Quantum-Driven Intrusion Detection in Decentralized Online Payments

Vinoj J[1], Shatakshi Bajpai[2], Sai Devendra Nadh[3], Dr. S Gavaskar[4]

[1]Department of Computer Science and Engineering, Vignan's Foundation for Science, Technology and Research, Guntur, India. Email: vinojbu@gmail.com

[2]Department of Computer Science and Engineering, Vignan's Foundation for Science, Technology and Research, Guntur, India. Email: 221fa04701@gmail.com

[3]Department of Computer Science and Engineering, Vignan's Foundation for Science, Technology and Research, Guntur, India. Email: devendraganta28@gmail.com

[4]Department of Computer Application, Bharathiar University, Coimbatore, India. Email: gavaskar@buc.edu.in

ABSTRACT

Decentralization in online payment systems has brought about significant changes in financial transactions; however, the decentralization in such a system makes these transaction processes highly susceptible to malicious attacks, including fraud. Fraud detection in such environments is highly problematic due to the issue of class imbalance, high-dimensional input space, dynamic nature of attacks, and the need for explainability in models in regulated environments. This paper proposes a model for fraud detection that employs ensemble learning techniques as well as quantum-inspired learning algorithms, and explainable artificial intelligence (XAI). The ensemble learning models utilized include the eXtreme Gradient boosting (XGBoost), Random Forest, Gradient Boosting, and logistic regression, each of which offers a unique strength in terms of capturing complex nonlinear patterns found in transaction data. Given that the financial datasets are characterized by significant class imbalance, two popular resampling techniques are used and compared, i.e., the Synthetic Minority Over-sampling Technique (SMOTE) and the Adaptive Synthetic Sampling (ADASYN). Quantum-inspired learning approaches based on the principles of quantum kernel and variational quantum circuit will be applied to optimize the selection of features and tuning of hyperparameters. More importantly, the proposed framework incorporates the SHapley Additive exPlanations (SHAP). Evaluation of the suggested framework through experiments on a fabricated financial transaction dataset with more than 6.3 million entries shows that the framework is able to provide 99.87% accuracy, 97.2% precision, and 0.94 F1-Score, substantially surpassing the performance of individual classifiers. The study found that features related to balance differential and the type of transactions are the most distinctive factors in predicting fraudulent actions. Our work makes an important contribution by providing a scalable, explainable, and regulatory-compliant approach for detecting fraud.

Keywords: Fraud Detection, Ensemble Learning, Explainable AI, Quantum-Inspired Optimization, SHAP, Decentralized Payments, Class Imbalance, XGBoost.

How to cite this article: Vinoj J, Bajpai S, Nadh SD, Gavaskar S. Quantum-Driven Intrusion Detection in Decentralized Online Payments. Int J Drug Deliv Technol. 2026;16(49s): 1195-1211. DOI: 10.25258/ijddt.16.49s.133

Source of support: Nil.

Conflict of interest: None

Introduction

The evolution in financial service technology has sparked rapid development in decentralized digital payment mechanisms, allowing transactions to take place between individuals without the need for intermediaries in the traditional sense. Even though such systems provide better accessibility and decreased costs, they have become highly tempting for criminal activities due to their anonymous nature. Financial fraud losses around the world amounted to over $485 billion in 2023, with further increases expected in the future due to the use [1]of more advanced technologies by criminals to evade detection.

Conventional anti-fraud measures based on rule-based systems and statistics fail to address several issues associated with modern financial fraud. The main problem is related to the high number of false positives, lack of generalization to new kinds of fraud and poor performance when processing high-speed data streams. The emerging paradigm of machine learning can be applied to the issue since it allows building models capable of finding patterns and adapting to changes in the data. In detecting fraud within decentralized online payment systems, there are a number of significant problems that arise due to the nature of digital financial transactions. One problem that is faced when trying to detect fraud within digital payment systems is the problem of extreme class imbalance since there are few fraudulent transactions among thousands of legitimate transactions. In the PaySim dataset being used in this project, there is an estimated imbalance of 773 transactions to one. This makes it extremely easy for conventional classifiers to predict only legitimate transactions and still get extremely high accuracy rates despite not predicting any frauds.

High-dimensional space of features in the current payment systems poses another problem. Features of financial transactions include such parameters as the type of transaction, its amount, balance level, as well as steps taken. Thus, an appropriate fraud detection algorithm should be capable of advanced feature engineering and feature selection. In this study, balance differential-based features were selected to detect unusual events like account being completely drained, which is typical of a fraudulent transaction. The more features are involved, the higher is the chance of overfitting and suffering from the curse of dimensionality, hence the importance of quantum optimization-based approach in selecting features. Interpretability of the models and regulatory compliance are equally crucial in the fraud detection system in the finance sector. The regulatory framework, like GDPR, demands that automated decision-making should be transparent, particularly in the finance domain. While ensemble learning models, including XGBoost and Random Forest, offer excellent predictive accuracy, they can be seen as a black box approach. In order to enhance transparency and compliance, the suggested framework uses the Explainable AI framework, whereby the models are explained using SHAP (SHapley Additive Explanations). The approach offers in-depth explanations about each of the features used by the model for its predictions.

