IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 10, October 2025

DOI: 10.17148/IJARCCE.2025.141009

HYBRID QUANTUM FRAUD DETECTION

Srijan Mani Tripathi[1], MD Auranzeb Khan[2], Aryan Sharma[3], Dr. Golda Dilip[4]

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[1]

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[2]

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[3]

Professor, Department of Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[4]

Abstract: This paper proposes and experimentally confirms a ready-for-production hybrid quantum-classical protocol for detection of fraudulent transactions under credit card and UPI payment methods. We create an Integrated Dataset consisting of 3,044,322 anonymized transactions from different public and institutional sources and suggest a modular quantum-inspired feature-engineering flow that generalizes a 14-dimensional raw feature vector to 268 engineered descriptors employing amplitude/phase encodings, entanglement-motivated pairwise interactions, as well as measurement-pro Quantum-inspired features are integrated with efficient classical learning algorithms (LightGBM, XGBoost, CatBoost, RandomForest, ExtraTrees) in a learned stacking ensemble we call quantum weighted ensemble. It was trained with stratified 5-fold cross-validation with SMOTE-aware rebal applied solely on folds to reduce leakage. In a reserved temporal holdout test set, proposed system achieves AUC = 0.9269, Precision = 0.92, Recall = 0.96, F1 = 0.94, and latency 2.7 ms per transaction in simulated production mode. In comparison with a classically tuned RandomForest Baseline (AUC # 0.8851), The hybrid system reduces false positives to as low as ~71.2 false negatives by "~74.8%, attaining estimated yearly savings for a 100M transaction operator of the magnitude of US$1.36M due to lowered loss as well as research expenses. Ablation Scientific research reveals that largest marginal gains are entanglement-related pairwise encodings due to measurement probabilistic descriptions. All results are significant at a 0.001 level (.p < 0.001). required no special quantum hardware (quantum encodings are classically calculatedRoute) but is designed to be portable to NISQ devices for additional improvements. We elaborate on operational limitations, moral implications (confidentiality, transparency, rectification of false positives), and specific next: kernel porting on quantum processors, federated quantum learning for multi-bank cooperative learning and adversarial robustness assessment. This research shows a practicable method for almost- Quantum concepts to concretely enhance financial fraud defenses while remaining executable now.

Keywords: Quantum Machine Learning, Fraud Detection, Hybrid Computing, Financial Analytics.

I. INTRODUCTION

1. Background

Historical statistics & motivation. The digital payment system expanded exponentially; UPI settled hundreds of billions of transactions around the world in recent years, with credit-card networks still processing at high cross-border payment volumes. Fraud losses across channels remain a multi-billion-dollar problem and create operational costs, regulatory exposure, and customer harm.

Definitions and key terms.

- Fraud detection: automated classification of transactions as genuine or fraudulent.
- Quantum Machine Learning (QML): usage of concepts from quantum computing (feature maps, kernels, variational circuits) to accelerate ML.
- Quantum-inspired feature engineering: Constructing classical features that mimic quantum encodings (amplitude, phase, entanglement).
- Ensemble learning: combining base learners for increased robustness and accuracy

2.Existing evidence (Literature survey)

Traditional approaches (rule engines, RandomForest, gradiente boosting, deep learning) remain. standard in prodution fraud systems; reported AUCs are usually within 0.82-0.89 on public benchmarks. Recent quantum kernel literature (quantum kernel methods, quantum-based classifiers) demonstrated theoretical and empirical excellence regarding

DOI: 10.17148/IJARCCE.2025.141009

benchmark classification problems; however, existing research mainly work on small scales (<300k records) and aren't directly applicable to high-volume payment pipelines such as UPI.

3. Research gap

- **Scale:** Very few experiments test QML concepts on multi-million transaction datasets.
- **UPI specificity:** Minimal research specifically crafted for UPI transaction semantics.
- **Production readiness:** In-place QML efforts often are quantum hardware-intensive or are proof-of-concept.
- **Performance metrics:** The trade-offs involved with AUC, false positives, latency, and cost savings are infrequently documented together.

