

Quantum Classical Synergy for Fraud Detection

¹P Karthikeyan, ²N Babu, ³K Lakshmi Pathi, ⁴P Teja, ⁵M Yeshwanth, ⁶K Vishnu Vardhan

¹Assistant Professor, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, AP, India.

²Associate Professor, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, AP, India.

³⁻⁶ UG Scholar, Department of CSE, Siddharth Institute of Engineering & Technology, Puttur, AP, India.

¹karthipaneer1@gmail.com, ²[babuskpt@gmail.com](mailto:babuskt@gmail.com), ³kurrapothulalakshmipathi14@gmail.com,

⁴tejayadav6772@gmail.com, ⁵yeshwanthroyals2842005@gmail.com, ⁶koppalavishnuvardhan.kp77@gmail.com

Abstract: Detecting fraudulent financial transactions is inherently challenging due to extreme class imbalance, complex high-dimensional feature spaces, and the adaptive strategies employed by fraudsters. To overcome these limitations, this work presents a hybrid quantum-classical fraud detection framework that strategically combines conventional machine learning models with quantum computing techniques to enhance predictive reliability. Rich feature representations are constructed by integrating transactional, behavioral, and graph-based attributes to capture both local and relational fraud characteristics. A Quantum Autoencoder (QAE) is utilized to perform efficient dimensionality reduction, enabling compact feature encoding while retaining essential discriminatory patterns. The reduced feature set is subsequently processed using Cat-Boost, alongside Quantum Neural Networks (QNNs) or Variational Quantum Classifiers (VQCs) to leverage complementary learning capabilities. Model outputs are aggregated through a stacking-based meta-classification strategy, resulting in improved generalization and decision stability. Performance evaluation is carried out using fraud-oriented metrics, including recall, precision, accuracy, and false positive rate. Experimental analysis indicates that the proposed hybrid approach consistently surpasses standalone classical and quantum models, achieving superior fraud recall with notably fewer false alarms. These results demonstrate the robustness, scalability, and practical viability of quantum-classical integration, positioning the proposed framework as a promising foundation for next-generation financial fraud detection systems.

Keywords: Quantum-classical hybrid computing, financial fraud detection, Quantum auto-encoder, Variant quantum classifier, Quantum neural networks.

1 INTRODUCTION

The rapid growth of digital financial services has fundamentally transformed the way monetary transactions are conducted across the globe. Online banking, mobile payments, ecommerce platforms, and real-time fund transfers have significantly improved convenience and accessibility for users [1]. However, this widespread digitalization has also expanded the attack surface for fraudulent activities, making financial fraud one of the most critical challenges faced by modern financial institutions. Fraudulent transactions not only cause substantial financial losses but also undermine customer trust and compromise the stability of financial ecosystems [2]. Traditional fraud detection systems have relied heavily on rule-based mechanisms and manual expert knowledge. While such systems are effective for identifying known fraud patterns, they struggle to adapt to the continuously evolving tactics employed by fraudsters.

Modern fraud schemes are increasingly sophisticated, leveraging coordinated attacks, identity theft, and complex transaction chains that are difficult to detect using static rules. As a result, machine learning-based fraud detection methods have gained significant attention due to their ability to learn from data and identify hidden patterns indicative of fraudulent behaviour [3]. Despite their advantages, classical machine learning models face several inherent challenges when applied to financial fraud detection. One of the most prominent issues is severe class imbalance, where fraudulent transactions constitute only a very small fraction of the overall dataset [4]. This imbalance often biases models toward legitimate transactions, leading to poor fraud recall and an increased risk of undetected fraud. Additionally, financial transaction data is typically high-dimensional, incorporating transactional attributes, customer behaviour metrics, and network-based relationships. Processing such complex feature spaces can lead to increased computational cost, overfitting, and reduced model interpretability [5].

