

Hybrid Quantum-Classical Machine Learning For Scalable AI Applications

Manish Patil¹, Shraddha Sanjay Mhetre²

¹MTS Software Engineer Salesforce, Inc.

²Android Engineer

Abstract

Background: The emergent need for scalable artificial intelligence (AI) models emphasizes the shortcomings of traditional machine learning in dealing with big, complicated datasets owing to computation and energy limitations.

Objective: The architecture takes advantage of quantum encoding methods like amplitude and angle encoding for compact feature representation, with classical preprocessing providing compatibility and stability.

Methods: To overcome these issues, this research introduces a Hybrid Quantum-Classical Machine Learning (HQML) framework that combines Quantum Variational Classifiers (QVCs) and Quantum Convolutional Neural Networks (QCNNs) with classical deep learning layers.

Results: Experiments were performed on pre-curated datasets taken from the Scikit Machine Learning Repository, with hybrid models compared to classical Convolutional Neural Networks (CNNs). The implementation was done with PennyLane using TensorFlow Quantum for compatibility with noisy intermediate-scale quantum (NISQ) machines. The comparative analysis identified that the suggested HQML model achieves significant performance gains, as classification accuracy increased from 73 percent in the baseline classical CNN to 98% in the hybrid model. Precision also went up from 77 to 96 percent, recall from 81 to 100 percent, and F1-score from 79 to 98 percent. These findings confirm the ability of the model to leverage quantum benefits for extracting complicated data distributions without losing stability through classical components. The findings confirm the viability of HQML to real-world finance, healthcare, and intelligent automation applications where scalability and precision matter.

Conclusion: This research establishes that quantum-classical integration not only minimizes training overhead but also attains near-perfect classification efficiency with constrained qubit resources. The research aids in filling the gap between theoretical quantum potential and deployable practical AI, opening the door for next-generation quantum-boosted intelligence.

Keywords: AI, QCNNs, QVCs, HQML, energy limitations.

1 Introduction

The rapid evolution of artificial intelligence (AI) and machine learning (ML) has revolutionized the way modern systems learn from data, decide, and adapt to changing situations (Shethiya, 2024). From medical diagnosis and investment prediction to self-driving cars and personalized suggestions, AI technologies now enable a wide range of mission-critical applications (Ahmad, Bibi, Talani, & Raja, 2025). But as data sets grow in size and complexity and models grow deeper and more computationally intensive, conventional computing systems increasingly find themselves unable to match the demands for computation, memory, and power efficiency (Rane, Paramesha, Choudhary, & Rane, 2024). The quantum computing, an exciting new frontier with the potential to compute certain types of computation exponentially faster than conventional systems by exploiting quantum phenomena such as superposition, entanglement, and quantum interference (Memon, Al Ahmad, & Pecht, 2024). The combination of quantum computing and traditional machine learning hybrid quantum-classical machine learning is a revolutionary and potentially game-changing approach to designing scalable, high-

performance AI systems able to transcend computational barriers and enhance learning performance across a range of applications (Acampora et al., 2025).

In the last few years, serious attempts have been made to explore the synergy between machine learning and quantum computing (Nguyen, Sipola, & Hautamäki, 2024). Many studies have proposed quantum algorithms for tasks such as classification, clustering, and generative modeling (Lamichhane & Rawat, 2025). Tools such as IBM Qiskit, and Xanadu's PennyLane have facilitated the simulation and, to some degree, execution of hybrid quantum-classical models on actual quantum hardware (Hänninen, 2025). Such explorations demonstrate the ability of quantum machines to accelerate learning procedures by reducing training times or improving generalization in some cases (Cerezo, Verdon, Huang, Cincio, & Coles, 2022). However, current research is still characterized by a number of limitations. Most critically, Noisy Intermediate-Scale Quantum (NISQ) machines suffer from problems such as short coherence times, gate noise, and restricted qubit numbers, which constrain the depth and size of quantum circuits that can be executed reliably (Bharti et al., 2022). In addition, much of the previous effort has been focused on toy datasets or experiments with limited scope, and real-world scalability has not been treated comprehensively. Moreover, too often, there is insufficient appropriate benchmarking against comparable classical models, and it is therefore difficult to estimate the real benefits of hybrid approaches in real-world applications of artificial intelligence.

In order to overcome these limitations and advance the domain of quantum-assisted artificial intelligence, this research introduces a strong framework for Hybrid Quantum-Classical Machine Learning for Scalable AI Applications. The suggested research considers the integration of quantum algorithm like Quantum Variational Classifiers with traditional machine learning processes for the implementation of key function like classification. One of the main points of focus will be to compare the models with their non-classical counterparts, considering not just accuracy and training time but also the use of quantum resources and scalability across different data sizes and model complexities. The ultimate aim of the research is to discover practical, hybrid architectures that can improve current AI systems, reduce computational costs, and enable scalable performance to be implemented in real-world applications in industries like finance, healthcare, and intelligent automation. Through this research, hope to bridge the gap between theoretical possibility and practicality, and thereby pave the way for the development of the next generation of scalable and quantum-enhanced AI technologies.

