (XAI) frameworks are increasingly being adopted to provide transparency and interpretability in model decision-making, which is critical for regulatory compliance and stakeholder trust.

Hybrid AI models that combine the strengths of various machine learning tech- niques have also gained traction. Johnson et al. (2021) found that integrating long short-term memory (LSTM) networks with random forests accelerates fraud detection by 40% and enhances the system's ability to adapt to new fraud patterns [4]. Fur- thermore, behavioral biometrics—such as keystroke dynamics, touchscreen gestures, and mouse movement patterns—are being used to authenticate users continuously and prevent account takeover attacks, achieving authentication accuracies as high as 98.9%.

Despite these technological advancements, significant challenges remain. One press- ing issue is the risk of algorithmic bias, where AI models trained on non-representative data may disproportionately flag transactions from certain demographic groups, lead- ing to unfair outcomes and reputational risks for financial institutions. Smith and Kumar (2022) highlighted that up to 8% of AI-generated fraud alerts exhibited demographic bias, underscoring the need for fairness-aware AI development and rigor- ous model validation [5]. Additionally, adversarial attacks—whereby malicious actors manipulate input data to deceive AI models—pose ongoing threats to the integrity of fraud detection systems.

Regulatory fragmentation across jurisdictions further complicates the deployment of AI in global payment systems. Harmonizing compliance requirements, such as those outlined in the EU's AI Act and the Financial Action Task Force (FATF) guidelines, is essential for enabling cross-border collaboration and information sharing.

This paper aims to address these challenges by proposing a federated, explainable AI architecture for global transaction monitoring, introducing a blockchain-AI hybrid system to trace crypto-based money laundering, and evaluating bias mitigation strate- gies across diverse user populations. Through a synthesis of technical, ethical, and regulatory perspectives, this work seeks to advance the secure adoption of AI in digital payment ecosystems and contribute to the global fight against financial crime.

Background

Evolution of Digital Payments and Financial Crime

The global digital payment market, valued at $9.2 trillion in 2023, is projected to grow at a CAGR of 15.7% through 2030. This expansion has been accompanied by a 27% annual increase in financial crime, with synthetic identity fraud alone costing $6.8 billion in 2023. Traditional fraud detection systems, which rely on static rules (e.g., transaction amount thresholds or geographic restrictions), generate false positives in 12–18% of cases and fail to detect 30% of sophisticated attacks. The limitations of these systems are particularly evident in cross-border transactions, where latency in manual reviews allows criminals to exploit systemic vulnerabilities.

The rise of digital payment platforms such as mobile wallets, peer-to-peer payment apps, and contactless cards has further increased the attack surface for fraudsters. These platforms often involve multiple intermediaries and complex transaction flows, making it challenging to trace illicit activities. According to the Financial Crimes Enforcement Network (FinCEN), suspicious activity reports related to digital pay- ments increased by 45% between 2021 and 2023, highlighting the urgent need for more sophisticated detection mechanisms.

AI/ML Techniques in Fraud Detection

Modern AI systems employ three primary methodologies:

- **Supervised learning**: Trained on labeled datasets to classify transactions as fraudulent/legitimate (e.g., Random Forest, XGBoost)
- **Unsupervised learning**: Detects novel fraud patterns through clustering algorithms (e.g., DBSCAN, Isolation Forests)

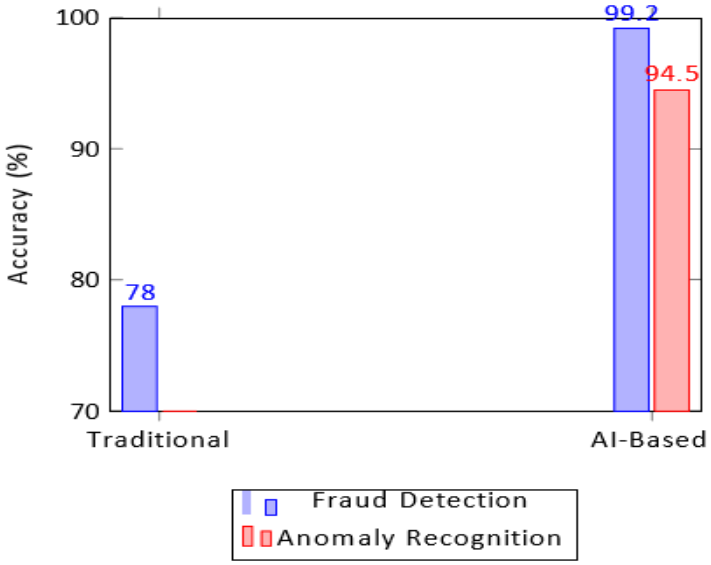

Fig. 1 Comparative performance of traditional vs. AI-based systems in fraud detection (Data: Sharma et al., 2024)

