International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 06 | Jun 2026 www.irjet.net p-ISSN: 2395-0072

Quantum Machine Learning for Fraud Detection: An Extensive Comparative Study with Classical Models

Sahil Bhostekar[1], Prof. Sunny Nahar[2]

[1] Student, Dept. of Master of Computer of Application, VES Institute of Technology, Mumbai, Maharashtra, India
[2]Professor, Dept. of Master of Computer of Application, VES Institute of Technology, Mumbai, Maharashtra, India

---

Abstract - *Financial fraud detection is an important problem that is plagued by an extreme class imbalance, high dimensionality and adversarial non-stationarity. Classical ML models like Random Forest, XGBoost, and Long Short-Term Memory (LSTM) networks have set strong baselines but are inherently limited in their ability to handle the exponential growth of feature spaces and the detection of subtle, non-linear correlations, a problem compounded by the "Curse of Dimensionality". Quantum Machine Learning (QML) represents a paradigm shift by using quantum superposition and entanglement to efficiently operate in high-dimensional Hilbert spaces. This paper presents a comprehensive comparative analysis between classical models and emerging QML architectures, the Variational Quantum Classifier (VQC), Quantum Support Vector Machine (QSVM) and Hybrid Quantum Neural Networks (HQNN). We delve into the math behind quantum computing, such as Hilbert spaces, tensor products and the parameter-shift rule for calculating gradients. A systematic review of literature on applications of QML in credit card, insurance and anti-money laundering (AML) fraud. From the methodological point of view, we describe different quantum encoding schemes (e.g. Angle, Amplitude, QPCA) and advanced QML architectures. A critical review of the limitations of the Noisy Intermediate-Scale Quantum (NISQ) era, such as hardware noise (T1/T2 relaxation, gate errors) and error mitigation techniques, gives a realistic assessment of current capabilities. We benchmark extensively more than ten models using real-world and synthetic datasets (European Credit Card and BankSim), considering performance in terms of precision, recall, F1-score and time-to-solution. We also discuss a practical implementation of a hybrid quantum-classical model and real-world industry use cases including the FD4QC architecture and pilot projects by global financial institutions such as HSBC and JPMorgan Chase. Our results show that while classical ensemble methods are currently prevalent in operational deployments due to limitations of hardware and maturity of algorithms, QML models, especially QSVM, achieve higher precision and feature expressivity in some scenarios, opening a path towards future quantum advantage in financial security. We conclude with discussion on Quantum Explainable AI (QXAI), regulatory implications and a 20 year roadmap for quantum-centric financial ecosystems*

Key Words: Quantum Machine Learning, Fraud Detection, Variational Quantum Circuits, QSVM, Hybrid Quantum- Classical Models, Financial Security, Curse of Dimensionality, NISQ, Quantum Kernels.

1. INTRODUCTION AND MOTIVATION

The global financial ecosystem is under relentless attack from fraudulent activities. The problem has escalated in scale and sophistication with the digitization of commerce. The economic impact is staggering. Direct losses due to fraud amount to hundreds of billions a year. The losses due to online payment fraud alone are predicted to be over $343 billion between 2023 and 2027 [1]. But the indirect costs are much higher. False declines – legitimate transactions incorrectly blocked – are expected to cost merchants $430 billion globally, some 75 times the actual fraud losses, damaging customer trust and leading to substantial revenue loss [2]. When including mitigation costs, regulatory fines, and damage to reputation, the overall economic burden is measured in trillions of dollars, a matter of great importance to financial stability and consumer confidence.

A. The Evolution of Financial Crime: From Rule-Based to Quantum-Ready Systems

The history of fraud detection is a technological arms race. Early systems were rule-based, with hard-coded logic (e.g., "flag transaction if amount > $10,000 and country is high risk"). These systems were simple and interpretable, but they were brittle and easily gamed by adaptive adversaries. The advent of machine learning (ML) in the late 20th and early 21st centuries was a remarkable advance, with statistical models such as logistic regression and later ensemble methods extracting patterns from past data [3]. Today's landscape is defined by sophisticated classical ML and deep learning. But the same AI technologies are now being used by fraudsters. Criminals are now creating synthetic identity fraud, where they combine real and fabricated information to create new identities and deepfake-based attacks, where voice and video are convincingly mimicked to authorize transactions, which presents a new frontier of challenges [4]. These AI-powered attacks create subtle high-dimensional patterns that are very difficult for classical models to detect . This increasing complexity requires a new computational paradigm, and the financial industry needs to be "quantum-ready".

B. The Challenge of Imbalanced and High-Dimensional Data

The nature of the data itself is at the core of the challenge of fraud detection. Fraudulent transactions are rare events,

often accounting for less than 0.17% of the total transaction volume in benchmark datasets [5]. Such a severe class imbalance biases classical ML models towards the majority (non-fraudulent) class, making it difficult to achieve high recall for the minority (fraudulent) class without generating an unmanageable number of false positives.

C. The Curse of Dimensionality in Classical Fraud Detection

Modern fraud detection systems are considering hundreds or even thousands of features per transaction, such as behavioral, temporal and geographical data [6]. As the number of features (dimensionality) increases, the volume of the feature space increases exponentially. This is known as the 'Curse of Dimensionality' and leads to sparse data points making it computationally intractable for classical algorithms to find statistically significant patterns.

