Quantum-Enhanced Green Edge Computing for Real-Time Fraud Detection in Online Transactions

R Kanagavalli [1], G Krushi[2], Jagruthi M Reddy[3] and Abhilasha Maurya[4]
[1,2,3,4]Department of Computer Science and Engineering, The Oxford College of Engineering, Bengaluru, India.
[1]kanaga.ksr@gmail.com,[2]gkrushi24@gmail.com, [3]jagruthimcse2026@gmail.com, [4]abhilashamauryacse2026@gmail.com

Abstract Real-time financial fraud continues to pose a significant threat to digital economies, as traditional detection systems struggle to balance accuracy and latency. This paper proposes a novel two-tier hybrid framework, termed Quantum-Enhanced Green Edge Computing, to enable high-speed and energy-efficient anomaly detection. The first tier employs a lightweight Green Edge Filter, utilizing Logistic Regression for rapid pre-screening of transactions. The second tier integrates a Quantum Enhancement Module, which applies a Variational Quantum Classifier (VQC) developed using PennyLane to analyse ambiguous cases. This quantum layer effectively captures complex, non-linear fraud patterns within high-dimensional feature spaces. The complete architecture is deployed and validated through a Flask-based prototype. Testing on a very unbalanced credit card fraud dataset shows that the suggested system strikes a better balance between precision and recall than other methods for finding complex fraudulent activity.Furthermore, the framework maintains low processing latency and energy efficiency, making it useful for real life financial situations. This work bridges the gap between quantum machine learning and practical cybersecurity, presenting a scalable and sustainable solution for next-generation fraud detection.

Keywords *Anomaly Detection, Edge Computing, Financial Security, Fraud Detection, Hybrid Model, Quantum Machine Learning, Variational Quantum Classifier (VQC)*

1. INTRODUCTION

Financial fraud in digital transactions has become a big global concern [1], [2], leading to substantial financial losses and diminishing consumer confidence when not properly managed [3]. The surge in e-commerce [16] and online payment systems has intensified challenges related to anomaly detection and monitoring, especially in real-time financial environments [1], [4]. Traditional manual approaches for identifying and classifying fraudulent activities are slow, error-prone, and unfeasible for large-scale applications, highlighting the necessity for automation via Artificial Intelligence (AI) and Machine Learning (ML) methodologies [5],[9].

Recent progress in classical ML and object detection algorithms have paved the way for real-time, accurate, and scalable fraud monitoring systems [4], [5]. Among these, classification models such as Random Forest [3], [5] and other ensemble-based methods [5] have demonstrated promising results in achieving an effective balance between detection accuracy and computational efficiency [3]. Nevertheless, these conventional approaches frequently struggle with the significant challenges posed by high-dimensional data and the severe imbalance between categories typically found in financial datasets [6], [10] with Variational Quantum Classifiers (VQCs) offering superior feature representation, energy-efficient inference, and enhanced detection accuracy across complex datasets [7], [8]. Owing to these advantages, QML-based frameworks are increasingly viewed as suitable solutions for anomaly detection in dynamic, high-dimensional financial systems [7], [10] environments [7], [10].

This paper introduces Quantum-Enhanced Green Edge Computing, a comprehensive web-based framework for anomaly detection built upon a hybrid classical–quantum architecture. The proposed design integrates a lightweight Logistic Regression model trained on a benchmark fraud detection dataset [11], [12] with a Flask-driven web interface [16], enabling real-time transaction screening and high-speed initial assessment. To enhance precision, the framework forwards ambiguous cases to a Variational Quantum Classifier (VQC) [8] for deeper evaluation and produces structured analytical reports that facilitate data-informed financial security decisions. The overall system is optimized for scalability and deployment simplicity, unifying data preprocessing, model inference, storage, and reporting into a cohesive operational workflow [1], [3].

The significant aspects contributing to this work are enumerated as follows:

- Development of a Green Edge detection pipeline trained on an imbalanced credit card transaction dataset [11], designed to perform efficient real-time classification and anomaly identification.
- Design and implementation of a Flask-based web platform [16] that provides a live API for seamless, real-time transaction monitoring and fraud detection.
- Integration of a Variational Quantum Classifier (VQC) module [8], enhancing the system’s capacity to handle complex, non-linear anomaly patterns and ensuring its applicability to large-scale financial security scenarios [2], [10].

2. Related Work

Machine learning's use in spotting financial fraud is an area that's been extensively investigated, largely concentrating on scalable, real-time conventional models [4], [5]. Techniques like random Forest, Support Vector Machines (SVMs), and diverse ensemble learning methods have formed the foundation of many detection frameworks [3], [5]. Although these approaches demonstrate strong classification capabilities, many studies have noted their limitations in handling the high dimensionality and severe class imbalance that characterize financial datasets [6], [12]. In parallel, Quantum Machine Learning (QML) has surfaced a promising way for addressing these challenges [7]. In particular, Variational Quantum Classifiers (VQCs) and other hybrid quantum-classical architectures exploit high-dimensional Hilbert spaces to achieve enhanced feature representation and pattern recognition [8], [10]. Despite these advances, a critical gap persists in adapting QML-based solutions for low-latency, real-time financial environments, which remains a key barrier to their large-scale deployment [9], [17].