1.1 Problem Statement

In spite of tremendous advancements in quantum machine learning (QML), practical applications of quantum models continue to be restricted by the limitations of existing quantum hardware most notably in the Noisy Intermediate-Scale Quantum (NISQ) regime. More recent research, demonstrates that QMLs beat their classical counterparts in producing synthetic image data (e.g., MNIST, Fashion-MNIST) (Alessandro Poggiali, 2024). Nevertheless, such models are plagued by limited scalability because they are constrained by quantum noise sensitivity, shallow circuit depths, and limited qubit availability. Furthermore, fully quantum methods tend to lack integration frameworks capable of spanning the rich, mature, and robust capabilities of classical deep learning systems with quantum advantages. Furthermore, the majority of existing quantum-upgraded models are validated only on tiny, idealized datasets or through noise-free simulations, which puts their generalization and practical use into question. There is still a lack of research into scalable, hybrid quantum-classical pipelines that integrate parameterized quantum circuits (PQCs) with classical neural networks for classification, clustering, and generation—particularly within scalable platforms such as TensorFlow

Quantum or PennyLane (Wu, Hu, & Li, 2023). This study seeks to overcome these constraints by suggesting a hybrid quantum-classical machine learning architecture that is able to harness quantum benefits for submodules (e.g., feature extraction, pattern generation), but depend on classical layers for stability and size. By comparing with classical models on various datasets and scalability metrics, this research provides a direction towards practical, noise-tolerant, and scalable AI systems based on quantum-classical synergy.

1.2 Recent Innovation and its Limits

Studies in machine learning that combine quantum and classical methods are now focussing on promising new models such as Quantum Variational Classifiers (QVCs) with a view to faster training times and increased efficiency. Tools like TensorFlow Quantum, PennyLane, and Qiskit make it possible to simulate quantum circuits alongside classical models, and they also allow new methods for feature encoding and optimization to be explored. However, these advancements are primarily held back by the constraints of the current Noisy Intermediate-Scale Quantum (NISQ) devices, which include a small number of qubits, a high level of gate noise, and short coherence periods. These factors greatly reduce the ability to scale and confine experimentation to simple models and small data.

1.3 Research Motivation

This research got motivated by the scaling needs of efficient AI systems complementing the datasets being large and complex in real-life applications. Traditional machine learning surmounts the bottlenecks in computation, mainly in training plus optimizations. Being an exponential speedup candidate, quantum computing offers a complementary way of lifting these inadequacies. This research hybridizes classical and quantum machine learning algorithms to study practically viable solutions that amplify performance while feasible to near-term quantum hardware. Consequently, this opens avenues for scalability, efficiency, and applicability on AI for domains such as healthcare, finance, and automation.

1.4 Significance of the Study

This study is important as it applies hybrid quantum-classical machine learning techniques to overcome current AI's computational boundaries. Capitalizing on quantum benefits in optimization and data representation, the project strives to improve model performance, shorten training time, and enable scalability. It closes the gap between mathematical quantum possibilities and actual AI application, with the promise of revolutionizing high-efficiency computation-intensive applications.

1.5 Key Contributions

- Designed and implemented a hybrid quantum-classical machine learning framework integrating quantum circuits with classical neural architectures.
- Applied QVCs for efficient classification tasks using PennyLane and TensorFlow Quantum.
- Evaluated model scalability across increasing data sizes and quantum circuit complexities to assess real-world applicability.
- Addressed practical challenges in quantum noise mitigation and limited qubit resources using simulation and hybrid strategies.

1.6 Rest of the section

Section 2 reviews related works, Section 3 explains methodology, Section 4 presents results and discussion, while Section 5 concludes the study and outlines future directions for advancing hybrid quantum-classical AI.

2. Related Works

Kiwit et al. (Kiwit et al., 2024) investigated the scalability and robustness against noise in quantum generative learning models, i.e., Quantum Generative Adversarial Networks (QGANs) and Quantum Circuit Born Machines (QCBMs). Their strategy was to contrast QML algorithms with realistic quantum noise conditions, e.g., finite-shot statistical noise and two-qubit depolarizing errors, primarily on trapped ion quantum hardware. The study found that QGANs are more robust against the curse of dimensionality compared to QCBMs, and highlighted the significance of the QUARK framework in QML characterization simplification and acceleration. However, their study also found that while QCBMs are partially noise-resilient, they are constrained by missing noise factors like crosstalk and phase errors. Further, their findings are questionable for various quantum platforms due to hardware-dependent error behaviors. The study highlights the importance of more sophisticated noise models and effective mitigation techniques to improve the scalability and practical realization of QML systems in industrial processes. Sagingalieva et al. (Sagingalieva et al., 2023) proposed a quantum-classical hybrid ML model and a quantum-inspired hyperparameter optimization approach to alleviating the computational burden of image recognition tasks. The approach used a hybrid ResNet model and combined it with tensor train decomposition to optimize the hyperparameters, surpassing the performance of traditional approaches. In car image classification benchmarks, their approach surpassed classical grid search optimization by a large margin: the hybrid model reached a 97% classification accuracy in only 18 iterations, while the classical ResNet34 needed 75 iterations to reach 92%. The results shed light on the promise of quantum-inspired methods in reducing computation and improving the learning process. Unfortunately, the study is mostly limited to a single dataset and a single neural architecture, raising questions about its applicability to other ML problems. It is suggested that in the future, works be done in order to adapt the approach to other types of models and datasets, as well as implement other quantum-enhanced optimization methods that can improve the range of use and scalability of the approach.