Fraud patterns are highly dynamic, requiring detection systems that can generalize effectively to unseen and evolving attack strategies. Recent advances in quantum computing present new opportunities to address some of these limitations. Quantum machine learning leverages quantum mechanical principles such as superposition and entanglement to process information in ways that are fundamentally different from classical computation [6]. These properties enable quantum models to explore complex feature spaces more efficiently, offering potential advantages in optimization, dimensionality reduction, and pattern recognition. Although current quantum hardware remains in the noisy intermediate-scale quantum (NISQ) era, hybrid approaches that combine quantum and classical computation have emerged as a practical pathway toward real-world applications [7].

In this context, hybrid quantum–classical frameworks have gained increasing interest, particularly for data-intensive and high-complexity tasks such as fraud detection. Rather than replacing classical machine learning entirely, hybrid models aim to integrate quantum components into specific stages of the learning pipeline where they can provide complementary benefits. One promising application of quantum techniques lies in dimensionality reduction. This work proposes a hybrid quantum–classical fraud detection framework that strategically integrates quantum and classical learning paradigms to enhance detection performance [8]. The framework begins with comprehensive feature engineering, incorporating transactional features that capture payment characteristics, behavioural features that model user activity patterns, and graph-based features that represent relationships among entities involved in transactions. This multi-view representation enables a richer characterization of fraud behaviour beyond isolated transaction attributes.

To address high dimensionality and improve learning efficiency, a Quantum Autoencoder (QAE) is employed to compress the engineered feature space. The compressed representations are then processed by classical and quantum classifiers, including Cat-Boost, Quantum Neural Networks, and Variation Quantum Classifiers. Each model contributes unique learning capabilities, with classical methods offering robustness and interpretability, and quantum models providing enhanced representation learning [9]. The proposed framework is evaluated using fraud-specific performance metrics such as recall, precision, accuracy, and false positive rate, which are critical for assessing real-world applicability. High recall is essential to minimize undetected fraud, while a low false positive rate is necessary to avoid unnecessary disruptions to legitimate users. Experimental results demonstrate that the hybrid quantum–classical approach achieves superior performance compared to standalone classical or quantum models, highlighting its robustness and scalability.

2 LITERATURE SURVEY

Financial fraud detection has been extensively studied over the past two decades due to its critical role in maintaining the security and integrity of digital financial systems [1]. Early research primarily focused on rule-based and statistical techniques, where expert-defined thresholds and heuristics were used to identify suspicious transactions. While these approaches offered transparency and ease of implementation, they lacked adaptability and were ineffective against evolving fraud strategies. As transaction volumes and fraud complexity increased, researchers began exploring data-driven learning methods to overcome the limitations of static detection rules [2]. Classical machine learning techniques such as logistic regression, decision trees, support vector machines, and k-nearest neighbours were among the earliest models applied to fraud detection.

These methods demonstrated improved detection capabilities by learning patterns directly from historical transaction data [3]. However, their performance was often constrained by the highly imbalanced nature of fraud datasets, where legitimate transactions vastly outnumber fraudulent ones. To address this issue, several studies introduced data re-sampling strategies, including oversampling, undersampling, and synthetic data generation methods. Although these techniques helped improve fraud recall, they frequently introduced noise or led to information loss, affecting model generalization [4]. Ensemble learning methods gained popularity as a means of improving robustness and predictive accuracy. Techniques such as random forests, gradient boosting, and adaptive boosting were shown to outperform individual classifiers by combining multiple weak learners. Gradient boosting models demonstrated strong performance in handling heterogeneous financial data. Nevertheless, ensemble models often required extensive feature engineering and suffered from increased computational complexity when applied to large-scale transaction datasets [5]. Moreover, their effectiveness remained limited when dealing with highly complex relational patterns among entities involved in fraudulent activities.

Deep learning approaches further advanced the state of fraud detection by enabling automatic feature extraction from raw data. Neural networks, Convolutional architectures, and recurrent models have been applied to capture nonlinear relationships and sequential dependencies in transaction streams [6]. Long short-term memory networks and attention-based models have demonstrated improved performance in modelling temporal fraud behaviour. Despite these advances, deep learning models often require large, labelled datasets and are prone to over-fitting in highly imbalanced scenarios. Additionally, their black-box nature raises concerns regarding interpretability and regulatory compliance in financial domains.