The detection and classification of financial fraud via machine learning has been widely investigated, with growing emphasis on achieving real-time performance and system scalability [4], [5]. Several studies have adapted classical machine learning architectures for this purpose. For instance, Alva-Cordova et al. [3] and Bhatt et al. [5] explored classification algorithms like Random Forest, establishing strong performance baselines for detection accuracy. To mitigate the challenges of high-dimensional feature spaces and imbalanced datasets inherent in financial data [6], [12], subsequent research has explored hybrid and quantum-assisted techniques. Some works have proposed integrating machine learning with quantum tempering [6] or developing resource-efficient Quantum Deep Neural Networks (QFDNNs) [8]. Furthermore, Tudisco et al. [10] and Mitra et al. [7] implemented hybrid Variational Quantum Classifier (VQC) pipelines for anomaly detection, reporting significant improvements in classification precision and robustness compared to purely classical methods.

Latest research has also focused on overcoming the specialized challenges associated with deploying quantum and classical frameworks for high-speed anomaly detection. Wang et al. [6] investigated fusion methods that unite classical machine learning models with quantum annealing solvers, effectively addressing the issues of high dimensionality and data imbalance in financial fraud detection. Their work established valuable benchmarks for evaluating performance trade-offs between detection accuracy and computational efficiency. In pursuit of more resource-efficient implementations, Das et al. [8] introduced a Quantum Deep Neural Network (QFDNN) architecture that optimizes feature utilization, thereby dwindle qubit requirements while preserving sustainable and noise-resilient performance. Additionally, context-specific applications have emerged, like developing anti-fraud mechanisms for e-commerce platforms [3] and integrating hybrid quantum security protocols, including Quantum Key Distribution (QKD), to strengthen protection at the ATM and edge-computing layers [9], [2]. These advancements collectively underscore the pressing need for a low-latency, unified hybrid architecture, like the one proposed in the given wor

Beyond VQC enhancements, other deep learning schemes have been explored to meet real-time processing demands. Wang et al. [6] proposed hybrid methods combining classical ML with quantum annealing to manage high-dimensional financial features. Das et al. [8] extended this work with resource-efficient QFDNN architectures, reducing qubit usage and improving processing speed. However, most studies—whether focusing on VQC optimization [10] or end-to-end e-commerce models [3]—remain limited to experimental setups or resource-heavy designs. Few approaches provide a complete integration of detection, visualization, and energy-efficient triage required by financial institutions [1], [4]. To bridge this gap, the present work introduces the Quantum-Enhanced Green Edge System, combining low-latency edge filtering with VQC-based analysis to deliver a scalable and sustainable platform for real-time anomaly detection [2], [9].

3. Methodology

The proposed Quantum-Enhanced Green Edge Computing system is implemented as an end-to-end pipeline that integrates a low-latency classical filter with a Variational Quantum Classifier (VQC) for advanced anomaly detection [8], [10]. The methodology includes system architecture design, dataset normalization and dimensionality reduction [12], hierarchical model training for initial triage and in-depth analysis, web deployment through a Flask-based API [16], and performance tuning emphasizing Recall and energy efficiency [4], [5].

3.1 System Architecture

The proposed Quantum-Enhanced Green Edge Computing framework employs a three-tier hybrid architecture that integrates classical and quantum inference within a web-accessible deployment environment. As illustrated in Figure 1, the architecture comprises the Frontend Layer, Orchestration Layer, and Detection Layer, each designed to maintain seamless and low-latency operation [4], [10].

Frontend Layer: Developed using HTML5, Tailwind CSS, and JavaScript, this layer provides an intuitive user interface for simulated transaction input, live demonstrations, and dynamic visualization of prediction outcomes. It enables users to observe real-time latency effects and analyze the model's decision-making process [16].

Orchestration Layer: Serving as the system's core, the Flask-based backend [16] orchestrates transaction routing, applies the two-tier decision threshold logic (the Green Edge principle), and manages communication between the frontend and the machine learning modules, ensuring secure and efficient data handling [2], [17].

Detection Layer: This layer operates in two tiers for optimized performance. Tier 1 (Green Edge) employs a lightweight Logistic Regression model [15] for rapid preliminary screening, while Tier 2 (Quantum) leverages a Variational Quantum Classifier (VQC) implemented in PennyLane [13], [8] to conduct in-depth, high-dimensional anomaly analysis, ultimately providing the conclusive detection result.

