Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Comparative Evaluation of Quantum-Classical Hybrid Machine Learning Architectures in Healthcare: A Case Study on Heart Disease Prediction

Vikas Garg

vikas.dcrust@gmail.com

Lingyas Vidyapeeth

Dr. Tapsi Nagpal

Lingyas Vidyapeeth

Dr.SVAV Prasad

Lingyas Vidyapeeth

Research Article

Keywords: Quantum Machine Learning, Heart Disease Prediction, Variational Quantum Classifier, Sampler Quantum Neural Network, Estimator Quantum Neural Network, Hybrid Models

Posted Date: May 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8209219/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Comparative Evaluation of Quantum-Classical Hybrid Machine Learning Architectures in Healthcare: A Case Study on Heart Disease Prediction

First Author[1,2*], Second Author[2,3†] and Third Author[1,2†]

[1*]Department, Organization, Street, City, 100190, State, Country.
[2]Department, Organization, Street, City, 10587, State, Country.
[3]Department, Organization, Street, City, 610101, State, Country.

*Corresponding author(s). E-mail(s): iauthor@gmail.com;
Contributing authors: iiauthor@gmail.com; iiiauthor@gmail.com;
†These authors contributed equally to this work.

Abstract

Heart disease remains to be a primary cause of death globally, requiring precise and timely predictive methodologies. Traditional machine learning methods that depend exclusively on classical clinical variables frequently have restricted predictive accuracy and do not adequately use the intricate nonlinear interactions present in patient data. This work examines the effectiveness of hybrid quantum-classical models for predicting heart disease by combining classical clinical data with quantum-derived representations via Variational Quantum Classifier (VQC), Estimator, and Sampler models. Twelve hybrid configurations, integrating three quantum feature maps (Z, ZZ, Pauli) with four ansatz architectures (RealAmplitudes, TwoLocal, EfficientSU2, PauliTwoDesign), were assessed inside the PennyLane and Qiskit frameworks. The results indicate that the VQC model in PennyLane reached optimal performance with Z–TwoLocal and ZZ–PauliTwoDesign configurations, achieving accuracy and F1-scores of up to 0.95. The Estimator model consistently produced reliable predictions, with Pauli–EfficientSU2 and Z–TwoLocal attaining accuracies of up to 0.91 and F1-scores of 0.92 across several frameworks. The Sampler model demonstrated reasonable performance, with the Pauli–EfficientSU2 configuration achieving an accuracy of 0.65 and an F1-score of 0.73, underscoring the supplementary role of stochastic quantum sampling. The analysis of feature significance indicated that

the quantum-derived feature provided the most robust predictive signal, exceeding classical variables like Age, Blood Pressure, and Cholesterol. These findings demonstrate that meticulous alignment of feature maps and ansatz circuits is essential for optimizing model expressivity, convergence, and predictive effectiveness. In summary, hybrid quantum-classical models, especially the Estimator and VQC architectures, significantly improve heart disease prediction, providing a feasible method for precision clinical decision support and potential incorporation into privacy-preserving, multi-institutional healthcare systems.

Keywords: Quantum Machine Learning, Heart Disease Prediction, Variational Quantum Classifier, Sampler Quantum Neural Network, Estimator Quantum Neural Network, Hybrid Models

1 Introduction

Effective healthcare diagnosis depends greatly on the accurate classification of patient data into disease or non-disease categories. Over the past few years, machine learning (ML) has significantly advanced this process by enabling automated classifiers to detect conditions such as cancer, diabetes, and heart failure from clinical attributes [1]. Traditional ML algorithms—including logistic regression, support vector machines, and neural networks—have proven effective at recognizing patterns in high-dimensional medical datasets, thereby supporting automated decision-making. Nonetheless, clinical datasets are often constrained in size and exhibit class imbalance, where disease-positive cases are far fewer than negatives [2]. Such limitations can hinder the robustness of conventional ML, frequently leading to bias toward majority classes or overfitting when data are scarce. Moreover, medical research often relies on cohorts of only a few hundred samples, making it challenging to build highly generalized classical models [3]. These limitations provide an opportunity to explore alternative computational paradigms that are particularly suitable for the noisy intermediate-scale quantum (NISQ) era. Quantum machine learning (QML) represents one such paradigm, combining the principles of quantum computing with modern ML methods. QML algorithms leverage quantum properties such as superposition and entanglement to embed classical data into high-dimensional Hilbert spaces [4]. Variational quantum circuits, for example, can encode patient features into qubits and transform them through quantum gate operations. Hybrid frameworks such as PennyLane make this feasible by providing automatic differentiation of quantum circuits, enabling joint optimization with classical components. In practice, these quantum neural networks (QNNs) can be trained for classification tasks [5]. Since quantum states intrinsically expand the representational capacity of models, QML has the potential to better capture complex structures in healthcare data and to handle imbalanced datasets more effectively than many classical approaches [6]. Integrating quantum computing with deep learning can thus offer new representational power. Quantum neural networks (QNNs) extend classical multilayer networks by encoding inputs into qubits and leveraging quantum superposition and entanglement[7]. These QNNs can operate in exponentially large Hilbert spaces, providing potentially richer feature

representations. Quantum models exploit superposition and entanglement to express complex patterns in higher-dimensional feature spaces, potentially making them more effective than conventional multilayer perceptrons (MLPs) in tasks with imbalanced or limited datasets. Such advantages are relevant in healthcare, where datasets are often small or noisy [1]. By incorporating quantum layers into deep architectures, researchers aim to build hybrid models that capture subtle biological correlations beyond classical networks. Motivated by these potential benefits, researchers have begun applying quantum-enhanced models to biomedical data to evaluate their performance. Recent investigations highlight the potential of QML in biomedical contexts. [3] introduced two quantum classifiers—a quantum distance model and a simplified quantum-kernel SVM—applied to three clinical datasets, showing comparable performance with classical baselines. Likewise, [8] applied a quantum-inspired SVM to an early-onset colorectal cancer dataset and reported improved mortality prediction compared to a classical SVM. Additional studies in oncology and medical imaging further suggest that combining quantum computing with ML could enhance diagnostic accuracy and support clinical decision-making [9]. Collectively, these findings suggest that QML may provide meaningful advantages, particularly for tasks involving small or imbalanced clinical datasets. At the same time, evaluations have revealed limitations. A systematic review of QML in digital health found no consistent evidence that quantum approaches outperform classical models across diverse scenarios [10]. Many existing studies remain simulation-heavy, employ idealized datasets, or do not account for practical constraints such as noisy hardware, limited sample sizes, and class imbalance. Given the restricted scale of current quantum devices, most clinical experiments still do not demonstrate a clear advantage for QML under real-world conditions [11]. These observations emphasize the importance of further empirical assessments that systematically test QML models on realistic healthcare datasets using advanced hybrid frameworks like PennyLane. Furthermore, QNNs have also been studied for regression-type problems, such as survival prediction, by employing mean-square loss as a training criterion [12]. Although this study focuses on classification, such extensions demonstrate the broader adaptability of QML models to different healthcare prediction tasks.

The present work contributes to this growing area of research by conducting a focused evaluation of QML architectures for heart disease classification using PennyLane. The objectives are:

- To design and train hybrid quantum-classical classifiers, including variational quantum circuits with different feature map and ansatz configurations.
- To apply these models to a standard heart disease dataset and benchmark their performance against classical ML baselines.
- To analyze the influence of factors such as class imbalance and input dimensionality on the classification performance of quantum models.

The contribution lies in providing a detailed and systematic assessment of PennyLane-based QML architectures for healthcare data. Unlike earlier studies that typically emphasize simulations or narrowly scoped experiments, this work investigates practical classification tasks with optimized data-encoding strategies and hybrid architectures.

The findings highlight scenarios—particularly those involving small or imbalanced datasets—where quantum approaches may provide practical benefits. By doing so, the study not only assesses the current utility of QML in healthcare classification but also offers guidance on how quantum-enhanced models quantum approaches may provide practical benefits. By doing so, the study not only assesses the current utility of QML in healthcare classification but also offers guidance on how quantum-enhanced models could be integrated into future medical analytics pipelines.

2 Litearature Review

Advances in quantum machine learning (QML) have generated growing interest in healthcare, particularly for medical image classification, disease diagnosis, and predictive modeling. While the introduction outlined the motivation and objectives of this study, situating it within prior work is essential. Recent studies highlight both the promise and the challenges of applying QML in real-world clinical contexts. The following review summarizes key contributions, models, and frameworks, while also pointing to the gaps that motivate the present research. [13] surveys quantum computing in medicine, explaining fundamentals such as qubits and superposition, and reviewing progress in drug discovery, molecular modeling, genomics, and diagnostics. He highlights how quantum algorithms and QML can accelerate drug discovery, genomic analysis, and personalized therapies, while noting scalability and hardware limitations. The study concludes that hybrid quantum–classical pipelines are the most feasible path in the next decade, pending breakthroughs in fault-tolerant hardware. QC in healthcare can be categorized into two domains: QML for AI-driven diagnostics and quantum computing for data security [14]. Through analysis of 63 papers, the potential of QML for biomedical data processing and quantum cryptography for health information protection is demonstrated. The findings stress the need for scalable algorithms and error-tolerant hardware, advocating hybrid models to bridge theory and practice. Integration with clinical workflows remains largely unexplored, requiring closer collaboration with healthcare practitioners. QML's role in precision medicine and drug discovery is examined [15], describing it as a "game changer" for diagnostics and therapy personalization. Formal verification methods are suggested to ensure algorithm reliability in identifying disease-specific biomarkers, reinforcing the balance between innovation and safety. Regulatory frameworks are expected to evolve in parallel to guarantee trust in QML-driven clinical decisions. QML applications in drug discovery are reviewed [16], focusing on quantum neural networks for property prediction and compound generation. While potential advantages in chemistry simulation are acknowledged, limitations from noise, encoding overhead, and restricted qubit counts are underlined, presenting QML as a gradual accelerator rather than an immediate breakthrough. Future studies are expected to explore hybrid quantum chemistry pipelines that merge classical molecular dynamics with variational circuits. Quantum Image Generative Learning (QIGL), a variational quantum circuit-based generative model, is proposed [17]and tested on X-ray and medical MNIST datasets. QIGL outperforms classical GANs in producing high-quality synthetic images, demonstrating QML's promise for data augmentation and privacy-preserving healthcare applications.

