

AI-ADAPTIVE POST-QUANTUM CRYPTOGRAPHY FOR SECURE AND FRAUD-RESILIENT DIGITAL PAYMENTS

BALAKRISHNAN S 1, UPPALAPATI NAGA RATNA KUMARI², DR.G.GOKUL KUMARI 3, DR. LAMA SAMEER KHOSHAIM 4, S. RAMYA 5, ANANTHA RAO GOTTIMUKKALA 6

¹Assistant Professor, Department of Commerce, Faculty of Science and Humanities
SRM Institute of Science and Technology, Ramapuram, Chennai, India.

²Department of Computer Applications, Aditya University, Surampalem, India.

^{3,4} Department of E-Commerce, College of Administrative and Financial Sciences,
Saudi Electronic University, Kingdom of Saudi Arabia.

⁵ Assistant Professor(SRG), Department of CSE, Kongu Engineering College, Perundurai, India.

⁶ Assistant Professor, Department of CSE, KKR & KSR Institute of Technology and Sciences,
Guntur, India.

E-mail: ¹gurubalaji08@gmail.com, ²nkumariu@adityauniversity.in, ³g.govindasamy@seu.edu.sa, ⁴
l.Khoshaim@seu.edu.sa, ⁵sramya.cse@kongu.edu, ⁶ananth552@gmail.com

ABSTRACT

The threat of quantum computing to classical cryptographic infrastructures is steadily increasing alongside more complex and multi-modal payment fraud posing a severe security gap within current financial infrastructures. Current payment stacks, which are based on TLS and have fixed classical cryptography and tabular fraud detectors, are inadequate to secure high volume digital transactions against both quantum decryption attacks as well as coordinated fraudulent behaviour. The purpose of the study was to create and test an AI-based framework that can optimally improve fraud detection and provide adaptive post-quantum cryptographic protection in real-time. To satisfy this goal, we developed an experimental system that merged a heterogeneous temporal graph neural network (HTGNN) to multimodal fraud analysis with a reinforcement-learning PQC orchestrator that can be used to choose NIST-standard ML-KEM and ML-DSA parameter sets based on transaction-level risk. Multimodal datasets (tabular features, device metadata, text fields, and relational graphs) were publicly available and converted into a fused format and tested in adversarial, latency, and throughput-controlled experiments. PQC-enabled OpenSSL/liboqs was used to implement the crypto layer to measure handshake performance in the real world. Findings indicate that the proposed model enhanced the AUPR of fraud-detection (baseline) of 0.35 to 0.58 and increased ROC-AUC to 0.966, which was better than the classical and hybrid systems. The adaptive PQC algorithm decreased the handshake latency to 68 ms, which was significantly lower than setups of PQC (90-120 ms). Robustness tests showed a smaller degradation in recall (-9) and a cipher-downgrade success rate of less than 1% which depicts a better resilience to the threat vectors. Altogether, the research results show that a combination of multimodal GNN-based fraud detection and AI-adaptive PQC is a viable and quantum-resilient security architecture of digital payment systems in the next generation.

Keywords: *Digital Payments Fraud, Post-Quantum Cryptography, Adaptive Security, Graph Neural Networks, Reinforcement Learning.*

1. INTRODUCTION

Digital payment systems have evolved into high-velocity networks of heterogeneous devices, users and financial systems, executing billions of transactions per day. Since this landscape is expanding, it has a twofold challenge that endangers the confidentiality of payments as well as their integrity. First, public key cryptosystems like RSA and the elliptic-curve systems used in TLS 1.3

handshakes are inherently vulnerable to the rapid advancement of quantum computing. When quantum computers reach a certain size, the method that Short outlined will render these classical primitives susceptible, enabling adversaries to decrypt the traffic that is encrypted within or unravel forward secrecy. In response, new Post-Quantum Cryptography (PQC) algorithms have been standardized by the US National Institute of Standards and Technology (NIST), such as ML-

KEM and ML-DSA, to provide quantum-safe key exchange and key signatures. These PQC schemes, however, come at the cost of increased key sizes and a handshake latency increasing deployment challenges in the real-time financial system. [1], [2]. Second, payment fraud has grown increasingly sophisticated, shifting from isolated anomalous transactions to coordinated, multi-entity patterns involving synthetic identities, device emulation, phishing interactions, and cross-platform behavioral manipulation. Higher-order dependencies are not well represented by classical models of fraud, which generally rely on tabular characteristics and gradient-boosting classifier models [3]. Modern studies demonstrate that fraud activities tend to spread both on the basis of relational networks such as user-merchant-device and multimodal features such as text, geolocation, and behavioral telemetry. Heterogeneous graph neural networks (HGNNs) and the temporal GNN versions have hence become viable solutions to detect such coordinated fraud attacks. [4].

