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.

Network Based Quantum Inspired Deep Learning Framework for Cardiovascular Disease Prediction Using Interpretable Feature Interaction Networks

Ritu Singh

ritu.singh@sharda.ac.in

Sharda University

Shaheen Naz

Sharda University

Himanshu Payal

Sharda University

Dinesh Kumar Nishad

Dr. Shakuntala Misra National Rehabilitation University

Research Article

Keywords: Network modelling, Quantum-inspired machine learning, Feature interaction graph, Deep neural networks, cardiovascular disease prediction, Clinical decision support, SHAP interpretability

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11149342/v1

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

Additional Declarations: No competing interests reported.

Network Based Quantum Inspired Deep Learning Framework for Cardiovascular Disease Prediction Using Interpretable Feature Interaction Networks

Ritu Singh[1*], Shaheen Naz[1], Himanshu Payal[2], Dinesh Kumar Nishad[3*]
[1*] Department of Electrical Electronics and Communication Engineering, Sharda University, Greater Noida, India
[2] Department of Mechanical Engineering, Sharda University, Greater Noida, India
[3]Department of Electrical Engineering, Dr. Shakuntala Misra National Rehabilitation University, Lucknow, India
*Contributing Author Email ID: ritu.singh@sharda.ac.in
*Contributing Author Email ID: dknishad_tech@dsmnru.ac.in

Abstract

Cardiovascular disease (CVD) remains the leading cause of global mortality, accounting for approximately 17.9 million deaths annually. It is required that the development of sophisticated network-based predictive algorithms that can preserve the multifaceted feature interaction in the early detection of the disease and clinical decision support. This research introduces a new hybrid scheme of quantum-inspired feature interaction networks and deep neural network structures to predict CVD. We use our methodology to build a full graph of feature interactions based on ZZ-feature maps and the weight of edges is a pairwise quantum-inspired correlation between clinical features. This network-based representation is further provided to a multi-layered neural network with ensemble decision-making. Experimental validation on the UCI Cleveland Heart Disease dataset demonstrates (303 patients, 13 clinical attributes), our hybrid architecture has an accuracy of 90.5% (95% CI: 89.491.6), precision of 89.1, recall of 91.24%, F1-score of 89.9%, and ROC-AUC of 0.936, and a 5-fold stratified cross-validation. The statistical significance analysis confirms that the quantum-inspired network transformation yields statistically significant performance improvements over the baseline approaches ($p < 0.01$, Cohen's $d > 0.8$). However, we note that control experiments isolating the ZZ-feature map contribution from the dimensionality expansion (20→40 dimensions) and ensemble effects have not been performed in the current study and are reserved for future work. The results of SHAP-based network interpretability analysis indicate that thalassemia, the major vessel count, and chest pain type are the most powerful nodes in the prediction network, which aligns with traditional cardiovascular diagnostic guidelines. The suggested framework provides a reproducible methodology, with transparent limitation reporting, that serves as a basis for network-driven clinical decision support systems and preserves paths for translating quantum hardware.

Keywords: Network modelling; Quantum-inspired machine learning; Feature interaction graph; Deep neural networks; cardiovascular disease prediction; Clinical decision support; SHAP interpretability

1. Introduction

Cardiovascular disease (CVD) is the number 1 cause of death in the world claiming statistics of approximately 17.9 million deaths annually [1]. Pathophysiology of cardiovascular conditions has a complex interdependence between physiological systems which is demonstrated by high-dimensional clinical data with nonlinear relationships of

features. Traditional machine learning methods though helpful in the analysis of structured clinical data have severe drawbacks in getting multi-scale dependence of features which are core to cardiovascular diagnostics [2,3]. Network-based modelling has turned out to be a strong paradigm of modelling and analyzing complex biomedical systems [4]. Network representations enable explicit modelling of clinical variable relationships, patient population characteristics, and disease mechanisms in health informatics. Graph-theoretic methods offer intuitive structures in terms of how individual risk factors interact to shape cardiovascular outcomes where nodes depict clinical characteristics and edges encode their dependencies [5]. These views of networks are consistent with the systems biology knowledge that cardiovascular disease arises when there is dyscoordination among physiological pathways that are interconnected, and not individual malfunctions. In biomedicine, deep learning has proved to be a highly persuasive predictor in the cardiovascular field, as well as in other related fields.