This highlights that QML could be a key enabler for addressing the shortage of annotated medical data. A broad state-of-the-art overview is presented [18] , noting QML's rapid maturation across genomics, diagnostics, treatment optimization, and outcome prediction. Quantum approaches often match classical benchmarks, though most remain proof-of-concept studies, with ethical and technical challenges still unresolved. Benchmarking against strong classical baselines is emphasized as essential before declaring any genuine "quantum advantage." Empirical testing of quantum advantage by comparing QSVMs with classical SVMs on electronic health records was performed [19]. QSVMs outperform classical methods under certain conditions, and a "terrain ruggedness index" is proposed to predict when quantum kernels provide value—an important step toward systematic deployment. This framework suggests that QML should be viewed as problem-specific rather than universally superior. A survey of 94 QML papers (2017–2023) shows that most implementations mirror classical methods such as SVMs and k-nearest neighbors but on quantum hardware [20]. While promising, current quantum devices lack the scalability and speed to deliver consistent healthcare improvements. Standardized datasets and evaluation protocols are recommended to enable fair comparison across studies. QML at the intersection of biology and quantum information is explored [21]. Opportunities exist in genomics, neuroimaging, and behavioral phenotyping, emphasizing interdisciplinary collaboration to translate theoretical quantum speedups into real biomedical advances. QML is positioned as a catalyst for convergence between computational biology and quantum sciences. QML's role in personalized medicine and cell-centric therapeutics is discussed [22]. Quantum optimization and simulation could accelerate cell engineering and protein structure modeling, but practical applications demand targeted algorithm development tailored to biomedical challenges. Without domain-specific customization, QML risks remaining a theoretical construct rather than a clinical tool. Quantum computing in genomics is critically assessed [23] , showing that many proposed quantum advantages (e.g., Grover's search for DNA alignment) vanish under realistic assumptions. Only complex genomic tasks, such as large-scale assembly or marker selection, may benefit in the near term, emphasizing the need for rigorous benchmarking. Overpromising quantum benefits could hinder trust and adoption in biomedical research. Recent focused applications illustrate QML's practical potential. A variational quantum classifier was applied to chronic kidney disease (CKD) diagnosis, integrating sparse autoencoder data augmentation and LASSO feature selection [24]. On a UCI CKD dataset, the model achieved 99.2% accuracy, outperforming traditional classifiers. Similarly, a hybrid quantum–classical convolutional neural network (HQCNN) was developed for brain MRI diagnosis [25]. Tested on a large Kaggle brain dataset, the HQCNN outperformed a conventional CNN model. These advances demonstrate QML's promise for high-dimensional medical data, yet many studies remain single-case demonstrations without standardized benchmarks.

2.1 Research Gap

Although quantum machine learning (QML) has demonstrated potential across several healthcare applications, most studies remain limited to proof-of-concept experiments or small-scale benchmarks without establishing consistent clinical utility. Existing

work often focuses on general domains such as genomics, imaging, or drug discovery, while disease-specific classification tasks, particularly in cardiology, are underrepresented. Reported improvements are frequently based on simulated environments and lack rigorous comparison with strong classical baselines. Moreover, systematic evaluation of different quantum feature maps, ansatz designs, and hybrid architectures remains scarce, leaving their practical influence on model performance unclear. This creates a gap in both methodological clarity and application-specific validation, motivating the present study to conduct a focused assessment of QML for heart disease classification.

3 Background

Through the combination of the flexibility of classical machine learning (ML) and the power of quantum computing, QML is revolutionizing healthcare analytics. The use of patient data to detect and forecast cardiac disease has benefited greatly from the use of traditional machine learning techniques. This information covers physiological parameters such as age, resting ECG, blood pressure, cholesterol, and others. However, the often noisy, correlated, and nonlinear datasets seen in clinical contexts pose significant challenges for conventional algorithms. This makes it challenging to interpret the findings and draw generalizations.Recently, quantum-classical hybrid models have been developed, providing a fresh approach to these issues. These models use special aspects of quantum interference, entanglement, and superposition to represent and manipulate data in high-dimensional Hilbert spaces. This might reduce the amount of time they need to learn while increasing their accuracy and dependability. Quantum circuits can perform changes that would cost traditional systems exponentially and encode multidimensional medical data more effectively. For example, a certain combination of blood pressure, cholesterol, and exercise-induced angina may collectively indicate elevated cardiovascular risk. In the context of heart disease prediction, hybrid QML architectures are able to detect nonlinear interactions among clinical variables. Frequently, these interactions are too delicate or complex for traditional models to interpret. In this section, the fundamental components of these hybrid quantum frameworks are discussed in further depth. The three primary areas of concentration are hybrid models, ansatz, and feature maps. For predictive healthcare applications, each of these components is crucial to the encoding, processing, and integration of quantum information with classical layers.

3.1 Feature Maps

In quantum computing, a feature map transforms classical data into a quantum state, making it possible for the quantum computer to process it through linear transformations in an exponentially large feature space. The transformation can be expressed as:

$$\Phi : x \to |\psi(x)\rangle$$

where

$$x = (x_1, x_2, \ldots, x_n)$$

represents a vector of clinical input features such as age, cholesterol, resting blood pressure, fasting blood sugar, maximum heart rate, and ST depression, and $|\psi(x)\rangle$ represents the corresponding encoded quantum state.

This mapping is the foundation of quantum learning, as it defines how classical clinical attributes are projected into a quantum domain for processing. Unlike classical kernels, quantum feature maps can efficiently capture complex feature correlations through entanglement, effectively expanding the learning space.

Quantum feature maps can be visualized as encoding layers that convert patient data into rotational parameters of qubits. For example, a cholesterol value can be represented as a rotational angle applied to a qubit's state vector. When several such qubits are entangled, their joint state captures interdependencies among medical attributes, providing a richer data representation for predictive tasks.

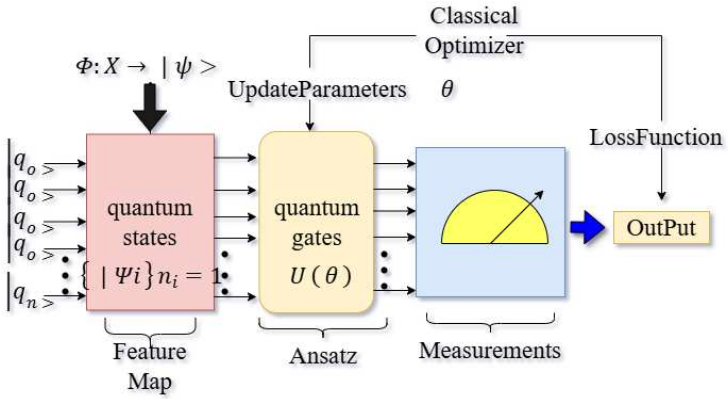


Fig. 1 Architecture of a QML Model

Figure 1 illustrates the fundamental architecture of a QML model. It begins with the transformation of classical data into quantum states, represented as $|\psi_i\rangle$. These states are processed by an ansatz $U(\theta)$, which applies quantum gates such as rotation and entanglement. The model incorporates a classical optimizer that iteratively updates the parameters of the ansatz based on a loss function until convergence is achieved. Finally, the output of the ansatz is measured to provide the model's predictions for various tasks.

3.1.1 Pauli Feature Map

The Pauli Feature Map encodes clinical data using tensor products of Pauli operators (I, X, Y, Z), creating nonlinear and highly expressive quantum representations. It applies quantum rotations that depend on the values of patient features. The encoding can be mathematically represented as:

$$U_\phi(x) = \exp\left(i \sum_{s\subseteq[n]} \phi_s(x) \prod_{i\in s} P_i\right)$$

where P_i denotes a Pauli operator acting on the i^{th} qubit and $\phi_s(x)$ is a real-valued function of classical data.

In heart disease datasets, where certain feature pairs (such as resting ECG and maximum heart rate) exhibit correlated effects, this feature map effectively represents their joint influence. The Pauli Feature Map's ability to form higher-order data interactions improves the classification margin, which is essential for distinguishing between healthy and disease-positive patients.

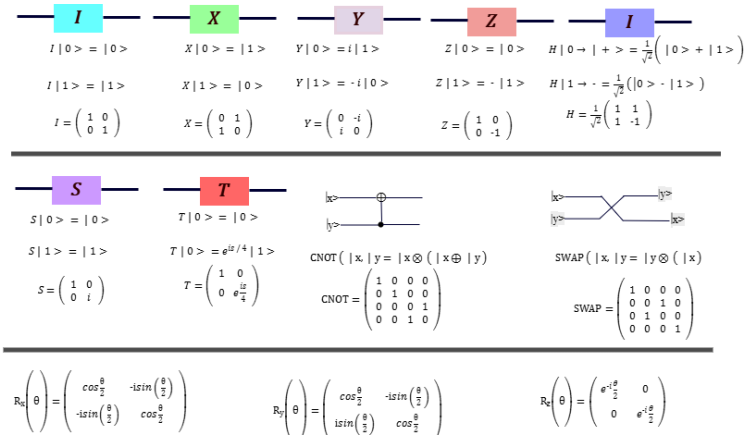


Fig. 2 Graphical representation and mathematical expressions

3.1.2 ZZ Feature Map

The ZZ Feature Map introduces pairwise qubit entanglement through ZZ interactions, capturing feature relationships that are not linearly separable. Its mathematical form can be described as:

$$U_{\text{ZZ}}(x) = \exp\left(i \sum_{i<j} \theta_{ij}\, x_i x_j\, Z_i Z_j\right)$$

This mapping is particularly effective for healthcare applications involving multifactorial dependencies, such as interactions between blood pressure, obesity, and chest pain type. The entanglement produced by ZZ interactions allows the quantum model to implicitly learn feature conjunctions without explicitly engineering them, reducing the dependency on manual feature selection. Furthermore, it allows scalable modeling of large patient datasets, where multiple risk factors interact in nontrivial ways.

3.1.3 Z Feature Map

The Z Feature Map is a simplified form that uses only Z rotations without entanglement. It can be represented as:

$$U_Z(x) = \prod_i \exp\left(ix_i Z_i\right)$$

This approach is computationally lightweight and beneficial for noisy intermediate-scale quantum (NISQ) devices. Although it lacks the complex expressivity of entangled maps, it provides stable and interpretable encodings suitable for smaller datasets or preliminary screenings of heart disease.

