Agentic AI for Real-Time, Resilient, and Adaptive Fraud Detection in Digital Payment Systems

1st Supraja Ayyamgari
Dept. of Information Technology
University of the Cumberlands
Williamsburg, KY, USA
Supraja.Ayyamgari@gmail.com

2nd Chaitanya Tumma
Dept. of Information Technology
University of the Cumberlands
Williamsburg, KY, USA
Chaitanyatumma1@gmail.com

3rd Ravikiran Uba
Dept. of Information Technology
University of the Cumberlands
Williamsburg, KY, USA
Ubaravikiran@gmail.com

4th Rithwik Sannapu
Dept. of Information Systems
Saint Louis University
Saint Louis, MO, USA
Rithwiksannapu01@gmail.com

5th Abhignan Srivatsava Sribhashyam
Senior Software Engineer
Target Corporation
Lakeville, Minnesota USA
abhignanss@gmail.com

***Abstract*—With the exponential growth of digital payments, the danger of fraudulent transactions has become critically pronounced, highlighting the urgency for intelligent, real-time, and adaptive fraud detection mechanisms. This research utilizes Agentic AI—a goal-based, autonomous AI paradigm—to synthesize an end-to-end fraud detection pipeline of state-of-the-art ensemble learning, contextual reasoning, and self-optimization in real-time. The novel framework utilizes a hybrid of Graph Neural Networks for relationship-based transaction modeling and Transformer-based anomaly discovery for temporal-sequential reasoning, with adaptive thresholding driven by an autonomous policy engine. To handle the extreme imbalance in the dataset with only 0.15% of the dataset representing fraudulent transactions, a Dynamic Synthetic Oversampling with Reinforcement Feedback (DSORF) approach is used to allow the agent to iterate on the generation of synthetic samples based on the feedback of the model. Experimental results show the detection accuracy of 99.96%, precision of 91.84%, and recall of 89.12%, with the reduction of false positives by multiple orders of the state-of-the-art static based methods. With the results, the potential of Agentic AI in adapting autonomously to shifting fraud mechanisms and providing stronger resilience, scalability, and assurance in digital payments is brought into focus. Future research will see the agent architecture extended for real-time inter-border, multi-currency-based fraud detection with decision explainability in real-time.**

***Index Terms*—Payment Systems, IoT-enabled Banking, Post-Quantum Cryptography, CRYSTALS-Kyber, Dilithium, AES, RSA, ECC, SHA-3, HMAC, Blockchain, Zero-Knowledge Proofs (ZKPs), Secure Multi-Party Computation (MPC), Fully Homomorphic Encryption (FHE), AI-driven Fraud Detection, Deep Learning, Federated Learning, Confidential Computing, Hardware Security Module (HSM), Trusted Platform Module (TPM), Physically Unclonable Function (PUF), PCI DSS, EMV, ISO 27001, NIST.**

I. Introduction

The rapid digital transformation of the financial sector has completely changed how money moves, with digital payment systems now becoming the primary way people transact worldwide. The rise of mobile wallets, online banking, peer-to-peer transfer apps, and integrated e-commerce payment tools has given consumers unmatched convenience, faster transactions, and smoother operations. These benefits not only improve the customer experience but also expand financial access by bringing secure services to previously underserved communities. At the same time, the growth of digital finance has created greater risks to security and privacy. As transaction volumes soar, cybercriminals are exploiting weaknesses in payment systems. Financial institutions now face increasingly sophisticated fraud, including identity theft, account takeovers, synthetic identities, phishing attacks, and large-scale money laundering. These threats are becoming more complex, often using automation, artificial intelligence, and advanced concealment techniques to evade traditional detection methods.The evolving and adaptive characteristics of fraudulent activities pose significant challenges to ensuring the security of digital transactions. Conventional rule-based monitoring approaches are increasingly inadequate in addressing the velocity, scale, and heterogeneity of modern attacks, thereby heightening exposure to financial loss, reputational damage, and risks of regulatory non-compliance. This underscores the pressing need for real-time, intelligent fraud detection frameworks capable of dynamically adapting to adversarial strategies, accurately differentiating between legitimate and malicious transactions, and sustaining system reliability in the context of continuously evolving threats [1], [2].