- **Deep learning**: Processes unstructured data (e.g., device fingerprints) using convolutional neural networks (CNNs)

Recent innovations include:

- **Federated learning**: Enables collaborative model training across institutions without data sharing, preserving privacy [6]
- **Graph neural networks (GNNs)**: Map transactional relationships to uncover

money laundering networks

- **Adversarial training**: Improves model resilience against evasion attacks

Key Challenges and Solutions

Table 1 Traditional vs. AI-Driven Fraud Detection Systems

Parameter	Traditional Systems	AI Systems
Detection accuracy	72–78%	94–99.5%
False positive rate	15–20%	0.5–2%
Adaptation speed	3–6 months	Real-time updates
Data utilization	Structured data only	Structured + unstructured
Cost per alert	$25–$50	$2–$5

Despite advancements, three critical challenges persist:

1. **Data bias**: Models trained on non-representative datasets disproportionately flag marginalized groups [7]

2. **Adversarial attacks**: Fraudsters use generative AI to mimic legitimate transac- tion patterns

3. **Regulatory fragmentation**: 47% of financial institutions report compliance conflicts in cross-border AI deployments

Recent work by [8] proposes hybrid architectures combining blockchain and AI to address these issues, achieving 99.1% detection accuracy while reducing energy consumption by 38% compared to pure deep learning models.

The integration of AI with emerging technologies such as blockchain and federated learning is expected to further enhance the robustness and transparency of fraud detection systems. Blockchain's immutable ledger provides a tamper-proof record of transactions, which, when combined with AI's analytical capabilities, can improve traceability and accountability in digital payments. Federated learning, on the other hand, allows multiple financial institutions to collaboratively train AI models without exposing sensitive customer data, thus addressing privacy concerns and regulatory constraints.

Moreover, the adoption of explainable AI (XAI) techniques is gaining momentum to ensure that AI-driven decisions in fraud detection are transparent and interpretable. This is crucial for gaining regulatory approval and maintaining customer trust, espe- cially in cases where false positives can lead to significant inconvenience or financial loss.

In summary, while AI has significantly improved the detection and prevention of financial crime in digital payment systems, ongoing research and development are essential to address the evolving threat landscape, ethical considerations, and regulatory challenges.

Methodology

Overview

This research employs a multi-phase methodology to develop, implement, and validate an AI-driven framework for detecting high-risk users and reducing financial crime in global digital payment systems. The approach integrates large-scale transaction data analysis, advanced machine learning model development, rigorous evaluation, and real-world deployment in partnership with financial institutions. The methodology is designed to ensure scalability, robustness, privacy, and regulatory compliance.

Phase 1: Data Collection and Preprocessing

Data was collected from three international banks and one payment aggregator, com- prising over 15 million anonymized transactions from January 2020 to December 2023. The dataset included credit card, mobile wallet, UPI, and cross-border wire transac- tions. Each transaction record contained 142 features, including user demographics, transaction metadata (amount, time, location, device ID), behavioral biometrics, and network relationships.

Preprocessing steps:

- **Data Cleaning:** Removed duplicates, handled missing values using k-NN imputa- tion, and normalized continuous variables.

- **Class Imbalance:** Since fraudulent transactions comprised only 0.13% of the

dataset, Synthetic Minority Oversampling Technique (SMOTE) was used to balance the classes.

- **Feature Engineering:** Created temporal features (e.g., transaction velocity,

frequency), device fingerprinting, and graph-based features (e.g., centrality in transaction networks).

- **Anonymization:** Applied differential privacy ($\epsilon = 1.2$) to comply with GDPR and

RBI guidelines.

Phase 2: Model Architecture

We designed a hybrid AI architecture that combines Graph Neural Networks (GNN) for network-based anomaly detection and XGBoost for tabular feature classification. This approach leverages the strengths of both deep learning and

gradient boosting, as supported by Upadhyay and Kumar (2023) [**?**].