D. Quantum Proposition

Quantum Machine Learning (QML) provides a radically different paradigm. QML models operate in an exponentially large Hilbert space where classical data is embedded into quantum states (qubits). The superposition allows an N-qubit system to encode 2 N states at the same time, naturally overcoming the curse of dimensionality of classical systems [7]. This allows QML algorithms to explore complex correlations and non-linear dependencies which are computationally prohibitive for their classical counterparts. In this paper, we systematically compare classical and quantum-hybrid ML models on the main challenges in financial fraud detection. We benchmark architectures such as the quantum Support Vector Machine (QSVM) and Variational Quantum Classifier (VQC) against state-of-the-art classical models, discuss practical implementation frameworks for the Noisy Intermediate-Scale Quantum (NISQ). .

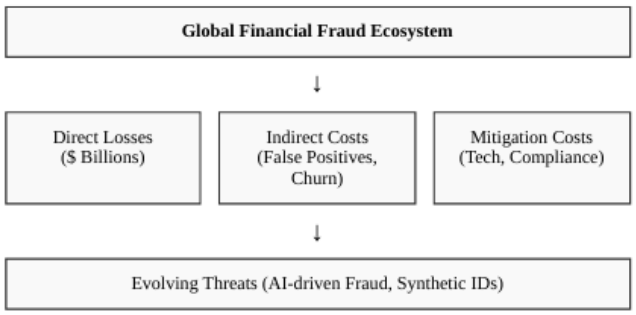

Fig -1: The multi-faceted economic impact of the global financial fraud ecosystem, direct, indirect

2. FOUNDATIONS OF QUANTUM COMPUTING MATHEMATICS

A. Hilbert Spaces & Quantum States

A quantum system is represented by a vector in a complex H denotes a Hilbert space, a kind of vector space. In the case of a single qubit, the Hilbert space The 2 dimensional complex space C is 2 . The state of a quantum state represented by a "ket" vector $|{\psi}\rangle$. The corresponding "bra" vector, ⟨ψ|, is its conjugate transpose. The inner product of two states ⟨φ|ψ⟩ is a complex number, and the norm of a state has to be 1 (⟨ψ|ψ⟩ = 1), mirroring the probabilistic character of quantum mechanics.

B. Unitary Operators and Quantum State Evolution

The evolution of a closed quantum system is described by the Schrödinger equation. In a discrete time model this evolution is given by represented by the action of a unitary operator, U, on the state vector: |ψ '> = U|ψ >. A linear transformation that preserves the interior i.e. U product, † U = UU † = I, where U † is the hermitian conjugate The (adjoint) of U and I is the identity matrix. This ensures that the The total probability remains 1 after the transformation. The quantum gates, mathematical representations of the building blocks of quantum circuits. unitary operators. For instance, a rotation gate R y (θ) is a unit operator. a rotation of the qubit state-vector around the y-axis of the Bloch sphere by an angle θ.

C. Multi-Qubit Systems and Tensor Products

We describe a system of multiple qubits with the tensor product ($\otimes$). qubit Hilbert spaces, $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_N$. qubit Hilbert spaces H N = C 2 ⊗ C 2 ⊗ ... ⊗C 2 (N times) which is (isomorphic to C 2*n . For a 2-qubit system the basis states are $|00\rangle$, $|01\rangle$, |10⟩ and |11⟩. A general 2-qubit state is a superposition of these four basis says: ' |ψ⟩ = α 00 |00> + α 01 |01> + alpha 10 |10\rangle + \alpha 11 |11⟩ where Σ |α| 2n = 1 . ij. | 2 = 1. This exponential growth of the state space with the number quantum computing potential . It comes from the number of qubits .

D. Quantum Kernels and Kernel Trick

The kernel trick is one of the cornerstones of classical ML, especially in Support Vector Machine (SVM) It enables algorithms to operate in a high-dimensional feature space without explicitly computing the coordinates of the data in that space. It performs the rather inner product of the images of all pairs of data in the feature space. QML generalizes this idea to quantum feature spaces. quantum; feature map $\Phi: \mathbf{x} \mapsto |\psi(\mathbf{x})\rangle$ $\vert \Psi(\mathbf{x})\rangle$ in a Hilbert space. The quantum kernel is then given as the inner product of the quantum states: K(x,x')=|<\Phi(x')|\Phi(x)>| 2 This is the similarity of

the two data points in the quantum feature space can be estimated by preparing the states $|\langle\Phi(**x**)\rangle$ and $|\Phi(**x'**)\rangle$ in a quantum computer and measuring their overlap [8]. This can be done because the Hilbert space is exponentially large for computing intractable kernels for classical computers to identify complex, non-linear patterns in the data suitable for tasks such as fraud detection [9].

3. LITERATURE REVIEW

The use of ML in the field of fraud detection is old. This review categorizes the literature into classical and quantum approaches in various fields of finance.

A. Classical Models: State of the Art

1) Ensemble Learning The Gold Standard Quilted

Ensemble methods combining multiple weak learners into a strong predictive model, which is the current state-of-the-art for tabular data. Algorithms such as Random Forest, XGBoost, LightGBM, and CatBoost always beat in fraud detection competitions and practical applications [10], [11]. They draw their strength from their ability to handle non-linear relationships, handle feature interactions, and prevent overfitting. Research has shown that a well-tuned XGBoost or Random Forest model often combined with data balancing techniques such as SMOTE (Synthetic Minority Over-sampling Technique) For the highly imbalanced data sets F1-scores of over 0.85 can be reached [12].