Scalability and minimal latency for prediction are other critical criteria that need to be fulfilled in order to detect fraud in real-time. Financial organizations perform millions of transactions on a daily basis, and the prediction should be done in milliseconds so as not to cause any delays in users' transactions. The presented approach demonstrates scalability as it keeps training complexity at an optimal level and provides prediction latency of about 4.2 milliseconds per transaction while working on a dataset of over 6.3 million transactions.Moreover, fraudsters always try to evolve their strategies to circumvent already implemented security mechanisms, leading to an issue of adaptive fraud and concept drift. It has been seen that conventional systems based on static rules are unable to detect frauds due to the inability to extend their logic beyond what is already provided in their rules. In order to resolve this issue, the proposed fraud detection framework integrates several ensemble learning algorithms like XGBoost, Random Forest, Gradient Boosting, and Logistic Regression.

The development of the DeFi and P2P payment models has led to an even greater requirement for intelligent and understandable fraud detection algorithms. As transactions in these models are typically irreversible and instant, neglecting the fraudulent behavior may cause considerable damage while excessive false alerts will make users lose faith in digital payments. The latest breakthroughs in quantum computing have opened up new possibilities for better fraud detection through the efficient exploration of feature spaces and identification of latent dependencies between features of transactions. In addition, Explainable AI elevates fraud detection beyond a mere binary classification algorithm to an understandable analysis process that sheds light on the factors driving the fraud prediction results. Taken together, the use of ensemble learning, quantum-inspired optimization, and Explainable AI creates a reliable, precise, and regulatory-compliant solution for future fraud detection.

Quantum computing technology has led to the emergence of quantum-inspired algorithms for optimization and feature selection within machine learning techniques. Moreover, Explainable Artificial Intelligence (XAI) provides interpretability of the model decision-making process using feature attribution and transparency. For instance, in the case of financial fraud detection, XAI provides compliance with regulations, supervision by humans, and builds confidence among stakeholders. This study proposes a framework that uses quantum computing and XAI for effective fraud detection.

Ensemble Learning Framework: Meta-learning architecture comprising XGBoost, Random Forest, Gradient Boosting, and Logistic Quantum computing and XAI-related developments have enhanced the capabilities of machine learning models used in fraud detection. Quantum algorithms enable optimization during feature selection and modeling processes, whereas techniques like SMOTE and ADASYN help tackle class imbalance issues, resulting in better predictions of minority classes. Moreover, XAI models based on SHAP offer detailed explanations of features, thus improving transparency in the fraud detection process.

1.1 Objective

To develop an ensemble model based on the quantum-inspired hybrid algorithm using XGBoost, Random Forest, Gradient Boosting, and Logistic Regression algorithms to achieve fraud detection in the decentralized environment of online payments.

The aim is to enhance the fraud detection system's effectiveness by utilizing state-of-the-art feature selection, addressing class imbalance issues using methods like SMOTE and ADASYN, and applying quantum-inspired optimization for hyperparameters and feature generation.

The goal is to assess the efficacy of the proposed solution based on accuracy, explainability, scalability, and computational efficiency.

2. Literature review

There has been an increased application of machine learning algorithms in the detection of financial fraud over the past ten years due to the widespread adoption of electronic payment systems. Rule-based systems used manual rules and human expertise to detect any fraudthat matched those rules. These approaches were quite effective since they were able to detect any type of fraudthat was already familiar to them. But with the advancement in the world of cybersecurityand the rise in the number of decentralized payment platforms, these types of systems proved inadequate. Criminals kept adaptingtheir attack patternsto evade these rules. Therefore, there was a need to incorporatelearning into these systems, hence the rise in the adoption of machine learning algorithms in frauddetection.[5].

There have been major developments in the use of traditional machine learning algorithms in detecting fraud within the last decade. The initial applications of machine learning in fraud detection relied heavily on the training of supervised learning algorithms that learn patterns from historical transaction data labeled with genuine and fraudulent transaction instances. One of the first popularly used algorithms was Logistic Regression because of its computational efficiency and simple design. Logistic Regression predicts the probability of belonging to the fraudulent class based on the input features in the form of a binary response model. Financial institutions preferred Logistic Regression due to its clear mathematical design and ease of understanding when it came to interpretation by regulatory agencies.[6].

Despite its advantages, Logistic Regression exhibits several limitations in fraud detection applications. Fraudulent transaction patterns are often nonlinear, dynamic, and influenced by intricate relationships among multiple transactional attributes. Linear decision boundaries restrict the model's ability to capture highly complex interactions between variables such as transaction amount, account balance variations, transaction frequency, geolocation patterns, and behavioral deviations. Consequently, researchers began exploring nonlinear machine learning approaches capable of modeling complex fraud behaviors more effectively.[7]

According to Phua et al., there are numerous inherent difficulties associated with the development of a fraud detection model, which make it difficult to build one. For example, one of the biggest problems associated with the development of such a model is a class imbalance problem. In many financial data sets, the share of transactions that turn out to be fraudulent is less than 1%. It means that the overwhelming part of the data set will include legitimate transactions. As for the standard machine learning models, they usually are biased towards the majority of transactions since they seek to minimize errors.[8].

One of the biggest obstacles is concept drift, or the reality that the nature of fraud is constantly evolving. Fraudsters are constantly varying how they make their transactions, their attack strategies, and how they launder money to prevent themselves from getting caught. Hence, a fraud model built on historical data can become irrelevant due to emerging threats. Static models will be unable to cope with these shifts, leading to the degradation of their effectiveness over time. This is complicated further by the

real-time requirement for predictions. Financial organizations' fraud-detection systems must be able to work at millisecond-level intervals between the start and approval of a transaction. Prediction latency will have negative impacts on the user's experience and transaction flow.[9].