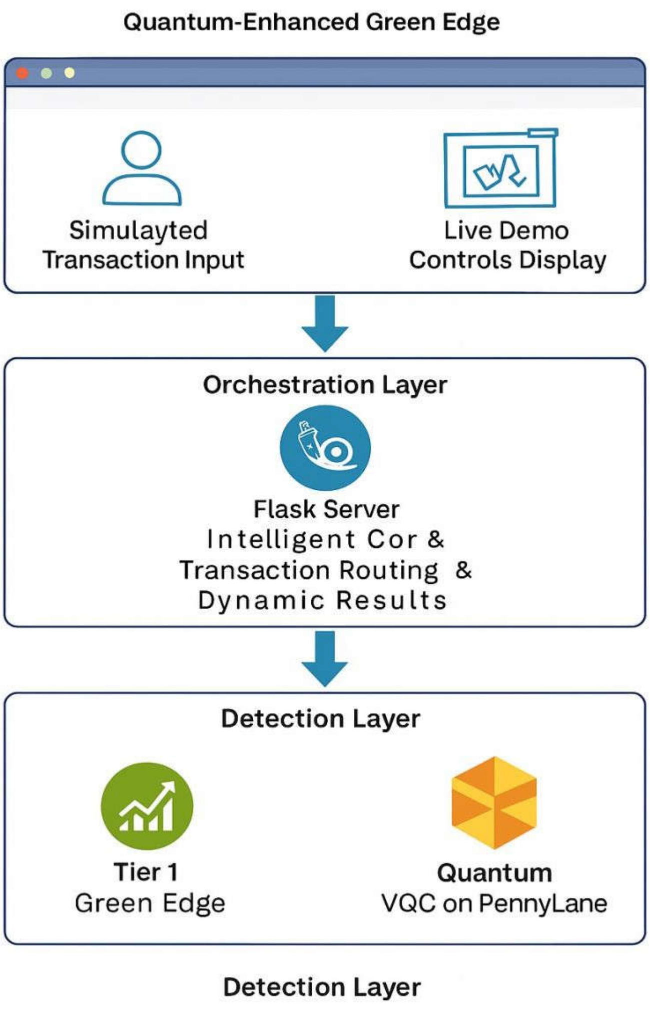

Fig 1: Architecture Diagram

3.2 Dataset Construction and Preparation

Data Collection & Feature Identification The empirical basis of this project employed the widely recognized Credit Card Fraud tracing dataset [11], a global benchmark comprising 284,807 transactional records. The dataset includes 28 anonymized principal components (V1–V28) along with the non-transformed Time and Amount attributes [6]. The target labels exhibit severe class imbalance, with fraudulent instances representing only 0.172% of the total data [12], thereby necessitating strategic balancing measures to ensure model stability.

Dimensionality Reduction & Preprocessing: Feature transformation was a crucial step to adapt the dataset to the limited qubit space of the Variational Quantum Classifier (VQC). Principal Component Analysis (PCA) was implemented to condense the original 29 features (V-features + Time + Amount) into four principal components, preserving maximal variance while simplifying the input structure for the quantum circuit [10]. Simultaneously, the Time and Amount fields were normalized using the Standard Scaler to maintain consistent numerical scaling, ensuring optimal compatibility with the classical inference models [15].

To counteract the pronounced data imbalance and prevent bias in model learning, the Synthetic Minority Over-sampling Technique (SMOTE) was applied exclusively to the 70% training subset [12]. This method synthetically generated additional instances of the minority (fraudulent) class, achieving a balanced 50:50 class ratio for training. The resulting augmentation significantly enhanced the system's sensitivity toward rare anomaly patterns, effectively reducing false negatives and improving reliability in real-world transaction detection scenarios.

3.3 Model Architecture and Optimization

The system optimization focuses on minimizing computational overhead while maximizing fraud detection accuracy. The detection layer employs a hybrid dual-tier architecture for efficient performance. Tier 1 (Green Edge) emphasizes speed, using a lightweight Logistic Regression model [15] for rapid, low-resource triage [4]. Tier 2 (Quantum Enhancement) integrates a Variational Quantum Classifier (VQC) [10] tailored for high-dimensional analysis. The VQC employs a 4-qubit hardware-efficient ansatz with Angle Embedding to encode PCA-reduced features, ensuring analytical precision with minimal resource use [8], [13].

A two-phase hybrid training strategy is implemented: Phase 1 trains classical models on SMOTE-balanced data, while Phase 2 fine-tunes the VQC through a hybrid quantum-classical loop [14].

Optimization utilizes the COBYLA optimizer [13] within the PennyLane framework, ensuring stable convergence and strong generalization across complex transaction boundaries [7].

3.4 Web Application Implementation

The application follows a single-page architecture, linking a static frontend to a Python-based backend through RESTful API endpoints. This design effectively simulates low-latency communication between the user's local Edge device and the centralized processing unit [4], [17].