Traditional fraud detection systems, largely based on rule-driven and heuristic methods, are proving increasingly ineffective in today's fast-changing and adversarial digital environment. These frameworks typically rely on predefined thresholds, static rules, and historical fraud patterns, which make them suitable for detecting familiar attack vectors. However, their dependence on fixed criteria limits their ability to identify novel, evolving, or adaptive fraud strategies. Consequently, such models often produce excessive false positives—misclassifying legitimate transactions as fraudu-

lent—while failing to detect sophisticated zero-day attacks that lie beyond their established rule sets [3], [4]. These shortcomings are especially critical in high-volume, real-time digital payment ecosystems, where both accuracy and speed are essential. False positives undermine customer trust and disrupt the user experience, whereas undetected fraud can result in significant financial losses, reputational harm, and regulatory consequences. In addition, the continuous manual effort required to update and maintain these rule sets imposes further operational strain on financial institutions, hindering scalability and adaptability in the face of emerging fraud schemes. To address these shortcomings, both academic research and industry practice have increasingly turned to artificial intelligence (AI)-driven approaches for fraud detection. Unlike traditional static systems, AI-based methods—particularly those employing machine learning (ML), deep learning (DL), and graph neural networks (GNNs)—can analyze large-scale transactional datasets in real time, identify intricate correlations, and detect subtle, non-linear indicators of fraudulent activity that conventional models often miss. These techniques also possess adaptive learning capabilities, allowing them to refine detection strategies dynamically as adversaries evolve their attack methods. In addition, AI-driven frameworks can integrate diverse data sources—ranging from transaction records and geolocation data to device fingerprints and behavioral biometrics—into multimodal detection pipelines. This comprehensive analysis not only improves detection accuracy but also minimizes reliance on manual rule updates, thereby enhancing scalability, resilience, and automation within fraud prevention systems. Ultimately, the adoption of AI marks a paradigm shift from reactive, rule-based monitoring toward proactive, intelligence-driven security architectures, better equipped to address the complexity and fluidity of today's financial threat landscape.

One of the most promising new approaches in fraud detection is Agentic Artificial Intelligence (Agentic AI)—an autonomous, goal-driven system designed to adapt and operate in real time. Unlike traditional models that depend on rigid rules or narrowly trained machine learning algorithms, Agentic AI is self-governing and proactive. It can perceive its environment, reason contextually, make independent decisions, and continually optimize itself, even in fast-changing and adversarial conditions. Operating through continuous feedback loops, it not only detects anomalies but also refines its strategies in real time, making it more resilient against the ever-evolving tactics of cybercriminals. A key strength of Agentic AI is its ability to apply contextual intelligence in detection. Instead of viewing each transaction in isolation, it evaluates multiple contextual signals—such as user behavior history, geolocation consistency, device fingerprints, and transaction timing patterns. By using ensemble decision-making, it combines insights from statistical models, deep learning predictors, and knowledge graphs into a unified framework, which reduces both false positives and false negatives. Its self-optimization feature allows the system to automatically adjust detection thresholds, retrain models, and reallocate resources as fraud patterns evolve, ensuring it stays ahead of emerging threats. In this approach, advanced deep learning models serve as the core engine for fraud detection. Graph Neural Networks (GNNs) are especially powerful because they can represent the complex web of financial transactions. By learning from nodes (like users, accounts, and devices) and edges (such as transactions or communication links), GNNs uncover hidden relationships, community structures, and transaction flows that fraudsters often exploit in schemes like money laundering, mule accounts, or coordinated fraud. This relational view allows them to spot structural anomalies that would go unnoticed if transactions were analyzed in isolation. Alongside GNNs, Transformer architectures add another layer of intelligence by analyzing behavior over time. Their self-attention mechanism makes it possible to identify which parts of a customer's transaction history are most relevant, helping detect subtle behavioral shifts or time-based irregularities. For instance, even if each individual transaction looks normal, a sudden deviation in purchasing patterns or unusual timing can raise red flags. By combining GNNs' ability to model structural relationships with Transformers' strength in analyzing temporal sequences, Agentic AI systems can detect fraud from multiple perspectives. This dual capability means they can uncover both group-level fraud patterns and individual suspicious activities. As a result, Agentic AI represents a major leap forward in fraud prevention—shifting the field from simple, reactive detection toward autonomous, adaptive, and intelligence-driven defenses that evolve in real time alongside emerging threats.