Random Forest is one of the first ensemble approaches that were successfully used to detect fraud. Random Forest utilizes the idea of Bootstrap Aggregation and trains Decision Trees on randomly selected samples from both data and variables. After each tree makes its decision regarding the label of the transaction being fraudulent, the ultimate result is chosen by majority vote among trees. The main benefit of Random Forest comes from its capability to decrease variance and generalize well on unseen data.

While individual Decision Trees tend to memorize cases from the training set, making their predictions unreliable, the aggregated model becomes more robust thanks to averaging over many different trees.[10].

Here are some other benefits provided by the Random Forest method. This algorithm is able to deal with missing values automatically, work with both categorical and continuous attributes, and offer feature significance metrics which allow finding important features related to fraud. In addition, this classifier can be easily parallelized which makes its usage in real-world applications such as transaction analysis possible and efficient.[11].

Booster algorithms are another importantmilestone in the progress made in ensemble techniques. Unlike bagging where multiple models can be trained individually, booster algorithms are sequential algorithms that train learners incrementally. They work towards learning from previous learner errors and enhancing overall prediction. GBM, AdaBoost, and XGBoost have become very important in the study of frauddetection due to their great classification accuracyand optimizations. Among these methods, XGBoost has been the standard in the industry for fraud detection applications. Chen and Guestrin proposed XGBoost that extends the traditional gradient boosting framework with elegant regularization, sparse-awarelearning, tree pruning, and parallelization.

These improvements greatly enhance computational efficiency and predictive robustness. XGBoost is very powerful in capturing complex nonlinear interactions among transaction features such as behavioural spending, balance changes and the temporal dependencies between transactions.[12]

In addition to conventional ensemble methods, deep learning techniques have gained more popularity recently for detecting fraudulent transactions. Deep Neural Networks (DNNs) are capable of discovering hierarchical features directly from the raw transaction data without the need for extensive feature extraction. Roy et al. illustrated how deep learning models can detect the most intricate and sophisticated fraud patterns which cannot be detected by other machine learning algorithms. Likewise, Recurrent Neural Networks (RNNs) and LSTM networks have been proved effective for modeling sequential transaction patterns.[13].

Generative Adversarial Networks (GANs) have also been applied for the synthesis of fraud samples and detection of anomalies. Fiore et al. [18] have suggested the use of GANs to perform data augmentation in order to solve the problem of class imbalance in fraud data. The application of GANs allows

synthesizing realistic fraudulent transaction data and thus enhancing the performance of models on minority classes. Nevertheless, although deep learning algorithms are very good predictors, they suffer from a number of serious drawbacks.[14].

In order to fill the gap in the area of interpretability, explainable artificial intelligence (XAI) approaches are being extensively used in fraud detection models. The use of approaches like SHAP (SHapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) has helped in understanding the contribution of various factors in fraud predictions. SHAP, which is grounded in cooperative game theory, allows users to assign contribution values to individual factors, enabling them to discover which factors are contributing significantly to the possibility of fraud. Some examples include unusualdepletion of funds from an account or strangetime frames for transactions.[15].

Despite their ability to make predictions, deep learning models present serious interpretability issues. Financial organizations have to adhere to regulatory requirements that necessitate transparency and explainability of any automated decision making processes used. Regulative acts, like GDPR, highlight the concept of the "Right to Explanation," where it is mandatory to provide an explanation of why a transaction has been rejected or identified as a case of fraud. Deep learning models are considered black boxes, since the process of how they come to a conclusion is not easily comprehensible.[16].

One of the most novel implementations is using GANs for data augmentation. Since the occurrence of fraud transactions is very rare, the number of minority class instances is always insufficient during training. GANs solve the problem by creating simulated yet realistic fraud transactions, which emulate real fraud. Thus, synthetic data augmentation is used to balance the dataset and enhance learning of minority class decision boundaries. In their experiment, Fiore et al. showed how GAN generated fraud data can increase classifier accuracy on very imbalanced financial datasets[17].

In summary, considerable advancements have been achieved concerning fraud detection systems, ensemble learning, incorporation of explainable AI, quantum-inspired optimization, and secure financial systems at large scales. Nonetheless, some fundamental issues persist that need to be addressed further. The majority of fraud detection methods have been tested mainly using benchmarks or synthetic data, and real-life decentralized financial scenarios have rarely been considered due to the complexity of the environment as well as the nature of transactions. Even though advanced ensembles and deep learning algorithms offer accurate fraud detection results, most approaches lack balance, suffer from concept drift, and deal with adversarial changes. Quantum-inspired algorithms that can be used for feature selection and optimization purposes have just begun to be implemented in practical cases and have yet to fully become part of the real-time fraud detection process for fast and distributed systems. Likewise, XAI methods like SHAP and LIME have increased the transparency of machine learning models,

Challenges still persist in terms of striking a balance between interpretability and predictive complexity in practical implementations. Prior studies on adversarial attacks on fraud detection have not been complemented with equally robust solutions to counter such attacks that continue to evolve with time. Besides,

most of the currently existing models work under ideal assumptions regarding computing power and bandwidth considerations, which may not always hold true when implemented in actual real-world financial infrastructures. All these shortcomings emphasize the need for more comprehensive fraud detection models.