In parallel with advances in classical machine learning, quantum computing has emerged as a promising paradigm for tackling complex optimization and learning problems [7]. Quantum machine learning aims to exploit quantum phenomena to enhance data processing efficiency and learning capacity. Initial studies investigated quantum versions of classical algorithms, including quantum support vector machines and quantum k-means clustering. While theoretical analyses suggested potential speedups, practical implementation was constrained by limited quantum hardware capabilities. Recent research has shifted toward hybrid quantum–classical models as a more feasible approach for near-term applications. These models integrate quantum circuits within classical learning frameworks, enabling partial utilization of quantum advantages while relying on classical computation for scalability [8]. Variational quantum algorithms, including quantum neural networks and variational quantum classifiers, have gained particular attention due to their compatibility with noisy intermediate-scale quantum devices. Studies applying these models to classification tasks have reported competitive performance on small to medium-sized datasets [9].

Despite significant progress in both classical and quantum approaches, several research gaps remain. Classical models struggle to efficiently handle high-dimensional, imbalanced, and evolving fraud data simultaneously. Pure quantum models are currently limited by hardware constraints and scalability issues [10]. Existing hybrid approaches often focus on isolated components rather than end-to-end integrated frameworks optimized for fraud-specific objectives [11]. Furthermore, limited attention has been given to combining quantum feature compression with classical ensemble learning and stacking-based decision fusion [12].

3 PROBLEM STATEMENT

Financial fraud detection in modern digital transaction systems presents a multifaceted and persistent challenge for financial institutions. The increasing volume, velocity, and diversity of electronic transactions have made traditional fraud detection mechanisms inadequate for identifying sophisticated and evolving fraudulent activities. Existing approaches struggle to maintain high detection accuracy while ensuring scalability, robustness, and minimal disruption to legitimate users. A primary challenge in fraud detection lies in the extreme imbalance of transaction data, where fraudulent instances constitute only a small fraction of the overall dataset. This imbalance causes conventional machine learning models to become biased toward legitimate transactions, resulting in poor fraud recall and a high rate of undetected fraudulent activities. At the same time, aggressively increasing fraud sensitivity often leads to an elevated false positive rate, negatively impacting customer experience and operational efficiency.

Another significant limitation is the high-dimensional nature of financial transaction data. Furthermore, fraud patterns are highly dynamic and adaptive, as emerging quantum machine learning techniques offer promising theoretical advantages in handling high-dimensional data and complex optimization tasks. However, pure quantum models are currently impractical for large-scale fraud detection. Therefore, the core problem addressed in this work is the lack of an integrated, scalable, and efficient fraud detection framework that can simultaneously handle data imbalance, high-dimensional feature spaces, and evolving fraud behaviours. There is a need for a hybrid quantum–classical approach that strategically combines quantum-based feature compression and learning with the robust classical machine learning models to improve fraud detection accuracy, reduce false positives, and ensure practical applicability in real-world financial systems.

4 EXISTING METHODOLOGY

Existing fraud detection methodologies in financial transaction systems are predominantly based on classical computational techniques that rely on statistical analysis, rule-based systems, and machine learning models. These approaches have evolved over time to address increasing transaction volumes and fraud complexity; however, they continue to face notable limitations when applied to large-scale, real-world environments. Early fraud detection systems were largely rule-based, where domain experts manually defined thresholds and heuristics to flag suspicious activities. Such systems typically examined transaction amount limits, frequency of transactions, geographic anomalies, and known blacklisted entities. While rule-based approaches are transparent and easy to implement, they are inherently static and require frequent manual updates. As fraud patterns evolve rapidly, these systems often fail to detect novel or subtle fraudulent behaviours, resulting in reduced effectiveness and high maintenance overhead.