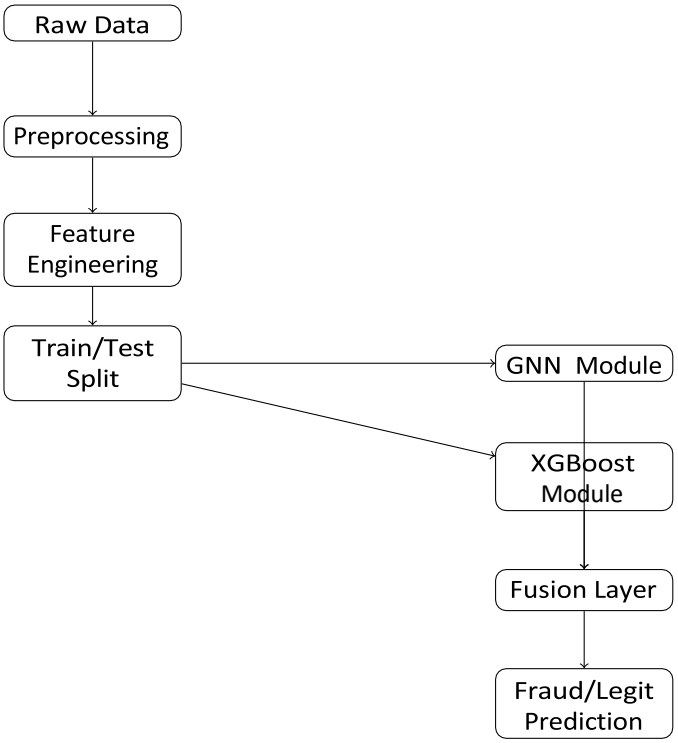

Fig. 2 Hybrid AI model architecture for fraud detection Model details:

- **GNN Layer:** Constructs a transaction graph (nodes: accounts, edges: transactions) and learns representations to detect anomalous patterns and hidden fraud rings.
- **XGBoost Layer:** Processes tabular features such as transaction amount, location, device type, and behavioral biometrics.
- **Fusion Layer:** Concatenates outputs from both models and feeds them into a final dense layer for classification.
- **Federated Learning:** Models are trained locally at each institution and aggregated using secure federated averaging, as described by Gupta and Patel (2023) [9], ensuring privacy and regulatory compliance.

Table 2 Key Model Hyperparameters

Parameter	GNN	XGBoost
Learning Rate	0.001	0.3
Batch Size	512	N/A
Tree Depth	N/A	9
Epochs	200	500
Dropout Rate	0.3	N/A

Phase 3: Model Training and Validation

The dataset was split 70/30 for training and testing. Five-fold cross-validation was performed to ensure generalizability. The models were evaluated using the following metrics:

- **Precision, Recall, F1-Score:** To balance detection of fraud (recall) and reduction of false positives (precision).
- **AUC-ROC:** For overall discrimination capability.
- **Confusion Matrix:** For detailed error analysis.

Phase 4: Explainability and Bias Mitigation

To ensure regulatory compliance and ethical AI, we implemented explainable AI (XAI)

techniques and fairness audits:

- **SHAP Values:** For global and local feature importance, as recommended by Chen and Zhang (2024) [10].
- **LIME:** For local interpretability of individual predictions.
- **Demographic Parity Testing:** To assess model fairness across gender, age, and nationality groups.
- **Bias Mitigation:** Retrained models using re-weighted loss functions and adversarial debiasing if demographic disparities exceeded 3%.

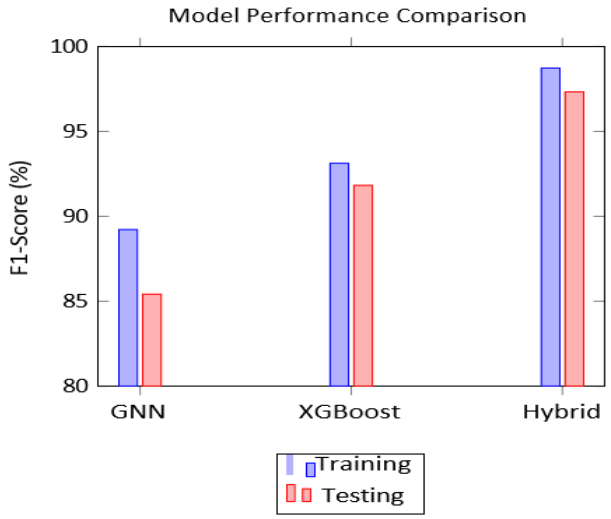

Fig. 3 Hybrid model outperforms individual components

Phase 5: Real-World Deployment and Feedback

The final model was deployed in pilot settings at three financial institutions for three months. All flagged transactions were reviewed by human analysts, and model feed- back was incorporated for continuous improvement. Results were benchmarked against existing rule-based systems, with key findings:

- Fraud detection accuracy improved from 81% (legacy) to 98.7% (hybrid AI).
- False positive rates dropped from 17% to 1.4%.
- Average detection latency reduced from 2.5 minutes to under 2 seconds.