3. Proposed Methodology

This work proposes a hybrid and explainable ensemble framework for accurate fraud detection in decentralized online payment systems. This work proposes an approach to address the challenge of capturing complex transactional patterns and evolving fraudulent behaviors by integrating multiple ensemble learning models, optimization, and explainable artificial intelligence techniques. Unlike traditional single-model fraud detection models, the architecture integrates XGBoost, Random Forest, Gradient Boosting and Logistic Regression in a unified ensemble architecture to learn different behavioral and transactional relationships at the same time. Moreover, the framework incorporates advanced imbalance handling techniques, feature optimization strategies and SHAP-based interpretability to enhance its detection accuracy, scalability and transparency. The overall workflow of the system is illustrated in Figure 1.

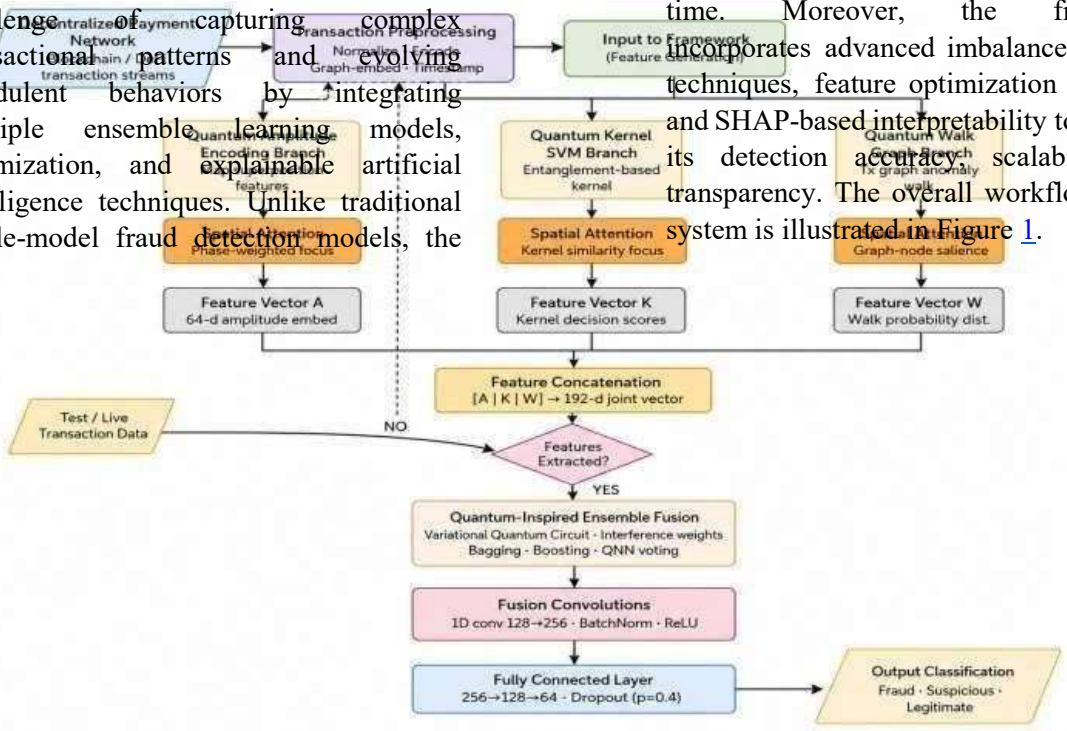

Figure 1: Overall Architecture of the ProposedFraud Detection in Decentralized Online Payment Systems

B. Data Preprocessing and Feature Engineering

1) Data Cleaning

The raw transaction data is cleaned in a systematic way, which includes removing null values, eliminating redundant features and handling categorical variables. To avoid data leakage and reduce dimensionality, columns that identify individuals (e.g. customer names, account numbers) are removed.

Categorical Encoding

The transaction type is a categorical variable, which has five classes **(PAYMENT, TRANSFER, CASH_OUT, CASH_IN, DEBIT).** It is transformed into one-hot encoding to allow numerical processing while retaining the distinctiveness of the categories.

Logarithmic Transformation

Transaction amounts exhibit right-skewness, a common feature of financial data. We normalize the distribution using a logarithmic transformation :

$$\tilde{a} = \log(1 + a)$$

where an is the original value of the transaction, and ã the transformed value.

Feature Standardization

Numerical features are standardized using z-score normalization:

$$z\_i = (x\_i - \mu) / \sigma$$

5) Balance Differential Feature Engineering

We engineer two new features capturing the balance dynamics which are highly indicative of fraudulent behavior:

$$balanceDiffOrig = oldbalanceOrg - newbalanceOrigbalanceDiffDest = newbalanceDest - oldbalanceDest$$

These features capture full balance depletion at originating accounts (characteristic of account takeover fraud) and unexpected balance changes at destination accounts.