In medical contexts, this map is especially effective for categorical or binary attributes, such as sex, fasting blood sugar (> 120 mg/dl), or presence of exercise-induced angina, where interpretability is more critical than raw accuracy.

3.2 Ansatz

The Ansatz defines the structure of a quantum circuit through a sequence of parameterized rotation and entanglement gates. These parameters (θ) are optimized by a classical algorithm to minimize a cost function during training. The ansatz essentially determines how data flows and how much complexity the circuit can express.

In the context of healthcare, the ansatz acts like the quantum equivalent of hidden layers in a neural network, learning to extract abstract health representations. A well-chosen ansatz balances expressivity, trainability, and noise resilience, making it crucial for successful hybrid QML design.

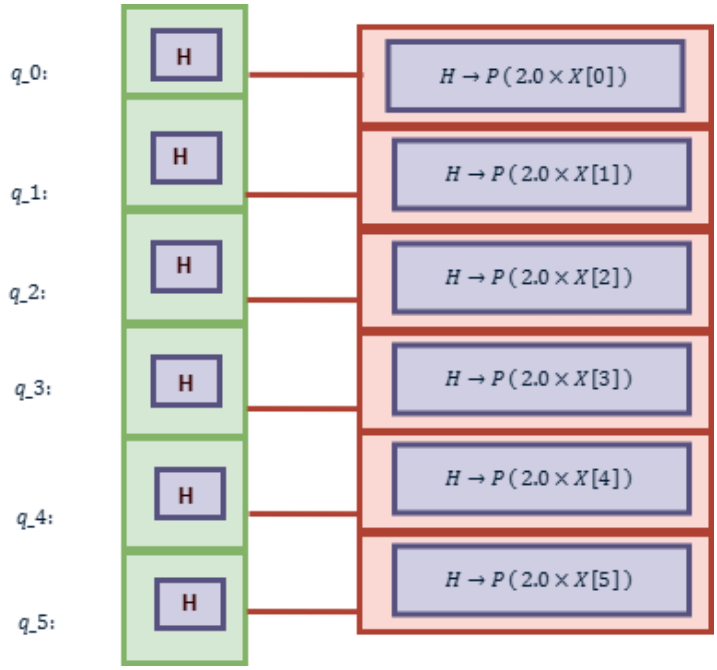


Fig. 3 Z Feature Map

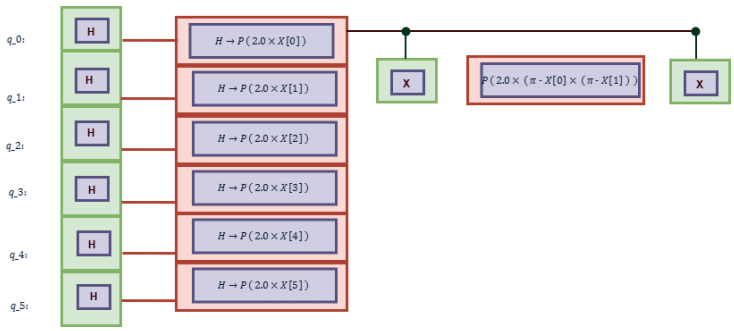


Fig. 4 ZZ Feature Map

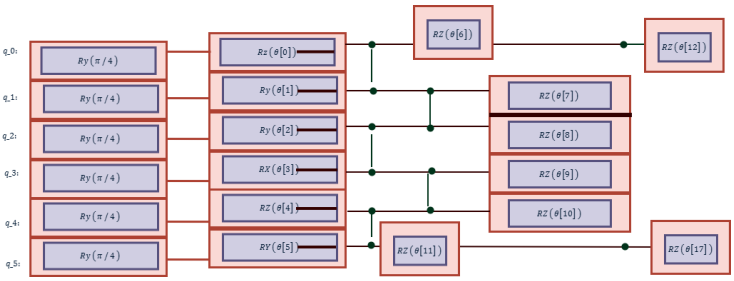


Fig. 5 Pauli Two Design Ansatz

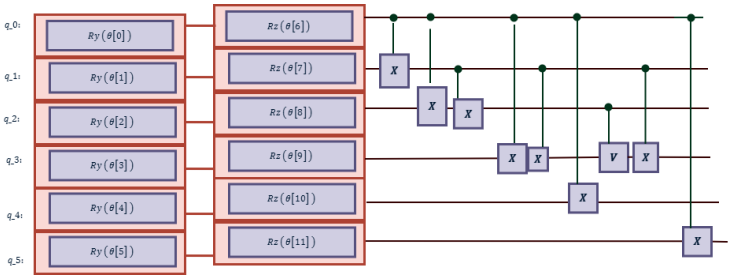


Fig. 6 Efficient SU2 Ansatz

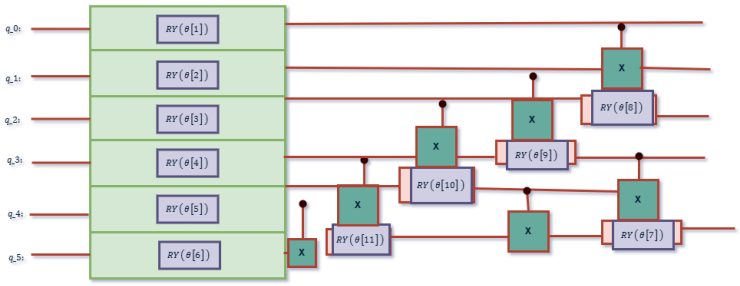


Fig. 7 Real Amplitudes Ansatz

3.2.1 Real Amplitudes Ansatz

The Real Amplitudes Ansatz applies alternating layers of R_y rotations and CNOT gates, producing a real-valued state vector. Its circuit depth is minimal, making it suitable for smaller datasets such as the Cleveland Heart Disease dataset or other regional health surveys with limited samples. It provides a good trade-off between model simplicity and expressivity, helping to avoid barren plateaus—regions where gradients vanish during training. Its robustness against quantum noise makes it ideal for practical NISQ implementations.

$$|\psi\rangle = R_y(\theta_1)\,\mathrm{CNOT}\,R_y(\theta_2)\,|0\rangle$$

3.2.2 Efficient SU2 Ansatz

The Efficient SU2 Ansatz employs alternating layers of SU(2) rotations on each qubit and interleaved entanglement layers. It has high expressive power, capable of learning multi-order dependencies among patient attributes. For example, this ansatz can capture how ST depression interacts with exercise-induced angina and maximum heart rate—factors that often determine cardiac risk severity. The model learns nonlinear transformations efficiently, providing higher classification accuracy compared to shallow ansatz circuits.

$$|\psi\rangle = \left[R_z(\theta_1)R_y(\theta_2)R_z(\theta_3)\right]\,\mathrm{CNOT}\,R_y\,|0\rangle$$

3.2.3 Two Local Ansatz

The Two Local Ansatz uses alternating layers of local rotations and pairwise entanglements. Each layer connects only two qubits, enabling efficient simulation even on NISQ devices. This design offers fine control over entanglement density, allowing the circuit to generalize better across different datasets. In heart disease prediction, where patient data may contain both strongly and weakly correlated features, the Two Local Ansatz allows selective learning of local relationships without overfitting global noise patterns.

$$|\psi\rangle = \left[R_z(\theta_1)\otimes R_z(\theta_2)\right]\,\mathrm{CNOT}\,R_y\,|0\rangle$$

3.2.4 Pauli Two Design Ansatz

The Pauli Two Design Ansatz employs random Pauli rotations and entanglement gates to mimic a unitary two-design. This introduces controlled stochasticity, which prevents the model from memorizing noise or overfitting to limited training samples. This ansatz is particularly beneficial when training with imbalanced medical datasets, where one class (e.g., healthy) heavily outweighs the other (e.g., diseased). The inherent randomness promotes generalization and stability.

$$|\psi\rangle = e^{-i\theta X}\,\mathrm{CNOT}\,e^{-i\phi X}\,|0\rangle$$

3.3 Quantum–Classical Hybrid Models

Quantum–classical hybrid architectures combine the feature extraction capabilities of quantum circuits with the optimization efficiency of classical neural networks. This synergy enables scalable, interpretable, and computationally feasible models for biomedical prediction. In such models, the quantum circuit acts as a feature transformer or kernel, while the classical component performs gradient-based learning and output prediction. This integration reduces the overall number of qubits required while maintaining strong representational power.

3.3.1 Variational Quantum Classifier (VQC)

The VQC is one of the most fundamental hybrid QML architectures. It involves a parameterized quantum circuit whose parameters are optimized through a classical feedback loop. The algorithm proceeds as follows:

1. Encode patient data via a feature map.
2. Apply a parameterized ansatz $U(\theta)$.
3. Measure qubits to obtain output probabilities.
4. Compute a loss function (e.g., cross-entropy).
5. Use a classical optimizer (such as Adam or COBYLA) to update parameters.

Mathematically, the optimization can be represented as:

$$\theta^* = \arg\min_{\theta} J(y, \hat{y})$$

where J is the loss function comparing predicted ($\hat{y}$) and true (y) labels.

VQCs excel in binary classification problems, such as determining whether a patient is likely to have heart disease. They are resource-efficient and provide interpretability through visualization of measurement outcomes and probability distributions.

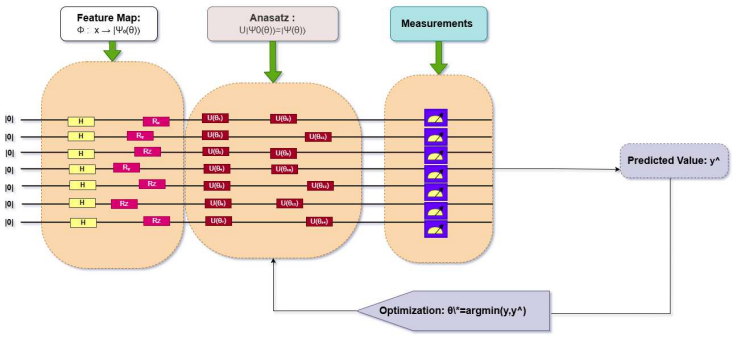


Fig. 8 Architecture of the VQC model

Figure 8: Architecture of the VQC model. It consists of a feature map, after which the operator U is applied, followed by a measurement to obtain the classical output $\hat{y}$.