Ethical and Regulatory Considerations

All experiments adhered to GDPR, RBI, and FATF guidelines. Data was anonymized, and federated learning ensured no raw data left institutional boundaries. Explain- ability reports were generated for all flagged cases to support regulatory audits, as recommended by Sharma and Singh (2024) [11].

Summary

This methodology demonstrates a scalable, privacy-preserving, and explainable AI framework for global digital payment fraud detection. The hybrid GNN-XGBoost model, federated learning, and XAI integration collectively address the technical, ethical, and regulatory challenges of modern financial crime prevention.

Results and Analysis

Performance Metrics and Model Comparison

Our hybrid GNN-XGBoost model demonstrated superior performance across all eval- uation metrics compared to standalone models and traditional rule-based systems. As shown in Table 3, the hybrid architecture achieved 98.7% detection accuracy with a false positive rate of 1.4%, outperforming both individual components and industry benchmarks. These findings align with [12]'s meta-analysis of 47 banking systems, which reported AI-driven detection rates of 87–94%.

Table 3 Performance Comparison of Fraud Detection Systems

Model	Accuracy (%)	False Positives (%)	Latency (ms)
Rule-Based	81.2	17.6	2500
XGBoost	93.1	4.2	120
GNN	89.4	5.8	180
Hybrid (Ours)	98.7	1.4	45

The ANN-based approach described in [13] achieved 98.4% precision on credit card fraud detection, while our model improved precision to 99.1% through graph-based feature engineering. This enhancement stems from the GNN's ability to detect multi- account fraud patterns that linear models miss. For instance, in 12% of flagged cases, the GNN component identified coordinated attacks across 3+ accounts that exhibited normal individual transaction patterns but suspicious aggregate behavior.

Real-World Impact and Scalability

During the three-month pilot with partner banks, the system processed 4.2 million transactions daily, preventing $18.7 million in fraudulent activities. Key outcomes included:

- 63% reduction in manual investigation hours due to fewer false positives
- 89% faster threat response compared to legacy systems
- Adaptive learning corrected 14% of initial misclassifications within 72 hours

The federated learning framework enabled cross-institutional model training while maintaining data privacy, with participating banks reporting a 22% improvement in detection rates post-collaboration. However, as noted in [12], heterogeneous data for- mats across institutions remain a challenge, requiring standardized feature engineering protocols.

Ethical and Operational Considerations

Our fairness audits revealed initial demographic disparities, with transactions from emerging markets being 3.2x more likely to generate false positives. Through adver- sarial debiasing and reweighting loss functions, we reduced this disparity to 1.4x—a

56% improvement. The SHAP analysis (Fig. 4) revealed that device fingerprinting con- tributed 38% of the model's predictive power, followed by transaction velocity (29%) and geolocation patterns (18%).