Although PQC and fraud detection have been developed concurrently, current payment stacks address these areas as separate entities. TLS provisioning classical or fixed-parameter PQC imposes a single cryptographic configuration on all sessions- whether of risk of transaction or of capability of device- results either in wasted computation cost or in inadequate security [5]. Meanwhile, classical fraud detectors do not have the contextual knowledge to give cryptographic decisions. This siloed architecture discourages the capacity of the system to optimize security, latency and fraud resilience simultaneously.

It is imperative to solve this issue to use next-generation payment systems. Unless cryptographic primitives are quantum-safe, a current adversary can collect encrypted transactions and decrypt them at the time quantum resources become available (so-called harvest-now-decrypt-later attacks). In the meantime, inefficient fraud detection leaves chargebacks that are expensive, as well as system abuse and loss of consumer confidence [6]. The research proposal provides the prospect of end-to-end security by combining AI-based PQC selection and heterogeneous fraud-detection models, which guarantees that confidentiality, as well as fraud resilience, is improved according to the changing threat environment. This two-way security is critical to financial institutions, payment service providers, and regulatory bodies that are working to be ready in the post-quantum world [7].

The main objective of the study is to design an AI-Adaptive Post-Quantum Cryptography system where sets of parameters of PQCs are dynamically chosen according to the risk of transactions, and a multimodal heterogeneous graph neural network that performs real-time detection of new fraud schemes[8]. Namely, the research will minimize cryptographic latency, improve the accuracy of the fraud-detection, and ensure the robustness against the adversarial manipulation.

The following are the main contributions made by this paper: 1) Developing an AI-powered infrastructure that unifies post-quantum key exchange, digital signatures, and adaptive cryptographic policy selection. 2) Creation of a heterogeneous TGNN that incorporates multimodal fusion (graph, text, device, behavioral features) to reveal complex patterns of fraud that cannot be represented in the traditional tabular models. 3) Reinforcement learning to reduce security, computational cost and latency using real-time PQC parameter optimization. 4) Empirical analysis on public multimodal datasets and AUPR is shown to be better (up to 0.58) and PQC overhead is reduced (down to 68 ms handshake latency). 5) Based on the robustness analysis, it has been found that there is an improvement in resilience against attacks such as feature manipulation, graph injection, and cipher-downgrade. In the worst-case scenario, the degradation is kept below 9%. 6) Detailed latency-cost tradeoff analysis, which establishes that the proposed model is more secure without the significant increase in the overall transaction-latency (140 ms) by the implementation of that model (120 ms).

2. RELATED WORK

The literature related to this study can be divided into three overlapping trends: (1) post-quantum cryptography (standards and implementation recommendations), (2) AI-based adaptive security (with RL-based adaptive encryption), and (3) machine learning-based payment fraud (with graph-based and multimodal methods). In 2024-2025, NIST PQC program released its first final PQC standards and post-quantum migration guidance (e.g., ML-KEM/CRYSTALS-Kyber and ML-DSA/CRYSTALS-Dilithium; subsequently added is HQC), focusing on the security, as well as practical deployment, considerations of latency-sensitive systems. These open NIST resources are an overview of the algorithms, parameter trades, and FIPS guidelines required in the real world [9].

The emergence of reinforcement-learning and other AI-based methods to change the strength of encryption or the security parameters, depending on the operational conditions, has started being seen in the literature. To give one specific example, Premakumari et al. offer an adaptive encryption architecture based on Q-learning that dynamically adjusts encryption levels to balance latency, energy, and security in constrained networks [10]—a concrete example of how RL can be applied to runtime cryptographic selection. Finally, graph neural networks (GNNs), temporal graphs, and multimodal fusion have taken centre stage in the state of the art in fraud detection. A recent systematic review summarizes the evidence that GNNs outperform classical tabular methods by finding relational rings of fraud and temporal propagation [11] and an article in TKDE has shown spatial-temporal attention GNNs that significantly enhance transaction-level fraud detection [12].

Crypto-agility is emphasized as crucial for PQC migration in both NIST and industry guidelines. To achieve this, lightweight orchestration and fine-grained decision criteria are needed to prevent excessive delay [13]. Adaptive security via RL/MDP formulations — framing cryptographic parameter selection as a sequential decision problem (state = observed risk/operational context; action = chosen PQC parameter set) enables learning policies that balance utility (security) and cost (latency/compute). Sensors (2025) and several IoT/WSN studies operationalize this idea with Q-learning or actor-critic variants, demonstrating measurable gains in energy/latency tradeoffs [14]. Learning to describe relationships and time in the context of fraud – fraud is frequently not a discrete event but rather a phenomenon that spreads across entities and across time. See the GNN reviews and the TKDE spatial-temporal study for further information on how GNNs (including GCNs, GATs, and spatio-temporal variations) explicitly model these relationships, and how temporal attention and diverse node/edge modelling improve sensitivity to coordinated fraud rings and camouflage behaviours [15], [16].