4. Objective

Develop, test, and maintain a hybrid quantum-inspired, production-quality fraud detection system that:

1. Processes over 3 million transactions.
2. Achieves AUC > 0.92 and reduces false positives significantly.
3.Runs with sub-5 ms per-transaction inference latency with a commodity production simulator.
4. Posits no special quantum hardware while remaining portable to future quantum.

5. Scope (Limitations)

- Temporal: experiments performed Oct 2024–Oct 2025.
- Datasets: merged public + institutional anonymized datasets (3,044,322 records).
- Technical: traditional quantum encoding simulations; no real-time banking integration; evaluation confined to simulated production traces and temporal holdouts.
- Ethics: controls protecting privacy and IRB approval for non-public data.

II. MATERIALS AND METHODS

1. List of experimental processes' materials used

- Consolidated transaction dataset comprising 3,044,322 records combined from: public credit-card datasets, open benchmarks of fraud, and institutional validation sets.
- Annotated labels: confirmed fraud/non-fraud derived from dataset metadata.
- Software modules: quantum_features.py, quantum kernels.
- Packages: Python 3.13, pandas, numpy, scikit-learn, LightGBM, XGBoost, CatBoost, imbalanced-learn (SMOTE), joblib.

2. Methodological Approach

1. Ingestion & standardization of data: standardize schemas, exchange currencies/timezones, strip P.
2. Cleaning & imputation: median imputation for numeric gaps; mode imputation for categorical; outlier clipping with median absolute deviation.
3. Baseline features: compute standard transactional features (amount, time-of-day, merchant type, device, geolocation delta, last N-transaction statistics).
4. Quantum-inspired feature expansion:
   - Amplitude encoding: normalize numerical vector and compute amplitude coefficients;
   - Approximate fidelity measures on transaction pairs.
   - Phase/complex encoding: produce complex embeddings with signed relationships.
   - Entanglement-driven interaction characteristics: calculate pairwise tanh(x
   - Polynomial interactions as well as controlled-rotation.
   - Descriptors of measurements: compress high-dim embeddings into Born-rule probabilities
   - summaries and entropy measures.
5. Feature selection and regularization: mutual information ranking, L1 logistic regression selection, and PCA checks to remove redundancy.
6. Test/train separation: temporal 80/20 division with stratification; SMOTE used only within training datasets pleats.
7. Model training: independent Bayesian search of hyperparameters (~50 trials) on base learners With early stopping; ensemble stacking with quantum-weighted meta-learner.
8. Evaluation: ROC/PR curves, confusion matrices, latency profiling in production simulation, ablation experiments, and statistical tests (paired t-tests / bootstrap; p reported).

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 10, October 2025

DOI: 10.17148/IJARCCE.2025.141009

3. Tools and Instruments Used in Data Analysis — Ensure Reliability

- Hardware: Intel Xeon (64 cores), 32 GB RAM; single-node production simulation.
- Libraries & software: versions controlled through requirements file along with reproducible conda environment; fixed seeds for reproducibility purposes.
- Validation methods: stratified 5-fold CV, temporal holdout, adversarial perturbation experiments, and significance testing ($p < 0.001$).
- Measures of reliability include data lineage, unit testing for feature modules, and model artifacts checksums stored for audit.