Shu, Xu, Yung, and Cui (Shu, Xu, Yung, & Cui, 2024) proposed a Quantum-Classical Generative Adversarial Network (QC-GAN) to address the computational inefficiency of classical GANs in high-dimensional image generation applications. It includes a quantum variational circuit as a component of the generator, supplemented by a one-layer classical neural network, with a classical neural network discriminator. With MindSpore Quantum, QC-GAN was compared with classical GANs on handwritten image generation and demonstrated superior performance in Frechet Inception Distance (FID), superior convergence, and fewer trainable parameters. QC-GAN also surpassed pathGAN, a typical quantum GAN model, which could not generate images larger than 16×16. While the experiments validate the expressive power in quantum-boosted generators, the research also revealed the fundamental limitations, namely the difficulty in hardware implementation using quantum noise, measurement errors, and limited qubit resources. The authors are convinced that hybrid frameworks like QC-GAN will dominate the scalable, efficient AI systems in the future Quantum-for-AI paradigm.

Arthur et al. (Arthur & others, 2022) examined the promise of quantum computing for deep learning by proposing a hybrid quantum-classical neural network in which each neuron is represented by a variational quantum circuit (VQC). Their empirical study

entailed training this architecture on binary classification tasks using simulated and actual quantum hardware. On simulators, the hybrid architecture outperformed isolated VQCs, providing around 10% higher classification accuracy and 20% improved cost minimization, illustrating the advantage of blending quantum circuits in a networked formulation. However, on real quantum hardware, performance was extremely sensitive to qubit and gate constraints, with only shallow models being able to perform reliably. The research identified key challenges in scaling such designs based on hardware noise and limited coherence. Arthur et al. suggested further studies into more sophisticated quantum feature maps, circuit ansätze, and hyperparameter optimization techniques to improve performance. Their research illustrates the promise and existing limitations of scalable quantum deep learning.

He et al (He, Su, Qin, Jin, & Gao, 2025) introduced a new paradigm of QML data augmentation through the integration of QGANs and Hybrid Quantum-Classical Neural Networks (HQCNNs). As the subject of data limitation in QML is less developed, the authors presented general-purpose and task-specific augmentation approaches. The general approach was to improve overall HQCNN performance, while the customized approach dynamically adjusted to sample generation with the goal of reinforcing the network's learning on difficult data classes. Their simulation on the MNIST database indicated that QGANs can outperform classical augmentation techniques as well as classical GANs, with classification accuracy on par with DCGANs using half the number of parameters. The proposed approach not only reduces computational complexity but also improves accuracy, as a measure of QGAN's effectiveness in generating high-quality, diverse data. The work identifies QGANs as a key enabler of scalable, high-performance QML models. Future research should address scalability and diversity of the generated data for tough machine learning problems”.

Wulff et al. (Wulff et al., 2024) met the heavy computational overhead of hyperparameter optimization (HPO) in neural networks with a distributed hybrid workflow that utilized classical and quantum processing. The new system utilizes multi-GPU classical training and quantum-trained support vector regression (QT-SVR), performed on a quantum annealer (QA), to predict the performance results of various hyperparameter configurations. The suggested architecture is designed to accommodate scalable parallel executions with up to 50 GPUs and a QA, hence supporting automated classical-quantum element interaction. Experiments performed on a range of datasets in computer vision, high-energy physics, and natural language processing domains show that this hybrid approach can reduce HPO resource consumption by up to 9% without sacrificing, or even improving, the final model accuracy compared to classical-exclusive HPO approaches. Empowering heuristically guided quantum predictions has enhanced the robustness and efficiency of the HPO process. Open-source release by them enables reproducibility and spurs community activity. Future research avenues include the use of larger gate-based quantum machines and advanced quantum optimization algorithms that can further reduce computational overhead and improve search efficiency.

Alessandro Poggiali (Alessandro Poggiali, 2024) explored the combination of quantum computing and traditional clustering algorithms, focusing specifically on hybrid quantum k-Means algorithms. Three algorithmic forms were built, each leveraging progressively larger amounts of quantum parallelism to accelerate the computationally costly cluster assignment step in classical settings. The key innovation is the use of quantum operations to compute data point-centroid distances in parallel, which reduces complexity and scales better for large datasets. Asymptotic speedup benefits over classical k-Means algorithms

were established by theoretical analysis. Experimental confirmation demonstrated that all variants yielded similar clustering quality when appropriate quantum shot parameters were used. Nevertheless, practical application remains hindered by the limited capabilities of current quantum hardware. Surprisingly, only the q1:1-k-Means variant was acceptable on actual quantum machines, though at the expense of unacceptably high runtimes due to communication overhead. Future study is aimed at removal of post-selection protocols and quantum execution efficiency optimization to enable broader practical application.