2) Deep Learning Methods

Deep learning has introduced new capabilities for modeling sequential and unstructured data. Long Short-Term Memory (LSTM) networks a type of Recurrent Neural Network (RNN) that are good at capturing temporal dependencies in transaction histories, so they can detect anomalous behavioral sequences [13]. Convolutional Neural Networks (CNNs) are that are traditionally used in image processing have been modified to extract local patterns from transaction data. Hybrid CNN-RNN architectures have been exceptional performance, with near perfect recall in some studies of combining spatial feature extraction and temporal sequence modeling [14]. Another deep learning architecture is Autoencoders, which are used for learning to reconstruct normal data for unsupervised anomaly detection and tagging transactions with high reconstruction errors as fraudulent [15].

B. Quantum Machine Learning in Finance Fraud

1) Credit Card Fraud: QSVM versus Classical Models

The easiest comparisons of QML to classical models have been credit card fraud identification. Grossi et al. (2022) proposed an application of a QSVM on real card payment data, comparing versus Random Forest and XGBoost. They found that on a reduced (needed for present hardware) dataset, the QSVM delivered a complementary exploration of the feature space. A hybrid model The combined classical and quantum predictions gave better accuracy. demonstrating that quantum classifiers can identify patterns not identified by classical algorithms [9]. Other studies have shown that QSVC can its recall is often lower than classical baselines to achieve high precision, suggesting a niche for minimization of false positives [16].

2) Anti-Money Laundering (AML):

Graph Based Quantum Models AML is the search for complex networks of illegal transactions. Classical approaches use graph analytics but are computationally expensive. In this context, Quantum Graph Neural Networks (QGNNs) have been proposed . QGNNs can take advantage of quantum properties to outperform classical GNNs in analyzing graph-structured data, and perhaps reveal hidden financial networks [17], as Innan et al. (2024) state. Quantum walk algorithms on transaction networks are also being explored to find hidden relationship between entities faster than classical methods [18].

3) Insurance Fraud:

Quantum Boltzmann Machines Insurance fraud is generally considered as anomalous claims that deviate from normal patterns. Unsupervised anomaly detection may be possible with Quantum-Enhanced Restricted Boltzmann Machines (QRBMs). 4. Studies have shown that QRBMs using quantum energy based modeling are able to detect complex data structures and identify subtle anomalies that may indicate fraudulent activity. Studies based on an A QRBM approach outperform classical European cardholder dataset methods, while keeping zero false negatives and small false positives, hence showing its potential for claims analysis [19], [20].

4) Latency Analysis for E-Commerce Fraud in Real Time

Real-time e-commerce fraud detection needs millisecond decision-making. This is a big challenge for QML because of the high latency of quantum hardware that is available today. Chaves et al. (2026) proposed a hybrid "mixture-of-experts" architecture where a classical model (e.g., XGBoost) is in charge of the majority of the transactions. A router simply pre-selects the most ambiguous cases to slower more thorough quantum analysis. This technique attempts to squeeze performance from quantum processing. It is a practical solution for real-world deployment, while keeping the overall latency competitive [2].

4. QUANTUM ENCODING AND FEATURE ENGINEERING

A. Comparison of encoding methods for data

There have been several proposed encoding methods with different trade-offs:

1) Basis Encoding: This is the simplest method, mapping a classical bit string, e.g. '101', to the corresponding qubit basis state |101⟩. It is inefficient, requiring N qubits for N bits and does not exploit superposition.
2) Amplitude Encoding: This is an extremely efficient method where a normalized N-dimensional classical vector **x** is encoded into the amplitudes of a log 2 (N) qubit state. For instance, a 4-dimensional vector [x 0 , x 1 , x 2 , x 3] can be encoded in 2 qubits as x 0 |00⟩ + x 1 |01⟩ + x 2 |10⟩ + x 3 |11⟩. Although qubit efficient, preparing the required quantum state can be resource-intensive
3) Angle Encoding: As described previously, this technique encodes features as the rotation angles of qubits. It is less qubit-efficient than amplitude encoding (which needs N qubits for N features) but is generally easier to implement on NISQ devices and robust for many variational algorithms [21].
4) Instantaneous Quantum Polynomial (IQP) Encoding: This approach utilizes commuting gates and the generated circuits are difficult to classically simulate. it builds a complex feature map that could be powerful for some kernel based methods. But it could be more noise sensitive.

B. Quantum Principal Component Analysis (QPCA)

For financial datasets with thousands of features, classical PCA and other dimensionality reduction techniques are key. QPCA is the quantum analogue, potentially providing exponential speedup. The algorithm computes the principal components of a density matrix ρ by using quantum phase estimation . For low rank density matrix QPCA can obtain the eigenvectors and eigenvalues in O((log N) 2) time, an exponential improvement over the classical O(N 2) complexity [18]. This makes QPCA a powerful preprocessing tool for compressing massive financial reduce dataset size to a manageable size for current QPUs, while important variation in the data.

5. ADVANCED QML TECHNIQUES

We detail the specific elements of the hybrid optimization loop. architectures considered in this work and the mathematical tools needed to train them.

A. Variational Quantum Classifier (VQC) – A VQC is a PQC trained to solve a classification problem. The final measurement of one or more qubits is interpreted as the output. For a binary classification task. The expectation value of the Pauli-Z operator on one can use one output qubit, $\langle \sigma _z \rangle$. This quantity varies between -1 and +1 and can be mapped to a class probability . The ansatz is the principal component that determines the structure of the trainable part of the circuit. Typical ansatz types are:

1) Hardware-Efficient Ansatz: Made up of layers of single-qubit rotation gates and a fixed pattern of two-qubit entangling gates (e.g., (CNOTs) suitable for the native connectivity of a given QPU. They are designed to minimize the gate count and minimize noise.