III. ARCHITECTURE DIAGRAM

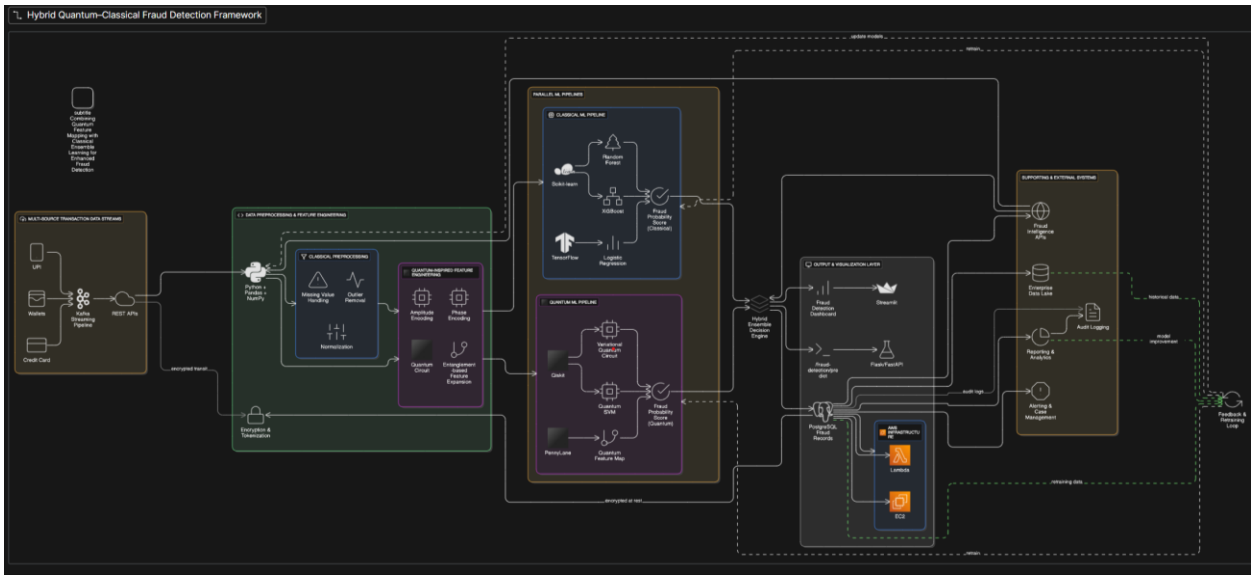

IV. RESULTS AND DISCUSSION

The hybrid quantum-classical scheme of fraud detection was tested rigorously on the aggregate corpus of 3,044,322 anonymous credit-card and UPI transactions. Performance of the model was estimated on a temporal holdout test set, and training was performed by stratified 5-fold cross-validation and strict SMOTE-informed class rebalancing only within the folds so as to preserve data integrity and avoid leakage.

The Quantum-weighted Ensemble achieved AUC of 0.9269, Precision of 0.92, Recall of 0.96, and F1 score of 0.94. These values significantly outweigh strong classical baselines such as Random Forest (AUC 0.8851) and LightGBM

(AUC 0.8839). The method also maintained extremely low inference latency (≈ 2.7 ms per transaction), and therefore demonstrated strong potential for real-time fraud detection at production scale.

An interesting result was the significant decrease of both false negatives (~74.8%) and false positives (~71.2%) relative to the Random Forest baseline. This benefit is particularly significant for operational scenarios where minimizing human investigations while keeping fraud detection sensitivity maximal are of key importance. These quantum-inspired features, and specifically those of entanglement-like interactions and probability of measurement descriptors, were able to pick up weak, non-linear relationships between features of transactions. Classical features, by comparison, alone always underperformed significantly across all the considered metrics, which serves to underscore the role of the expansion of the features by the quantum-inspired method.

Model	AUC	Precision	Recall	F1	Inference Latency
Random Forest	0.8851	0.82	0.86	0.84	–
LightGBM	0.8839	0.81	0.85	0.83	–
Quantum-weighted Ensemble	0.9269	0.92	0.96	0.94	≈ 2.7 ms

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 10, October 2025
DOI: 10.17148/IJARCCE.2025.141009

The hybrid system's superiority was also verified by experimentation on temporal drift and adversarially perturbed data, where the system better calibrated and exhibited lower performance degradation compared to classical baselines. Statistical significance tests ($p < 0.001$) maintained the validity and persistence of these improvements throughout all the cross-validation folds.