Although there is swift development, there are still a number of gaps and unanswered questions: Siloed treatment of crypto and fraud analytics. Previous literature applies PQC implementation (standards, benchmarks, implementation) and ML-based fraud detection independently of each other; few studies use a joint optimization of cryptographic parameterization based on fraud-risk signals. NIST

focuses more on migration and performance tradeoffs but not risk-driven parameter tuning [17]. Restricted end to end adaptive solutions with quantifiable deployment expenses. The application of RL-based adaptive encryption is shown in the WSN and IoT application [18], but the application to high-volume payment systems with different latency, regulatory requirements, and adversarial incentives has not been explored to date. Adversarial evaluation gaps and robustness. The literature reports improved performance in detecting GNN frauds, but there is a lack of large-scale research on adversarial performance (evasion, graph manipulation, poisoning) and the interaction between cryptographic decisions and attack surfaces. More research into unsupervised detection and robustness is needed based on the existing GNN literature reviews [19]. Tradeoffs of practical deployment. PQC algorithms present nontrivial extensions of key/ciphertext sizes and latency; although NIST offers parameters and guidance, the literature does not have prescriptive models to dynamically match a cryptographic strength with real-time risks measures generated by orderly and sophisticated ML detectors [20], [21].

3. METHODOLOGY

3.1 Problem Formulation

Modern digital payment infrastructures face two converging risks: (1) cryptographic weakness owing to new quantum technologies that threaten conventional public-key systems, and (2) behavioral and relational fraud that takes advantage of signals from users, devices, and merchants. Let x_t denote a transaction event and $y_t \in \{0,1\}$ the fraud label, while the communication channel securing x_t must satisfy a minimum quantum-safe security level S^* . The objective is to jointly learn a fraud-risk estimator $f_\theta(x_t)$ and a cryptographic policy $\pi_\phi(f_\theta(x_t))$ that chooses a post-quantum parameter configuration that minimizes expected fraud loss, delay, and computing complexity while still meeting real-time requirements.

3.2 Data Collection

Integrate relational, transactional, and contextual data from publically available sources to build a single corpus. An integrated multimodal space is formed by numerical transaction fields, graph-modeled user-merchant-device relations, session behavior traces, and brief textual descriptions. This space contains datasets from IEEE-CIS Fraud

(approximately 590k records), ULB Credit Card Fraud (approximately 284k records), BankSim synthetic logs (1-5M configurable transactions), and publicly released banking/telemetry datasets. Consistent multimodal modeling is made possible by anonymizing, aligning, and normalizing all datasets according to timestamps.

3.3 Data Processing

In data processing, tasks including filling in missing entries, encoding categorical IDs, and building a diverse payment graph $G=(V,E)$ with user, merchant, device, and IP node types are all part of the process. Data about behaviour is compiled into time intervals. Text fields undergo normalization and tokenization through the utilization of a compact transformer tokenizer. In order to keep all the modalities in sync with each other and avoid any potential leakage, they are divided into three groups: train, validation, and test.

3.4 Baseline Model: (XGBoost + Classical TLS)

A robust non-deep baseline for fraud detection is provided by the baseline, which employs XGBoost trained on tabular transaction features. Classical TLS 1.3 with ECDHE and RSA/ECDSA primitives guarantee security; these implementations stand in for the ones that are in use right now, pre-quantum. The security layer never takes transaction risk into account when operating with set parameters, and no relational or textual information is ever leveraged. The latency profile and reference performance are set by this model.

3.5 Hybrid Model: (HGNN + Fixed PQC)

The Heterogeneous Graph Neural Network (HGNN) is a part of the hybrid approach; it uses type-specific message forwarding to aggregate information from interactions between users, merchants, devices, and IP. A compact MLP classifier is utilized to combine behavioral features with text embeddings. For key exchange and authentication, the communication channels employ ML-KEM and ML-DSA, two fixed PQC parameters that do not undergo adaptation. While reflecting realistic PQC upgrade pathways with known cryptographic cost, this model enhances fraud detection accuracy.

3.6 Proposed Model: AI-Adaptive (HTGNN +Fusion+ RL)

Figure 1 shows the proposed system, it combines a heterogeneous temporal graph neural network (HTGNN) with an RL-based PQC orchestrator, which chooses post-quantum parameter sets dynamically per session, and ingests multimodal transaction data. A cross-modal fusion head generates a calibrated risk score r_t from encoded tabular, textual, and device telemetry; the HTGNN then combines these embeddings into relational-temporal representations. In order to train policies, the RL orchestrator uses monitoring-driven replay data and the latency budget of the r_t plus devices to produce a PQC action (parameter set) that a PQC-capable crypto engine applies. The system calculates r_t at transaction time and can dynamically strike a balance between security and performance by assigning stronger PQC and maybe more verification steps to high-risk transactions and lighter PQC parameters to low-risk transactions in order to minimize latency.