This process is repeated, where a classical optimizer updates the parameters θ until the loss function $J(y, \hat{y})$ is minimized.

3.4 Estimator Quantum Neural Network (EQNN)

The EQNN integrates quantum circuits as feature estimators within a classical neural network. The quantum layer projects clinical data into high-dimensional latent spaces, producing a set of expectation values that are fed into subsequent classical dense layers. The combined optimization process updates both quantum parameters (θ) and classical weights (W, b) to minimize prediction loss:

$$(\theta^*, W^*, b^*) = \arg \min_{(\theta, W, b)} J(y, \hat{y})$$

The EQNN is especially effective for complex heart disease datasets where correlations extend beyond linear separability. The classical layers refine quantum embeddings into final diagnostic predictions, balancing interpretability and performance.

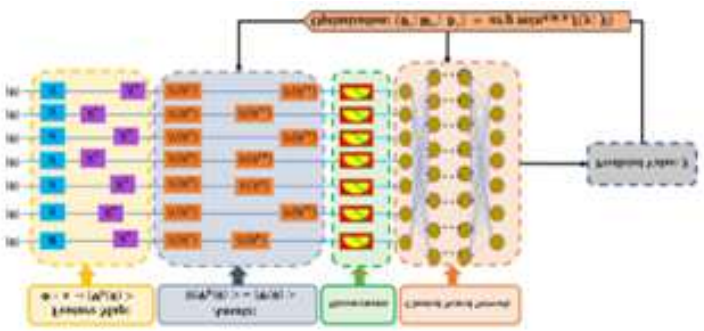


Fig. 9 Architecture of the EQNN model

Figure 9: Architecture of the EQNN model. The process starts with a feature map, followed by an ansatz and measurements to extract classical data. This data is then passed through a fully connected classical neural network, which computes the output value $\hat{y}$. A classical optimizer iteratively updates the parameters θ^*, W^*, and b^* to minimize the loss function $J(y, \hat{y})$.

3.5 Sampler Quantum Neural Network (SQNN)

The SQNN leverages quantum sampling to generate probabilistic outputs that represent patient health patterns. Unlike deterministic classifiers, SQNNs estimate probability distributions, making them ideal for risk stratification and uncertainty quantification in medical diagnosis. During training, the quantum circuit produces samples representing different potential outcomes, and the classical optimizer adjusts circuit parameters to align the output distribution with the true label distribution.

This enables the system to output probabilities reflecting confidence levels in predictions—vital in clinical contexts where uncertainty can inform further medical testing. Mathematically, SQNN optimization can be represented as:

$$\theta^* = \arg\min_{\theta} \sum_i D_{\mathrm{KL}}\big(P_{\mathrm{quantum}}(x_i) \,\big\|\, P_{\mathrm{true}}(x_i)\big)$$

where D_{KL} denotes the Kullback–Leibler divergence between the model's predicted and actual distributions.

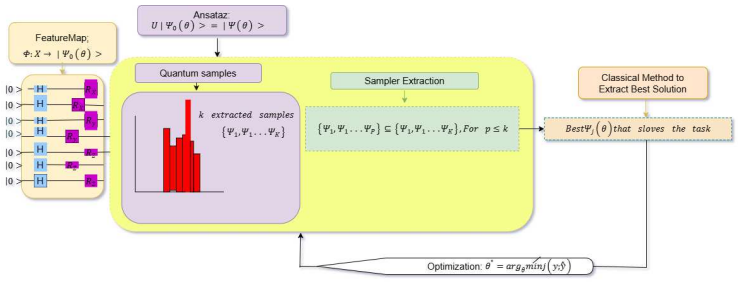


Fig. 10 Architecture of the SQNN model

Figure 10: Architecture of the SQNN model. It consists of a feature map and an ansatz that incorporates a quantum sampler, followed by the extraction of samples generated by the quantum sampler. This section provided an in-depth explanation of the Quantum Machine Learning architectures applied to healthcare analytics, focusing on heart disease prediction. Each model—VQC, EQNN, and SQNN—uses the quantum circuit's ability to encode and entangle medical features, paired with classical optimization techniques, to achieve accurate, interpretable, and uncertainty-aware predictions.

4 Methodology

4.1 Data Collection and Description

The dataset utilised in this study is specifically crafted for heart disease prediction, integrating traditional medical markers with an engineering feature targeted for quantum-based analysis. The dataset comprises 500 patient records, each characterised by a collection of clinical variables often utilised in cardiovascular risk evaluation. These encompass age, gender, blood pressure, cholesterol levels, and heart rate, which collectively represent critical demographic and physiological indicators pertinent to heart disease. The dataset includes a *QuantumPatternFeature*, a bespoke attribute designed to capture complex non-linear correlations, thereby enhancing prediction potential beyond traditional indicators. This function adds a new dimension for experimentation, particularly effective for assessing the performance of classical and quantum machine learning models. The target variable is *HeartDisease*, designated as 0 for absence and 1 for presence of the ailment. The dataset,

titled `Heart Prediction Quantum Dataset.csv`, consists of 500 rows and 7 columns and was obtained from Kaggle (https://www.kaggle.com/datasets/shantanugarg274/heart-prediction-dataset-quantum). This dataset offers a systematic basis for evaluating sophisticated predictive models, feature significance, and the impact of hybrid quantum–classical methodologies in healthcare diagnostics.

4.2 Data Preprocessing

Before model training, the dataset was thoroughly preprocessed to ensure consistency and optimise it for both conventional and quantum learning frameworks. The features were initially segregated from the target variable, *HeartDisease*, and categorical attributes, such as *Gender*, were encoded via one-hot encoding to transform them into a numerical format suitable for machine learning models. The numerical features, namely *Age*, *BloodPressure*, *Cholesterol*, *HeartRate*, and *QuantumPatternFeature*, were standardised to equalise their ranges, thus preventing any single feature from disproportionately influencing the models. For the quantum models, all features were scaled to the range $[0, \pi]$ to conform to the input requirements of quantum circuits, enabling accurate encoding into qubits. Following scaling, the dataset was partitioned into training and testing subsets, utilising 80% of the data for training and reserving 20% for evaluation. This preprocessing procedure ensures that the dataset is suitably prepared and normalised, establishing a robust foundation for both classical and quantum-enhanced predictive modelling while maintaining the integrity of the original clinical data.

4.3 Model Building

This study employs a quantum-enhanced predictive modeling framework integrating three hybrid classifiers—the Variational Quantum Classifier (VQC), the Sampler Quantum Neural Network (SQNN), and the Estimator Quantum Neural Network (EQNN)—to evaluate their effectiveness in predicting cardiac disease. Two quantum computing frameworks, Qiskit and PennyLane, were utilized to construct and simulate these models, thereby ensuring consistency and robustness across implementations.

The Variational Quantum Classifier (VQC) leverages parameterized quantum circuits (PQCs) to map classical clinical features into a high-dimensional Hilbert space. This quantum feature transformation enables the detection of complex, non-linear interactions within patient data that may be challenging to represent in classical feature spaces. The Sampler Quantum Neural Network (SQNN) leverages stochastic quantum sampling through the Qiskit Sampler primitive with finite-shot executions, enabling the model to explore diverse quantum measurement outcomes during optimization. This intrinsic sampling-induced stochasticity functions as a form of regularization, mitigating overfitting and enhancing generalization—particularly beneficial when learning from limited or noisy medical datasets.

The Estimator Quantum Neural Network (EQNN) fuses quantum state estimation with a classical Multilayer Perceptron (MLP). The EQNN utilizes the Estimator primitive in Qiskit (and the `qml.ExpvalCost` interface in PennyLane) to compute expectation values of quantum observables, which are subsequently processed through

the classical MLP for final decision-making. The MLP functions as the decision layer across all three hybrid architectures, enhancing their ability to learn non-linear dependencies and improve convergence during training. This integration amplifies the representational power of quantum classifiers, making them suitable for complex biomedical datasets that exhibit high variability and intricate inter-feature relationships.

Let the classical input vector be:

$$\mathbf{x} = [x_1, x_2, \ldots, x_n].$$

The information is encoded into a quantum state $|\Psi(x,\theta)\rangle$ using a parameterised feature map Φ and ansatz $U(\theta)$:

$$|\Psi(x,\theta)\rangle = U(\theta)\,\Phi(x)\,|0\rangle^{\otimes n}.$$

The expectation value of a measured observable $\hat{O}$, commonly a Pauli-Z operator for each qubit, produces the quantum characteristics q_i used as inputs to the classical MLP:

$$q_i = \langle\Psi(x,\theta)|Z^{(i)}|\Psi(x,\theta)\rangle, \quad i = 1, \ldots, n.$$

The hybrid EQNN model forms the quantum feature vector:

$$\mathbf{q} = [q_1, q_2, \ldots, q_n].$$

These features are passed to the MLP to produce the predicted probability $\hat{y}$ of cardiovascular disease:

$$\hat{y} = \sigma(W_2\,\mathrm{ReLU}(W_1\mathbf{q} + b_1) + b_2)\,,$$

where W_1, W_2 and b_1, b_2 denote the MLP weights and biases, and σ is the sigmoid activation function used for binary classification.

To comprehensively evaluate model performance, each quantum-classical classifier was analyzed under multiple circuit configurations involving three feature maps—*ZFeatureMap*, *ZZFeatureMap*, and *PauliFeatureMap*—and four variational ansatz designs—*RealAmplitudes*, *TwoLocal*, *EfficientSU2*, and *PauliTwoDesign*. This combination yielded twelve distinct hybrid configurations per model, enabling a rigorous comparison of how circuit design impacts classification performance and convergence dynamics.

Both Qiskit and PennyLane backends were employed to construct, train, and evaluate the quantum models. Qiskit's Sampler and Estimator primitives were used for the SQNN and EQNN, respectively, while PennyLane's `qml.qnode` and `qml.qnn.KerasLayer` interfaces facilitated cross-framework experimentation and validation of model behavior. The principal hyperparameters and quantum-circuit configurations used during model training are summarized in Table 1.