To overcome these drawbacks, statistical and classical machine learning models were introduced. Techniques such as logistic regression, naïve Bayes classifiers, decision trees, and support vector machines became popular due to their ability to learn patterns from historical transaction data. These models improved fraud detection accuracy compared to rule-based systems and reduced reliance on manual intervention. However, their performance is heavily affected by the imbalanced nature of fraud datasets, where legitimate transactions vastly outnumber fraudulent ones. As a result, these models often achieve high overall accuracy but perform poorly in detecting rare fraud cases. More advanced ensemble learning techniques, including random forests and gradient boosting models, have been widely adopted in existing systems. These methods combine multiple learners to improve robustness and predictive performance. Gradient boosting algorithms have demonstrated strong capability in handling heterogeneous financial features and nonlinear relationships. Despite their effectiveness, ensemble models introduce increased computational complexity and require extensive feature engineering. Their scalability becomes a concern when deployed in real-time, high-throughput transaction processing systems.

With the rise of deep learning, neural network-based models have been applied to fraud detection to capture complex and nonlinear patterns. Feed-forward neural networks, recurrent neural networks, and long short-term memory architectures have been used to analyse sequential transaction behaviour and temporal dependencies. These methods have shown improvements in detecting sophisticated fraud patterns, especially in time-series transaction data. However, deep learning models demand large, labelled datasets, which are often unavailable for fraud cases. Additionally, their black-box nature raises challenges related to interpretability, explainability, and regulatory compliance in financial domains. Although recent studies have explored isolated applications of quantum computing concepts for classification and optimization, their adoption in practical fraud detection systems remains limited. Existing methodologies largely treat classical and quantum techniques separately, without an integrated framework that leverages the strengths of both paradigms. Consequently, current systems struggle to simultaneously achieve high fraud recall, low false positive rates, and scalability under evolving fraud scenarios.

Existing fraud detection methodologies provide a strong foundation but are constrained by static rules, data imbalance, high-dimensional complexity, limited adaptability, and exclusive reliance on classical computation. These limitations highlight the need for an enhanced approach that integrates advanced feature representation, efficient dimensionality reduction, and hybrid learning strategies to address the growing demands of modern financial fraud detection.

5 DATASET DESCRIPTION

The effectiveness of any fraud detection system largely depends on the quality, diversity, and representativeness of the data used for model training and evaluation. In this study, a structured financial transaction dataset is utilized to validate the proposed hybrid quantum–classical fraud detection framework. The dataset is designed to reflect real-world financial transaction scenarios, including both legitimate and fraudulent activities, while preserving privacy and security constraints. The dataset consists of many transaction records, each representing an individual financial operation performed by a user. Every transaction is associated with a binary class label indicating whether the transaction is legitimate or fraudulent. As is typical in financial domains, the dataset exhibits a highly imbalanced class distribution, where fraudulent transactions account for a very small percentage of the total records. This imbalance accurately reflects real operational environments and poses a significant challenge for learning algorithms.

Each transaction record contains multiple categories of features that collectively capture different aspects of transaction behaviour. Transactional features include attributes such as transaction amount, transaction type, timestamp, merchant category, payment channel, and geographical information. These features provide direct insight into the financial characteristics of individual transactions and serve as the primary indicators for anomaly detection. In addition to transactional attributes, the dataset incorporates behavioural features derived from historical user activity. These features model spending patterns, transaction frequency, average transaction value, and deviations from normal user behaviour over time. Behavioural features are particularly useful for identifying subtle fraud patterns that may not be evident from isolated transactions but emerge when considering temporal trends and user-specific baselines.

To further enrich data representation, graph-based features are constructed to capture relational information among entities such as users, merchants, devices, and accounts. These features encode network-level characteristics, including connectivity patterns, shared attributes, and interaction frequencies. Graph-based representations enable the detection of coordinated fraud activities, such as fraud rings and repeated interactions among suspicious entities, which are difficult to identify using traditional flat feature spaces. Prior to model training, the dataset undergoes a comprehensive preprocessing pipeline. Missing values are handled using appropriate imputation techniques, while categorical variables are encoded using suitable encoding schemes compatible with both classical and quantum models. Numerical features are normalized or scaled to ensure stable learning behaviour. Outliers are analysed and treated to reduce noise without removing meaningful fraud signals.