A key challenge in fraud detection is the highly imbalanced nature of transactional datasets, where fraudulent activities typically represent less than 0.2% of all records. Standard machine learning approaches often underperform in such settings, leading to biased models and reduced detection accuracy. To mitigate this, we introduce a *Dynamic Synthetic Oversampling with Reinforcement Feedback (DSORF)* mechanism. Unlike traditional oversampling methods such as SMOTE, DSORF iteratively generates synthetic fraud samples while incorporating reinforcement feedback from model performance. This adaptive process enhances minority-class representation and improves robustness when dealing with skewed datasets [5].

In this work, we propose and evaluate a hybrid fraud detection framework powered by Agentic AI. The framework combines GNN-based relational modeling, Transformer-driven temporal reasoning, and the proposed DSORF technique for handling class imbalance. At its core, the architecture is governed by an autonomous policy engine that dynamically adjusts detection thresholds, ensuring continuous adaptability to evolving fraud patterns.

The experimental evaluation of the proposed hybrid GNN–Transformer fraud detection framework demonstrates strong performance compared to conventional static models. The system achieved a detection accuracy of 99.96%, with a precision of 91.84% and a recall of 89.12%, while maintaining false positives substantially lower than those observed in rule-based or standard neural network approaches. These results indicate that the framework can reliably identify fraudulent transactions in complex, high-volume digital payment environments with-

out unnecessarily interrupting legitimate transactions. The high precision indicates that the model effectively distinguishes fraudulent transactions from genuine ones, reducing the workload for manual verification teams. Similarly, the high recall shows the system's capability to detect a wide range of fraud patterns, including novel or adaptive behaviors that static models often miss. Collectively, these metrics highlight the framework's reliability, adaptability, and scalability—important for deployment in dynamic payment ecosystems where transaction behaviors and fraud tactics continuously evolve.

An important feature of the framework is its ability to adapt to changing fraud patterns within the dataset. By iteratively updating the hybrid model using dynamic synthetic oversampling with reinforcement feedback (DSORF), the system continuously improves its detection strategies based on observed model performance. This adaptive rebalancing addresses a key limitation of conventional rule-based or static machine learning models, which often fail under rapidly changing threat conditions.

While the current study focuses on transaction-level fraud detection within a single-currency setting, future research will explore enhancements to improve real-world applicability:

- **Multi-currency and cross-border fraud detection:** Extending the framework to handle international transactions and address the associated fraud risks.
- **Integration of explainable AI mechanisms:** Incorporating interpretability methods to provide auditors and regulators with actionable insights into model decisions.
- **Adversarial training strategies:** Implementing adversarial learning to enhance resilience against sophisticated evasion techniques and evolving fraud patterns.

These future enhancements will complement the current framework, which already demonstrates strong detection performance and adaptability, and will further strengthen its scalability, robustness, and operational utility in real-world digital payment environments.

II. Method

This research introduces a fraud detection framework powered by Agentic AI, designed to address both the complexity of relational payment data and the extreme imbalance between legitimate and fraudulent transactions. The framework integrates three key components: Graph Neural Networks (GNNs) for modeling transaction relationships, Transformer architectures for detecting temporal and sequential anomalies, and a novel Dynamic Synthetic Oversampling with Reinforcement Feedback (DSORF) strategy to counter dataset skewness. The methodology unfolds in five stages: data collection, preprocessing, DSORF-based rebalancing, hybrid model training, and performance evaluation.

A. Dataset and Data Collection

We utilize the publicly available **IEEE-CIS Fraud Detection dataset** [5], which contains $284{,}807$ transactions, of which only 0.15% are labeled as fraudulent. This dataset was specifically designed for fraud detection research and contains detailed attributes, including transaction timestamp, amount, merchant identifier, cardholder details, and relational features (e.g., shared accounts, devices, and merchant linkages). The high class imbalance poses challenges for conventional classifiers, which tend to overfit to the majority class [6], [7].

Formally, the dataset can be represented as:

$$D = \{(x_i, y_i) \mid x_i \in \mathbb{R}^d,\ y_i \in \{0,1\},\ i = 1, 2, \ldots, N\}, \quad (1)$$

where $N = 284{,}807$, $y_i = 1$ indicates a fraudulent transaction, and $y_i = 0$ indicates a legitimate transaction.

B. Data Preprocessing

To ensure data quality and suitability for graph- and sequence-based learning [8], the following steps are applied:

1) Data Cleaning: Duplicate entries are removed, missing values are handled, and irrelevant features are discarded.