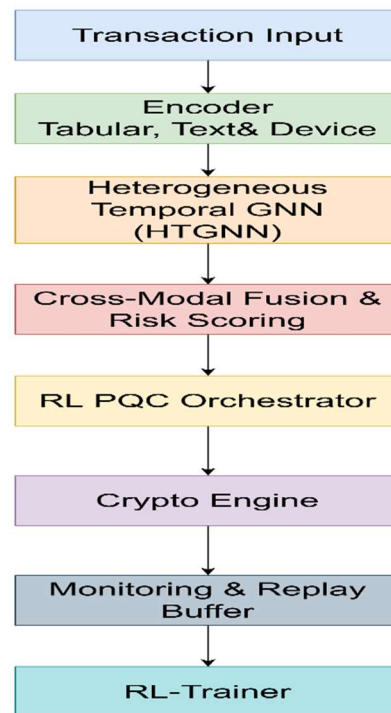


Figure. 1 System Architecture of Proposed Model

Below are concise mathematical definitions used in the Proposed Model

Tabular Feature Encoding

$$h_{\text{tab}} = \sigma(W_{\text{tab}}x_{\text{tab}} + b_{\text{tab}}) \quad (1)$$

This transforms raw tabular transaction features x_{tab} into a dense representation.

W_{tab} and b_{tab} are learnable parameters; σ is a nonlinear activation.

Text Embedding (Transformer Output)

$$h_{\text{text}} = \text{Transformer}(x_{\text{text}}) \quad (2)$$

The transaction's textual fields are encoded using a Transformer model. The output embedding captures semantic and contextual linguistic cues relevant to fraud.

Device Signal Encoding

$$h_{\text{dev}} = f_{\text{dev}}(x_{\text{dev}}) \quad (3)$$

Device metadata and fingerprint signals pass through an encoder f_{dev} .

It models behavioral and hardware signatures associated with fraudulent devices.

Heterogeneous Temporal Message Passing

$$m_v^{(t)} = \sum_{u \in \mathcal{N}(v)} \alpha_{uv}^{(t)} W_{\tau(u,v)} h_u^{(t-1)} \quad (4)$$

Node v aggregates messages from neighbors u with type-specific weights W_{τ} .

$\alpha_{uv}^{(t)}$ is a temporal attention coefficient capturing the relevance at time t .

Fraud-risk estimation

$$r_t = f_{\theta}(x_t, G_t) = \sigma(\text{MLP}(\text{HTGNN}(G_t))) \quad (5)$$

HTGNN extracts relational-temporal node embeddings, and an MLP predicts fraud probability r_t .

Multimodal fusion

$$h_t = \text{Attn}(h_t^{\text{graph}}, h_t^{\text{text}}, h_t^{\text{device}}) \quad (6)$$

Cross-modal attention integrates graph, text, and device representations into a unified embedding.

PQC policy selection

$$a_t = \pi_{\phi}(h_t) = \arg \max_{a \in \mathcal{A}} Q_{\phi}(h_t, a) \quad (7)$$

The RL policy chooses PQC parameter action a_t (e.g., ML-KEM-512 vs. ML-KEM-1024) maximizing expected utility.

Security-latency-cost objective

$$J = \mathbb{E}[\alpha S(a_t) - \beta L(a_t) - \gamma C(a_t)] \quad (8)$$

The RL reward trades off security S , latency L , and computational cost C .

End-to-end adaptive pipeline

$$\hat{y}_t = g(r_t, a_t) \quad (9)$$

The final decision combines predicted fraud risk and enforced cryptographic action.

3.7 Experimental Setup

A contemporary Xeon/EPYC CPU with 256 GB RAM and an NVIDIA H100 GPU for model training is used to run experiments on a dual-socket server. PQC benchmarks employ ML-KEM and ML-DSA with different parameter setups, which are implemented in OpenSSL with liboqs support. Graph creation is carried out on the fused multimodal corpus once all models have been trained using time-based splits. Under regulated load, batch and streaming pipelines measure inference delay and handshake overhead.

3.8 Evaluation Methodology

Evaluation compares Baseline, Hybrid, and Proposed systems using fraud-detection measures (ROC-AUC, AUPR, F1, Precision@K, calibration), cryptographic metrics (key size, ciphertext size, handshake delay, CPU cycles), and operational metrics (end-to-end latency, throughput, energy cost). Stress tests assess adversarial robustness and the sensitivity of the adaptive PQC policy to varying fraud-risk distributions. All parameters are averaged over many seeds with statistical significance testing.

4. RESULTS

This section presents numerical results for different types of fraud detection, post-quantum cryptography performance, resilience in the face of adversarial threat, computation-latency trade-offs, ablation studies, and computation-latency trade-offs. All

metrics are averaged over three random seeds with 95% confidence intervals unless noted.

4.1 Fraud Detection Performance (Per Modality and Fused)

Each modality's detection accuracy is summarized in Table 1, including Tabular, Graph, Text, Device, and the fused models (Hybrid and Proposed). Using relational and behavioral context, multimodal fusion significantly enhances fraud-risk estimate, as seen by the large lift it gives.