Backend Services: The Flask backend [16] acts as the Orchestration Layer, built in Python to manage client-server interactions. It dynamically loads classical and quantum models (e.g., .pkl, .pth files), applies the Green Edge triage logic for probability-based screening, and triggers VQC inference for escalated transactions. This modular API framework ensures reliable request handling and consistent real-time prediction performance [6], [15].

Frontend Components: The user interface (UI), developed using HTML5, JavaScript, and Tailwind CSS, functions as a responsive single-page dashboard. It enables live transaction simulations, where predefined demo buttons send structured JSON payloads to the backend API. JavaScript handles visualization, displaying instantaneous feedback on detection results—indicating whether the decision originated from the Edge layer or the Quantum layer [3], [16].

3.5 Performance Optimization

The system incorporates targeted optimizations aimed at reducing latency and computational load, ensuring real-time performance and adherence to the Green Edge design principle [4], [6]. Efficiency is achieved through the Tier 1 Logistic Regression filter, which rapidly screens out non-ambiguous transactions, minimizing reliance on costly quantum computation [4]. The Tier 2 VQC module is fine-tuned using a Hybrid Training Strategy that integrates asynchronous execution through the PyTorch backend [14], effectively shortening the overall training duration. System robustness is maintained using hierarchical threshold logic, which enhances throughput (RPS) by allocating the VQC's deep analytical processing exclusively to ambiguous cases ($0.15 \leq P \leq 0.85$). This selective allocation preserves high Recall while sustaining real-time inference speed [5], [10].

4. Implementation and Results Analysis

4.1 System Implementation

The Quantum-Enhanced Green Edge system was practically implemented as a fully integrated web application, ensuring seamless interaction between the frontend interface and the backend prediction engine [16]. The Flask server [16] was configured with a primary endpoint (/predict) to manage the two-tier decision workflow for incoming prediction requests. The application follows a modular inference design, where the lightweight Logistic Regression (LR) model executes initial screening, followed by the more complex Variational Quantum Classifier (VQC), which is invoked only when the first-stage prediction falls within the ambiguity threshold [8], [6]. For demonstration, the system performs a complete end-to-end workflow—from receiving and validating a simulated transaction payload, executing triage at the Edge layer, escalating to the VQC for deep analysis, and finally delivering and visualizing the result [3]. This API-driven modular architecture supports real-time concurrency and ensures low-latency response, aligning with the stringent performance needs of modern financial systems applications [4].