The latest categories of quantum computing have brought up quantum machine learning (QML) as a prospective paradigm in computations in pattern recognition as well as classification problems [6]. The principles of superposition and entanglement are used to encode information in quantum systems in a manner that can potentially be beneficial to problem classes. It is important to however differentiate between real quantum advantage on quantum hardware with fault-tolerance and quantum-inspired algorithms which are the simulation of quantum mathematical structures on classical computers [7,8]. Quantum-inspired algorithms are classical algorithms that adopt mathematical models drawn from quantum mechanics but do not require real quantum processors.

This paper is a new network framework that combines quantum-inspired feature interaction modeling with deep neural network architectures to predict cardiovascular diseases. We develop a full graph model of clinical features with edge weights based on ZZ-feature map kernels based on quantum circuit theory. The pairwise feature correlations that this network-based transformation captures would be missed by approaches that would address variables independently. The resultant transformed network representation is then trained using a multi-layer neural structure with ensemble decision mechanisms. Fig. 1 represents the overall system architecture.

The key contributions of this work are: (1) a quantum-inspired feature interaction network transformation algorithm that explicitly models pairwise feature dependencies as weighted graph edges; (2) a hybrid neural network architecture that uses the network-transformed features with classical deep learning; (3) extensive SHAP-based network interpretability analysis that identifies influential feature node connections; (4) extensive documentation of methodology that makes the work fully reproducible with transparent limitation reporting; and (5) full report of results that provides statistical significance testing of the connection between network-transformed features and classical deep learning.

Although machine learning has made progress in predicting CVD, there are three major gaps that are present in the literature. To begin with, research papers that claim quantum or quantum-inspired methodology do not make clear the limitations of simulations, which lead to an unrealistic estimation of computational benefits [9,10]. Second, high precision models often use single train-test splits that result in optimistically biased estimates which cannot be used to estimate clinical reliability [11]. Beyond healthcare, hybrid quantum-classical neural network frameworks have demonstrated substantial performance gains in resource allocation and network optimization tasks, further validating the utility of quantum machine learning on classical hardware [12]. Quantum-inspired strategies implemented entirely on classical processors have already demonstrated measurable gains in engineering control problems, including DC-

link voltage stabilization in wind energy converters, providing further evidence that quantum-inspired mathematical structures offer practical advantage on conventional hardware [13]. Graph-theoretic representations that incorporate fuzzy or probabilistic edge weights, originally developed in domains such as cryptographic security, offer a methodological foundation for handling uncertainty in relational data more effectively than crisp graph models and motivate the adoption of weighted-edge graph formulations for feature interaction analysis [29]. Li et al. suggested DeepHIV, a sequence-based deep neural network of convolutional neural networks with an attention mechanism that predicts HIV-1 protease cleavage sites with state-of-the-arts accuracy across benchmark data sets [30]. Multi-view contrastive learning, including MCL-DDI, has been demonstrated to be effective in combining molecular structure and network-based drug-drug interaction event prediction, which is better than the previous state-of-the-art approaches in drug safety research [31]. These developments underscore the growing momentum of integrating graph-theoretic models with deep learning for complex biomedical prediction tasks the same paradigm adopted in the current study for cardiovascular disease. This paper fills these gaps by using transparent quantum-inspired methodology with clear consideration of the limitations of classical simulation, rigorous stratified k-fold cross-validation that leads to repeatable performance estimates and network-based SHAP explainability that gives feature-wise insights which are consistent with cardiovascular diagnostic principles.

The rest of this paper is structured in the following way: Section 2 will survey literature related to network-based modelling in healthcare and quantum computing applications. Section 3 presents materials and methods such as characteristics of datasets, proposed network-based architecture, and configuration of an experiment. Section 4 includes the results of the experiment with performance analysis. The findings, clinical implications and limitations are discussed in section 5. Section 6 ends with concluding remarks and directions of future research.