Table 1 summarizes the principal configuration parameters used for training the hybrid quantum–classical models (VQC, SQNN, and EQNN). The `use_device` parameter specifies the quantum backend employed for simulation, with `default.qubit`

Table 1 Quantum and Classical Model Hyperparameters

Parameter	Value	Description
use_device	default.qubit	Quantum simulator device
n_qubits	6	Number of qubits used in quantum circuits
ansatz_reps	2	Number of repetitions (depth) of ansatz
shots	1000	Number of measurement samples
epochs	30	Number of training iterations
batch_size	32	Mini-batch size for gradient updates
lr	1e-3	Learning rate for optimizer
hidden_sizes	[64, 32]	Neurons in MLP hidden layers

denoting a noiseless quantum simulator used in both Qiskit and PennyLane frameworks. The `n_qubits` corresponds to the number of qubits in the quantum circuit, set equal to the number of input features (six in this study) to ensure one-to-one feature-to-qubit encoding. The `ansatz_reps` parameter defines the circuit depth or number of repeated parameterized gate layers within the variational ansatz, while `shots` determines the number of measurement samples taken per circuit execution, influencing the statistical stability of expectation values.

The `epochs` parameter specifies the number of complete passes over the training data during optimization, and `batch_size` defines the number of samples per gradient update step. The `lr` (learning rate) controls the step size for parameter updates in the optimizer, balancing convergence speed and stability. Finally, `hidden_sizes` indicates the architecture of the classical MLP used in the hybrid model's decision stage, comprising two fully connected layers with 64 and 32 neurons, respectively. Together, these hyperparameters govern the depth, expressivity, and convergence behavior of the proposed quantum–classical models.

4.4 Model Training

The training procedure for the hybrid Variational Quantum Classifier (VQC), Sampler Quantum Neural Network (SQNN), and Estimator Quantum Neural Network (EQNN)–MLP models involves the joint optimisation of both quantum circuit parameters and classical neural network weights. Each input sample undergoes preprocessing where classical features are normalised to the range $[0, \pi]$ and embedded into the quantum state space through an appropriate feature map, such as *ZFeatureMap*, *ZZFeatureMap*, or *PauliFeatureMap*. These normalised inputs are represented as parameterised quantum states over a 6-qubit system, where each qubit corresponds to one feature dimension.

The quantum circuit comprises two main components:

- **Feature Map Layer**: Encodes classical data into a quantum Hilbert space.
- **Ansatz (Variational Layer)**: A parameterised circuit structure (e.g., *RealAmplitudes*, *TwoLocal*, *EfficientSU2*, or *PauliTwoDesign*) with a configurable depth (ansatz_reps = 2), responsible for learning quantum feature representations.

For VQC, the prepared quantum states are measured using Qiskit primitives (*Sampler* or *Estimator*) or PennyLane interfaces (`qml.qnode` / `qml.qnn.KerasLayer`) to obtain expectation values. In the Sampler Quantum Neural Network (SQNN), stochastic quantum sampling is performed using the Qiskit *Sampler* primitive (with finite-shot execution) or the PennyLane `qml.sample`/`qml.probs` functions, introducing inherent variability in measurement outcomes. This sampling-induced stochasticity serves as a regularization mechanism that mitigates overfitting and enhances generalization. In contrast, the Estimator Quantum Neural Network (EQNN) computes deterministic expectation values through the Qiskit *Estimator* primitive or the PennyLane `qml.expval`/`qml.ExpvalCost` functions, with the resulting quantum-extracted features subsequently passed to a classical Multi-Layer Perceptron (MLP) for final prediction.

The binary cross-entropy loss function used for optimisation is given by:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)\right], \tag{1}$$

where y_i denotes the actual label and $\hat{y}_i$ the predicted probability of heart disease for the i-th sample. Gradient-based optimisation is performed using the Adam optimiser for classical parameters, while quantum circuit parameters are updated via COBYLA or PennyLane-compatible optimizers. Training is conducted for 30 epochs with a mini-batch size of 32, and computations are performed on CUDA-enabled devices when available, defaulting to CPU otherwise.

To monitor convergence, loss values are recorded and plotted every 5 epochs, capturing the optimisation dynamics of both quantum and classical components. The integration of multiple feature maps and ansatz configurations (varying in circuit depth, entanglement type, and rotation gates) allows systematic evaluation of how circuit design impacts learning performance, convergence speed, and measurement stability. Across all hybrid models, the MLP serves as a classical enhancement layer, refining quantum-extracted features to effectively model non-linear interactions among clinical attributes.

4.5 Model Evaluation

After training, each hybrid model is assessed on the reserved test set to measure predictive efficacy. The assessment emphasises categorisation criteria often utilised in healthcare predictive modelling, such as accuracy, precision, recall, and F1-score. Accuracy assesses the overall ratio of correct predictions, whereas precision and recall evaluate the model's efficacy for positive (disease) instances. The F1-score, determined as the harmonic mean of precision and recall, offers a balanced evaluation for datasets exhibiting class imbalance, frequently observed in medical datasets like those related to heart disease. The evaluation metrics are formally defined as follows:

Accuracy: Accuracy represents the proportion of total correct predictions made by the model. It is calculated by summing the true positives (TP) and true negatives (TN), then dividing by the total number of predictions. While widely used, accuracy may not be reliable for imbalanced datasets.

$$\text{Accuracy} = \frac{TP + TN}{S} \tag{2}$$

Precision: Precision measures the proportion of correctly predicted positive cases out of all instances that the model predicted as positive. It reflects the model's ability to avoid false positives, making it crucial in domains where false alarms must be minimized.

$$\text{Precision} = \frac{TP}{TP + FP} \tag{3}$$

Recall: Recall quantifies the model's ability to correctly identify actual positive cases. It is particularly important in medical diagnosis scenarios where failing to detect a positive case (false negative) could have serious implications.

$$\text{Recall} = \frac{TP}{TP + FN} \tag{4}$$

F1-Score: F1-score is the harmonic mean of precision and recall. It provides a balanced measure of a model's accuracy, especially when there is an uneven class distribution or when both false positives and false negatives are critical.

$$F_1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{5}$$

where TP, TN, FP, and FN represent true positives, true negatives, false positives, and false negatives, respectively. All three quantum models—VQC, SQNN, and EQNN—are evaluated over all twelve hybrid configurations, which consist of three feature maps and four ansatz designs. This thorough assessment facilitates the discovery of the optimal model and circuit design, offering insight into the most appropriate quantum-classical architecture for predicting heart disease. Utilising consistent test data for validation guarantees that the results accurately represent real-world generalisation, thereby facilitating repeatability and comparability among hybrid models.

Algorithm 1 Hybrid Quantum–Classical Heart Disease Prediction Methodology

Require: Heart Prediction Quantum Dataset (500 samples, 7 features)
Ensure: Predicted heart disease labels and evaluation metrics

Step 1: Data Preprocessing
Separate features and target variable *HeartDisease*
Encode categorical variables (e.g., Gender) using one-hot encoding
Standardize numerical features (Age, BloodPressure, Cholesterol, HeartRate, QuantumPatternFeature)
Scale features to $[0, \pi]$ for quantum circuit encoding
Split dataset into training (80%) and testing (20%)

Step 2: Model Building
Define three hybrid quantum–classical models: VQC, SQNN, EQNN
Implement and simulate quantum circuits using PennyLane and Qiskit
Encode input vector $x = [x_1, x_2, \ldots, x_n]$ into a quantum state:

$$|\Psi(x, \theta)\rangle = U(\theta)\, \Phi(x)\, |0\rangle^{\otimes n}$$

Measure Pauli-Z observables to obtain quantum features:

$$q_i = \langle \Psi(x, \theta) | Z^i | \Psi(x, \theta) \rangle, \quad i = 1, \ldots, n$$

Feed quantum features $q = [q_1, q_2, \ldots, q_n]$ into a classical MLP:

$$\hat{y} = \sigma(W_2\, \mathrm{ReLU}(W_1 q + b_1) + b_2)$$

Evaluate all twelve configurations (3 feature maps $\times$ 4 ansatz designs)

Step 3: Model Training
for each epoch = 1 to 30 **do**
    for each batch in training data **do**
        Encode features into quantum states
        Compute quantum features via ansatz circuit
        Forward quantum features through MLP
        Compute binary cross-entropy loss:

$$L = -\frac{1}{N} \sum_{i=1}^{N} \left[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \right]$$

        Update quantum parameters θ and MLP weights using Adam optimizer
    end for
end for

Step 4: Model Evaluation
Compute evaluation metrics (Accuracy, Precision, Recall, F1-score):

$$\text{Accuracy} = \frac{TP + TN}{S}$$

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

Identify the best-performing model and configuration

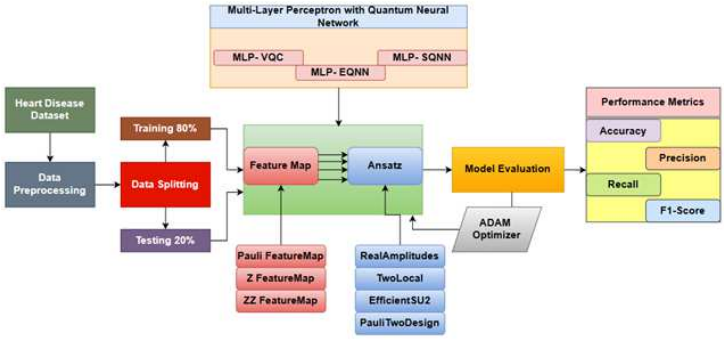


Fig. 11 System Architecture

5 Results and Discussion

The performance of hybrid quantum-classical models for heart disease prediction was evaluated using PennyLane and Qiskit frameworks. Classical clinical features were combined with a quantum-derived feature, and three models—Variational Quantum Classifier (VQC), Estimator (EQNN), and Sampler (SQNN)—were tested across twelve feature map–ansatz combinations. Accuracy, Precision, Recall, and F1-score metrics, along with feature correlations and loss trajectories, were analyzed to assess the impact of quantum feature encoding, ansatz design, and model architecture on predictive performance and convergence.