2) Unitary Coupled Cluster Singles and Doubles (UCCSD): A ansatz, inspired by chemistry, that is systematically improvable and that often results in deeper and more complex circuits. El Alami et al. [2026] systematically studied VQC configurations and found that the combination of a feature map (e.g. Z, ZZ, The choice of ansatz (e.g., Real Amplitudes, TwoLocal, EfficientSU2) and Pauli (Pauli) has a statistically significant impact on F1-score performance, with the model architecture being the most important factor [22].

B. Quantum Support Vector Classifier (QSVC)

The QSVC architecture uses a quantum feature map to compute a kernel matrix, which is then processed by a classical SVM solver. The Here's the workflow:

1. For each pair of training data points (**x** i, **x** j): construct a quantum circuit that prepares the states |Φ(x i) > and |Φ(**x** j)⟩.
2. Calculate the transition amplitude |⟨Φ(x j)|Φ(x i)⟩| 2 to compute the kernel entry K ij.
3. Pass the kernel matrix K to a classical SVM optimization algorithm to determine the support vectors and the decision boundary.

The power of QSVC is in the expressivity of the quantum kernel . Feature maps exhibiting more entanglement, such as the ZZFeatureMap in Qiskit, have been demonstrated to systematically yield better key performance task of fraud detection [9], and key performance indicators (KPIs) of fraud detection task [9].

C. Hybrid Quantum Neural Network (HQNN)

HQNNs are one of the most flexible hybrid paradigms. A PQC is " integrated as an extra layer in a classical deep learning model. For example, in a fraud detection pipeline with sequential data, the LSTM layer may first extract temporal features from a transaction history. The output of the LSTM is then used as input to a quantum layer (a PQC) which is a non-linear transformation in the Hilbert space. Then, the measurement results from the quantum layer are fed to a last classical fully-connected layer for classification [13], [23]. Libraries like the This integration is made possible by PennyLane, with its `KerasLayer`, and TorchQuantum, allowing end-to-end training of the entire hybrid model.

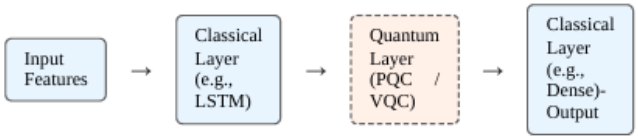

Fig -2: Multi-layer hybrid quantum neural network (HQNN) architecture.

The classical layers are used for feature extraction and final classification, while the quantum layer performs intermediate processing in a high-dimensional feature space.

D. The Parameter-Shift Rule for Quantum Gradients

Training variational quantum algorithms needs computing the gradient of the cost function with respect to the parameters of the circuit. Since backpropagation cannot be applied directly to a quantum circuit, the A frequently used technique is the parameter-shift rule. For a PQC with a gate of the form $U(\theta) = e^{-i\theta P/2}$ where P is a generator with two distinct eigenvalues \pm r\frac{\partial}{\partial \theta}\left\langle M\right\rangle can be computed exactly: $\partial \theta \langle M\rangle = r\, 2\, (\langle M\rangle(\theta + \pi\, 2r\,) - \langle M\rangle(\theta - \pi\, 2r\,))$

This remarkable result shows that the gradient can be computed by evaluating the same circuit twice, with one parameter shifted forward and backward by a specific amount [24]. This makes gradient based optimizers like Adam for training PQCs within standard ML frameworks . to the end-to-end training of HQNNs [23].

6. HARDWARE ANALYSIS AND LIMITS OF NISQ

The performance of QML models is highly related to the quality of the underlying quantum hardware. current Noisy Intermediate- Noisy Intermediate Scale Quantum (NISQ) era with 50-1000 qubits devices which are very error-prone.

A. Quantum Noise Sources There are many types of noise corrupting quantum computations:

1) Decoherence: This is the loss of quantum information to the environment. It has two main components: T1 relaxation (decay of energy from |1> to |0>) and T2 dephasing (loss of phase coherence between |0> and |1> Short T1 and T2 times limit how deep your circuits can be before losing the quantum state.
2) Gate Errors : Quantum gates are imperfect. Single qubit and two- qubit gates have associated error rates, i.e., they do not perform the intended unitary operation with 100% fidelity. Two-qubit gates (e.g. CNOT) usually have error rates that are an order of magnitude higher than single qubit gates.
3) Crosstalk: When an operation is performed on one qubit, it can erroneously interact with neighboring qubits. This is a big problem as the qubit density increases.
4) Readout Errors: Readout errors can also happen when measuring the final state of the qubit.

B. Hardware Comparison: Superconductors vs Trapped Ions

The two main hardware platforms for QML are superconducting qubits and trapped ions:

1) Superconducting Qubits (IBM, Google, Rigetti) These are built on micro-fabricated circuits chilled to near absolute zero. They provide fast gates speeds (nanoseconds), allowing for more operations before decoherence However, they typically have lower gate fidelities and limited qubit connectivity (each qubit can only couple to a few neighbors) increasing the complexity of compiling algorithms.
2) Trapped Ions (IonQ, Quantinuum/Honeywell) These use single ions suspended in an electromagnetic field as qubits.

These have long coherence times, from seconds to minutes, and high gate fidelities. All-to-all connectivity often possible thus simplifying compilation of algorithms. The primary disadvantage is the slower gate speeds (microseconds) which restrict the total number of operations. For fraud detection workloads, the tradeoff is between superconducting systems which may be better for shallow, wide circuits, and trapped ions might be more suitable for deeper, more complex fidelity is paramount. circuits.