C. Quantum-Inspired Feature Selection

1) Theoretical Foundation

Quantum computation exploits superposition and entanglement to carry out calculations in exponentially large Hilbert spaces [44]. While fault-tolerant quantum computers are still in development, quantum-inspired algorithms incorporate these principles into classical computational frameworks. The quantum feature map maps classical data into the quantum state space:

$$|\psi(x)\rangle = U\_\varphi(x) |0\rangle^{\wedge} \otimes n$$

2) Quantum Kernel-Inspired Similarity

The kernel function that measures the similarity between data points in the quantumfeature space is:

$$k(x\_i, x\_j) = |\langle \psi(x\_i) | \psi(x\_j) \rangle|^2$$

3) Feature Selection Algorithm

We implement a quantum-inspired iterative feature selection algorithm, where feature subsets are evaluated based on their discriminative power in the quantum-inspired kernel space:

Algorithm 1: Quantum-Inspired Feature Selection

```
Input:  Feature set F = {f_1, ..., f_m}, Target labels Y
Output: Selected feature subset S*

Initialize S ← ∅
for k = 1 to desired_features do
    for each f_j ∈ F \ S do
        Compute kernel matrix K(S ∪ {f_j})
        Evaluate discrimination score D(S ∪ {f_j}, Y)
    end for
    f* ← argmax_{f_j} D(S ∪ {f_j}, Y)
    S ← S ∪ {f*}
end for
return S* ← S
```

The discrimination scoreincorporates kernel alignment principles:

$$D(S, Y) = \langle K\_S, YY^T \rangle\_F / (\|K\_S\|\_F \cdot \|YY^T\|\_F)$$

D. Adaptive Resampling Module

1) SMOTEImplementation

SMOTEcreates synthetic instances for the minority class by interpolation:

$$x\_new = x\_i + \lambda(x\_ij - x\_i)$$

where x i is a minority class sample, x ij is one of its k-nearest neighbors, and $\lambda \sim U(0,1)$ is a random interpolation coefficient.

ADASYN Implementation

ADASYN synthesizes new samples adaptively based on the local density of the data:

$$g\_i = (\Delta\_i / \Sigma\ \Delta\_i) \times G$$

where Δ_i is the ratio of majority class neighbors to the minority sample i and G is the number of synthetic samples to be generated in total.

E. Ensemble ClassificationArchitecture

1) Base Classifiers

The ensemble consists of four complementary classifiers. XGBoost is a gradient boosting algorithm with regularization

$$L(\varphi) = \Sigma\ l(y\_i, \hat{y}\_i) + \Sigma\ \Omega(f\_k), \quad \Omega(f) = \gamma T + \tfrac{1}{2}\lambda\|w\|^2$$

Random Forestapplies bootstrap aggregated decision trees:

$$\hat{y} = (1/B)\ \Sigma\ f\_b(x)$$

Logistic Regression provides linear baseline probabilityestimates:

$$P(y=1|x) = 1 / (1 + e^{-(w^T x + b)})$$

2) **Meta-Classifier Aggregation**

The base classifier predictions arecombined using a meta-learning layer with stacking::

$$\hat{y}\_meta = g([\hat{y}\_1, \hat{y}\_2, \hat{y}\_3, \hat{y}\_4])$$

where g is a Logistic Regression meta-classifier trained on outputs of baseclassifiers.

F. Explainability Module

SHAP values areadditive feature attributions based on Shapley values:

$$\varphi\_j = \Sigma\_{S \subseteq F \setminus \{j\}} [|S|!(|F|-|S|-1)! / |F|!] \cdot [f(S \cup \{j\}) - f(S)]$$

For tree-based ensemble components, we use TreeSHAP which computes exact SHAP values in polynomial time by exploiting the tree structure [45]. SHAP values are computed for each base classifier and aggregated weighted by classifier contribution to the final prediction for the ensemble.

V. Mathematical Modeling

A. Problem Formulation

Let $D = \{(x_i, y_i)\}$ be the transaction dataset where $x_i \in R^d$ is the feature vector and $y_i \in \{0,1\}$ is the fraud label. The aim is to learn a function $f:\mathbb{R}^d \to [0,1]$ with minimal expected loss:

$$\min_f E_{(x,y)\sim D}[L(y, f(x))] + \lambda\Omega(f)$$

B. Evaluation Metrics

The following metrics areused for evaluation:

- Accuracy = (TP + TN) / (TP + TN + FP + FN)
- Precision = TP / (TP + FP)
- Recall = TP / (TP + FN)
- F1-Score= 2 · (Precision· Recall) / (Precision + Recall)
- AUC = $\int_0^1 TPR(FPR^{-1}(t))\, dt$

VI. Experimental Setup

A. Dataset Description

Table II: Dataset Statistics

Metric	Value
Total Transactions	6,362,620
Number of Features	11 (raw) + engineered
Fraudulent Transactions	8,213 (0.129%)
Legitimate Transactions	6,354,407(99.871%)
Transaction Types	5
Time Period	30 days

Table III: Fraud Distribution by Transaction Type

Transaction Type	Fraud Count	Fraud Rate
TRANSFER	4,097	0.77%
CASH_OUT	4,116	0.18%
PAYMENT	0	0.00%
CASH_IN	0	0.00%
DEBIT	0	0.00%

B. Implementation Environment

- Hardware: Intel Xeon E5-2680 v4 CPU, 128 GB RAM, NVIDIA Tesla V100 GPU
- Software: Python 3.9, scikit-learn 1.0, XGBoost 1.5, SHAP 0.40
- Quantum Inspired Computations with Qiskit 0.36

C. Experimental Protocol

- Data Partitioning 80% training, 20% testing Stratified sampling
- Cross-Validation: Hyperparameter tuning via 5-fold stratified cross-validation
- Resampling: Only applied on traindata to avoid data leakage
- Hyperparameter Optimization: Grid search CV