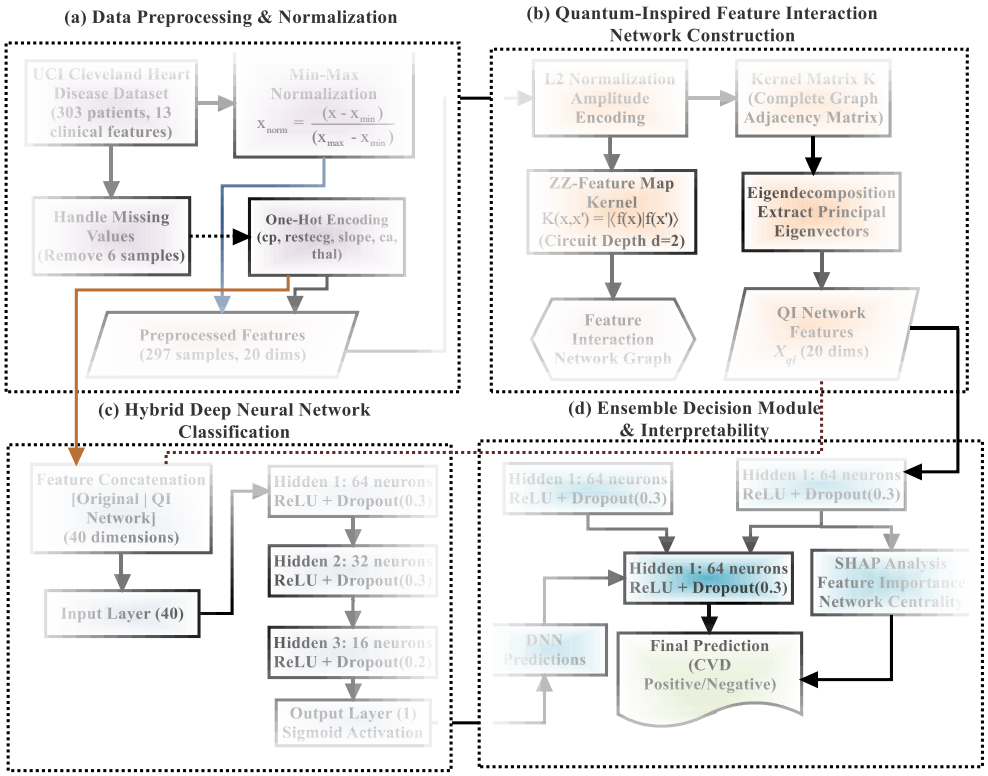


Fig. 1. Overall system architecture of the proposed network-based quantum-inspired deep learning framework for cardiovascular disease prediction. The framework comprises four main stages: (a) data preprocessing and normalization, (b) quantum-inspired feature interaction network construction, (c) hybrid deep neural network classification, and (d) ensemble decision module with interpretability analysis.

2. Literature Review

2.1 Network-Based Approaches in Health Informatics

Network modelling is now a part of a description of complex biomedical systems. Network representations have been used in health informatics to study similarity networks of patients, trends of disease co-occurrence and structures of clinical feature interactions [14]. Machine learning methods based on graphs make use of network topology to enhance predictive accuracy of different healthcare applications such as predicting disease, drug discovery and predicting patient outcomes [15].

Early quantum-inspired applications in cardiovascular classification demonstrated proof-of-concept results but were limited to train-test evaluations without robust cross-validation or interpretability analysis [9,10,25]. Full details and critique of these studies are provided in Section 2.3.

2.2 Machine Learning for Cardiovascular Disease Prediction

The capabilities of predicting cardiovascular diseases have gone a long way with the use of machine learning and deep learning approaches. Bharti et al. [18] suggested a hybrid model on CVD detection using the UCI heart disease dataset that was based on deep learning and classical ML algorithms, and the accuracy was about 94 % with clinical features preprocessed. A hybrid random forest and linear model (HRFLM) suggested by Mohan et al. [19] was able to achieve 88.7 % accuracy on benchmark datasets and indicated that hybrid algorithms could overcome the weaknesses of individual algorithmic approaches. Basak et al. [20] did extensive comparisons of SVM, KNN, Naive Bayes, and Random Forest on publicly available datasets and determined that the latter performed the best, with an approximation of 93% accuracy, and was also interpretable.