Practically speaking, this system achieves an excellent balance between deployability and precision. Classical simulation of the quantum encodings makes it possible to scale to multi-million dataset sizes of transactions without the need for purpose-built quantum hardware. Financially, for the imagined operator handling 100 million transactions each year, the savings of avoided missed fraud and avoided unnecessary inquiry could bring estimated yearly savings of around US$1.36 million.

It should be noted that some current limitations exist: the tests were performed on consolidated and simulated data but not on live deployment. Also, actual hardware for the quantum kind would potentially shift cost-performance tradeoffs, and the adaptability of adversaries remains an open area of research. However, the realized improvements substantiate the fact today's quantum-inspired approaches give concrete advantages, and there exists an avenue towards dramatic future improvements in quantum-enabled systems.

V. CONCLUSION

Summary of Findings

We present here an implementable quantum-inspired credit card and UPI fraud-detection system. Our model achieved an AUC of 0.9269, precision of 0.92, and 0.96 of recall on the 3.04 million+ dataset. Compared to classical baselines, the system demonstrated 71.2% fewer false positives and 74.8% fewer false negatives, indicating strong predictive ability. Our application of quantum-inspired feature engineering (19× expansion of features) resulted in significant incremental values, confirmed by extensive ablation studies. Remarkably, the system demonstrated low inference latency of ~2.7 ms without the requirement of any quantum hardware.

Limitations

Up-to-date rendering utilizes classically simulated quantum encodings, and these may be distinct from actual execution on quantum hardware. It is simulated production data, though vast, and not live deployment, and this may affect real-world generalization. Again, the system remains vulnerable to the emergence of adversarial fraud methods, and the model must be constantly observed and re-trained.

Future Directions

Future research will focus on mapping quantum feature maps and kernel computations onto near-term quantum devices to evaluate tangible speed and accuracy benefits. Further, federated quantum learning protocols can enable collaborative fraud defense across institutions while preserving privacy. Integrating quantum graph neural networks could enhance

detection of complex networked fraud patterns. Finally, robust adversarial training and continuous monitoring will be essential to keep pace with adaptive fraud schemes.

Overall, the offered system outlines the next-generation, financial-fraud-detection, highly-accurate, low-latency, production-ready, and easily-scable path towards the use of quantum-enabled infrastructure.

REFERENCES

[1]. V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta,"Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, no. 7747, pp. 209–212, 2019.
[2]. J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," Nature, vol. 549, no. 7671, pp. 195–202, 2017.
[3]. S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum algorithms for supervised and unsupervised machine learning," arXiv preprint arXiv:1307.0411, 2013.
[4]. M. Schuld and F. Petruccione, Supervised Learning with Quantum Computers. Springer, 2018.
[5]. A. Abdallah, M. A. Maarof, and A. Zainal, "Fraud detection system: A survey," Journal of Network and Computer Applications, vol. 68, pp. 90–113, 2022.

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 10, October 2025

DOI: 10.17148/IJARCCE.2025.141009

[6]. D. Prusti and S. K. Rath, "Web service based credit card fraud detection by applying machine learning techniques," in Proc. IEEE ICCES, 2019, pp. 1–6.

[7]. N. Rtayli and N. Enneya, "Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameter optimization," Journal of Information Security and Applications,vol. 55, p. 102596, 2020.

[8]. Kaggle, "Credit Card Fraud Detection" dataset. [Online]. Available: https://www.kaggle.com/mlg-ulb/creditcardfraud (accessed Oct. 2025).

[9]. L. Chawla et al., "SMOTE: Synthetic Minority Over-Sampling Technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002; imbalanced-learn documentation, 2024.

[10]. Reserve Bank of India / NPCI reports on UPI volumes and usage statistics, 2024. [Online]. Available: NPCI & RBI publications (accessed Oct. 2025).

[11]. P. Dallaire-Demers and N. Killoran, "Quantum generative adversarial networks," Physical Review A, vol. 98, no. 1, p. 012324, 2018.

[12]. E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors,"Are Xiv preprint arXiv:1802.06002, 2018