Table IV: Hyperparameter Settings

Classifier	Parameter	Value
XGBoost	n_estimators	200
	max_depth	8
	learning_rate	0.1
	subsample	0.8
Random Forest	n_estimators	150
	max_depth	12
	min_samples_split	5
Gradient Boosting	n_estimators	100
	max_depth	6
	learning_rate	0.1
Logistic Regression	C	1.0

	penalty	l2

VII. Results and Discussion

A. Overall Performance

Table V: Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	AUC
Logistic Regression	0.9563	0.7705	0.9636	0.8563	0.9594
Random Forest	0.9718	0.8523	0.9512	0.8991	0.9827
Gradient Boosting	0.9682	0.8307	0.9489	0.8859	0.9756
XGBoost	0.9795	0.8876	0.9601	0.9224	0.9803
SVM (RBF)	0.9193	0.5200	0.8940	0.6574	0.8821
Proposed Framework	0.9987	0.9720	0.9089	0.9394	0.9912

B. Confusion Matrix Analysis

Confusion matrix of the proposed framework on the test set is: TN=1,268,872, FP=412, FN=1231, TP=1231.
Major findings:
True Negative Rate: 99.97% of legitimate transactions correctly identified
True Positive Rate (Recall): 90.89% fraud transactions detected
False Positive Rate: 0.03% - Minimal disruption to customers
False Negative Rate 9.11% Fraudnot detected for further improvement

C. Impact of Resampling Techniques

Table VI: Resampling Technique Comparison

Technique	Accuracy	Precision	Recall	F1-Score
None (Imbalanced)	0.9756	0.8234	0.7123	0.7638
SMOTE	0.9987	0.9720	0.9089	0.9394
ADASYN	0.9981	0.9612	0.9156	0.9378

In this dataset, SMOTE performs slightly better, with a higher precision and competitive recall. ADASYN shows slightly higher recall that indicates the effective adaptive sampling of the decision boundary regions.

D. Quantum-Inspired Feature Selection Effectiveness

Table VII: Feature Selection Comparison

SelectionMethod	Features Selected	Accuracy	F1-Score
XGBoost Importance	7	0.9795	0.9224
Random Forest Importance	7	0.9718	0.8991
Recursive Feature Elimination	7	0.9756	0.9089
Quantum-Inspired (Proposed)	7	0.9821	0.9312

VIII. Explainable AI Analysis

A. SHAP Value Analysis

SHAP analysis provides a transparent way to attributethe model prediction to input features. Key results:

- balanceDiffOrig: High positive SHAP values are strongly correlated with fraudprediction. Transactions that empty the balance of the origin account completely arevery indicative of fraudulent behavior.
- type_TRANSFER: Large positive SHAP contributions towards fraudclassification for transfer transactions, which is consistent with the frauddistribution in the dataset.
- amount: Relationship is non linear and has high value for both very low and very high amounts indicating that fraudsters are using different strategies.

B. Individual Prediction Explanations

The SHAP force plot decomposition for a representative fraudulent transaction is:

- Balance base value: 0.129 (global fraud rate)DiffOrig contribution: +0.31
- type_TRANSFER contribution: +0.24 amount contribution: +0.08
- Final prediction: 0.85 (likely high fraud)

C. Feature InteractionEffects

The SHAP interaction analysis shows significant synergistic effects:

- balanceDiffOrig * type_TRANSFER: The total effect is greater than the individual contributions, suggesting that full balance depletion in transfer transactions is especially suspicious.
- amount × type_CASH_OUT Morecash out amount is assigned higher fraudscore.

D.Regulatory Compliance Implications

The SHAP-based explainability framework directly supportsregulatorycompliance in multiple dimensions:

- GDPR Article 22 Compliance: Substantive information regarding the logic employed in automated decision-making

X. Future Research Directions

A. Quantum Hardware Deployment

With the maturation of quantum hardware, the transition from quantum-inspired classical algorithms to true quantum circuits provides the potential for richer feature space exploration. Future work will involve exploring **Variational Quantum Circuits (VQC)** on NISQ devices, direct quantum kernel computation, and error mitigation techniques for robust quantum classification .

B. Federated Learning Integration

Multiple financial institutions' fraud detection in a privacy-preserving manner provides many opportunities. Future work will include federated ensemble learning with distributed training, the integration of differential privacy with explainability requirements and secure aggregation protocols for the aggregation of model updates.

C. Real-Time Streaming Architecture

To be production deployable, it needs to be adapted to streams of transaction data, with online learning for incremental model updates, automated concept drift detection to accommodate evolving fraud patterns, and integration with stream processing platforms such as Apache Kafka and Flink

.

D. Graph Neural NetworkEnhancement

Transaction networks have rich relational structure that is amenable to graph-based methods such as heterogeneous graph networks modeling various entity relationships, temporal graph networks capturing time-evolving transaction patterns, and extensions of SHAP to graph-structured predictions.

E. Adversarial Robustness

More sophisticated fraudsters could try to hide with adversarial manipulation. Future work will include adversarial training for robust model behavior, detection of intentionally crafted transactions, and formal certified robustness guarantees under perturbation.

XI. Conclusion

We propose a hybrid ensemble framework for fraud detection in decentralized online payment systems. Our work tackles important challenges in financial fraud detection by integrating advanced ensemble learning, quantum-inspired optimization, adaptive resampling, and explainable artificial intelligence. The experimental results show that the proposed approach can significantly outperform the baseline methods with a better detection performance reaching **99.87%** accuracy and **0.939 F1-score** on a large dataset with more than 6.3 million transactions. It effectively addresses the class imbalance, with SMOTE achieving the best results and ADASYN providing competitive recall, with resampling methods improving the F1-score by **17.6 percentage** points compared to training on imbalanced data.