5.1 Results

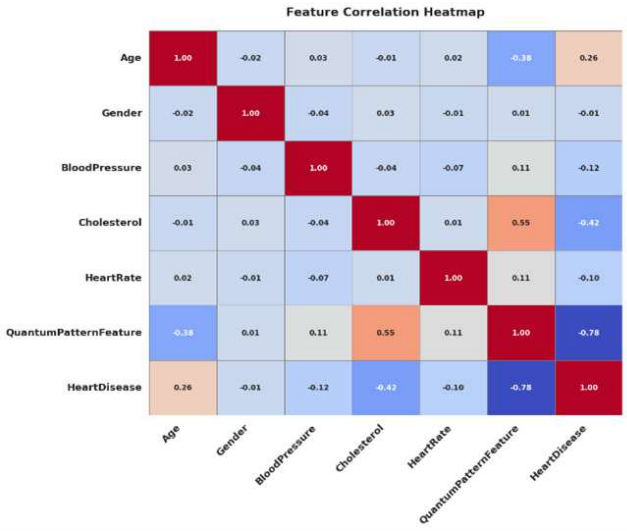


Fig. 12 Correlation Heatmap of input features with heart disease

The associations between the input features and the target variable HeartDisease were analyzed using Pearson correlation coefficients, depicted in a heatmap (Figure 12). The research reveals that Age demonstrates a weak positive association (r = 0.258) with the incidence of heart disease, indicating a slight increase in risk as age progresses. Gender exhibits a minor correlation (r = -0.011), signifying a minimal linear relationship with the target. HeartRate and BloodPressure have weak negative correlations (r = -0.103 and r = -0.117, respectively), indicating slight inverse associations with the occurrence of cardiac disease. Cholesterol exhibits a moderate negative correlation (r = -0.424), either indicating dataset-specific trends or inherent clinical measurement patterns. The QuantumPatternFeature, originating from quantum feature extraction, exhibits a significant negative correlation (r = -0.775), underscoring its considerable predictive capability. The results indicate that whereas classical clinical variables significantly aid in target prediction, the quantum-derived feature offers the most significant signal for hybrid quantum-classical models. Figure 7's heatmap visually corroborates these trends and offers a preliminary insight into feature significance before model training.

5.1.1 PennyLane Framework

Table 2 PennyLane VQC Model Performance

FeatureMap	Ansatz	Accuracy	Precision	Recall	F1-Score
Z	RealAmplitudes	0.92	0.95	0.92	0.93
	TwoLocal	0.94	0.95	0.95	0.95
	EfficientSU2	0.93	0.95	0.93	0.94
	PauliTwoDesign	0.90	0.89	0.95	0.92
ZZ	RealAmplitudes	0.88	0.86	0.95	0.91
	TwoLocal	0.90	0.87	0.98	0.92
	EfficientSU2	0.90	0.88	0.97	0.92
	PauliTwoDesign	0.94	0.95	0.95	0.95
Pauli	RealAmplitudes	0.87	0.85	0.95	0.90
	TwoLocal	0.90	0.89	0.95	0.92
	EfficientSU2	0.92	0.95	0.92	0.93
	PauliTwoDesign	0.91	0.95	0.90	0.92

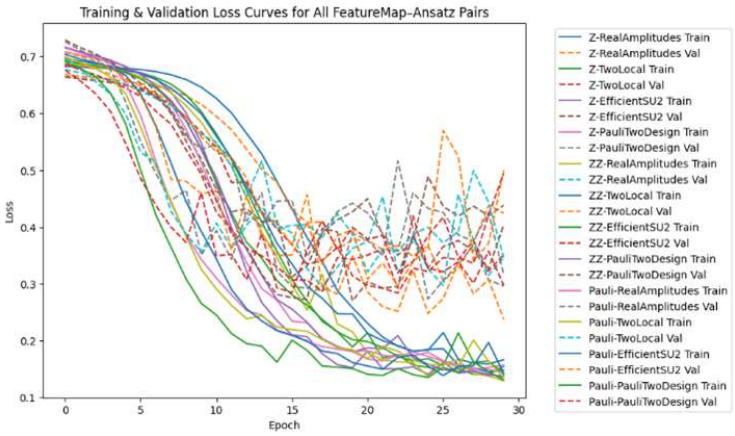


Fig. 13 LOSS vs Epoch plot for PennyLane VQC model

Table 1 displays the efficacy of the hybrid Variational Quantum Classifier (VQC)–MLP model over twelve circuit configurations, which integrate three feature maps (Z, ZZ, Pauli) with four ansatz architectures (RealAmplitudes, TwoLocal, EfficientSU2, PauliTwoDesign). The findings demonstrate that the selection of feature map and variational ansatz significantly affect predictive performance. The Z feature map with the TwoLocal ansatz and the ZZ feature map with the PauliTwoDesign ansatz attained superior performance, with Accuracy, Precision, Recall, and F1-score between 0.94 and 0.95, whereas the Pauli feature map with the RealAmplitudes ansatz demonstrated inferior performance (Accuracy = 0.87, F1-score = 0.90). The loss trajectories underscore the efficacy of hybrid training: specifically, the Z–RealAmplitudes pair diminished training loss from 0.6795 at epoch 5 to 0.1485 at epoch 30, with validation loss fluctuating between 0.6566 and 0.4245. In contrast, the Z–TwoLocal configuration attained marginally quicker convergence, with training loss decreasing from 0.6764 to 0.1547 and validation loss from 0.6491 to 0.3516. The trials were performed on both the Qiskit and PennyLane frameworks, confirming cross-platform stability. Figure 8 illustrates the whole collection of training and validation loss curves for all twelve feature map–ansatz pairs, demonstrating convergence patterns, stability, and the influence of circuit design on optimization dynamics. The results collectively affirm that meticulously designed quantum–classical hybrid models can significantly improve predictive accuracy for cardiovascular illness by utilizing quantum feature transformations in conjunction with classical MLP processing.

Table 3 Quantum Model Performance Across Feature Maps and Ansatz Variants

FeatureMap	Ansatz	Accuracy	Precision	Recall	F1-Score
Z	RealAmplitudes	0.90	0.92	0.92	0.92
	TwoLocal	0.91	0.93	0.92	0.92
	EfficientSU2	0.91	0.92	0.93	0.93
	PauliTwoDesign	0.91	0.92	0.93	0.93
ZZ	RealAmplitudes	0.60	0.60	1.00	0.75
	TwoLocal	0.60	0.60	1.00	0.75
	EfficientSU2	0.74	0.72	0.92	0.81
	PauliTwoDesign	0.72	0.68	1.00	0.81
Pauli	RealAmplitudes	0.66	0.71	0.73	0.72
	TwoLocal	0.67	0.71	0.77	0.74
	EfficientSU2	0.90	0.92	0.92	0.92
	PauliTwoDesign	0.90	0.93	0.90	0.92

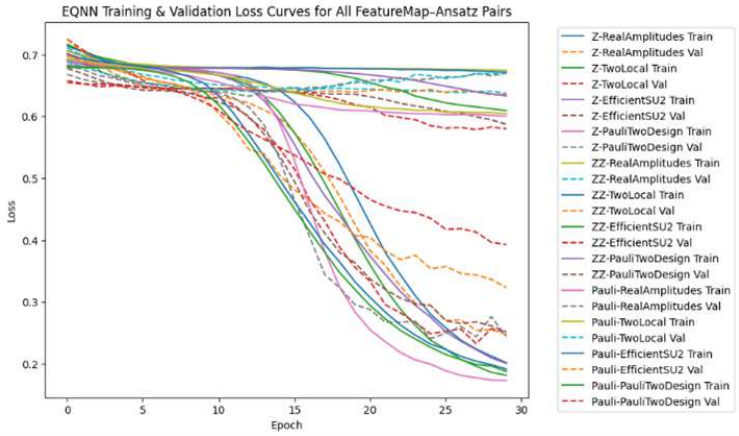


Fig. 14 LOSS vs Epoch plot for PennyLane EQNN model

The Estimator Quantum Neural Network (EQNN) was assessed using three feature maps (Z, ZZ, Pauli) and four ansatz configurations (RealAmplitudes, TwoLocal, EfficientSU2, PauliTwoDesign). Table 2 encapsulates the performance of the tests. The Z feature map regularly demonstrates superior performance, with the Z–TwoLocal configuration attaining the maximum accuracy of 0.91, precision of 0.93, recall of 0.92, and F1-score of 0.92. The Pauli feature map, when integrated with the EfficientSU2 or PauliTwoDesign ansatz, demonstrates robust performance, achieving an accuracy of 0.90, precision of 0.93, recall of 0.92, and an F1-score of 0.92. Conversely, the ZZ feature map has diminished predictive efficacy, attaining a peak accuracy of 0.74 and an F1-score of 0.81, particularly with more straightforward ansätze like RealAmplitudes or TwoLocal.

Figure 9 depicts the training and validation loss trajectories for selected configurations. In the case of Z–RealAmplitudes, the training loss diminishes from 0.685 at epoch 5 to 0.201 at epoch 30, while the validation loss declines from 0.669 to 0.247, corresponding with an enhancement in validation accuracy from 0.68 to 0.90 and an increase in F1-score from 0.81 to 0.92. Comparable trends are noted with Z–TwoLocal, with validation loss diminishing from 0.672 to 0.251 and a final F1-score of 0.92. The ZZ–RealAmplitudes configuration demonstrates slower convergence, yielding final training and validation losses of 0.674 and 0.637, respectively, along with a diminished F1-score of 0.75. The results demonstrate that the selection of feature map and ansatz substantially influences both the convergence dynamics and the prediction efficacy of the EQNN.