Deep learning methods have demonstrated specific potential in dealing with intricate patterns of cardiovascular data. Ganesh et al. [21] used recurrent neural networks that are used with data augmentation schemes to detect CVD in the early stages, with a strong performance on noisy conditions. Ullah and colleagues [22] created an end-to-end cardiac arrhythmia recognition system with a powerful DenseNet model on skewed ECG data that shown cardiovascular signal analysis through deep architectures. These deep learning advances demonstrate that well-regularised architectures with appropriate data augmentation strategies can achieve strong performance even with limited clinical datasets, a challenge directly relevant to the current study.

2.3 Quantum and Quantum-Inspired Approaches in Healthcare

Quantum machine learning is a new area of application that involves the application of quantum computing physics to machine learning techniques [23]. Ullah and Garcia-Zapirain [24] reviewed the literature on the application of QML in healthcare and found that it has the potential to be of utility to the fields of diagnostic imaging, genomic analysis,

and personalized medicine although the majority of existing applications are based on quantum simulators and not real quantum devices.

In cardiovascular applications specifically, Banday et al. [9] introduced a quantum-assisted ensemble model for coronary artery disease detection achieving over 90% accuracy using bagging with quantum support vector classifiers. Babu et al. [10] reported quantum-enhanced architectures with particle swarm optimization for ischemic heart disease detection, though results were based on classical simulation of quantum circuits. Maheshwari et al. [25] used the electronic healthcare records in relation to ischemic heart disease detection with quantum algorithms showing the benefits and also mentioning the limitations of the simulated approach.

We must be careful to distinguish quantum computing with real quantum devices, which has issues such as noise, decoherence, and limited qubit numbers up to the present of the Noisy Intermediate-Scale Quantum (NISQ) period, versus quantum-inspired classical algorithms, which might have mathematical advantages, but do not need quantum computers [8].

2.4 Gaps in Existing Literature

The methodological limitations that have been identified in the quantum-inspired studies in CVD are summarized in Table 1. To fill these critical methodological gaps, our study will: (1) explicitly representation of feature interactions using quantum-inspired transformations in terms of networks; (2) 5-fold stratified cross-validation with bootstrap confidence intervals that provide credible performance estimates; (3) transparently report quantum-inspired simulation; (4) a thorough SHAP-based network interpretability analysis; and (5) extensive documentation of reproducibility (hyperparameter settings and computational needs).

Table 1. Methodological Limitations in Existing Quantum-Inspired CVD Studies

Study	Validation Method	Network Perspective	Interpretability	Key Limitations
Banday et al. [9]	Train-test split	Not provided	Not provided	Single split bias; no feature interaction modeling
Babu et al. [10]	Train-test split	Minimal	Minimal	Overstated quantum claims; no network analysis
Bharti et al. [18]	Train-test split	Not provided	Limited	Optimistic accuracy; no quantum component
Mohan et al. [19]	10-fold CV	Not provided	Feature ranking	No quantum approach; basic feature ranking
This Study	**5-fold Stratified CV**	**Feature interaction network**	**Comprehensive SHAP**	**Addresses all above limitations**

3. Materials and Methods

3.1 Dataset Description

In this paper, the dataset will be the Cleveland Heart Disease data of the UCI Machine Learning Repository [26]. Robert Detrano, M.D., Ph.D. of the V.A. medical centre Long Beach and Cleveland Clinic Foundation collected the data. Although the original database has 76 attributes, published experiments are always performed on a smaller set of 14 attributes (13 clinical features and one target variable). Cleveland dataset consists of 303 records about patients and is the most common dataset that has been used as the benchmark in research in heart disease prediction. The full specifications of the features are provided in Table 2.