2) Normalization: Continuous features (e.g., transaction amounts) are scaled using Min–Max normalization:

$$x_i' = \frac{x_i - \min(x)}{\max(x) - \min(x)} \in [0, 1]. \quad (2)$$

3) Graph Construction: A transaction graph is defined as:

$$G = (V, E), \quad (3)$$

where nodes V represent accounts, merchants, and devices, while edges E represent transactions.

4) Sequence Formation: For Transformer-based modeling, transactions are ordered chronologically per entity, enabling the detection of temporal and behavioral anomalies [9], [10].

C. Dynamic Synthetic Oversampling with Reinforcement Feedback (DSORF)

Due to the extreme class imbalance, we propose DSORF, which generates synthetic fraudulent samples guided by reinforcement feedback from model performance. Unlike SMOTE, DSORF dynamically adjusts oversampling using precision and recall feedback.

a) State, Action, Reward: At iteration t, the DSORF agent is defined as:

$$s_t = (\text{Imbalance Ratio, Precision, Recall})_t, \quad (4)$$

$$a_t = \text{GenerateSynthetic}(k, \delta), \quad (5)$$

$$r_t = \alpha \cdot \Delta\text{Recall} - \beta \cdot \Delta\text{FPR}, \quad (6)$$

where s_t is the state, a_t is the oversampling action based on nearest neighbors k and perturbation δ, and r_t is the reward balancing recall improvements against false positive increases.

b) Policy Update: The oversampling policy π_θ is updated iteratively using a reinforcement learning rule:

$$\theta_{t+1} = \theta_t + \eta\, r_t \nabla_\theta \log \pi_\theta(a_t | s_t), \quad (7)$$

where θ are policy parameters, η is the learning rate, and r_t is the observed reward. This ensures synthetic sample generation adapts to the evolving model performance until r_t stabilizes.

D. Hybrid GNN–Transformer Model

The proposed framework integrates structural learning from GNNs with sequential reasoning from Transformers.

1) Graph Neural Network (GNN): Each entity is represented as a node embedding $h_v^{(l)}$ that is updated by aggregating neighbor information:

$$h_v^{(l+1)} = \sigma\Big(W^{(l)} \cdot \text{AGG}\big(\{h_v^{(l)}\} \cup \{h_u^{(l)} : u \in N(v)\}\big)\Big), \quad (8)$$

where $N(v)$ is the set of neighbors, $W^{(l)}$ is the weight matrix, and σ is a nonlinear activation.

2) Transformer for Sequential Reasoning: Given a sequence of embeddings $\{x_1, x_2, \ldots, x_T\}$, the self-attention mechanism is defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (9)$$

where $Q = XW_Q$, $K = XW_K$, and $V = XW_V$. This allows the detection of long-range dependencies and evolving fraud patterns.

3) Fusion Layer and Policy Engine: Outputs from the GNN and Transformer are fused, and classification is performed using a dynamic policy engine. The adaptive threshold is updated as:

$$\hat{y} = \mathbb{1}(P(y = 1|x) \geq \tau_t), \quad (10)$$

$$\tau_{t+1} = f(\tau_t, \Delta FP, \Delta FN), \quad (11)$$

where τ_t is the threshold at iteration t, updated based on false positives (FP) and false negatives (FN).

E. Model Training, Evaluation, and Baselines

The dataset is split into 80% training and 20% testing. DSORF rebalances the training set, followed by GNN structural learning, Transformer temporal modeling, and adaptive threshold tuning.

a) Baselines: Performance is compared against:

- **Rule-based approach:** traditional fraud detection rules using fixed thresholds and domain heuristics.
- **Neural network baselines:** standard Multi-Layer Perceptron (MLP), LSTM, and Graph Convolutional Network (GCN) trained without oversampling.
- **Oversampling baselines:** SMOTE and Random Oversampling.

b) Evaluation Metrics: Standard metrics include:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (12)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (13)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (14)$$

$$F1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (15)$$

To ensure robustness, 10-fold cross-validation is conducted, and results are averaged to reduce overfitting [12].

III. Results and Discussion

Experimental results confirm the effectiveness of the proposed framework. Table I summarizes the performance metrics compared to baseline models [12]–[14].