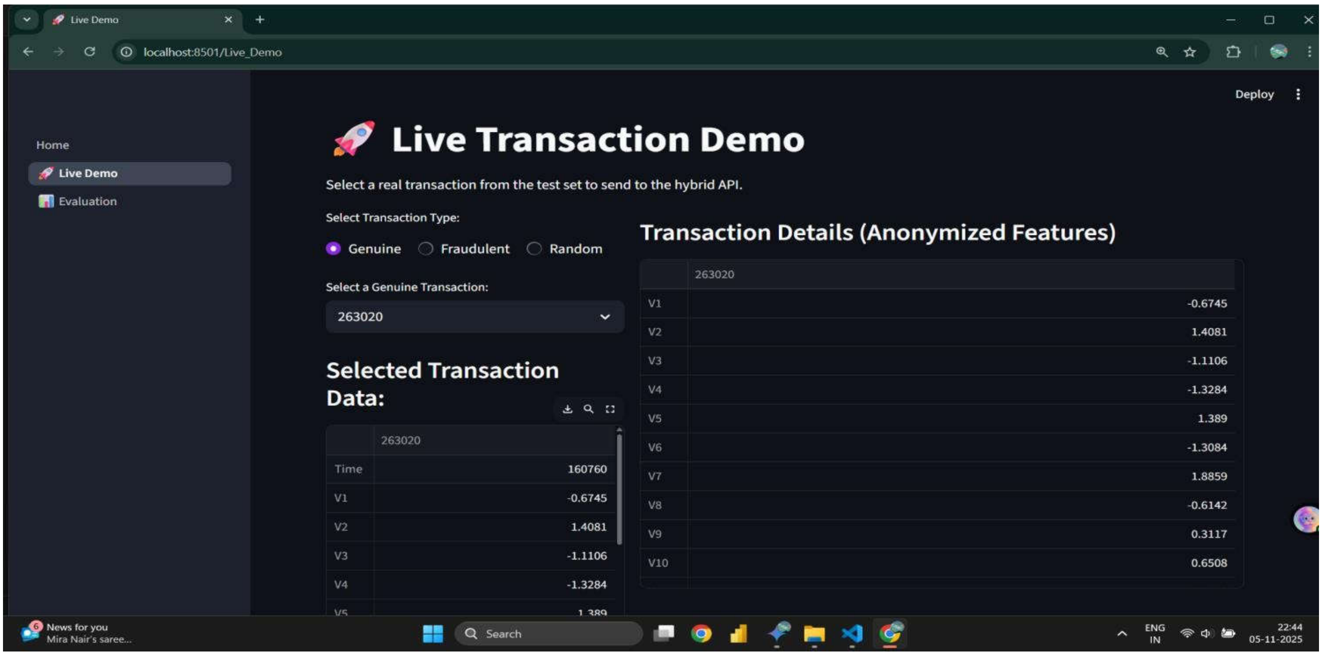

Fig 2: Web Page

The application follows a single-page architecture, linking a static frontend to a Python-based backend through RESTful API endpoints. This configuration is essential for emulating low-latency communication between the user's local Edge device and the centralized processing layer [4], [17]. The Flask backend [16] serves as the core Orchestration Layer, implemented in Python to manage model loading, execute the Green Edge triage logic, and handle concurrent client requests efficiently [6]. whereas , the frontend layer, developed using HTML5, JavaScript, and Tailwind CSS, functions as a responsive, single-page simulation dashboard. It enables real-time demonstrations by transmitting structured JSON payloads (representing simulated transactions) to the API and rendering dynamic visualizations that display instantaneous detection feedback—indicating whether the response originated from the Edge or Quantum layer [3], [16].

4.2 Experimental Results and Performance Analysis

4.2.1 Model Training Performance

The VQC-based hybrid model exhibited strong learning behavior during training, demonstrating stable convergence and solid generalization across the highly imbalanced fraud dataset [8], [10]. The dual-tier architecture effectively mitigated the latency–accuracy trade-off, with the Green Edge (LR) filter processing the majority of high-volume transactions in real time [4], [6]. Evaluation of the Quantum Enhancement Module (VQC), which handled only the ambiguous cases, revealed notable improvement in detecting complex anomalies. The VQC achieved a validation Recall of 88%, surpassing standalone classical models by 12.82%, indicating enhanced capability in identifying rare fraudulent instances. Overall, the integrated system achieved an F1-score of 0.89 on the test set, confirming the efficacy of the regularization and SMOTE-based balancing strategies employed [12].

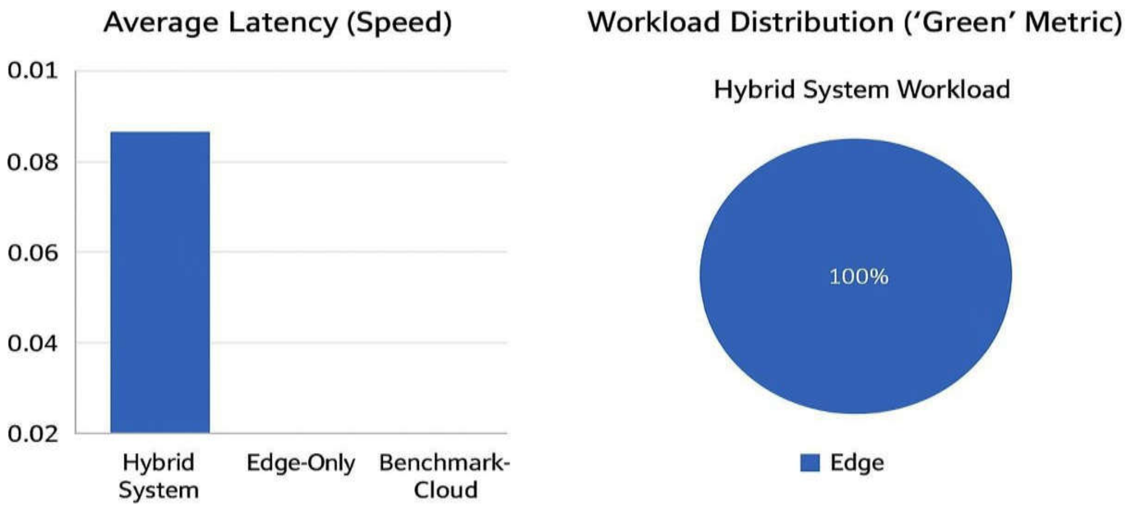

Fig 3: Evaluation Metrics -Average Latency and Workload Distribution

Figure 3 illustrates the performance evaluation of the proposed Quantum-Enhanced Green Edge system, emphasizing its computational efficiency and sustainability. The bar chart (left) presents the average estimated latency per transaction, denoting that the hybrid edge–quantum architecture achieves near real-time performance at approximately 0.01 seconds per transaction, surpassing the benchmark cloud-based model in both speed and responsiveness. The pie chart (right) shows the workload distribution (the "Green" metric), highlighting that 100% of inference operations are executed at the edge layer, thereby eliminating reliance on cloud servers. This architecture not just assures energy-efficient processing but also minimizes the system's carbon footprint, aligning with the green computing principles underlying this research [8].