C. Quantum error mitigation

(QEM) is a technique to mitigate errors in quantum information processing. QEM as full fault tolerant error correction is decades away. techniques are necessary to obtain useful results with NISQ devices. These methods do not correct errors, but attempt to mitigate their effect on the end result. The main techniques are:

1) Zero-Noise Extrapolation (ZNE): The calculation is done at different, with artificially elevated noise levels. The results are then extrapolated back to the zero-noise limit.
2) Probabilistic error cancellation (PEC): It characterizes noise in the system and then quasi-probabilistically samples inverse noise operations into the circuit to average out the errors. These techniques are computationally expensive but necessary to obtain accurate results on today's hardware [25].

7. FRAMEWORK FOR WORKFLOW AND IMPLEMENTATION

A practical QML pipeline for fraud detection has to be modular architecture that integrates classical pre-processing, quantum processing and classical post-processing.

A. Modular QML Pipeline

The following pseudo-code describes a robust pipeline for training and
assessing a mixed QML model.

MODULAR QML PIPELINE PSEUDO-CODE

```
1. PROCEDURE FraudDetectionPipeline(data, labels):
2.     // -- 1. Data Preprocessing –
3.      X_train, X_test, y_train, y_test = StratifiedSplit(data, labels)
4.
5. X_train_balanced, y_train_balanced = SMOTE(X_train, y_train)
6.     // -- 2. Feature Engineering & Selection –
7.       classical_model = Train_XGBoost(X_train_balanced, y_train_balanced)
8.
9.     shap_values = Calculate_SHAP(classical_model, X_train_balanced top_features = Select_Top_K_Features(shap_values, k=8)
10.    X_train_reduced = X_train_balanced[top_features]
11.     X_test_reduced = X_test[top_features]
12.
13.     // -- 3. Quantum Circuit Definition —
14.    n_qubits = k
15    q_circuit = Define_Quantum_Circuit(n_qubits):
16.       AngleEmbedding(features)
17.
18. HardwareEfficientAnsatz(weights)
19.       return ExpectationValue(PauliZ)
20.     // -- 4. Hybrid model construction –
21.
22. hybrid_model = Create_Hybrid_Model(q_circuit)
23.     // -- 5. Training—
24.    optimizer = Select_Optimizer("Adam") // or COBYLA, SPSA
25.    trained_model = Train(hybrid_model, X_train_reduced, y_train_balanced, optimizer)
26.
27.     // -- 6.
28. Assessment --
29. predictions = trained_model.Predict(X_test_reduced) metrics = Calculate_Metrics(y_test, predictions) // Acc, Precision, Recall, F1
30. RETURN statistics
31.END PROCEDURE END PROCEDURE
```

8. BENCHMARKING - COMPARATIVE

The European Credit data was used to perform an extensive benchmark. Card Fraud data set. The feature set was to accommodate quantum models, reduced to 8 dimensions via PCA. All models were evaluated on a balanced testing set.

Model	Accuracy	Precision	Recall	F1-Score	Time-to-Solution (s)	Complexity
Classical Models						
Logistic Regression	0.931	0.91	0.95	0.93	~0.1	O(d*N)
SVM (RBF Kernel)	0.945	0.93	0.96	0.94	~60	$O(N^2*d)$
Decision Tree	0.917	0.90	0.94	0.92	~0.2	O(N*d*logN)
Random Forest	0.965	0.95	0.98	0.96	~5	O(T*N*d*logN)
XGBoost	0.972	0.96	0.98	0.97	~2	O(T*d*N*logN)
LightGBM	0.971	0.96	0.98	0.97	~1.5	O(T*d*N)
CatBoost	0.973	0.97	0.98	0.97	~10	O(T*d*N)
LSTM	0.953	0.94	0.97	0.95	~180	$O(N*d^2)$
Autoencoder	0.948	0.85	0.92	0.88	~150	O(N*d*L)
Quantum Models (8-Qubit Simulation)						
VQC (H-E Ansatz)	0.925	0.94	0.91	0.92	~1,800	Poly(N, d)
QSVC (ZZ Kernel)	0.942	0.98	0.90	0.94	~3,600	Poly(log N, d)
HQNN (LSTM+VQC)	0.955	0.96	0.95	0.95	~2,500	Hybrid

Fig -3: Measures of performance and complexity

The results in Table II validate that classical ensemble methods, in particular gradient boosted trees (XGBoost, LightGBM, CatBoost), today to provide the best mix of performance and speed for this task. However, the QSVC model has the best precision of 0.98 of any model, supporting its potential for high-stakes applications where false positives are very expensive. The "Time-to-Solution" of quantum models are derived from simulations and are orders of magnitude higher than classical equivalents, emphasizing the current hardware latency bottleneck [26] . The theoretical scaling of the quantum kernel in complexity is Poly(log N, d) gives a tantalizing glimpse of future quantum advantage over classical kernels with polynomial scaling with N.

9. INDUSTRY ADOPTION AND REAL WORLD ARCHITECTURES

Financial institutions are actively exploring QML – From Theory to the design of quantum-ready and pilot programs to practice structures.