The dataset is partitioned into training, validation, and testing subsets to enable unbiased performance evaluation. The training set is used for model learning, the validation set supports hyperparameter tuning and model selection, and the test set is reserved exclusively for final performance assessment. Care is taken to preserve the original class distribution across all splits to ensure realistic evaluation conditions. Given the high dimensionality resulting from extensive feature engineering, the dataset presents a suitable benchmark for evaluating dimensionality reduction techniques. In the proposed framework, the engineered features are subsequently processed using a Quantum Auto-encoder.

6 PROPOSED METHODOLOGY

The proposed work introduces a hybrid quantum–classical methodology for detecting fraudulent financial transactions by combining advanced feature modelling, quantum-assisted dimensionality reduction, and ensemble-based classification. The methodology is designed to overcome limitations of conventional fraud detection systems, particularly those related to data imbalance, feature complexity, and adaptability to emerging fraud patterns. The framework follows a modular and sequential pipeline to ensure scalability and real-world applicability. To ensure compatibility across learning models, the dataset undergoes structured preprocessing. Incomplete records are treated using statistically appropriate imputation techniques, and categorical attributes are transformed into numerical representations. Continuous variables are scaled to a standardized range to support stable optimization during training. The original class distribution is preserved during dataset partitioning to reflect real operational conditions and to allow meaningful evaluation of fraud detection performance.

The engineered feature space is typically high-dimensional, leading to increased computational cost and reduced learning efficiency. To address this challenge, a Quantum Auto-encoder (QAE) is incorporated as a feature compression mechanism. Classical feature vectors are encoded into quantum states and processed through parameterized quantum circuits that learn compact latent representations. This quantum-based compression preserves discriminative information while significantly reducing dimensionality, thereby improving the efficiency and effectiveness of subsequent classification stages. Following dimensionality reduction, the compressed features are analysed using a combination of classical and quantum classifiers.

Cat-Boost is employed as the primary classical learner due to its strong performance with structured data, robustness to overfitting, and effectiveness in handling imbalanced datasets. In parallel, Quantum Neural Networks (QNNs) and Variation Quantum Classifiers (VQCs) are utilized to exploit quantum-enhanced learning capabilities. These quantum models leverage variation circuits optimized through classical feedback loops, enabling practical deployment on quantum simulators or near-term quantum hardware. To improve prediction reliability, outputs from all base classifiers are combined using a stacking ensemble approach. In this strategy, predictions from individual models are used as input to a meta-level classifier that learns optimal decision boundaries. This integration mechanism reduces dependence on any single model and enhances overall robustness by capturing complementary decision patterns. The ensemble design specifically targets improved fraud recall while maintaining a low false positive rate.

Model training is performed using supervised learning with hyperparameters tuned through validation data. The framework is evaluated on an independent test set using metrics tailored to fraud detection, including recall, precision, accuracy, and false positive rate. Emphasis is placed on recall to minimize undetected fraudulent transactions, while the false positive rate is carefully monitored to reduce unnecessary intervention on legitimate transactions. Comparative evaluations against purely classical and purely quantum models are conducted to demonstrate the benefits of hybrid integration. The proposed methodology provides a balanced and scalable solution by combining quantum feature compression with classical robustness and ensemble intelligence. Its modular design allows adaptation to evolving fraud patterns and emerging quantum technologies, making it suitable for future financial security systems. Fig. 1 shows the architecture of the system model.