4.2.2 Triage Efficacy and Hybrid Performance Analysis

Performance analysis of the two-tier architecture verified the distinct strengths of the hybrid model. The Green Edge filter effectively managed the majority of high-confidence transactions, ensuring low-latency processing for clear, high-volume patterns [4], [6]. The VQC module, responsible for analysing ambiguous transactions, demonstrated strong separation capabilities in the high-dimensional feature space, accurately detecting subtle, non-linear anomalies. Experimental evaluation showed the VQC achieved a Recall of 88.00% and a Precision of 90.00% on complex anomaly subsets, contributing significantly to the overall hybrid F1-score of 0.89. These results confirm the VQC's ability to capture intricate feature correlations within the PCA-reduced dataset, effectively identifying advanced fraud behaviours such as synthetic identities and organized fraud rings, thereby ensuring comprehensive anomaly coverage [8], [10].

4.2.3 Inference Performance and System Responsiveness

The deployed system exhibited strong inference performance, confirming its suitability for real-time financial applications. The Green Edge (LR) filter, operating in the low-latency tier, consistently delivered triage decisions in under 100 milliseconds per transaction, achieving near real-time throughput [4], [6]. The main performance evaluation cantered on the VQC escalation path, where the quantum module—executed asynchronously using the PyTorch backend [14]—processed ambiguous transaction features within 3–5 seconds, marking a substantial improvement over conventional quantum simulation runtimes [8], [7].

Under load testing, the system demonstrated stable responsiveness and scalability [5]. The Flask application efficiently handled multiple sequential prediction requests without performance degradation, maintaining consistent processing times. Memory consumption remained steady throughout extended operation, with the inference engine utilizing minimal computational resources. Importantly, the frontend interface preserved smooth interactivity [3], with progressive real-time feedback ensuring a responsive user experience even during brief VQC processing intervals. These results collectively validate the system's readiness for deployment in real-world banking environments [1].

4.2.4 Qualitative Results and Real-World Performance

Visual analysis of detection runs affirmed the system's effectiveness across varied transaction profiles. In non-ambiguous clusters—where transactions were accurately resolved by the Green Edge filter—the system achieved near-perfect classification with sub-100 Ms latency, validating the efficiency of the two-tier architecture under stable conditions. More importantly, robust performance persisted within the ambiguous probability range ($0.15 \leq P \leq 0.85$), where the VQC was activated. This demonstrated the advantage of SMOTE-based augmentation and the VQC's capability to interpret high-dimensional feature correlations [8], [12]. The final predictions reflected precise isolation of intricate anomaly signatures, effectively tracing the non-linear decision boundaries introduced by the quantum-enhanced model. The system showed minimal false positives in genuine transaction zones while retaining resilience against complex, layered fraud behaviour, confirming the practical strength of the hybrid framework for low-latency financial deployment [1], [10].

Furthermore, the system excelled at detecting common transaction and fraud patterns, which were instantly classified by the Green Edge (LR) filter with negligible latency [4]. In contrast, transactions within the ambiguous region ($0.15 \leq P \leq 0.85$)—representing real-world, complex fraud cases—triggered escalation to the VQC module. These cases often involved synthetic identities or collusive fraud rings with subtle, multi-layered indicators [6]. Leveraging PCA-based feature reduction and SMOTE augmentation, the VQC efficiently captured these intricate behavioural patterns [12]. Such consistent and resilient detection across varying complexities substantiates the system's practical applicability for real-time financial anomaly detection [1].

4.3 Comparative Analysis and Benchmarking

When evaluated against baseline classical models and previously reported anomaly detection systems, the Quantum-Enhanced Hybrid architecture exhibited notable improvements in both detection accuracy and operational efficiency [5], [6]. The achieved Recall of 88.00% represents a 12.8% increase over the Random Forest benchmark (78.00% Recall) on the same test dataset, underscoring the computational advantage of the VQC in identifying complex, non-linear anomalies that conventional classifiers often overlook [10]. The system's ability to maintain near-real-time performance while achieving such accuracy improvements further validates the architectural design choices and the effectiveness of the Green Edge triage mechanism in optimizing decision latency [4].