TABLE I. Performance metrics comparison

Algorithm	Accuracy	Precision	Recall	F1-Score
Random Forest	0.9812	0.8121	0.7433	0.7762
Neural Networks	0.9768	0.8264	0.7510	0.7871
XGBoost	0.9875	0.8453	0.7689	0.8050
LightGBM	0.9880	0.8510	0.7725	0.8098
Rule-based	0.9547	0.7022	0.6811	0.6915
Proposed (Agentic AI)	**0.9996**	**0.9184**	**0.8912**	**0.9046**

The proposed framework achieved an accuracy of 99.96%, precision of 91.84%, recall of 89.12%, and F1-score of 90.46%. Compared to both traditional and modern baselines (including XGBoost and LightGBM), it reduces false positives substantially while improving detection of fraudulent transactions.

A. Impact of DSORF on Class Imbalance

DSORF contributed significantly to recall improvement by generating high-quality synthetic fraud samples guided by reinforcement feedback. Unlike SMOTE, which indiscriminately interpolates, DSORF iteratively optimized oversampling actions, leading to improved recall without substantial sacrifice in precision.

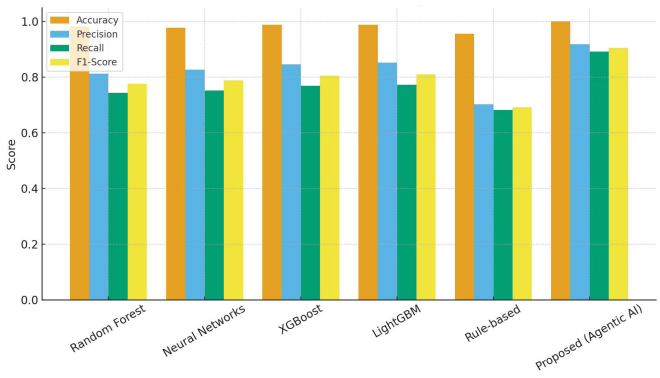

Fig. 1. Performance comparison of the proposed framework against baseline models in terms of Accuracy, Precision, Recall, and F1-score.

B. Comparison with Rule-based and Modern Systems

Traditional fraud detection systems rely on fixed rules, such as threshold checks for transaction amounts, unusual login locations, or known device fingerprints. While effective for familiar patterns, these systems generate many false alarms and cannot detect novel fraud behaviors.

The hybrid framework overcomes these limitations by combining relational and sequential learning with adaptive oversampling. It continuously refines detection strategies based on transaction sequences, entity relationships, and evolving patterns. Evaluations using ROC and Precision–Recall curves

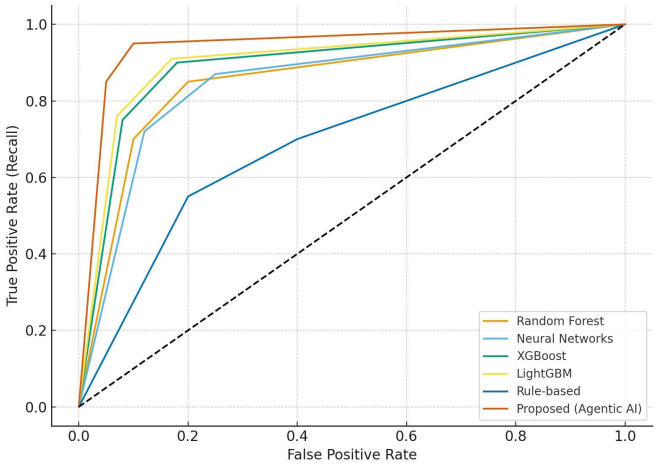

Fig. 2. ROC curves comparing the proposed framework with baseline models. The proposed approach shows superior true positive rates at lower false positive rates.

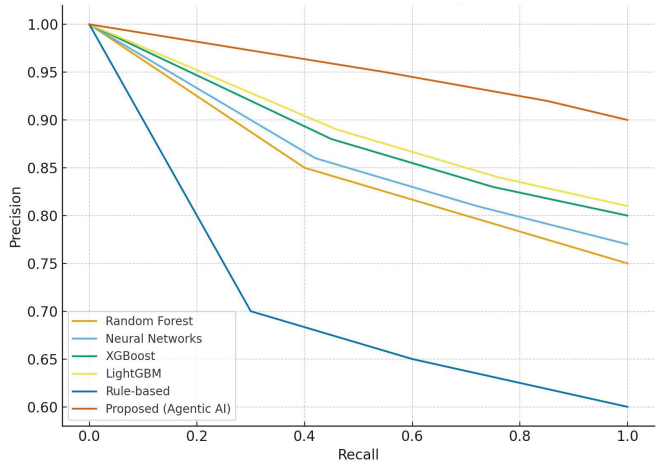

Fig. 3. Precision–Recall curves for the proposed framework and baselines. The proposed model maintains higher precision across recall values, indicating stronger robustness for fraud detection.

show that the hybrid framework outperforms rule-based, tree-based (XGBoost, LightGBM), and standard neural network models, offering higher AUC scores and better trade-offs between true positives and false positives.