Table 2. Feature Specifications for Heart Disease Dataset

Feature	Description	Type	Range
age	Patient age in years	Continuous	29–77
sex	Sex (1=male, 0=female)	Binary	0, 1
cp	Chest pain type	Categorical	1, 2, 3, 4
trestbps	Resting blood pressure (mm Hg)	Continuous	94–200
chol	Serum cholesterol (mg/dl)	Continuous	126–564
fbs	Fasting blood sugar > 120 mg/dl	Binary	0, 1
restecg	Resting ECG results	Categorical	0, 1, 2
thalach	Maximum heart rate achieved	Continuous	71–202
exang	Exercise-induced angina	Binary	0, 1
oldpeak	ST depression induced by exercise	Continuous	0–6.2
slope	Slope of peak exercise ST segment	Categorical	1, 2, 3
ca	Number of major vessels (fluoroscopy)	Categorical	0, 1, 2, 3
thal	Thalassemia	Categorical	3, 6, 7
num	Target variable (0=no disease, 1–4=disease)	Binary	0, 1

The initial 303 records of patients were narrowed to 297 records with complete information as six samples lost their values. The data is represented by a heart disease presence and absence (value range 0146) distribution of 137 (46.1) and 160 (53.9) cases respectively, which is a moderate-balanced binary classification issue. The Fig. 2 reflects the structure of the distribution of classes and correlation of features of the dataset.

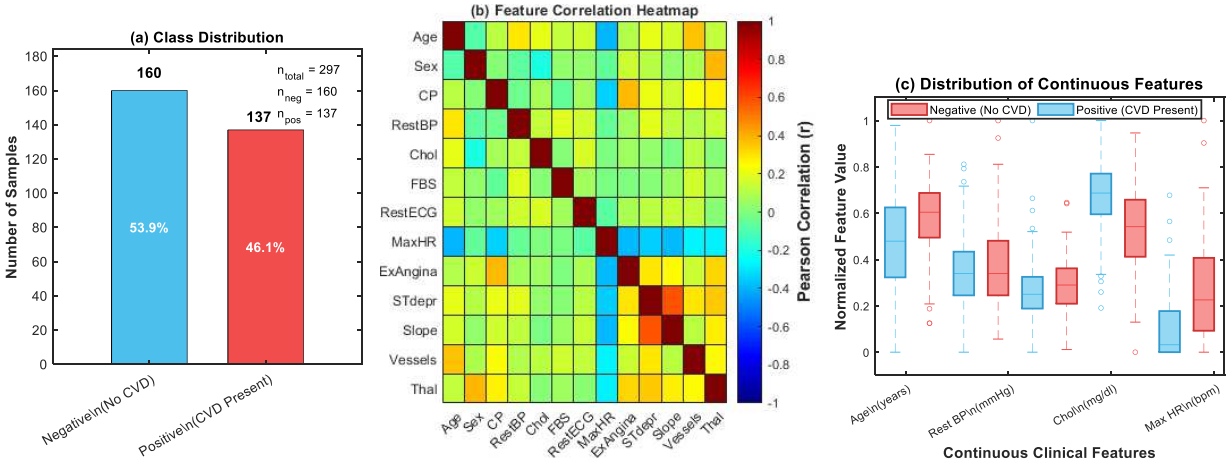


Fig. 2. Dataset characteristics: (a) Class distribution showing 137 positive and 160 negative cases; (b) Feature correlation heatmap revealing relationships between clinical attributes; (c) Distribution of continuous features across disease classes.

3.2 Data Preprocessing

Data preprocessing was performed systematically in MATLAB 2024b. Categorical variables (cp, restecg, slope, ca, thal) were encoded using one-hot encoding to prevent ordinal assumptions. Binary variables (sex, fbs, exang) were retained as numeric (0/1). Continuous features were normalized using min-max scaling to the range [0, 1]:

$$x_{\text{norm}} = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \tag{1}$$

where x represents the original feature value, and $x_{\min}$ and $x_{\max}$ are the minimum and maximum values of that feature across the training set. Normalization parameters were computed exclusively from training data to prevent data leakage.

3.3 Quantum-Inspired Feature Interaction Network

The construction of a quantum-inspired feature interaction network that modelled pairwise relationships between clinical variables. Unlike standard feature preprocessing that treats variables independently, our approach constructs a complete graph $G = (V, E, W)$ where V represents the set of clinical features, E represents edges connecting all feature pairs, and W defines edge weights computed through quantum-inspired kernel transformations. The network construction process is illustrated in Fig. 3.