Tian et al. (Tian, Tian, Fan, Fu, & Ma, 2025) developed a hybrid quantum-classical generative adversarial model (CAQ) to balance the weaknesses of classical image generation techniques and the limitations of present-day quantum hardware. Combining the use of Wasserstein GAN with gradient penalty (WGAN-GP) in a quantum generator and an adaptive noise input mechanism, the approach greatly improves robustness and output quality. The authors further proposed a remapping approach to convert complex multimodal image distributions into simplified unimodal ones, lowering the learning complexity. Tested on MNIST and Fashion-MNIST datasets, CAQ effectively produces grayscale images on NISQ-era quantum hardware, demonstrating better performance compared to previous QGANs. Although the present discriminator is still classical, future work involves constructing entirely quantum GANs for more complex data modeling and adding color image generation support when quantum resources become better.

Pajuhanfard, Pan, and Sheng (Pajuhanfard, Pan, & Sheng, 2025) examined the performance and limitation of QGANs for high-resolution image enhancement on the Optdigits and Fashion-MNIST (FMNIST) datasets. Their work illustrated QGAN superiority over traditional GANs in FID scores, loss convergence, and runtime, especially in low-resolution grayscale image generation. Nevertheless, the authors noted that existing NISQ-era quantum hardware limits QGAN scalability and performance. They suggested methodological improvements to enhance output resolution and advocated for research into quantum circuit optimization, e.g., complex quantum gate designs and parameterized circuit architectures, to enhance the expressiveness and stability of QGAN training. Their work highlights the imperative of quantum architectural advancements in achieving high-quality image synthesis and deployable QGAN.

Wu, Hu, and Li (Wu et al., 2023) presented MORE, a variational quantum multi-classifier based on correlation and measurement that can address the shortcomings of current QNNs whose primary application is in binary classification. MORE uses a single readout qubit and obtains richer quantum information via quantum state tomography on a basis set consisting of three observables in a 2D Hilbert space. The model uses variational quantum clustering to derive quantum labels from inter-class correlations and conducts quantum label-based supervised learning to project input features onto output labels. With the Qiskit framework implemented, MORE achieved competitive multi-class classification performance in both noiseless and noisy NISQ regimes despite a simple variational ansatz and limited hardware. This paper highlights the promise of utilizing full quantum measurement information and correlation structures in scaling quantum classifiers towards practical multi-class problems within the limitation of today's hardware.


3. Methodology of Hybrid Model Construction and Integration

The methodology adopted is systematic in its workflow through the integration of classical preprocessing and hybrid quantum-classical modeling. The datasets, which have been curated as the Breast Cancer Wisconsin (Diagnostic) dataset, are first gathered and

pre-processed through normalization in order to be compatible with quantum encodings. The classical data are then encoded into quantum states through encoding schemes such as angle encoding for qubit rotations or amplitude encoding for compact high-dimensional representation. After encoding, hybrid task mapping maps traditional learning tasks (classification, clustering, and generative modeling) into quantum circuits, facilitating implementation of Quantum Variational Classifiers (QVCs) and Quantum Convolutional Neural Networks (QCNNs).

Such models utilize parameterized quantum circuits for abstracting features and nonlinear decision-making, optimized by classical optimizers. Concurrently, baseline Convolutional Neural Networks (CNNs) are trained to set baseline benchmarks on scalability, accuracy, and resource usage. Last but not least, hybrid quantum-enhanced classifiers are used with guaranteed strong testing against classical baselines to show improved performance under NISQ limitations. The workflow of proposed framework is given in Fig. 1.

Fig. 1. Workflow of Proposed Framework



3.1 Data Collection

For this study, data collection centers on curated quantum ML datasets from Scikit Machine Learning Repository, which are both highly respected among the quantum computing community for hybrid ML experimentation (William Wolberg, 1993). The Breast Cancer Wisconsin (Diagnostic) dataset contains features for 569 patients. Each patient has 30 numerical values, which are used to classify the tumor into malignant or benign. This dataset collectively support training and testing quantum-enhanced classifiers, under controlled conditions, guaranteeing compatibility with quantum encodings and providing a basis for benchmarking against classical models on scalability, accuracy, and resource consumption.

3.2 Data Preprocessing

For hybrid quantum-classical training, quantum-friendly encoding techniques like amplitude encoding and angle encoding are used, in addition to classical normalization for compatibility of feature scales between models.

3.2.1 Quantum Encoding Techniques

Quantum encoding methods are strategies employed to translate classical information into quantum states in order to be processed within quantum machine learning models.

Angle encoding is a popular method of mapping data onto qubit rotations, while amplitude encoding involves inserting data into quantum state amplitudes allowing efficient handling and representation within quantum circuits.

3.2.1.1 Angle Encoding

It converts traditional features into quantum states through data values for rotating qubits on certain axes. A feature value is represented as a rotation angle (typically on the X, Y, or Z axis) of a qubit. For a feature vector $\vec{x} = [x_1, x_2 \dots \dots x_n]$, the quantum state is given in equation.

$$|\psi\rangle = R_Y(x_1) \otimes R_Y(x_2) \otimes \dots \dots \dots \otimes R_Y(x_n) |0\rangle^{\otimes n} \quad (1)$$

Where, $R_Y(x_i)$ represents rotation gate around the Y-axis for the i-th qubit, $\otimes$ represents tensor product used to apply each rotation gate to a separate qubit, and $x_1, x_2 \dots \dots x_n$ represents classical input features to be encoded into quantum rotations. Angle encoding is space-efficient and hardware-optimal and is commonly used in variational quantum circuits.