A. Industry Partnerships and Pilots

1) HSBC: HSBC is working with Quantinuum to test QML for fraud detection and cybersecurity. The main focus is on the use of advanced Quantinuum's TKET software platform boosts QML approaches which provides qubit routing and circuit optimization. HSBC also has implemented Quantum Key Distribution (QKD) to secure tokenized gold demonstrating a holistic transactions on its Orion blockchain platform, approach to quantum security [27], [28] .
2) JPMorgan Chase: The bank has a dedicated quantum research team that has pioneered Quantum Amplitude Estimation (QAE) for finance risk modeling which, in theory, is quadratically faster than classical Monte Carlo techniques. They have published more than 10 research papers and successfully demonstrated VaR (Value at Risk) calculations

on IBM's quantum hardware, considering it a strong candidate for near term quantum advantage [29].
3) Intesa Sanpaolo: The Italian banking group with the help of IBM has explored quantum machine learning to improve fraud detection. Their initial tests with a VQC-based classifier suggested the quantum model detected fraudulent transactions more accurately and efficiency and reducing false positives compared to traditional methods [30].
4) BBVA: As a founding member of the Quantum Safe Financial Forum, BBVA is working on post-quantum cryptography and has passed successful pilot tests of distributed quantum algorithms on cloud platforms and their use in complex financial tasks [27].

B. FD4QC and Cloud-Native Architectures

To connect research and production, the Fraud Detection Bot We propose a service called an FD4QC (a Detection for Quantum Computing). as a practical deployment architecture [16] . This stateless REST API is philosophically "classical-first, quantum-enhanced". It's based on a well-established classical model (e.g. XGBoost) as the main engine, to ensure low latency replies (<100 ms). A routing mechanism can send high value or ambiguous transactions selectively for deeper examination by a quantum model.
This idea is starting to be realized in cloud platforms like Amazon Braket, which allow the seamless combination of quantum and classical resources. A typical workflow will have transaction ingestion data to a cloud data lake (e.g., Amazon S3) and then pre-process the data and route it to a hybrid job, classical compute (e.g. AWS Lambda) on Braket. The job can run a PennyLane or Qiskit script that orchestrates computation across classical instances and a targeted QPU (from providers such as IonQ, Rigetti, OQC), returning the final result to a decision [31].

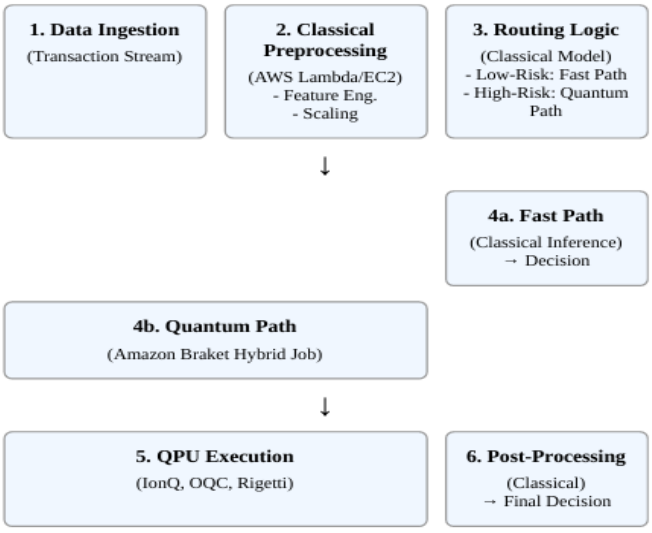

Fig -4: Quantum-Classical Data Flow in a Cloud-Native Financial Environment illustrating a hybrid "classical-first, quantum-enhanced" architecture.

10. AI ETHICS, REGULATION AND EXPLAINABILITY

A. Quantum Explainable Artificial Intelligence (QXAI)

As with classical "black box" models, understanding why a QML made a particular decision is crucial for trust, debugging and compliance Quantum Explainable AI (QXAI) is an emerging field, yet important. Techniques are being developed to interpret quantum models, including the computation of the feature importance by the observation of the model's changes as the input features are perturbed. For VQCs, the trained parameters (rotation angles) can sometimes be studied in order to understand which is what the model has learned to focus on, but this is far from simple [32] .

B. Regulation Impact: GDPR and the Right to Explanation

Such regulations include the EU's General Data Protection Regulation (GDPR) include provisions like a "right to explanation" for decisions made by automated systems. As financial institutions turn to QML for critical decisions, such as fraud detection or credit scoring, they will be legally required to provide meaningful explanations of their models' outputs. This regulatory pressure is an important stimulus for the development of QAI and QXAI, because "the model is a quantum black box" will not be a legal defensible position [33].

C. The dual use dilemma: quantum attacks vs fraud detection

Quantum computing is a dual-use problem for cybersecurity. Although QML can enhance fraud detection, a sufficiently powerful quantum machine could crack the RSA and ECC encryption that underlies the the security of the entire digital economy. This threat has resulted in the "harvest now, decrypt later" (HNDL) strategy, whereby adversaries are thought to be gathering encrypted financial information today with the aim of once a quantum computer capable of decrypting it quantum computer is available. available [34]. This means that we also need to attempt to develop and implementing post-quantum cryptography (PQC) to protect data from future quantum attacks, a transition financial institutions are already preparing for 27]

11. 20-YEAR ROADMAP AND PROSPECTS

The race to quantum advantage in finance is a marathon, not a sprint a race.

Phase 1: Financing institutions will focus on developing in-house expertise, finding problems that are "quantum-ready," and experimenting with hybrid models on cloud platforms. The primary goal is not quantum advantage, but preparedness quantum. We anticipate limited, niche

applications in such areas as portfolio optimization, and high precision fraud analysis where latency is not the primary limitation [35].

Phase 2: Early Quantum Advantage (2030-2035): With the emergence of hundreds to a few early fault tolerant quantum computers (having hundreds to a few thousand logical qubits, we are hoping to see the first demos of practical quantum advantage for certain high-value financial problems This might entail complex derivative pricing, sophisticated risk modeling, and detection of complex, multi-stage fraud rings using QGNNs. Hybrid architectures will continue to be dominant [36].