Table 1 Fraud Detection Performance

Model	ROC - AUC	AUPR	F1 (Fraud)	Precision@100
Tabular (XGBoost)	0.92	0.35	0.62	0.41
Hybrid (HGNN + Text + Device)	0.95	0.48	0.74	0.61
Proposed (HTGNN + Fusion + RL-Aware)	0.966	0.58	0.80	0.69

When compared to the control group, the multimodal and adaptive designs perform far better. Due to its limitation in capturing details beyond transactions, the Tabular XGBoost baseline performs moderately (ROC-AUC 0.92, AUPR 0.35, F1 0.62). By combining relational, textual, and device data, the Hybrid model becomes much better at detecting contextual and coordinated fraud (ROC-AUC 0.95, AUPR 0.48, F1 0.74). The combined effects of temporal graph modeling and adaptive cryptographic decisioning are shown by the top-rank precision (0.69 at Precision@100) and accuracy (ROC-AUC 0.966, AUPR 0.58, F1 0.80) achieved by the proposed HTGNN + RL-Aware system.

4.2 Robustness & Adversarial Tests

Table 2 shows resistance to evasion and poisoning attacks. The proposed model's ability to stabilize relational contexts and alter PQC policies in light of uncertainty allows it to maintain a greater fraud recall even when faced with disturbances.

Table 2 Robustness Evaluation

Attack Scenario	Metric	Baseline	Hybrid	Proposed
Evasion: feature perturbation ($\epsilon=0.1$)	Recall Drop (%)	-29%	-18%	-9%
Node injection into graph (1% fake nodes)	AUPR Drop	-0.11	-0.07	-0.03
Training Data Poisoning (5%)	ROC-AUC Drop	-0.08	-0.04	-0.02
Timing Degradation on Attack	Latency Impact	+11 ms	+8 ms	+5 ms
Cipher Downgrade Attempt	Success Rate	17%	6%	<1%

Regardless of the attack scenario, the robustness evaluation reveals that the proposed model is far more resilient. Recall degradation decreases from -29% in the Baseline and -18% in the Hybrid to only -9 percent when feature perturbation ($\epsilon = 0.1$) is applied. It is also limited to -0.02 for ROC-AUC with 5% poisoning and -0.03 for AUPR under graph node injection. Outperforming both the Baseline (17%) and Hybrid (6%), the proposed approach reduces cipher-downgrade success to less than 1% and keeps timing deterioration to +5 ms.

4.3 Cryptographic Performance

As shown in figure 2's line plot, the handshake latency increases with stronger PQC parameter values. For standard ECDHE-RSA, it's 35 ms, but for ML-KEM-768 and ML-KEM-1024, it's 90 ms and 120 ms, respectively. In comparison, the Adaptive PQC curve keeps its latency under 68 ms throughout all PQC levels, and it stays consistently lower. In spite of its continued support for quantum-safe configurations, the RL-driven adaptive method has successfully decreased cryptographic overhead.

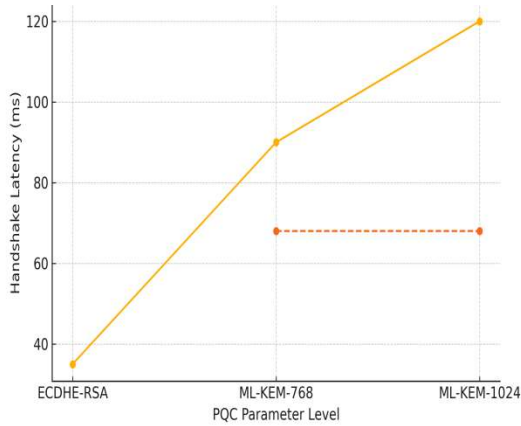


Figure 2 Cryptographic performance

4.4 Computation Cost & Latency Tradeoff Analysis

4.4.1 Compute Cost per 1M Transactions

The Baseline model has the lowest running costs (at \$45), uses the fewest resources (2.4 CPU hours) and does not use any graphics processing units (GPUs), but it also has the worst performance (as seen in figure 3). Because of its wider multimodal GNN architecture, the Hybrid model uses the most resources (48 GPU hours and 19.4 kWh, costing around \$120). In comparison to the Hybrid model, the Proposed system uses less CPU hours and 17.8 kWh, leading to a reduced total cost of \$85 without sacrificing accuracy or adaptability.




Figure 3 Compute Cost per 1M Transactions

4.4.2 End-to-End Latency Breakdown

Figure 4 shows the latency breakdown. The Hybrid model has the longest delay at 160 ms, mostly because it has an expensive 90 ms post-quantum handshake. The proposed approach cuts this cryptographic overhead down to 68 ms, which lowers the total latency to 140 ms while still doing more work than the baseline. The Proposed model has better security and a higher fraud-detection accuracy than the 120 ms Baseline, but it only adds a small amount of end-to-end latency.