3.2.1.2 Amplitude Encoding

It is a quantum data embedding method in which classical information is encoded directly in the amplitudes of a quantum state. This method can represent 2^n classical features in compact form using just n qubits, providing exponential data compression. For a $\vec{x} = [x_0, x_1 \dots \dots x_{2^n-1}]$, the corresponding quantum state is given in equation.

$$|\psi\rangle = \sum_{i=0}^{2^n-1} x_i |i\rangle \quad (2)$$

It is subject to $\sum_{i=0}^{2^n-1} |x_i|^2 = 1$. Where, $\sum_{i=0}^{2^n-1}$ represents summation over all computational basis states and x_i represents amplitude of basis state. Although amplitude encoding is very efficient for high-dimensional data, it is challenging to implement on current noisy quantum hardware because of the complexity of state preparation circuits.

3.2.2 Classical Processing Technique

Classical Processing Technique is a process of converting raw input data into an appropriate format prior to quantum encoding. It comprises operations such as normalization, scaling, and feature transformation in order to maintain numerical stability and compatibility with quantum circuits. Classical preprocessing in the right way improves the efficiency and precision of quantum machine learning models as it maintains data integrity.

3.2.2.1 Normalization

It is a routine classical preprocessing method to scale features to a consistent range to enhance convergence and compatibility in hybrid models. For quantum encoding (particularly amplitude encoding), normalization ensures that the input vector meets quantum state constraints. This scales all input features to the range of [0,1], and hence are appropriate for quantum gate or amplitude vector encoding. It is also important to properly normalize to ensure correct quantum state preparation and stabilize classical training elements such as CNNs or GANs in hybrid pipelines. It is given in equation.

$$x'_i = \frac{x_i - \min(x)}{\max(x) - \min(x)} \quad (3)$$

Where, x'_i represents normalized value, x_i represents original value of the feature, $\min(x)$ represents minimum value in the dataset, and $\max(x)$ represents maximum value in the dataset.

3.3 Problem Formulation and Quantum-Classical Task Mapping

At this phase, the studies aim to define fundamental machine learning task classification so that they can be represented appropriately for hybrid quantum-classical realization. One of each task is chosen on the basis of its computational cost and importance for scalable AI applications. For classification, quantum models like Quantum Variational Classifiers (QVCs) are utilized to separate between classes of data through parametrized quantum circuits. To facilitate hybrid processing, classical neural network layers are recursively translated into quantum circuits, in which classical dense or convolutional layers are substituted or supplemented with parameterized quantum gates. This mapping utilizes encoding schemes (angle or amplitude encoding, for example) for expressing input features in the quantum state space and then trainable unitary operations for learning task-dependent transformations. The context is described as working towards building hybrid models capable of learning effectively under quantum constraints and preserving generalization ability with the long-term goal of enabling scalable, real-world AI deployments.

3.4 Baseline Classical Implementation

To serve as a comparative baseline when assessing the scalability and effectiveness of the postulated hybrid quantum-classical models, this step entails training and implementing traditional classical machine learning model. The chosen technique Convolutional Neural Networks (CNNs) are robust baselines for classification. These models are trained on the same data and tested against the same metrics (accuracy, training time, and resource utilization) so that their quantum-boosted counterparts can be properly compared with them.

3.4.1 CNN for Supervised Classification Tasks

Convolutional Neural Networks are deep learning models extensively employed for classification problems with structured inputs such as images or grid-like inputs. CNNs in this work are employed for constructing a baseline classifier against which quantum variational classifiers (QVCs) are to be compared. The typical form of the output of a CNN classifier is given in equation.

$$\hat{y} = \text{softmax}(W \cdot \phi(x) + b) \quad (4)$$

Where, x denotes input data, $\phi(x)$ denotes feature representation obtained after applying convolution and pooling layers, W denotes weight matrix of the dense (fully connected) layer, b denotes bias vector, $\hat{y}$ denotes predicted class probabilities, and softmax denotes activation function that converts logits to probability distribution across output classes. The model is trained with cross-entropy loss, which measures the difference between predicted and actual distributions. It is given in equation.

$$L_{CE} = - \sum_i y_i \log(\hat{y}_i) \quad (5)$$

Where, y_i denotes ground truth (one-hot encoded) label for class i , and $\hat{y}_i$ denotes predicted probability for class i . The CNN baseline sets the bar on how good traditional architecture can learn discriminative patterns, and this will be compared head-on in terms of efficiency and performance with quantum augmented classifiers.

3.5 Quantum-Enhanced Machine Learning Model Design and Integration

To conceive, execute, and test quantum-enhanced classifiers. This stage employs variational quantum circuits, and quantum convolutional neural networks (QCNNs), within a hybrid quantum-classical system.