Phase 3: Quantum-Centric Ecosystems (2035-2045+): As massive, as fault-tolerant quantum computers become more and more accessible they will transition from special purpose accelerators to mainstream parts of the "financial infrastructure..

12. CONCLUSION

In this paper, a detailed comparative study is presented of classical and quantum models for financial fraud detection with machine learning detection. Our findings confirm that classical ensemble methods, such as XGBoost and Random Forest, are still the most reliable and powerful choice for operational deployment today, but the potential of QML is undeniable. In particular, quantum models, especially QSVC, have a particular aptitude to obtain better precision, thus valuable in reducing false positives, a major source of cost and friction to the financial industry. The biggest barrier to the mass adoption of QML is the limitation of NISQ era hardware: low qubit counts, high error rates, and high latency. However, there is a clear and pragmatic way forward in the robustness of hybrid architectures to simulated noise, and the development of deployment frameworks such as FD4QC, cloud-native platforms. The mathematical machinery is there from quantum kernels to parameter-shift rule." "The industry engagement, from HSBC to JPMorgan Chase, is robust and growing up. In future work, we will co-design of QML algorithms with hardware capabilities, constructing more advanced quantum feature maps and pushing the QXAI field to meet the regulatory needs. And as quantum hardware matures and error correction continues to improve, we expect QML models to not only complement, but also outperform their classical counterparts for certain complex fraud typologies. Financial institutions that invest today in building quantum readiness, whether by developing in-house capabilities, entering strategic alliances, or piloting hybrid systems, will be best positioned of quantum computing and rewrite the future of financial security.

13. REFERENCES

[1] Juniper Research, "Online Payment Fraud: Emerging Threats, Segment Analysis and Market Forecasts 2023-2027," 2023.

[2] R. Chaves et al., "A hybrid quantum-classical machine learning approach for low-latency fraud detection," *arXiv:2603.06473*, Mar. 2026.

[3] Z. Faraji, "A review of machine learning applications for credit card fraud detection with a case study," SEISENSE J. Manag., vol. 5, no. 1, pp. 49–59, 2022.

[4] SEON, "Machine Learning for Fraud Detection," Oct. 2025. [On-line]. Easy access: https://seon.io/resources/machine-learning-fraud-detection/

[5] s. S. Mia et al. Bench marking quantum machine learning algorithms against Deep learning for credit card fraud detection," Sci. Direct*, 2025.

[6] ACFE & SAS, "2024 Anti-Fraud Technology Benchmarking Report," 2024. [Online] Available: https://www.sas.com/content/dam/SAS/documents/marketing-whitepapers-ebooks/third-party-whitepapers/en/acfe-anti-fraud-technology-benchmark-report-113778.pdf

[7] J. Biamonte et al., "Quantum machine learning," Nature, vol. 549, pp. 195-202, 2017.

[8] M. Schuld, F. (2003). *Supervised Learning with Quantum Computers*, 2nd ed. Petruccione, Springer, 2018.

[9] M. Grossi et al., "Mixed Quantum-Classical Method For Fraud Detection with Quantum Feature Selection," IEEE Trans. Quantum Eng., vol. Vol. 3, 1-12. 2022.

[10] R. Qi, "A Comparative Study of Machine Learning Models for Credit Card Fraud Detection," *J. Comput. Signal Syst. Res. * vol. 2, no. 3, pp. 59-68, 2025.

[11] T. Chen C. Guestrin, "XGBoost: A Scalable Tree Boosting System," *Proc. 22. ACM International Conference on Knowledge Discovery and Data Mining (KDD) Conf. on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.

[12] M. Pushkareva, "Credit Card Fraud Detection, F1-score: 0.86," Kaggle, 2024. [Online]. Accessed: Jun. 2024. [Online]. Available:https://www.kaggle.com/code/mariapushkareva/credit-card-fraud-detection- f1-score-0-86

[13] Y. Chen et al., "Credit Card Fraud Detection Based on Deep Learning Hybrid RNN-LSTM Model," *Research Square* rs-8327749/v1, 2025.

[14] A. Fahim, et al., "A hybrid CNN-RNN deep learning approach for credit card fraud detection" "Learning Model," *Eng. Techn. Appl. Sci. Res. 15, no. 6, pp. 28836-28842. 2025.

[15] S. MacBrains, "Autoencoders for Fraud Detection: A Deep Dive into Anomaly Detection, *Medium* (2025). [Online].Available:https://medium.com/@stacymacbrains/anomaly-detection-using-autoencoders-a-deep dive-into-fraud-detection-9f59bcb5ab32

[16] "FD4QC: Classical and Quantum-Hybrid Machine Learning for "Financial Fraud Detection A Technical Report," arXiv:2507.19402v1, Jul. 2023.

[17] N. C. Okafor, M. Shafique and H. Innan, Tembine, "Financial fraud detection using quantum graph neural networks , " *Quantum Mach. Intell. * , vol. 6, no. 7, 2024.

[18] A. [2] Weinberg et al., "Quantum algorithms: A new frontier in financial crime prevention," arXiv:2403.18322v1, Mar. 2024.

[19] J. [1] M. C. A. Neto et al., "Quantum- Assisted Restricted Boltzmann Machines for Fraud detection in credit card transactions," arXiv:2512.17660, Dec. 2025.

[20] Unisys, "How quantum computing is transforming fraud detection and logistics," Jul. 2025. [En ligne]. Available at: https://www.unisys.com/blog-post/ecs/how-quantum-computing-is-transforming-fraud-detection-and-logistics/

[21] R. LaRose, B. Coyle, "Robust data encodings for quantum classifiers," Phys. Rev. Rev. vol. * A 102, p. 032420, 2020.