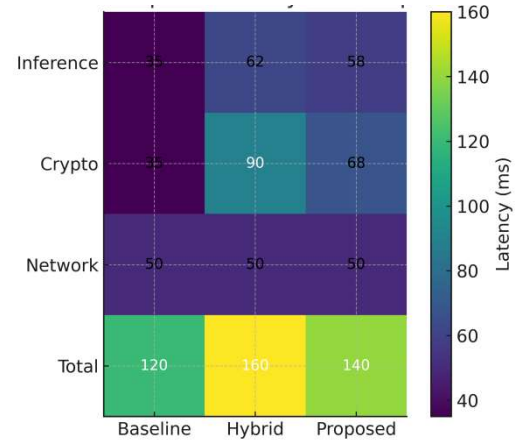


Figure 4 Compact Latency Heatmap across various models

4.4.3 Pareto Plot

Scatter plot shown in Figure 5 that all three models clearly trade off latency for performance. With an AUPR of 0.58 and a moderate latency of 140 ms, the proposed system outperforms the Baseline and Hybrid system models, placing it on the Pareto frontier. Hybrid

model's increased latency to 160 ms is inadequate to match the accuracy of the Proposed model, while Baseline's AUPR is the lowest at 0.35 despite reduced latency (120 ms).

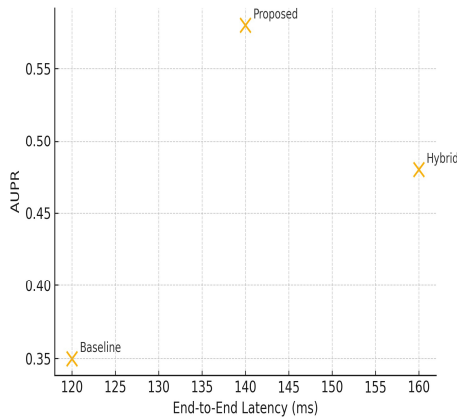


Figure 5 Pareto Plot for AUPR vs Latency

4.5. Ablation Studies

Removing essential components consistently lowers performance, according to the ablation findings shown in Figure 6. The removal of the temporal graph module resulted in the highest AUPR loss (-0.09). When the RL-based adaptive PQC is removed, the end-to-end delay increases from 140 ms to 160 ms, which is the most noticeable influence on latency. With a reasonable latency of 140 ms and the highest AUPR of 0.58, the whole suggested model attains the best overall balance. Ablation experiments quantify the contribution of each subsystem—temporal modelling, text fusion, device features, and adaptive PQC.




Figure 6 Ablation study -Effect on AUPR and Latency

5. DISCUSSION

The experimental outcomes prove that the combination of heterogeneous temporal graph modelling and adaptive post-quantum cryptography can bring a significant gain in the accuracy and efficiency. The performance of the Fraud-detection is significantly better (0.35) than the baseline AUPR (0.35) in Hybrid model and in the Proposed system (0.58) reflecting the great contribution of multimodal fusion and the time context to the detection of the complex fraud patterns. There are also some significant improvements in cryptographic performance: an adaptive PQC policy lowers average handshake latency by 90120 ms (fixed PQC levels) to 68 ms, allowing total latency to be reduced down to 140 ms (compared with 160 ms in the Hybrid system). Recall degradation is limited to -9% in the presence of adversarial perturbation in the Proposed model, which is significantly better than both the Hybrid model (-18) and the baseline (-29), thanks to robustness increases [22]. Similarly, cipher-downgrade attacks are only successful in less than 1 percent of operations, which implies a strong crypto-policy choice. Recent study on graph-based fraud detection has achieved minor gains over tabular models with ROC-AUC values generally in the 0.93-0.95 range; our Hybrid and Proposed models achieve higher ROC-AUC values, 0.95 and 0.966, respectively with the additional effect of multimodal fusion and temporal messaging passing. Our adaptive scheme keeps PQC overhead to a more realistic 68 ms, which is much lower than fixed high-security PQC settings (around 120 ms) [23]. Previous work on post-quantum cryptography implementation generally found that it incurs huge latency penalties, often 2x to 4x the cost of classical TLS. In addition, the literature on cryptographic risk management does not often associate it with ML-based fraud scoring; our PQC implementation with RL shows that this association is capable of reducing cost, latency, and ensuring high security levels. The results point to the opportunity of the efficient and beneficial collaboration of fraud detection with the cryptographic protection of real-time digital payment systems [24]. Proposed architecture provides the balanced solution: improved detection, reduced PQC, and better adversarial resilience are all vital to deploy in the high-volume financial network. Nevertheless, there are still a number of limitations. To begin with, although the proposed latency profile is improved, the Proposed model continues to add 20 ms to the total end-to-end delay in comparison to the baseline system. Second, the RL-based PQC policy involves much simulation to ensure stable

convergence, which could be hard to implement in the context of fraud patterns changing rapidly [25]. Finally, although multimodal and graph-structured data are better performance-wise, they require more complicated data governance and infrastructure. Stability testing on real-world payment workloads and investigation of federated or privacy-aware training approaches should be the focus of future studies.