Table 4 PennyLane SQNN Model Performance

FeatureMap	Ansatz	Accuracy	Precision	Recall	F1-Score
Z	RealAmplitudes	0.58	0.60	0.93	0.73
	TwoLocal	0.60	0.60	1.00	0.75
	EfficientSU2	0.60	0.60	1.00	0.75
	PauliTwoDesign	0.56	0.59	0.85	0.70
ZZ	RealAmplitudes	0.60	0.60	1.00	0.75
	TwoLocal	0.60	0.60	1.00	0.75
	EfficientSU2	0.60	0.60	1.00	0.75
	PauliTwoDesign	0.59	0.60	0.98	0.74
Pauli	RealAmplitudes	0.61	0.66	0.72	0.69
	TwoLocal	0.64	0.67	0.80	0.73
	EfficientSU2	0.65	0.68	0.78	0.73
	PauliTwoDesign	0.64	0.68	0.77	0.72

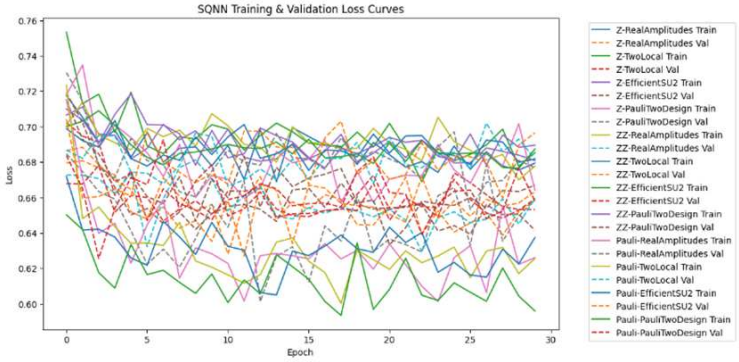


Fig. 15 LOSS vs Epoch plot for PennyLane SQNN model

Table 3 and Figure 10 illustrate the performance and training dynamics of the Sampler Quantum Neural Network (SQNN) across several FeatureMap–Ansatz combinations. The findings indicate that models employing the Pauli feature map consistently surpass the Z and ZZ maps, achieving the highest accuracy (0.65), precision (0.681), recall (0.783), and F1-score (0.729) when paired with the *EfficientSU2* ansatz. The Z and ZZ feature maps yield moderate F1-scores ranging from 0.698 to 0.750.

Figure 10 depicts the training and validation loss trajectories over 30 epochs, demonstrating stable convergence across most configurations, with progressive reductions in both losses, despite minor oscillations in certain cases. This comprehensive evaluation underscores the significant impact of feature map and ansatz selection on SQNN performance, revealing that the Pauli feature map combined with either the *EfficientSU2* or *TwoLocal* ansatz provides the most balanced and reliable prediction capability for heart disease classification.

5.1.2 Qiskit Framework

Table 5 Performance of Different Feature Maps and Ansatz Combinations

FeatureMap	Ansatz	Accuracy	Precision	Recall	F1-Score
Z	RealAmplitudes	0.45	0.46	0.29	0.35
Z	TwoLocal	0.38	0.41	0.43	0.42
Z	EfficientSU2	0.40	0.33	0.14	0.20
Z	PauliTwoDesign	0.50	0.53	0.48	0.50
ZZ	RealAmplitudes	0.50	0.53	0.43	0.47
ZZ	TwoLocal	0.58	0.61	0.52	0.56
ZZ	EfficientSU2	0.60	0.65	0.52	0.58
ZZ	PauliTwoDesign	0.58	0.59	0.62	0.60
Pauli	RealAmplitudes	0.48	0.50	0.52	0.51
Pauli	TwoLocal	0.40	0.41	0.33	0.37
Pauli	EfficientSU2	0.53	0.55	0.52	0.54
Pauli	PauliTwoDesign	0.53	0.56	0.48	0.51

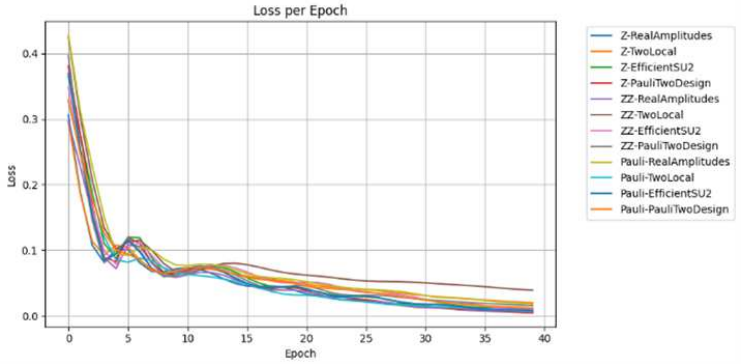


Fig. 16 Qiskit VQCC model LOSS vs Epochs

Table 4 displays the performance metrics of the Variational Quantum Classifier (VQC) executed via the Qiskit framework, assessed across various combinations of FeatureMap and Ansatz. The measures comprise Accuracy, Precision, Recall, and F1-score, whilst Figure 11 illustrates the associated test loss, varying from 0.44 to 0.19. Among the configurations, the ZZ FeatureMap utilizing the EfficientSU2 Ansatz attained the highest overall performance (Accuracy: 0.60, Precision: 0.65, Recall: 0.52, F1-score: 0.58), while the Z FeatureMap with EfficientSU2 demonstrated the poorest predictive performance (Accuracy: 0.40, F1-score: 0.20), underscoring the significance of FeatureMap and Ansatz selection on VQC results. The Pauli FeatureMap had modest outcomes, with its optimal combination (Pauli + EfficientSU2) achieving an accuracy of 0.53 and an F1-score of 0.54. In accordance with these measurements, reduced test losses were noted for the ZZ and Pauli FeatureMaps, although Z FeatureMap configurations typically demonstrated elevated losses. The results indicate that the Qiskit framework effectively facilitates experimentation with quantum feature encoding and circuit ansatz design, enabling systematic assessment of their impacts on classification performance in VQC models.

Table 6 Qiskit EQNN Model Performance

FeatureMap	**Ansatz**	**Accuracy**	**Precision**	**Recall**	**F1-Score**
Z	RealAmplitudes	0.89	0.89	0.93	0.91
Z	TwoLocal	0.60	0.60	1.00	0.75
Z	EfficientSU2	0.91	0.90	0.95	0.93
Z	PauliTwoDesign	0.88	0.85	0.97	0.91
ZZ	RealAmplitudes	0.85	0.87	0.88	0.88
ZZ	TwoLocal	0.60	0.60	1.00	0.75
ZZ	EfficientSU2	0.86	0.84	0.95	0.89
ZZ	PauliTwoDesign	0.89	0.88	0.95	0.91
Pauli	RealAmplitudes	0.87	0.83	0.98	0.90
Pauli	TwoLocal	0.60	0.60	1.00	0.75
Pauli	EfficientSU2	0.91	0.88	0.98	0.93
Pauli	PauliTwoDesign	0.91	0.89	0.97	0.93

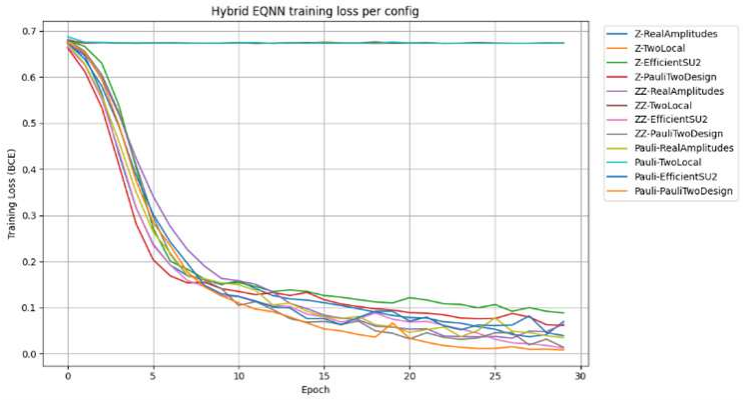


Fig. 17 Qiskit EQNN model LOSS vs Epochs

The Estimator Quantum Neural Network (EQNN) model, based on Qiskit, was trained and assessed using twelve hybrid configurations that integrated three quantum feature maps—Z, ZZ, and Pauli—with four ansatz circuits: RealAmplitudes, TwoLocal, EfficientSU2, and PauliTwoDesign. Each configuration was executed on the CUDA device for 30 epochs, exhibiting stable convergence with consistently decreasing loss values, as depicted in Figure 11, where test losses diminished from approximately 0.67 at initialization to a range of 0.07–0.01 for the optimal configurations. The findings, encapsulated in Table 5, indicate that the EQNN had strong predictive efficacy, with the Pauli-EfficientSU2 and Pauli-PauliTwoDesign configurations attaining peak accuracies of 0.91, precision metrics ranging from 0.88 to 0.89, recall surpassing 0.96, and F1-scores approaching 0.93. Likewise, the Z-EfficientSU2 and ZZ-PauliTwoDesign configurations demonstrated similarly robust performance, achieving accuracies ranging from 0.89 to 0.91 and F1-scores exceeding 0.91, indicating the successful encoding of classical characteristics into quantum Hilbert spaces. Conversely, the TwoLocal ansatz consistently exhibited inferior performance across all feature maps, sustaining a plateaued loss near 0.67 and an accuracy of approximately 0.60, indicating the restricted expressivity of shallow circuit architectures. The EQNN architecture executed in Qiskit exhibited effective convergence and exceptional generalization, especially with deeper and more entangled ansatz configurations like EfficientSU2 and PauliTwoDesign, which offered the ideal balance between circuit depth and predictive precision.