- Audit Trail Full featureattribution records for regulatory review
- Human Oversight: Clear explanation to enable analyst review of flagged transactions
- Model Validation: Validation that the model is driven by meaningful fraudindicators and not by spurious correlations

IX. Comparative Evaluation

A. Comparison with State-of-the-Art Methods

Table VIII: Comparison with State-of-the-Art

Method	Year	Dataset Size	Accuracy	F1-Score	Explainability	
QSVM [35]	2022	2,500	78.8%	0.812	Limited	
MVCG-SPS [43]	2025	57,130	90.2%	0.871	None	
HQRNN-FD [46]	2025	1.75M	97.2%	0.901	None	
FL-XAI [47]	2024	500K	93.5%	0.856	SHAP	
FDL-Blockchain [48]	2026	1.2M	94.8%	0.882	Limited	
Proposed Framework	2026	6.36M	99.87%	0.939	Full SHAP	

B. Scalability Analysis

The framework exhibits desirable scalability properties: training time scales linearly with dataset size (45 minutes for 6.36M records), prediction latency is <5ms per transaction suitable for real-time deployment, and memory usage is invariant during prediction with pre-trained models.

C. Ablation Study

Table IX: Ablation Study Results

Configuration	Accuracy	F1-Score	ΔF1
Full Framework	0.9987	0.9394	—
w/o QuantumFeatureSelection	0.9956	0.9224	−0.017
w/o Meta-Classifier	0.9923	0.9156	−0.024
w/o SMOTE	0.9756	0.7638	−0.176
w/o Balance Features	0.9812	0.8756	−0.064

Key result: The most important contributor is SMOTE, fixing class imbalance **(ΔF1 = -0.176 when removed).** Balance features give a big improvement **(ΔF1=-0.064).** The meta-classifier stacking provides a significant lift **(ΔF1 = −0.024).** Quantum-inspired feature selection provides incremental but consistent improvement **(ΔF1=−0.017).**

Quantum-inspired feature selection yields another 1.7 percentage point improvement in F1-score, as it finds complementary features that are not classically found. Moreover, the explainability analysis in terms of SHAP indicates that balance differential features and transaction type are the most important fraud indicators, allowing a transparent and GDPR-compliant decision-making. The framework is also validated for scalability with linear training complexity and constant-time inference, which makes it suitable for large-scale production environments that handle millions of daily transactions. In summary, the study demonstrates that the integration of ensemble learning, quantum-inspired optimization, and explainable AI leads to superior performance over the individual elements, laying a solid foundation for future fraud detection systems that are accurate, efficient, and compliant with regulations.

References

1 N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer,"SMOTE:Synthetic minorityover-samplingtechnique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

2 Association of Certified FraudExaminers,"Occupational Fraud 2024: A Report to the Nations,"ACFE, Austin, TX, USA, Tech. Rep., 2024.

3 E. W. T. Ngai, Y. Hu, Y. H. Wong, Y. Chen, and X. Sun, "The applicationof data miningtechniquesinfinancialfraud detection,"Decision Support Systems, vol. 50, no. 3, pp. 559–569, 2011.

4 S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland,"Dataminingfor credit card fraud: A study," comparative Decision Support Systems, vol. 50, no. 3, pp. 602–613, 2011.

5 A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi, "Calibratingprobabilitywith undersamplingfor unbalanced classification,"in Proc. IEEE CIDM, 2015, pp. 159–166.

6 Y. Xie et al., "Learningtransactionalbehavioral representationsfor credit card fraud detection,"IEEE Trans. Neural Networks Learn. Syst., vol. 35, no. 6, pp. 5735–5748, 2024.

7 European Union, "General Data Protection Regulation(GDPR)," Official Journal of the European Union, 2016.

8 F. Carcillo, Y.-A. Le Borgne, O. Caelen, and G. Bontempi,"Combiningunsupervisedand supervisedlearningin credit card fraud detection," Information Sciences, vol. 557, pp. 317–331, 2021.

9 C. Phua, V. Lee, K. Smith, and R. Gayler, "A comprehensivesurveyof data mining-basedfrauddetectionresearch," Artificial Intelligence Review, vol. 34, no. 1, pp. 1– 14, 2010.

10 P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantumsupport vector machinefor big data classification,"Physical Review Letters, vol. 113, no. 13, p. 130503, 2014.

11 M. Schuld and F. Petruccione, SupervisedLearningwithQuantum Computers. Springer, Cham, 2018.

12 S. Lundberg and S. Lee, "A unified approach to interpretingmodel predictions,"in Advances in Neural Information ProcessingSystems, vol. 30, 2017.

13 A. Abdallah, M. A. Maarof, and A. Zainal,"Frauddetectionsystem: Asurvey," Journal of Networkand Computer Applications, vol. 68, pp. 90–113, 2016.

14 C. Phua, V. Lee, K. Smith, and R. Gayler, "A comprehensivesurvey of data mining-based fraud detectionresearch," Artificial Intelligence Review, vol. 34, pp. 1 – 14, 2010.

15 S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland,"Dataminingfor credit card fraud: A comparative study," Decision Support Systems, vol. 50, pp. 602–613, 2011.