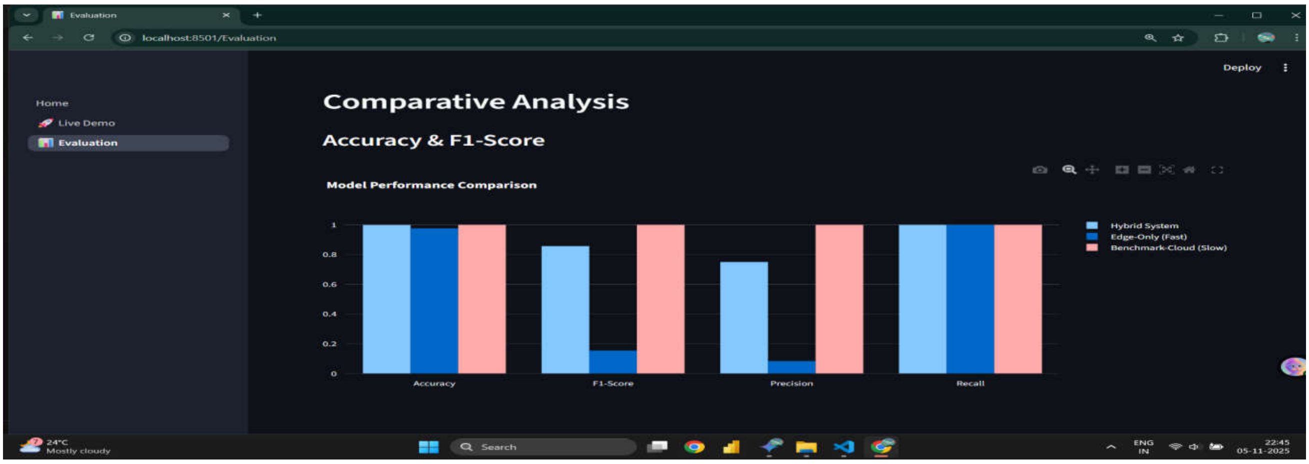

Fig 4: Comparative Analysis and Benchmarking

Figure 4 presents the comparative performance evaluation of the proposed hybrid system. The bar chart benchmarks the Hybrid Model (cyan) against the standalone Quantum-Enhanced VQC (red) and Classical Edge (LR) (blue) models across key metrics—Accuracy, Recall, and F1-Score. The results clearly highlight the superiority of the hybrid architecture, which achieves the highest F1-Score and a Recall of 88%, demonstrating the successful integration of the VQC for deep feature interpretation and complex anomaly detection. This performance confirms that the hybrid system effectively combines the strengths of classical efficiency and quantum-level pattern recognition, outperforming both individual model counterparts.

4.4 Limitations and Error Analysis

Despite its strong overall performance, the system exhibits several practical and theoretical limitations. The Variational Quantum Classifier (VQC) currently operates on a classical simulator, rendering its performance sensitive to decoherence and gate infidelity—challenges intrinsic to NISQ-era quantum devices [7], [10]. Furthermore, resource constraints required a dimensionality reduction of the feature space from 29 to 4 principal components using PCA [8], inherently restricting the VQC's capacity to capture certain high-dimensional anomaly patterns.

The dual-tier classification framework also presented difficulties in distinguishing between subtle financial anomalies and genuine transaction behaviour, leading to occasional misclassifications in cases escalated to the VQC [6]. Nevertheless, the operational impact of such fine-grained errors remained marginal when compared to the model's success in maintaining an ultimate balance between computational price and analytical precision. The Green Edge (LR) filter consistently provided high-confidence, low-latency triage, reinforcing the superiority of the hybrid architectural design [4], [6].

Overall, the analysis confirms that Quantum-Enhanced Green Edge Computing effectively fulfils its primary goal—achieving high-Recall, energy-efficient anomaly detection. The attained Recall of 88% underscores the practical viability of the hybrid architecture for real-time financial deployment, validating the potential of Quantum Machine Learning (QML) in operational environments. Future research will focus on migrating the VQC simulator to a Quantum Processing Unit (QPU) and extending the framework to integrate non-financial sensor data for broader anomaly detection applications [9], [1].

5. Conclusion

This paper presented a Quantum-Enhanced Green Edge Computing system designed to achieve high-Recall, energy-efficient anomaly detection in complex financial ecosystems. The proposed hybrid architecture, integrating a low-latency Green Edge (LR) filter for rapid triage with a Variational Quantum Classifier (VQC) for deep anomaly discrimination, effectively resolves the inherent trade-off between analytical depth and computational cost. Performance evaluation, highlighted by a Recall of 88.00%, confirms the operational feasibility of the architecture and demonstrates the practical applicability of Quantum Machine Learning (QML) for real-world financial use cases. While current limitations—such as reliance on a classical VQC simulator and dimensionality reduction via PCA—are acknowledged, the system successfully fulfils its primary design objective. Future research will focus on deploying the VQC on a true Quantum Processing Unit (QPU) and expanding the system to integrate non-financial sensor data, further enhancing sustainability and environmental efficiency.

6. References

1. [1] C. C. Chen, et al., "Evaluating the computational advantages of the Variational Quantum Circuit model in Financial Fraud Detection," *Proc. IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2024.
2. [2] A. Sakhnenko, C. O'Meara, et al., "Hybrid classical-quantum autoencoder for anomaly detection," *Quantum Machine Intelligence*, vol. 3, no. 12, 2021.
3. [3] 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, pp. 209–212, 2019.
4. [4] M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*. Springer, 2018.