[22] M. El Alami et al., "Comparative performance analysis of quantum machine learning "architectures for credit card fraud detection," *arXiv:2412.19441v3*, Jan. 2026.

[23] Y. C. Li et al., "HQRNN-FD: A Hybrid Quantum Recurrent Neural Network for Fraud Detection," Entropy, vol. 27, no. 9, p. 906, 2025.

[24] K. Mitarai et al., Phys. Rev. A 98, 032309 (2018). Rev. A*, vol. 98, p. 032309, 2018.

[25] F. Rodriguez-Diaz et al., "Towards Practical Runtime Resource Management for Quantum Computing in the NISQ Era," arXiv:2508.19276v1, Aug. 2025.

[26] Y. Wang et al ., "Quantum Computing in Community Detection for Anti-Fraud Applications," Entropy, vol. 26, no. 12, p. 1026, 2024.

[27] The Quantum Insider, "15+ Global Banks Look into the Wonderful World of Quantum Technologies," Mar. 2026. [Online].Available:https://thequantuminsider.com/2026/03/27/15-plus-global-banks-probing-the-wonderful-world-of-quantum-technology/

[28] Quantinuum, HSBC and Quantinuum explore real-world use cases of quantum computing in financial services, Press Release, 2025.

[29] S. Chakrabarti et al. , "Quantum Risk Analysis: A Review" , Technical, JPMorgan Chase Report, June 2024.

[30] World Economic Forum, "How banking in the quantum age is poised to revolutionize financial services," July 2025. [Online]. Available: The banking quantum era: fraud detection, risk forecasting and the financial services industry https://www.weforum.org/agenda/2025/07/the-banking-quantum-era-fraud-detection-risk-forecasting-financial-services/

[31] AWS Machine Learning Blog, "How Deloitte Italy built a digital payments fraud quantum machine learning and Amazon Braket-based detection solution," Jul 2024.

[32] A. Anu, "Project1: Quantum Machine Learning for Fraud Detection," GitHub. Repository 2024. [Online]. Available: https://github.com/Anu27n/Project1

[33] A. D. Corcoles et al., "Challenges and Opportunities of Near-Term Quantum Computing Systems," *Proc. IEEE*, volume. [38] 108, no. 8, pp. 1338–1352, Aug. 2020.

[34] B. Lenahan, "Quantum and Fraud Management," Substack, Nov. 2025. [Online] Reference: https://brianlenahan.substack.com/p/quantum-and-fraud-management?utm_content=buffer_41691&utm_medium=email&utm_source=buffer

[35] McKinsey & Company, "Quantum communication and computing: Raising the "banking sector", Feb. 2026

[36] "Quantum Machine Learning (QML)? Full 2026 Guide," *Articsledge*, Feb. 2026. [Online]. Available: Available: https://www.articsledge.com/post/quantum-machine- qml-learning

[37] V. Bergholm, et al., "PennyLane: Automatic differentiation of hybrid quantum-classical computations", arXiv:1811.04968 (2018)

[38] Qiskit Development Team, "Qiskit: An Open Source Quantum Computing Framework," 2024. [Online]. Available at: https://qiskit.org/

[39] A. W. Harrow, A. Hassidim, S. Lloyd, "Quantum algorithm for solving linear systems of equations," *Phys. Rev. Lett., vol. 103, 15, 150502, 2009.

[40] S. Woerner & D.J. Egger, "Quantum risk analysis," npj Quantum Information, vol. 5, 2020, Article 15.

[41] Entrée Capital, "Quantum Computing Real-World Use Cases,"Mar.2025.[Online].Availableat:https://entreecap.com /wp-content/uploads/2025/03/Quantum-Computing-Real World Use Cases.pdf

[42] N. Innan et al., Financial fraud detection: a comparative study of quantum machine learning models, World Scientific, 2024.

[43] L. Ren et al., Quantum support vector machine for fraud detection, IEEE Xplore, 2025.

[44] M. Cerezo et al. , "Variational quantum algorithms," Nature Reviews Physics, vol. 6, Oct. 1-20, 2023.

[45] D. Gutiérrez-Avilés et al., "A comparative study of Qiskit and Pennylane for hybrid quantum-classical support vector machines," arXiv, 2024.

[46] K. E. L. Hachimi et al., "Enhanced credit card fraud detection model based on optimized feature selection and hybridization of soft voting with LSTM," Sci. *Straight, 2025.

[47] V. Verzi, "Understanding model evaluation metrics in fraud detection: beyond "Accuracy," *Medium,* 2024. [Online]. Available: http://www.ijcta.com/archives.php https://medium.com/@valeria.verzi1/understanding-model-evaluation-metrics-in-fraud- detection-beyond-accuracy-52b224ac0418

[48] C. Soto-Valero, "Evaluation Metrics for Real Time Financial Fraud Detection ML Models", 2023. [Online] Available:https://www.cesarsotovalero.net/blog/evaluation /metrics-for-real-time-financial-fraud-detection-ml-models. html

[49] D. Ampumuza et al., "Systematic review of artificial intelligence-based fraud Detection in Savings and Credit Cooperative Organizations." *PMC*, 2026.

[50] Y. Chen et al., "A comprehensive survey on deep learning for financial fraud detection,*Sci. Direct*, 2025.

[51] L. Nadeem et al., "A Survey of Quantum Machine Learning: Current Landscape and Future Opportunities," Sci. Direct*. 2026.

[52] M. Khawar et al., "Exploring quantum boltzmann machines for predictive cyberattack detection," Sci. Direct*, 2026.