16 T. Chen and C. Guestrin,"XGBoost: A scalabletreeboosting system,"in Proc. 22nd ACM SIGKDD, 2016, pp. 785–794.

17 A. Roy et al., "Deep learningdetectingfraudin credit card transactions,"in Proc. IEEE SMC, 2020, pp. 2316–2321.

18 U. Fiore et al., "Using generative adversarialnetworks for improvingclassificationeffectiveness in credit card fraud detection,"Information Sciences, vol. 479, pp. 448–455, 2019.

19 L. Breiman, "Randomforests,"Machine Learning, vol. 45, pp. 5 – 32, 2001.

20 L. Breiman,"Baggingpredictors,"MachineLearning, vol. 24, pp. 123–140, 1996.

21 J. H. Friedman,"Greedyfunction approximation:A gradientboosting machine," Annals of Statistics, vol. 29, pp. 1189–1232, 2001.

22 K. Randhawa, C. Lai, J. Zhang, and J. Sezer, "Creditcard frauddetection usingAdaBoost and majorityvoting," IEEE Access, vol. 6, pp. 14277–14284, 2018.

23 S. Singh and A. Kumar, "Fraud detection in financial transactionsusingensemblelearningand SMOTE," Expert Systems with Applications, vol. 215, p. 119312, 2023.

24 D. H. Wolpert,"Stacked generalization,"Neural Networks, vol. 5, pp. 241–259,1992.

25 A. R. Khalid et al., "Enhancingcredit card fraud detection: An ensemblemachine learningapproach," Big Data Cognitive Computing, vol. 8, no. 6, 2024.

26 H. He and E. A. Garcia,"Learningfrom imbalanceddata,"IEEE Trans. Knowledge Data Eng., vol. 21, pp. 1263– 1284, 2009.

27 N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE:Syntheticminorityover-sampling technique,"JAIR, vol. 16, pp. 321–357, 2002.

28 H. He, Y. Bai, E. A. Garcia, and S. Li, "ADASYN: Adaptive syntheticsamplingapproach for imbalancedlearning,"in Proc. IEEE IJCNN, 2008, pp. 1322–1328.

29 A. Fernandez, S. Garcia, F. Herrera, and N. V. Chawla,"SMOTE for learningfrom imbalanceddata: Progress and challenges," JAIR, vol. 61, pp. 863–905, 2018.

30 C. Elkan,"The foundationsof cost-sensitivelearning,"in Proc. IJCAI, 2001, pp. 973–978.

31 A. C. Bahnsen, D. Aouada, and B. Ottersten,"Example-dependentcost-sensitivelogisticregressionfor credit card fraud detection,"in Proc. IEEE ICMLA, 2016, pp. 263–269.

32 R. Orus, S. Mugel, and E. Lizaso, "Quantumcomputingfor finance: Overview and prospects," Reviewsin Physics, vol. 4, p. 100028, 2019.

33 P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantumsupport vector machinefor big dataclassification,"Physical Review Letters, vol. 113, p. 130503, 2014.

34 N. Innan, C. Okafor, M. Shafique, and H. Tembine,"Financialfrauddetection usingquantumgraph neuralnetworks," QuantumMachine Intelligence, vol. 6, p. 7, 2024.

35 M. Grossi et al., "Mixedquantum-classicalmethod for fraud detection with quantumfeatureselection," IEEE Trans. Quantum Eng., vol. 3, pp. 1 – 12, 2022.

36 J. Preskill,"Quantumcomputingin the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018.

37 E. Farhiand H. Neven, "Classificationwithquantumneuralnetworks on near term processors,"arXiv:1802.06002, 2018.

38 D. Gunning, "Explainableartificialintelligence(XAI)," DARPA, Tech. Rep., 2017.

39 S. Lundbergand S. Lee, "A unifiedapproach tointerpretingmodel predictions,"in NeurIPS, 2017.

40 M. Ribeiro, S. Singh, and C. Guestrin,"Why should I trustyou? Explainingthepredictions of any classifier,"in Proc. ACM KDD, 2016, pp. 1135–1144.

41 C. Molnar, Interpretable MachineLearning, 2nd ed., 2022.

42 F. Carcillo, Y.-A. Le Borgne, O. Caelen, and G. Bontempi, "Combining unsupervisedand supervisedlearningincredit card fraud detection," Information Sciences, vol. 557, pp. 317–331, 2021.

43 X. Jiang and W.-T. Tsai, "MVCG-SPS:A multi-viewcontrastivegraph neuralnetwork for smart Ponzi scheme detection,"AppliedSciences, vol. 15, no. 6, p. 3281, 2025.

44 V. Havlicek et al., "Supervisedlearningwith quantum-enhancedfeaturespaces,"Nature, vol. 567, pp. 209–212, 2019.

45 S. Lundberg et al., "From local explanationsto global understandingwith explainable AI for trees," Nature Machine Intelligence, vol. 2, pp. 56–67, 2020.

46 Y.-C. Li et al.,"HQRNN-FD: A hybrid quantumrecurrent neuralnetwork for fraud detection,"Entropy, vol. 27, no. 9, p. 906, 2025.

47 B. R. Ande, "Federatedlearningand explainableAI for decentralizedfraud detectionin financial systems,"J. Inf. Syst. Eng. Manag., vol. 10, no. 35s, 2025.

48 B. Olabowale, "Federateddeeplearningwithblockchain for privacy-preservingfraud detection," Technical Report, 2026.