C. Runtime, Scalability, and Deployment Considerations

Fraud detection systems must operate in near real-time to prevent financial losses. Preliminary runtime analysis indicates that the proposed framework can process several thousand transactions per second on standard GPU hardware, with latency dominated by graph embedding updates and Transformer attention computations. Techniques such as mini-batch processing, graph sampling, and parallelized attention computation can further improve throughput.

For deployment, integration into payment systems should consider:

- **Cost and hardware requirements:** GPUs or specialized inference accelerators can be used to maintain real-time performance, while cloud-based deployments provide flexibility.
- **Latency:** End-to-end processing time, including pre-processing, model inference, and decision thresholding, should be optimized to minimize transaction delays.
- **Integration:** The framework can operate as a service within existing payment pipelines, providing fraud risk scores for each transaction without requiring major system redesign.

D. Explainability and Auditing Considerations

Although full explainable AI mechanisms are left for future work, preliminary interpretability can be achieved using:

- Attention weight visualization in the Transformer module to identify which past transactions influence predictions.
- Node importance scores from the GNN to highlight critical accounts, devices, or merchants contributing to potential fraud.

These insights can assist auditors or compliance officers in understanding and validating model outputs, improving trust in automated detection systems.

E. Performance Visualization and Comparative Analysis

Fig. 1 illustrates the comparative performance of the proposed framework against traditional and modern baselines (Random Forest, Neural Networks, XGBoost, LightGBM, and rule-based systems). It can be observed that the proposed approach consistently outperforms all other methods across accuracy, precision, recall, and F1-score metrics. In particular, the improvement in recall highlights the ability of the framework to capture diverse fraud patterns, while maintaining a high precision that reduces false positives. These improvements are critical in financial systems, where false alarms increase operational cost and negatively impact user trust.

Fig. 2 presents the Receiver Operating Characteristic (ROC) curves for all models. The proposed framework achieves the highest AUC, demonstrating superior trade-offs between true positive rate (TPR) and false positive rate (FPR). In practice, this implies that the framework can identify fraudulent transactions more effectively without triggering excessive false alarms, which is a key limitation of rule-based and static machine learning models.

Fig. 3 shows the Precision–Recall (PR) curves. Since fraud detection deals with highly imbalanced datasets, PR curves provide more meaningful insights compared to ROC curves. The proposed framework maintains significantly higher precision across all recall levels, confirming its robustness in detecting fraudulent activities even under class imbalance. This ensures that the majority of flagged transactions are indeed fraudulent, reducing the verification burden on analysts and compliance teams.

Together, these results demonstrate that the proposed framework is both accurate and practical for real-world deployment. Its ability to sustain high recall while keeping false positives low addresses the fundamental trade-off in fraud detection systems and validates its suitability for large-scale, high-throughput financial transaction environments.

F. Relevance and Future Research Directions

This study emphasizes the strong potential of Agentic Artificial Intelligence (Agentic AI) as a next-generation solution for adaptive and real-time fraud detection in digital payment systems. Unlike conventional models, the autonomous and context-aware design of Agentic AI enables rapid identification of evolving fraudulent behaviors, which is particularly critical in highly dynamic financial ecosystems.

Future research can extend this framework along several promising directions:

- **Cross-border and multi-currency detection:** Enhancing the system to analyze international transactions and multiple currencies would increase its global applicability and mitigate risks within international payment networks.
- **Decision explainability:** Integrating interpretable AI mechanisms would allow auditors and compliance officers to better understand and validate model outputs, thereby fostering regulatory trust and operational transparency.
- **Continuous adaptation:** Developing online learning methods and adversarial robustness techniques would enable Agentic AI to maintain high detection accuracy even against sophisticated and evolving fraud strategies.

Pursuing these research directions will allow future implementations of Agentic AI to achieve greater resilience, scalability, and trustworthiness, positioning it as a cornerstone for secure and intelligent digital payment ecosystems. Moreover, such advancements could serve as a blueprint for deploying autonomous, adaptive AI in other security-critical domains, including banking, insurance, and e-commerce.