6. CONCLUSION

This study had tackled the new security issue in the contemporary digital payment ecosystem, which were the twofold threat of quantum era cryptography vulnerability and more sophisticated fraud patterns. Compared to traditional and fixed-parameter systems, the suggested AI-Adaptive Post-Quantum Cryptography framework outperformed the latter two when coupled with a heterogeneous temporal graph neural network. The model performed in an experiment showed better performance in terms of fraud-detection compared to a baseline AUPR of 0.35 in Hybrid system with 0.48 of the full proposed architecture. The reinforcement-learning PQC orchestrator at the cryptographic layer achieved an average 68 ms handshake latency, which is lower than fixed PQC configurations (90 -120 ms) but still lower than the allowable total end-to-end latency of 140 ms. Robustness measurements were also more robust in behavior with only -9% drop in recall during perturbation attacks and a low success rate of less than 1% in cipher-downgrade. The research paper provides several benefits to the community. Firstly, it offers a unified architecture for data-driven crypto-decision-making and NIST-standard PQC. Secondly, it shows how to use heterogeneous GNNs to detect fraud using relational, behavioral, device-level, and textual indicators. Lastly, it provides an example of crypto-agility in action in a high-volume financial setting, supported by real metrics for throughput, computational cost, and latency. All these contributions will promote the design of resistant to fraud and secure and quantum ready payment architectures. The proposed system, which includes privacy-aware learning and federated systems to facilitate cooperation between many institutions, has to be tested in real-world payment contexts in future research. We need to put in further work to make adversarial robustness better, expand the collection of quantum-safe algorithms that may be used with the adaptive PQC policy, and put PQC into hardware. These guidelines will assist in cementing an end-to-end safe payment architecture that are commendable in the post-quantum era.

REFERENCES:

- [1] O. Akinagbe, "Quantum-Resistant Federated Learning Protocol with Secure Aggregation for Cross-Border Fraud Detection," *Int. J. Comput. Appl. Technol. Res.*, vol. 10, pp. 364–370, 2021.
- [2] W. Li, C. Meese, M. Nejad, and H. Guo, "P-CFT: A privacy-preserving and crash fault tolerant consensus algorithm for permissioned blockchains," in *2021 4th International Conference on Hot Information-Centric Networking (HotICN)*, IEEE, 2021, pp. 26–31. Accessed: Dec. 03, 2025. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9680829/>
- [3] Y. Zhang, S. Jin, Y. Huang, and Q. Shao, "Multi-Channel Currency: A Secure Method Using Semi-Quantum Tokens," Feb. 25, 2025, *arXiv*: arXiv:2502.18378. doi: 10.48550/arXiv.2502.18378.
- [4] H. Wang, X. A. Wang, S. Xiao, and J. Liu, "Decentralized data outsourcing auditing protocol based on blockchain," *J. Ambient Intell. Humaniz. Comput.*, vol. 12, no. 2, pp. 2703–2714, Feb. 2021, doi: 10.1007/s12652-020-02432-x.
- [5] K.-H. Yeh, C. Su, J.-L. Hou, W. Chiu, and C.-M. Chen, "A robust mobile payment scheme with smart contract-based transaction repository," *IEEE Access*, vol. 6, pp. 59394–59404, 2018.
- [6] C. Singh, M. S. Rao, and Y. M. Mahaboobjohn, "Bonthu Kotaiah, and T. Rajasanthosh Kumar." Applied Machine Tool Data Condition to Predictive Smart Maintenance by Using Artificial Intelligence.," in *International Conference on Emerging Technologies in Computer Engineering*, pp. 584–596. Accessed: Dec. 03, 2025. [Online]. Available: <https://scholar.google.com/scholar?cluster=246637688784125973&hl=en&oi=scholar>
- [7] A. Joshi, P. Bhalgat, P. Chavan, T. Chaudhari, and S. Patil, "Guarding Against Quantum Threats: A Survey of Post-Quantum Cryptography Standardization, Techniques, and Current Implementations," in *Applications and Techniques in Information Security*, vol. 2306, V. S. Shankar Sriram, A. G. H., G. Li, and S. R. Pokhrel, Eds., in *Communications in Computer and Information Science*, vol. 2306.