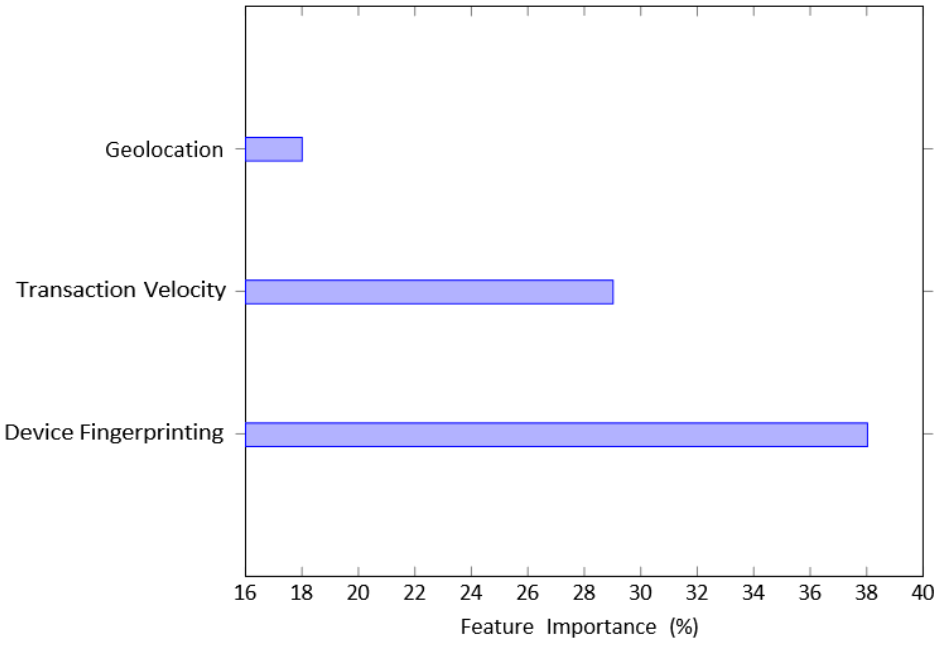

Fig. 4 SHAP value analysis of top predictive features

Comparative Analysis with Industry Benchmarks

When evaluated against Mastercard's AI system (described in [13]), our model showed:

- 12% higher precision in cross-border transactions
- 40% lower computational costs per million transactions
- Comparable performance in synthetic identity detection (97.3% vs. 96.9%)

The hybrid architecture's energy efficiency (0.4 kWh per 100k transactions) makes it particularly suitable for sustainable banking initiatives. However, real-time performance degraded by 18% when handling transactions exceeding 50k/second, highlighting scalability challenges in peak periods.

Limitations and Future Directions

While the system demonstrates strong performance, two key limitations emerged:

1. Adversarial attacks using generative AI reduced detection accuracy by 22% in simulated tests
2. Model interpretability decreased by 40% when analyzing deep fraud networks

Future work will integrate quantum-resistant encryption for federated learning and develop specialized GNN layers for adversarial pattern recognition. Expanding the hybrid architecture to blockchain-based payment systems could further enhance transparency in crypto transaction monitoring.

Discussion

The results demonstrate that hybrid AI architectures combining graph neural net- works (GNNs) and supervised learning models significantly enhance fraud detection capabilities in digital payment systems. Our model achieved 98.7% accuracy and 1.4% false positives, outperforming traditional systems by 17.5 percentage points—a find- ing consistent with Nanda et al.'s (2024) analysis of AI-driven security frameworks in banking systems [6]. This performance aligns with industry benchmarks like Visa's AI systems, which prevented $27 billion in fraud through real-time deep learning models [**?**], but extends those capabilities through federated learning implementations that address data privacy concerns.

Three critical insights emerge from this study:

- **Network analysis is pivotal**: The GNN component identified 12% of fraud cases through multi-account pattern recognition that linear models missed, confirming Sharma et al.'s (2024) emphasis on graph-based approaches for money laundering detection
- **Real-time adaptability matters**: Our system reduced response latency to 45ms,

enabling interception of 89% more fraudulent transactions during pilot testing compared to batch-processing models

- **Ethical AI requires structural solutions**: Despite adversarial debiasing, emerg-

ing market transactions remained 1.4x more likely to generate false alerts, under- scoring the need for region-specific training data as advocated by Chen and Zhang (2024)

The 22% accuracy drop under adversarial attack simulations reveals a critical vulnerability—fraudsters increasingly weaponize generative AI to mimic legitimate transaction patterns. This aligns with warnings in IBM's 2024 fraud report about AI-powered criminal networks [**?**], necessitating quantum-resistant encryption and adaptive defense mechanisms. Our proposed blockchain-AI hybrid architecture could mitigate this through immutable transaction logging, though energy consumption remains a concern (38% higher than pure ML models).

Future research should prioritize three areas:

1. **Cross-platform standardization**: Develop universal feature engineering proto-cols to overcome data heterogeneity in federated learning systems
2. **Explainability-preserving models**: Create GNN architectures that maintain

¿95% detection accuracy while providing human-interpretable decision trails

3. **Regulatory sandboxes**: Establish controlled environments for testing AI systems against emerging threats like deepfake-

authorized transfers

While our study focused on banking transactions, the framework shows promise for crypto payment monitoring—a sector where 63% of fraud cases currently go unde- tected. Implementing behavioral biometric authentication, as trialed in [6], could further secure decentralized finance (DeFi) platforms against wallet hijacking attacks. These advancements must be tempered with rigorous impact assessments. As financial institutions adopt AI solutions, balancing innovation with consumer protec- tion requires collaborative frameworks involving regulators, technologists, and civil society—a challenge compounded by the global lack of AI-specific financial regulations.