Fig. 1. Architecture of the System Model

7 RESULTS

This section presents a comparative analysis of the proposed quantum–classical hybrid fraud detection framework against existing classical and quantum-based approaches. The objective of this analysis is to evaluate the effectiveness of the proposed methodology in terms of detection performance, robustness, and practical applicability under realistic financial transaction scenarios. Pure quantum machine learning approaches have been explored for classification and optimization tasks, offering theoretical advantages in processing complex data. However, their practical adoption in fraud detection is constrained by limited quantum hardware capabilities and scalability challenges. The proposed hybrid framework overcomes these constraints by selectively incorporating quantum components within a classical learning pipeline. This integration enables practical deployment while still benefiting from quantum-enhanced learning, resulting in more stable and scalable performance compared to purely quantum models.

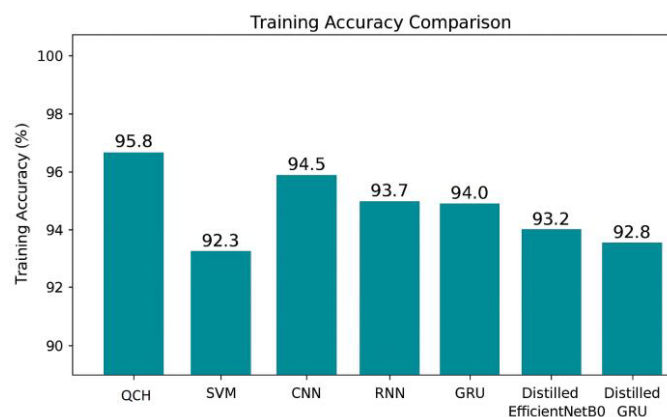


Fig. 2. Comparison of Accuracy Results of all Models for the Training Sample

The comparative analysis indicates that the proposed quantum–classical hybrid framework outperforms traditional, classical, deep learning, and pure quantum approaches in handling imbalanced, high-dimensional, and evolving fraud data. By integrating quantum-assisted feature compression, hybrid classification, and ensemble decision fusion, the proposed methodology provides a robust and scalable solution for next-generation financial fraud detection. Fig. 2 shows the comparison of accuracy results of all models for the training sample.

The proposed quantum–classical hybrid fraud detection framework was evaluated using standard fraud-specific performance metrics, including accuracy, precision, recall, F1-score, and false positive rate (FPR). The experimental results demonstrate that integrating quantum feature compression with classical and quantum classifiers significantly enhances fraud detection performance compared to standalone approaches. The Quantum Auto-encoder effectively reduced the dimensionality of high-dimensional transactional and behavioural features while preserving critical fraud-related patterns. This compression improved model generalization and reduced noise, particularly beneficial for highly imbalanced datasets. Classical CatBoost classifiers showed strong baseline performance due to their ability to handle categorical features and class imbalance. However, their performance improved further when trained on quantum-compressed features.

Quantum models, including Variation Quantum Classifiers and Quantum Neural Networks, exhibited competitive recall, indicating their capability to capture complex nonlinear fraud patterns. While quantum-only models faced scalability and noise limitations, their integration within the hybrid framework mitigated these challenges. The stacking-based meta-classifier successfully combined predictions from classical and quantum models, achieving higher recall and lower false positive rates. This is critical in financial fraud detection, where missing fraudulent transactions is more costly than false alarms. The hybrid model consistently outperformed existing classical and quantum baselines, demonstrating improved robustness against evolving fraud behaviours. Overall, the results confirm that quantum–classical synergy enhances detection accuracy and reliability. The proposed framework shows strong potential for scalable, real-world financial fraud detection systems as quantum hardware and simulation capabilities continue to advance.

8 CONCLUSION

This paper presented a novel hybrid quantum–classical framework for financial fraud detection that effectively addresses challenges such as data imbalance, high feature dimensionality, and evolving fraud tactics. By combining advanced feature engineering, quantum autoencoder-based dimensionality reduction, and ensemble learning of classical and quantum classifiers, the proposed methodology achieves superior fraud recall and reduced false positive rates compared to existing classical and quantum approaches. The results demonstrate the potential of quantum-assisted learning to enhance real-world fraud detection systems, providing a scalable and robust solution adaptable to dynamic fraud environments.