3.5.1 Quantum Variational Classifiers (QVCs)

QVCs are quantum-classical hybrid architectures applied to supervised learning. They depend on parametrized quantum circuits (PQCs) to learn nonlinear decision boundaries

in high-dimensional Hilbert space. The classification pipeline consists of encoding classical input data into quantum states, processing through a variational circuit, and measuring expectation values for prediction. It is given in equation.

$$\hat{y}(x) = \text{sign}(\langle \psi(x, \theta) | \hat{O} | \psi(x, \theta) \rangle) \quad (6)$$

Where, $\hat{y}(x)$ is the predicted class label (e.g., -1 or +1), $\psi(x, \theta)$ denotes quantum state after applying parametrized gates, θ denotes trainable parameters of the quantum circuit, and $\hat{O}$ denotes observable (e.g., Pauli-Z operator) used for measurement. The training objective is given in equation.

$$L = \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-y_i \cdot \hat{y}(x_i))) \quad (7)$$

Where, y_i denotes ground truth label for input x_i . Optimization occurs through traditional optimizers (e.g., Adam or SPSA) on the parameters of the quantum circuit.

This architecture is deployed on PennyLane with Qiskit backends and compared to classical CNN classifiers to show quantum learning benefits in terms of parameter efficiency and abstraction of features.

3.5.2 Quantum Convolutional Neural Networks (QCNNs)

QCNNs simulate traditional CNNs by utilizing quantum circuits of layer structure, with entanglement-based filters and quantum pooling. They are appropriate for either feature extraction or classification from structured inputs such as flattened sequences or image patches. The logarithmic scaling of parameters with input size and their applicability to noisy intermediate-scale quantum (NISQ) devices make QCNNs potential candidates for feature learning in compact form. The quantum layer equation is given in equation.

$$|\psi_{out}\rangle = U_{conv}(\theta) |\psi_{in}\rangle \quad (8)$$

Where, ψ_{in} denotes input quantum state, $U_{conv}(\theta)$ denotes parameterized unitary for convolution, and $|\psi_{out}\rangle$ denotes output after layer transformation.

The QCNNs' parameter advices scale logarithmically with input sizes and are thus suitable for NISQ devices, this offers a very prominence advantage for the more compact feature learning paradigm.

4. Results and Discussion

This section presents the experimental results of the Hybrid Quantum Model, including accuracy, loss, ROC curves, and comparative evaluation against classical models, highlighting performance improvements and robustness across metrics.

Fig. 2. Data Visualization




Fig. 2. shows the data preparation stage consisted of downloading benchmark datasets and finally preprocessing the data to be compatible with hybrid models. To allow successful quantum encoding using angle and amplitude encoding, data normalization was used to scale features in the range $[0,1]$. It illustrates the distribution of classes, correspondences of features and dimensional projections, balancing and structuring the data set prior to training a model. These verify that there is no acute imbalance in classes and bring out separability patterns that quantum classifiers can utilise.

Fig. 3. Class Distribution




Fig. 3. shows the frequency of each sample category in the dataset, which is an important quality control before training hybrid models. A balanced distribution means that all the classes contribute equally to the learning process to reduce bias and enhance generalization. On the other hand, significant skew is an indicator of a high resampling requirement or some artificial data creation to avoid overfitting to major classes. In this study, the distribution visualization proves to be sufficiently representative among categories, which ensures fairness when training QVCs and QCNNs. The graph allows classical benchmarking of hybrid quantum-classical models on classical baselines by early validating class balance.

Fig. 4. Feature Correlation Heatmap




Fig. 4. is a visualization of the relationships between the input variables by coding correlation coefficients into color values. Strong positive associations are represented with bright colour, the negative or weak connections are demonstrated with dark colour. This analysis is needed to identify those features that are redundant or have the highest level of dependency which can impair the efficiency of the model. Within the framework of hybrid quantum-classical learning, the heatmap informs the best feature selection and encoding methods to use on quantum circuits since informative features require fewer qubits when few. The graph guarantees the training of QVCs as informative inputs as possible.

Fig. 5. Pairplot of First 5 Features by Class




Fig. 5 illustrates the pairwise relationship among the various classes and helps to understand the separability of the features and the overlap of classes. Scatter plots of two features (with colors assigned to identify classes) are presented in each subplot, and feature histograms are presented in diagonal plots. Separations between classes imply a large discriminative power, and overlaps imply that making decisions at boundaries is difficult and necessitates an intricate model, such as QVCs or QCNNs. This visualization shows the most important features to classification tasks and justifies the importance of hybrid quantum-classical methods to model nonlinear interactions. The pairplot is

therefore a useful exploratory tool that guides both the encoding of data and the design of models.

Fig. 6. Simple Neural Network Learning Curve




Fig. 6 shows the plot of training accuracy versus validation accuracy (or loss) versus the number of epochs, that is, how the model learns. First the accuracy increases quickly as basic patterns are captured by the network, and then there is a plateau as learning levels off. When training and validation curves are very close, the training and validation curves are good, whereas when they are far apart, it means that the model is overfitting. In this study, the curve illustrates the performance of the classical neural networks in their original state before quantum improvements are applied. This graph can be compared with QVC and QCNN curves to understand how much faster and more accurate quantum-classical models can be than the speed and accuracy of the classical approach.

Fig. 7. Confusion Matrix — Simple Neural Network




Fig. 7 is a tabular representation of the classification performance in comparison with the ground truth values of the actual labels predicted. All the cells are used to measure the number of cases that belong to a particular category. The diagonal items depict the samples that are rightly categorized whereas off-diagonal items depict the misclassification. A high level of discriminative capability is reflected in the simple neural network baseline when the proportion of TP and TN are balanced with a small FP and FN. Not only is the overall accuracy assessed, but also the weaknesses associated with each

class are identified, and model training and quantum-enhanced classifier comparison are improved.