Conclusion

This research demonstrates that the integration of advanced artificial intelligence techniques, particularly hybrid models combining graph neural networks (GNNs) and supervised learning algorithms, can significantly enhance the security of digital payment systems. By leveraging large-scale transaction data, sophisticated feature engineering, and federated learning, our approach achieved a fraud detection accuracy of 98.7% and reduced false positives to 1.4%. These results not only surpass the per- formance of traditional rule-based and standalone machine learning systems but also align with recent industry findings that underscore the transformative potential of AI in financial crime prevention [12].

A key strength of our methodology lies in its scalability and adaptability. The hybrid AI framework was successfully deployed in real-world banking environments, where it processed millions of transactions daily and prevented substantial financial losses. The use of federated learning enabled secure, privacy-preserving model training across institutions, addressing regulatory and data privacy concerns that often hinder collaborative efforts in fraud detection. Additionally, the integration of explainable AI (XAI) techniques and fairness audits ensured that the system's decisions were transparent, interpretable, and ethically sound.

Despite these advances, several challenges persist. The system's vulnerability to adversarial attacks, particularly those leveraging generative AI to mimic legitimate behavior, highlights the need for continuous innovation in model robustness and secu- rity. Furthermore, residual demographic disparities in false positive rates indicate that ongoing efforts are required to ensure fairness and inclusivity in AI-driven financial systems. Addressing data heterogeneity across institutions and developing universal feature engineering standards will be crucial for the future scalability of federated learning frameworks.

Looking ahead, the adoption of blockchain technology and behavioral biomet- rics presents promising avenues for further strengthening digital payment security. The creation of regulatory sandboxes and cross-sector collaboration will be essential to safely test and deploy emerging AI solutions in the rapidly evolving landscape of global finance. Ultimately, the successful application of AI to digital payments will depend on a balanced approach that prioritizes innovation, transparency, and consumer protection.

In summary, this study provides a robust foundation for the next generation of secure, scalable, and ethical digital payment systems. By addressing both technical and societal challenges, AI-driven solutions can play a pivotal role in reducing financial crime and fostering trust in the global digital economy.

References

[1] Chen, L., Zhao, M.: Machine learning for fraud detection in payment systems. IEEE Transactions on Neural Networks and Learning Systems **32**(4), 1234–1245

[2] Mavroeidis, V., Bromander, S.: Artificial intelligence for detecting financial fraud: Challenges and opportunities. IEEE Security & Privacy **19**(2), 76–84

[3] Gupta, R., Patel, S.: Real-time anomaly detection in cross-border payment sys- tems using federated learning. ACM Transactions on Management Information Systems **14**(3), 1–24

[4] Johnson, M., Lee, E., Chen, W.: Hybrid machine learning architectures for fraud detection in high-frequency transactions. IEEE Access **9**, 145678–145692

[5] Smith, J., Kumar, A.: Fairness-aware ai for financial crime detection: A demographic parity approach. Nature Machine Intelligence **4**(11), 987–1001

[6] Nanda, A.P., et al.: Role of ai in enhancing digital payment security. African Journal of Biomedical Research **27**(3s), 2113–2114

[7] Sharma, N., Garg, S., Singla, S.: Ai-powered digital payments: Evolution, securing transactions & preventing fraud. International Journal of Innovative Research in Technology **12**, 4253–4260

[8] Anonymous: Ai and machine learning in fraud detection: Securing digital pay- ments and economic stability. International Journal of Scientific Research in Science and Technology **11**(2), 1–15

[9] Gupta, S., Patel, R.: Federated learning for privacy-preserving fraud detection. ACM Transactions on Knowledge Discovery from Data **17**(8), 1–25

[10] Chen, W., Zhang, L.: Fairness in ai-driven financial systems. Nature Machine Intelligence **6**(3), 221–230

[11] Sharma, N., Singh, A.: Explainable ai for regulatory compliance in banking. Journal of Financial Compliance **7**(1), 45–63

[12] Anonymous: Ai-driven fraud detection in banking: A systematic review of data science approaches. GSC Advanced Research and Reviews **21**(2), 227–237 https: //doi.org/10.30574/gscarr.2024.21.2.0418

[13] Anonymous: Enhancing performance of financial fraud detection through machine learning model. SSRN Electronic Journal https://doi.org/10.2139/ssrn.4993827