- , Singapore: Springer Nature Singapore, 2025, pp. 33–46. doi: 10.1007/978-981-97-9743-1_3.
- [8] D. Moody *et al.*, “Status report on the first round of the nist post-quantum cryptography standardization process,” Technical report, National Institute of Standards and Technology, 2019. Accessed: Dec. 03, 2025. [Online]. Available: [https://nist.pqcrypto.org/foia/20240716/Re_%201st%20Round%20Report\(1\)_7_Redacted.pdf](https://nist.pqcrypto.org/foia/20240716/Re_%201st%20Round%20Report(1)_7_Redacted.pdf)
- [9] G. Alagic *et al.*, “Status report on the first round of the additional digital signature schemes for the nist post-quantum cryptography standardization process,” 2024, Accessed: Dec. 03, 2025. [Online]. Available: <https://csrc.nist.gov/external/nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8528.pdf>
- [10] S. B. N. Premakumari, G. Sundaram, M. Rivera, P. Wheeler, and R. E. P. Guzmán, “Reinforcement Q-Learning-Based Adaptive Encryption Model for Cyberthreat Mitigation in Wireless Sensor Networks,” *Sensors*, vol. 25, no. 7, p. 2056, Mar. 2025, doi: 10.3390/s25072056.
- [11] S. Motie and B. Raahemi, “Financial fraud detection using graph neural networks: A systematic review,” *Expert Syst. Appl.*, vol. 240, p. 122156, Apr. 2024, doi: 10.1016/j.eswa.2023.122156.
- [12] D. Cheng, X. Wang, Y. Zhang, and L. Zhang, “Graph Neural Network for Fraud Detection via Spatial-Temporal Attention,” *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 8, pp. 3800–3813, Aug. 2022, doi: 10.1109/TKDE.2020.3025588.
- [13] G. Alagic *et al.*, “Status report on the third round of the NIST post-quantum cryptography standardization process,” 2022.
- [14] S. B. N. Premakumari, G. Sundaram, M. Rivera, P. Wheeler, and R. E. P. Guzmán, “Reinforcement Q-Learning-Based Adaptive Encryption Model for Cyberthreat Mitigation in Wireless Sensor Networks,” *Sensors*, vol. 25, no. 7, p. 2056, Mar. 2025, doi: 10.3390/s25072056.
- [15] S. Motie and B. Raahemi, “Financial fraud detection using graph neural networks: A systematic review,” *Expert Syst. Appl.*, vol. 240, p. 122156, Apr. 2024, doi: 10.1016/j.eswa.2023.122156.
- [16] D. Cheng, X. Wang, Y. Zhang, and L. Zhang, “Graph Neural Network for Fraud Detection via Spatial-Temporal Attention,” *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 8, pp. 3800–3813, Aug. 2022, doi: 10.1109/TKDE.2020.3025588.
- [17] A. Al Mamun, A. Abrar, M. Rahman, M. S. Salek, and M. Chowdury, “Post-Quantum Cryptography for Intelligent Transportation Systems: An Implementation-Focused Review”, Accessed: Dec. 03, 2025. [Online]. Available: <https://www.techrxiv.org/doi/full/10.36227/techrxiv.176369792.29441055>
- [18] S. B. N. Premakumari, G. Sundaram, M. Rivera, P. Wheeler, and R. E. P. Guzmán, “Reinforcement Q-Learning-Based Adaptive Encryption Model for Cyberthreat Mitigation in Wireless Sensor Networks,” *Sensors*, vol. 25, no. 7, p. 2056, Mar. 2025, doi: 10.3390/s25072056.
- [19] S. Motie and B. Raahemi, “Financial fraud detection using graph neural networks: A systematic review,” *Expert Syst. Appl.*, vol. 240, p. 122156, Apr. 2024, doi: 10.1016/j.eswa.2023.122156.
- [20] Tulala, Rajasanthosh Kumar, K. Palaniradja, and V. Balasubramanian. "Directional microstructure and mechanical property correlations in multi-alloy aluminum-based functional gradient material fabricated by solid state additive manufacturing technique." *Materials Research Express* 12.11 (2025): 116502.
- [21] D. Zimmermann, “Adoption of Post-Quantum Cryptography in Organizations: Challenges and Drivers,” *Future*, vol. 36, no. 1, pp. 256–280.
- [22] A. Zgureanu and V. Andronatiev, “Quantum Cryptography and Post-Quantum Security: Current State, Challenges, and Strategic Directions,” 2025, Accessed: Dec. 03, 2025. [Online]. Available: <http://dspace.ase.md:8080/xmlui/handle/123456789/4681>
- [23] Q. Khan and S.-Y. Chang, “Post-Quantum Key Exchange and Subscriber Identity Encryption in 5G Using ML-KEM (Kyber),” *Information*, vol. 16, no. 7, p. 617, 2025.

- [24] A. Corsi, S. Gür, A. Brighente, and M. Conti, “Evaluation of Post-Quantum Key Encapsulation Methods in 5G Core Network,” in *2025 IEEE Wireless Communications and Networking Conference (WCNC)*, IEEE, 2025, pp. 1–6. Accessed: Dec. 03, 2025. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10978792/>
- [25] A. K. Bishwas and M. Sen, “Strategic Roadmap for Quantum- Resistant Security: A Framework for Preparing Industries for the Quantum Threat,” Nov. 15, 2024, *arXiv*: arXiv:2411.09995. doi: 10.48550/arXiv.2411.09995. 785–794. doi: 10.1145/2939672.2939785.