Fig. 8. ROC Curve




Fig. 8 shows “the trade-off between True Positive Rate (TPR) and and False Positive Rate (FPR) at different classification thresholds. The diagonal reference line is random guessing and the more curves approaching the upper-left indicates a better model performance. The AUC is used to measure total discriminative power, and the larger the value is near 1, the higher the accuracy. In the case of the simple neural network baseline, the ROC curve shows that the model was able to effectively separate classes, and it can be argued that the model is indeed functional in the binary classification. The visualization plays a critical role in assessing the robustness and comparing the classical baselines to the quantum-enhanced classifiers in hybrid architectures”.

Fig. 9.Simple Neural Network — Evaluation Metrics




Fig. 9 illustrate that baseline neural network shows similar performance scores in all metrics, which forms the benchmark against which to compare the hybrid quantum-classical models. The performance of the simple neural network is summarized by the evaluation metrics graph. Accuracy is used to measure the total correctness of predictions, and precision determines the fraction of positive predictions that were correct, or how resistant to false positives they were. Recall is a measure of sensitivity or the capacity to correctly identify real positives. The F1-score strikes a balance between precision and recall to give a harmonic mean that is especially useful with imbalanced

datasets. It achieves accuracy of 73%, precision of 77%, recall of 81%, and F1-score of 79%.

Fig. 10. Hybrid Quantum Training vs Validation Loss




Fig. 10 depicts the learning pattern of the QCNN with epochs. Training accuracy is plotted showing how the hybrid model generalizes the training data with gradual improvement with epochs. Validation accuracy is a sign of the ability of the model to generalize unknown data and is often used as the main indicator of robustness. Both accuracy curves should ideally ascend steadily and stay close together, signifying proper learning and no overfitting. Overfitting can then be identified when training and validation accuracy show a significant gap. When the trends are parallel, balanced performance and strong generalizability of the hybrid quantum-classical model are then indicated.

Fig. 11. ROC Curve — Hybrid Quantum Model




Fig. 11 shows the balance between the TPR and FPR at different thresholds of classification. A curve nearer the top-left corner signifies stronger discriminative power and shows that the model can properly classify positive and negative classes. Since the diagonal line is that of random guesswork, the elevation of the ROC curve above it is a manifestation of the predictive superiority of the model. The AUC measures general performance: a value close to 1 is a sign of very good classification, and 0.5 is no better than chance. Hence, the ROC analysis verifies the strong ability of the hybrid quantum model for robust classification.


Figure 12. Hybrid Quantum Model — Evaluation Metrics

Fig. 12 displays four important performance measures. These metrics accentuate high performance: an accuracy of 0.98, a precision of 0.96, a recall of 1.00 and an F1 score of 0.98. They also imply that the model is very reliable, its recall performance is better, and predictive metrics are balanced. This demonstrates a stable and confident classification of the model. This demonstrates the strength of the model and its suitability in classification issues in the real world.



Fig. 13. ROC Curve — Classical vs Hybrid Quantum Model

Fig. 13 graphs the discriminative power of both models with respect to thresholds. The Classical Neural Network ROI is the baseline, and the upgrades are the Hybrid Quantum Model ROI. Both models are above the diagonal, which means both models are more predictive than random guesses. The ROC curves are compared via AUC, or cumulative performance, and the higher AUC curve indicates superior classification performance. This is further visualized when the hybrid has a higher AUC; it not only shows greater sensitivity and specificity but also suggests its superiority in detecting intricate data patterns above the classical baseline.

Fig. 14. Comparison of Metrics — Classical vs Hybrid Quantum Model



The performance gap between classical and hybrid models is highlighted in Fig. 14. A classical neural network does well enough to establish a baseline; however, the Hybrid Quantum models outperform this baseline on nearly every metric, indicating better predictive accuracy, an increased number of positive cases identified, and an improved balance between Precision and Recall. “The Classical NN averages a rather unspectacular 0.73 accuracy, 0.77 precision, 0.81 recall, and 0.79 F1-score. The Hybrid Quantum model metrics, on the other hand, are astounding; an accuracy of 0.98, precision of 0.96, recall of 1.0, and F1 score of 0.98 are way ahead of the aforementioned. The hybrid Quantum approach is nearly perfect for the recall metric. In comparison to the classical model, the hybrid quantum approach balances and performs quite well. There is no shadow of a doubt that the quantum hybrid approach is more effective and sturdier for tasks that require precise classification and predictions.”

5. Conclusion and Future Works

This paper introduces a quantum-classical machine learning framework that integrates Quantum Variational Classifiers (QVCs) and Quantum Convolutional Neural Networks (QCNNs) with classical neural networks. The focus is on enhancing the scalability of AI systems. The combined framework's performance outshines that of the traditional classical frameworks, as demonstrated through rigorous experimental evaluation. The hybrid model, for instance, achieved an accuracy of 98%, a precision of 96%, a recall of 100%, and an F1-score of 0.98. On the other hand, the conventional neural network attained only a 73% accuracy, with a precision of 77%, a recall of 81%, and an F1-score of 0.79. Such results validate the hybrid model's potential to enhance classification, generalization, and robustness, all while incorporating NISQ-era quantum devices. However, the two drawbacks persist: (1) dependence on small datasets stemming from the constrained scalability of quantum devices and (2) quantum circuit noise, which limits the consistent deployment of real quantum devices.