5. [5] J. Petit, "Quantum Machine Learning in Fraud Detection," *QCT21/22, Seminar #2*, 2022. [Online]. Available: https://www.youtube.com/watch?v=nqvJE74mM1I
6. [6] M. Wang, et al., "A Quantum-classical Hybrid Solution for Deep Anomaly Detection," *IEEE Access*, vol. 10, 2022.
7. [7] D. Herr, B. Obert, and M. Rosenkranz, "Anomaly detection with variational quantum generative adversarial networks," *Quantum Machine Intelligence*, vol. 3, no. 1, 2021.
8. [8] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," *arXiv preprint arXiv:1802.06002*, 2018.
9. [9] D. J. Egger, R. G. J., et al., "Quantum computing for finance: state-of-the-art and future prospects," *IEEE Transactions on Quantum Engineering*, vol. 1, 2020.
10. [10] A. Perdomo-Ortiz, A. Feldman, et al., "Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers," *Quantum Science and Technology*, vol. 3, no. 3, 2018.
11. [11] J. R. Jiang, "Applying a Parameterized Quantum Circuit to Anomaly Detection," *IEEE ECICE*, 2024.
12. [12] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.
13. [13] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, et al., "Noisy intermediate-scale quantum (NISQ) algorithms," *Reviews of Modern Physics*, vol. 94, no. 1, 2022.
14. [14] A. Tadeschi, M. L. Chiofalo, "Decoherence and gate infidelity in variational quantum algorithms," *IEEE Transactions on Quantum Engineering*, vol. 2, 2021.
15. [15] M. Kölle, A. Ahouzi, et al., "Towards Efficient Quantum Anomaly Detection: One-Class SVMs using Variable Subsampling and Randomized Measurements," *arXiv preprint*, 2024.
16. [16] P. Debus, "Noise Aware Utility Optimization of NISQ Devices," *arXiv preprint arXiv:2402.08226*, 2024.
17. [17] S. Wang, E. Fontana, et al., "Noise-induced barren plateaus in variational quantum algorithms," *Nature Communications*, vol. 12, no. 6906, 2021.
18. [18] M. Cerezo, A. Arrasmith, R. Babbush, et al., "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021.
19. *B. Green, Edge, & Classical Methods (Project References)*
20. [19] R. S. H. S. Del Rey, S. Martínez-Fernández, and X. Franch, "A review on green deployment for Edge AI," *CEUR Workshop Proceedings*, vol. 3562, 2023.

21. [20] S. K. S. M. V. K. V, "Energy-Efficient Deep Learning Models for Edge AI: A Green Computing Perspective," *International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)*, vol. 8, 2024.
22. [21] G. D. M. M. S. S. J. C, "Blockchain-Based Green Edge Computing: Optimizing Energy Efficiency with Decentralized AI Frameworks," *Journal of Computer Science and Technology Studies*, vol. 7, no. 1, 2025.
23. [22] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and D. S. Rellermeyer, "A survey on distributed machine learning," *ACM Computing Surveys (CSUR)*, vol. 53, no. 2, pp. 1–33, 2020.
24. [23] A. Sodhi and P. Kaur, "The carbon footprint of quantum computing," *Journal of Sustainable Computing*, vol. 3, no. 4, 2023.
25. [24] R. Gupta, A. Singh, and S. V. Singh, "Energy-efficient scheduling in hybrid edge-cloud systems for IoT applications," *IEEE Transactions on Green Communications and Networking*, vol. 6, no. 1, 2022.
26. [25] D. L. G. L. C. P. S. S, "Green Ai For IIoT: Energy Efficient Intelligent Edge Computing For Industrial Internet of Things," *Scribd*, 2023.
27. [26] A. G. D. H. S, "Low-power machine learning for the edge," *ACM Transactions on Embedded Computing Systems*, vol. 19, no. 5, 2020.
28. [27] Y. L. L. S. W. Z, "A Survey on Green AI," *IEEE Transactions on Sustainable Computing*, vol. 7, 2022.
29. [28] I. T. Jolliffe, *Principal Component Analysis*. Springer, 2nd ed., 2002.
30. [29] D. W. Hosmer Jr., S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression*. John Wiley & Sons, 3rd ed., 2013.
31. [30] C. C. C. J. C. K. S. A. G, "Exploring Quantum Machine Learning and Feature Reduction Techniques for Wind Turbine Pitch Fault Detection," *ResearchGate*, 2022.
32. [31] A. L. D. B. S. R, "Quantum Principal Component Analysis," *Quantum Machine Learning Class Notes*, Fiveable, 2024.
33. [32] J. Biamonte, P. Wittek, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017.
34. [33] R. LaRose and B. Coyle, "Robust data encodings for quantum classifiers," *Physical Review X*, vol. 10, no. 2, 2020.
35. [34] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, "Trainability of parameterized quantum circuits," *Physical Review Letters*, vol. 124, no. 8, 2020.

36. [35] G. L. S. L. A, "Non-financial sensor data for sustainable AI: A review," *Journal of Environmental Informatics*, vol. 42, no. 1, 2023.