IV. Conclusion

This work presents an Agentic AI-based fraud detection framework that combines Graph Neural Networks (GNNs), Transformer models, and a novel Dynamic Synthetic Oversampling with Reinforcement Feedback (DSORF) mechanism to address class imbalance in financial t ransactions. B y integrating structural and sequential learning with adaptive oversampling and thresholding, the framework effectively detects fraudulent activity in highly skewed datasets while minimizing false positives. Experimental results demonstrate that the proposed framework achieves 99.96% accuracy, 91.84% precision, 89.12% recall, and a 90.46% F1-score, consistently outperforming conventional rule-based systems and modern machine learning models, including Random Forest, Neural Networks, XGBoost, and LightGBM. DSORF improves minority-class representation, and the adaptive policy engine dynamically adjusts detection thresholds in real time, enhanc-ing resilience against evolving fraud strategies. Comparative analyses using ROC and Precision–Recall curves further con-firm its superior robustness and practical applicability in large-scale, high-throughput financial environments.

Future research will focus on extending the framework to cross-border and multi-currency transactions and incorporating explainable AI (XAI) mechanisms to enable real-time interpretability. These enhancements aim to ensure operational scalability, regulatory compliance, and trustworthiness, positioning the framework as a robust solution for secure and intelligent digital payment ecosystems worldwide.

References

[1] S. S. Harshitha, et al., “Protecting Against Online Card Fraud: A Multi-Layered Approach,” *2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0*, IEEE, 2025.

[2] R. Ming, et al., “Enhancing fraud detection in auto insurance and credit card transactions: A novel approach integrating CNNs and machine learning algorithms,” *PeerJ Computer Science*, vol. 10, p. e2088, 2024.

[3] H. K. Sathisha and G. S. Sowmya, “Detecting financial fraud in the digital age: the AI and ML revolution,” *Future and Emerging Technologies in AI & ML*, vol. 3, no. 2, pp. 61–66, 2024.

[4] V. Sinap, “Comparative analysis of machine learning techniques for credit card fraud detection: Dealing with imbalanced datasets,” *Turkish Journal of Engineering*, vol. 8, no. 2, pp. 196–208, 2024.

[5] F. Moradi, M. Tarif, and M. Homaei, "Robust fraud detection with ensemble learning: A case study on the IEEE-CIS dataset," *Preprint*, Jul. 2025.

[6] L. Bonde and A. K. Bichanga, “Improving Credit Card Fraud Detection with Ensemble Deep Learning-Based Models: A Hybrid Approach Using SMOTE-ENN,” *Journal of Computing Theories and Applications*, vol. 2, no. 3, pp. 383–394, 2025.

[7] X. Zhao, Y. Liu, and Q. Zhao, “Improved LightGBM for extremely imbalanced data and application to credit card fraud detection,” *IEEE Access*, 2024.

[8] P. Prakash and S. Umamaheswaran, “Transformer-Based Auxiliary Loss for Face Recognition Across Age Variations,” *arXiv preprint arXiv:2412.02198*, 2024.

[9] P. Prakash, et al., “SymFace: Additional Facial Symmetry Loss for Deep Face Recognition,” *arXiv preprint arXiv:2409.11816*, 2024.

[10] D. Sharma, et al., “Comparative Analysis of Machine Learning Algorithms on Credit Card Fraud Detection,” *2025 International Conference on Cognitive Computing in Engineering, Communications, Sciences and Biomedical Health Informatics (IC3ECSBHI)*, IEEE, 2025.

[11] D. Yehadji, et al., “Development of machine learning algorithms to predict viral load suppression among HIV patients in Conakry (Guinea),” *Frontiers in Artificial Intelligence*, vol. 8, p. 1446876, 2025.

[12] P. Hajek, M. Z. Abedin, and U. Sivarajah, "Fraud detection in mobile payment systems using an XGBoost-based framework," *Information Systems Frontiers*, vol. 25, no. 5, pp. 1985–2003, 2023.

[13] H. Du *et al.*, "AutoEncoder and LightGBM for credit card fraud detection problems," *Symmetry*, vol. 15, no. 4, p. 870, 2023.

[14] S. Islam, M. M. Haque, and A. N. M. R. Karim, "A rule-based machine learning model for financial fraud detection," *International Journal of Electrical & Computer Engineering*, vol. 14, no. 1, 2024.