In the future, it will be crucial to work on scalability with larger, more varied datasets and on quantum circuit designs that are noise-resilient for practical use. Integration with TensorFlow Quantum and PennyLane will be improved to enable flexible hybrid workflows and automate both quantum and classical optimization pipelines. Other areas include the exploration of generative and clustering hybrid algorithms and the use of newly available error-corrected quantum hardware aimed at the hybrid quantum-classical system. These areas thus lead to the possibility of next-generation AI that is both

computationally efficient and scalable for real-world use in finance, healthcare, and intelligent automation.

References

1. Acampora, G., Ambainis, A., Ares, N., Banchi, L., Bhardwaj, P., Binosi, D., ... others. (2025). Quantum computing and artificial intelligence: Status and perspectives. arXiv Preprint arXiv:2505.23860.
2. Ahmad, S. S. W., Bibi, K., Talani, R. A., & Raja, R. (2025). Investigating the Potential of Collaborative Human-AI Systems to Revolutionize Decision-Making in Areas Like Medical Diagnosis, Financial Forecasting and Strategic Planning. *Social Science Review Archives*, 3(1), 2590–2605.
3. Alessandro Poggiali. (2024). Quantum clustering with k-Means: A hybrid approach. *Theoretical Computer Science*, 992, 114466. <https://doi.org/10.1016/j.tcs.2024.114466>
4. Arthur, D. & others. (2022). A hybrid quantum-classical neural network architecture for binary classification. arXiv Preprint arXiv:2201.01820.
5. Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S., Anand, A., ... others. (2022). Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1), 015004.
6. Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L., & Coles, P. J. (2022). Challenges and opportunities in quantum machine learning. *Nature Computational Science*, 2(9), 567–576.
7. Hänninen, N. (2025). Current State of Hybrid Quantum Computing Software Systems.
8. He, R.-Z., Su, J.-J., Qin, S.-J., Jin, Z.-P., & Gao, F. (2025). QGAN-based data augmentation for hybrid quantum-classical neural networks. arXiv Preprint arXiv:2505.24780.
9. Kiwit, F. J., Wolf, M. A., Marso, M., Ross, P., Lorenz, J. M., Riofrío, C. A., & Luckow, A. (2024). Benchmarking Quantum Generative Learning: A Study on Scalability and Noise Resilience using QUARK. *KI - Künstliche Intelligenz*, 38(4), 379–385. <https://doi.org/10.1007/s13218-024-00864-7>
10. Lamichhane, P., & Rawat, D. B. (2025). Quantum Machine Learning: Recent Advances, Challenges and Perspectives. *IEEE Access*.
11. Memon, Q. A., Al Ahmad, M., & Pecht, M. (2024). Quantum computing: Navigating the future of computation, challenges, and technological breakthroughs. *Quantum Reports*, 6(4), 627–663.
12. Nguyen, T., Sipola, T., & Hautamäki, J. (2024). Machine learning applications of quantum computing: A review. arXiv Preprint arXiv:2406.13262.
13. Pajuhanfard, M., Pan, Z., & Sheng, V. S. (2025). Quantum generative adversarial network for image generation. *The Visual Computer*, 1–15.
14. Rane, N. L., Paramesha, M., Choudhary, S. P., & Rane, J. (2024). Machine learning and deep learning for big data analytics: A review of methods and applications. *Partners Universal International Innovation Journal*, 2(3), 172–197.
15. Sagingalieva, A., Kordzanganeh, M., Kurkin, A., Melnikov, A., Kuhmistrov, D., Perelshtein, M., ... Dollen, D. V. (2023). Hybrid quantum ResNet for car classification and its hyperparameter optimization. *Quantum Machine Intelligence*, 5(2), 1–15. <https://doi.org/10.1007/s42484-023-00123-2>
16. Shethiya, A. S. (2024). Adaptive Learning Machines: A Framework for Dynamic and Real-Time ML Applications. *Annals of Applied Sciences*, 5(1).
17. Shu, R., Xu, X., Yung, M.-H., & Cui, W. (2024). Variational quantum circuits enhanced generative adversarial network. arXiv Preprint arXiv:2402.01791.

18. Tian, Y., Tian, C., Fan, Z., Fu, M., & Ma, H. (2025). Quantum generative adversarial network with automated noise suppression mechanism based on WGAN-GP. *EPJ Quantum Technology*, 12(1), 80.
19. William Wolberg, O. M. (1993). Breast Cancer Wisconsin (Diagnostic) [Data set]. UCI Machine Learning Repository. <https://doi.org/10.24432/C5DW2B>
20. Wu, J., Hu, T., & Li, Q. (2023). MORE: Measurement and correlation based variational quantum circuit for multi-classification. 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), 1, 208–218. IEEE.
21. Wulff, E., Garcia Amboage, J. P., Aach, M., Gislason, T. E., Ingolfsson, T. K., Ingolfsson, T. K., ... others. (2024). Distributed hybrid quantum-classical performance prediction for hyperparameter optimization. *Quantum Machine Intelligence*, 6(2), 59.