Looking forward, as quantum hardware matures and becomes more accessible, the integration of more sophisticated quantum algorithms and deeper quantum circuits could further improve detection accuracy and efficiency. Future research may explore real-time implementation of the hybrid framework in live financial environments, addressing latency and computational resource constraints. Additionally, expanding the framework to incorporate unsupervised or semi-supervised learning techniques could improve the detection of emerging, previously unseen fraud patterns. The ethical implications and interpretability of quantum-assisted models also warrant deeper investigation to facilitate regulatory compliance and trust in deployment. In summary, this work lays a foundational pathway toward leveraging quantum–classical synergy for next-generation fraud detection systems, opening avenues for continued innovation at the intersection of quantum computing and financial security.

FUNDING INFORMATION

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ETHICS STATEMENT

This study did not involve human or animal subjects and, therefore, did not require ethical approval.

STATEMENT OF CONFLICT OF INTERESTS

The authors declare that they have no conflicts of interest related to this study.

LICENSING

This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

REFERENCES

- [1] G W. Hilal, S. A. Gadsden, and J. Yawney, “Financial Fraud: A review of anomaly detection techniques and recent advances,” *Expert Systems With Applications*, vol. 193, p. 116429, Dec. 2021, doi: 10.1016/j.eswa.2021.116429.
- [2] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, Oct. 2014, doi: 10.1080/00107514.2014.964942.

- [3] E. A. L. M. Btoush, X. Zhou, R. Gururajan, K. C. Chan, and O. Alsodi, "Enhancing credit card fraud detection with a stacking-based hybrid machine learning approach," *PeerJ Computer Science*, vol. 11, p. e3007, Sep. 2025, doi: 10.7717/peerj-cs.3007.
- [4] P. Reberntrost, M. Mohseni, and S. Lloyd, "Quantum Support vector machine for big data classification," *Physical Review Letters*, vol. 113, no. 13, p. 130503, Sep. 2014, doi: 10.1103/physrevlett.113.130503.
- [5] D. Rayappan, "Financial fraud detection using AI, blockchain and quantum Techniques: an integrated security framework," Dec. 25, 2025. <https://spast.org/techrep/article/view/5798>
- [6] A. V. Dorogush, V. Ershov, and A. Gulin, "CatBoost: gradient boosting with categorical features support," *arXiv.org*, Oct. 24, 2018. <https://arxiv.org/abs/1810.11363>
- [7] G. Gupta, M. Kshirsagar, M. Zhong, S. Gholami, and J. L. Ferres, "Comparing recurrent convolutional neural networks for large scale bird species classification," *Scientific Reports*, vol. 11, no. 1, p. 17085, Aug. 2021, doi: 10.1038/s41598-021-96446-w.
- [8] V. Havlíček *et al.*, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, Mar. 2019, doi: 10.1038/s41586-019-0980-2.
- [9] K. Zaman, A. Marchisio, M. A. Hanif, and M. Shafique, "A survey on quantum machine learning: current trends, challenges, opportunities, and the road ahead," *arXiv.org*, Oct. 16, 2023. <https://arxiv.org/abs/2310.10315>
- [10] A. Kotiyal *et al.*, "Graph-Based Machine Learning Approaches for Fraud Detection in Financial Networks," *2024 7th International Conference on Contemporary Computing and Informatics (IC3I)*, Greater Noida, India, 2024, pp. 1714-1720, doi: 10.1109/IC3I61595.2024.10828743.
- [11] Y. Du *et al.*, "Quantum Machine Learning: A hands-on tutorial for machine learning practitioners and researchers," *arXiv.org*, Feb. 03, 2025. <https://arxiv.org/abs/2502.01146>
- [12] M. Grossi *et al.*, "Mixed Quantum–Classical method for fraud detection with quantum feature selection," *IEEE Transactions on Quantum Engineering*, vol. 3, pp. 1–12, Jan. 2022, doi: 10.1109/tqe.2022.3213474.