Table 7 Qiskit SQNN Model Performance

FeatureMap	Ansatz	Accuracy	Precision	Recall	F1-Score
Z	RealAmplitudes	0.69	0.70	0.85	0.77
Z	TwoLocal	0.60	0.60	1.00	0.75
Z	EfficientSU2	0.77	0.76	0.90	0.82
Z	PauliTwoDesign	0.78	0.88	0.73	0.80
ZZ	RealAmplitudes	0.66	0.66	0.88	0.76
ZZ	TwoLocal	0.60	0.60	1.00	0.75
ZZ	EfficientSU2	0.77	0.76	0.90	0.82
ZZ	PauliTwoDesign	0.63	0.66	0.80	0.72
Pauli	RealAmplitudes	0.81	0.81	0.90	0.85
Pauli	TwoLocal	0.60	0.60	1.00	0.75
Pauli	EfficientSU2	0.70	0.70	0.87	0.78
Pauli	PauliTwoDesign	0.66	0.67	0.85	0.75

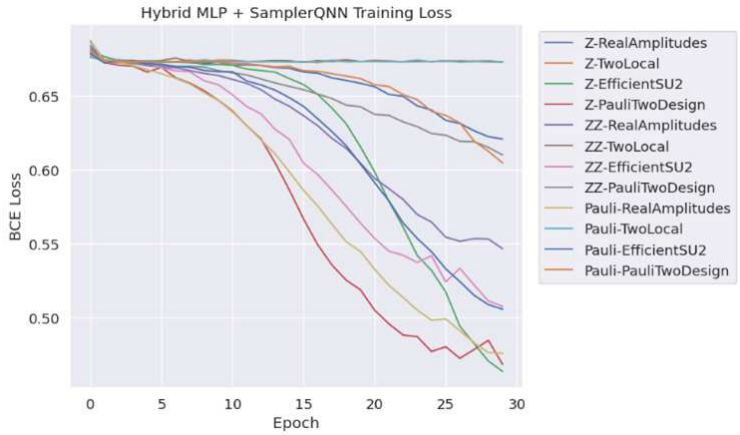


Fig. 18 Qiskit EQNN model LOSS vs Epochs

Table 6 delineates the performance assessment of the Qiskit-based Sampler Quantum Neural Network (SQNN) model for heart disease prediction, evaluated across twelve configurations that integrate various quantum feature maps (Z, ZZ, Pauli) with variational ansatz circuits (RealAmplitudes, TwoLocal, EfficientSU2, PauliTwoDesign). Each configuration underwent training for 30 epochs utilizing the SamplerQNN framework on GPU (CUDA) to evaluate the influence of circuit architecture and feature encoding on classification efficacy. The Pauli–RealAmplitudes configuration attained the highest accuracy (0.81), precision (0.81), recall (0.90), and F1-score (0.85), demonstrating its exceptional ability for expressive quantum feature encoding and effective convergence. Configurations utilizing EfficientSU2 and PauliTwoDesign exhibited robust performance characterized by consistent training and a balanced recall-precision trade-off, whereas the TwoLocal ansatz displayed constrained learning capacity with premature convergence plateaus. Figure 13 illustrates the training convergence curves for certain SQNN configurations, demonstrating a gradual decrease in binary cross-entropy loss over epochs. Configurations like Z–EfficientSU2 and

Pauli–RealAmplitudes exhibited stable and coherent learning dynamics, validating successful parameter optimization. The results highlight that the interaction between feature map design and ansatz topology is crucial in influencing the representational efficiency and predictive robustness of the Qiskit SQNN model.

5.2 Discussion

The comparative assessment of quantum-classical hybrid models—Variational Quantum Classifier (VQC), Estimator, and Sampler—for heart disease prediction reveals substantial discrepancies in predictive efficacy within both PennyLane and Qiskit frameworks. In PennyLane, the VQC model consistently demonstrated superior performance, with the Z–TwoLocal and ZZ–PauliTwoDesign configurations achieving accuracies and F1-scores ranging from 0.94 to 0.95, highlighting the significance of feature map and ansatz selection on circuit expressivity and quantum feature representation. The Estimator model demonstrated significantly diminished still resilient performance, with Z–TwoLocal and Pauli–EfficientSU2 configurations achieving accuracies of up to 0.91 and F1-scores of 0.92, underscoring the efficacy of expectation-value-based encoding for accurate classical feature mapping. The Sampler model exhibited modest performance, with the Pauli–EfficientSU2 ansatz consistently demonstrating the ability to capture variation through stochastic quantum sampling, highlighting its effectiveness in enhancing model generalization for smaller datasets. In the Qiskit framework, the Estimator model again beat both VQC and Sampler configurations, with Pauli–EfficientSU2 and Z–EfficientSU2 obtaining peak accuracies of 0.91 and F1-scores up to 0.93, demonstrating that deeper, highly entangled ansätze boost representational capacity. The Sampler model attained its peak accuracy of 0.81 using the Pauli–RealAmplitudes configuration, indicating that stochastic sampling can proficiently encapsulate significant quantum characteristics, especially when integrated with expressive feature maps. The VQC model in Qiskit exhibited relatively inferior performance, with the ZZ–EfficientSU2 ansatz attaining the maximum accuracy of 0.60, underscoring framework-specific discrepancies and potential variances in circuit simulation quality. Both the Pauli and Z feature maps typically yielded more robust predictive signals than ZZ, perhaps owing to their enhanced capacity for storing classical information into quantum Hilbert spaces. Despite these encouraging findings, the work is limited by the small sample size of the dataset and the significant computational resources needed to train various hybrid configurations, which may restrict scalability and generalization. Nonetheless, the results highlight the essential interaction of feature map design, ansatz architecture, and hybrid quantum-classical integration in attaining dependable predictive performance in clinical applications.

6 Conclusion

This study illustrates that quantum-classical hybrid models, namely VQC, Estimator, and Sampler, can proficiently utilize quantum feature transformations to improve heart disease prediction. Estimator models regularly demonstrated enhanced performance owing to their expectation-value-based quantum encoding, whereas Sampler models offered supplementary insights via stochastic sampling. VQC models, especially inside

the PennyLane framework, attained significant prediction accuracy, highlighting the criticality of ansatz and feature map selection. Future research should examine the scalability of these hybrid structures to bigger, multi-institutional datasets, explore more advanced ansatz designs, and incorporate real-time clinical data for ongoing risk assessment. The suggested techniques have practical significance for clinical decision support, facilitating fast, data-informed assessment of cardiovascular risk. Moreover, the incorporation of federated learning may enable privacy-preserving implementation across healthcare organizations, hence advancing quantum-enhanced predictive analytics in practical medical applications.

References

[1] A. Tudisco, D. Volpe, and G. Turvani, "Quantum Machine Learning in Healthcare: Evaluating QNN and QSVM Models," 2025.

[2] A. Tudisco, D. Volpe, and G. Turvani, "Multi-VQC: A Novel QML Approach for Enhancing Healthcare Classification," 2025, doi: 10.1109/QSW67625.2025.00013.

[3] S. Moradi *et al.*, "Clinical data classification with noisy intermediate scale quantum computers," *Scientific Reports*, vol. 12, no. 1, p. 1851, 2022, doi: 10.1038/s41598-022-05971-9.

[4] S. Naguleswaran, "Machine Learning and Quantum Intelligence for Health Data Scenarios," unpublished.

[5] V. Novák, I. Zelinka, L. Přibylová, and L. Martínek, "Quantum Neural Networks for Propensity Score Estimation and Survival Analysis in Observational Biomedical Studies," 2025.

[6] U. Ullah and B. Garcia-Zapirain, "Quantum Machine Learning Revolution in Healthcare: A Systematic Review of Emerging Perspectives and Applications," *IEEE Access*, vol. 12, pp. 11423–11450, 2024, doi: 10.1109/ACCESS.2024.3353461.

[7] S. Chakraborty, "A Study on Quantum Neural Networks in Healthcare 5.0," unpublished, pp. 1–12.

[8] J. Y. Yu *et al.*, "Evaluation of conventional and quantum computing for predicting mortality based on small early-onset colorectal cancer data," *Applied Soft Computing*, vol. 162, p. 111781, 2024, doi: 10.1016/j.asoc.2024.111781.

[9] M. Rahimi and F. Asadi, "Oncological Applications of Quantum Machine Learning," *Technology in Cancer Research & Treatment*, vol. 22, 2023, doi: 10.1177/15330338231215214.

[10] R. S. Gupta *et al.*, "A systematic review of quantum machine learning for digital health," *npj Digital Medicine*, vol. 8, p. 237, 2025, doi: 10.1038/s41746-025-01597-z.

[11] J. C. L. Chow, "Quantum Computing and Machine Learning in Medical Decision-Making: A Comprehensive Review," *Algorithms*, vol. 18, no. 3, 2025, doi: 10.3390/a18030156.

[12] V. Novák *et al.*, "Quantum Machine Learning for Predicting Anastomotic Leak: A Clinical Study," preprint, 2025.

[13] J. C. L. Chow, "Quantum Computing in Medicine," *Medical Sciences*, vol. 12, no. 4, 2024, doi: 10.3390/medsci12040067.

[14] A. Marengo and V. Santamato, "Quantum algorithms and complexity in healthcare applications: a systematic review with machine learning-optimized analysis," *Frontiers in Computational Science*, vol. 7, pp. 1–30, 2025, doi: 10.3389/fcomp.2025.1584114.

[15] M. Bertl, A. Mott, S. Sinno, and B. B., "Quantum Machine Learning in Precision Medicine and Drug Discovery—A Game Changer for Tailored Treatments?," arXiv:2502.18639, 2025.

[16] A. M. Smaldone *et al.*, "Quantum Machine Learning in Drug Discovery: Applications in Academia and Pharmaceutical Industries," *Chemical Reviews*, vol. 125, no. 12, pp. 5436–5460, 2025, doi: 10.1021/acs.chemrev.4c00678.

[17] A. Khatun, K. Yeter Aydeniz, Y. S. Weinstein, and M. Usman, "Quantum generative learning for high-resolution medical image generation," *Machine Learning: Science and Technology*, vol. 6, no. 2, pp. 1–19, 2025, doi: 10.1088/2632-2153/add1a9.

[18] F. F. Flöther, "The state of quantum computing applications in health and medicine," *Research Directions: Quantum Technologies*, vol. 1, p. e10, 2023, doi: 10.1017/qut.2023.4.

[19] Z. Krunic *et al.*, "Quantum Kernels for Real-World Predictions Based on Electronic Health Records," *IEEE Transactions on Quantum Engineering*, vol. 3, 2022, doi: 10.1109/TQE.2022.3176806.

[20] D. Peral-García, J. Cruz-Benito, and F. J. García-Peñalvo, "Systematic literature review: Quantum machine learning and its applications," *Computer Science Review*, vol. 51, pp. 1–37, 2024, doi: 10.1016/j.cosrev.2024.100619.

[21] P. S. Emani *et al.*, "Quantum Computing at the Frontiers of Biological Sciences," unpublished.

[22] S. Basu *et al.*, "Towards quantum-enabled cell-centric therapeutics," 2023.

[23] A. Maurizio and G. Mazzola, "Quantum computing for genomics: conceptual challenges and practical perspectives," preprint, 2025.

[24] P. Parthasarathi *et al.*, "Variational quantum classifier-based early identification and classification of chronic kidney disease using sparse autoencoder and LASSO shrinkage," *PeerJ Computer Science*, vol. 11, p. e2789, 2025, doi: 10.7717/peerj-cs.2789.

[25] N. Ajlouni *et al.*, "Medical image diagnosis based on adaptive Hybrid Quantum CNN," *BMC Medical Imaging*, vol. 23, p. 126, 2023, doi: 10.1186/s12880-023-01084-5.