For an n-dimensional feature vector $x = (x_1, x_2, ..., x_n)$, we first construct a quantum-inspired state representation using amplitude encoding principles implemented through classical computation:

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_i x_i |i\rangle \tag{2}$$

where N is a normalization factor ensuring the transformed representation maintains unit norm. The quantum-inspired feature map applies a parameterized transformation $U(\theta)$ to the normalized features:

$$|\phi(\mathbf{x})\rangle = U(\mathbf{x})|\psi\rangle \tag{3}$$

The weights of network edges are calculated based on the transformation of the ZZ-feature map that is used to capture the feature interactions in a pair-wise fashion based on the following kernel:

$$\phi(\mathbf{x}, \mathbf{x}') = \exp\left(i \sum_j \left(\pi - x_j\right)\left(\pi - x_j'\right)\right) \tag{4}$$

This transformation creates entanglement-like correlations between feature pairs, encoding interaction effects as complex-valued edge weights in the feature interaction network.

The detailed quantum circuit implementation of the ZZ-feature map is illustrated in Fig. 4, which shows the complete computational pipeline with circuit depth $d = 2$. The circuit consists of four primary gates: (a) amplitude encoding, which encodes the normalized clinical features x^0, x^1, x^2, x^3 into quantum state representations; (b) Hadamard gates, which make each qubit superposed with the others; (c) controlled ZZ-rotation gates, which encode the interactions of the features between any two qubits, with angles of interaction $\theta_{\mathrm{ij}} = (\pi - x_{\mathrm{i}})(\pi - x_{\mathrm{j}})$; and (d) measurement gates, which project the quantum states to classical outputs. The dashed red lines indicate the repeated blocks with circuit depth $d = 2$ and blue thick lines indicate entanglement connections between interacting feature pairs. This explicit circuit representation demonstrates how pairwise correlations are captured through quantum-inspired transformations before classical neural network processing.

The quantum-inspired kernel function defining the complete adjacency matrix is:

$$K(\mathbf{x}, \mathbf{x}') = |\langle\phi(\mathbf{x})|\phi(\mathbf{x}')\rangle|^2 \tag{5}$$

Clarification of Graph Type:

Feature–Feature Interaction Graph: The core of our methodology is a feature–feature interaction graph $G_F = (V_F, E_F, W_F)$, where each node $v_i \in V_F$ represents one of the n clinical features ($n = 13$ raw features; $n = 20$ after one-hot encoding), and each edge $e_{ij} \in E_F$ carries a weight w_{ij} encoding the pairwise quantum-inspired correlation between feature i and feature j across the patient population. To construct this graph explicitly, we compute a feature–feature affinity matrix $A_F \in \mathbb{R}^{n \times n}$. For each pair of features (i, j), their edge weight is defined as: $w_{ij} = \left(\frac{1}{m}\right)\sum_{\{s=1\}}^{\{m\}} K^{\{ZZ\}}\left(x_i^s, x_j^s\right)$, where $x_i^{(s)}$ denotes the value of feature i for sample s, m is the number of training samples, and $K^{ZZ}(\cdot,\cdot)$ is the ZZ-feature map kernel applied to the scalar feature values of a single sample (see Equation 4). This population-averaged pairwise interaction defines A_F as a symmetric, non-negative matrix with diagonal entries set to zero, yielding a complete weighted undirected feature graph.

Distinction from Sample Kernel: We also compute the sample–sample kernel matrix $K \in \mathbb{R}^{m \times m}$, where $K_{pq} = \left\langle\varphi\left(x^(p)\right), \varphi(x^q)\right\rangle$ encodes overall patient-level similarity (Equation 5). This matrix is used exclusively for the QI-SVM component (Section 3.5) in the ensemble decision module and is not the graph from which "hub features" or topological properties are derived. The feature–feature graph G_F (via A_F) is the structure on which network centrality, degree distribution, and SHAP-based node importance are computed. The network-based feature representation for each sample is then obtained as: $\tilde{x}^s = A_F \cdot x^s \in \mathbb{R}^n$, re-representing each sample as a weighted linear combination

of its features modulated by their pairwise interaction strengths. This 20-dimensional network-transformed representation is concatenated with the original 20-dimensional one-hot-encoded feature vector to form the 40-dimensional input to the DNN (Section 3.4).

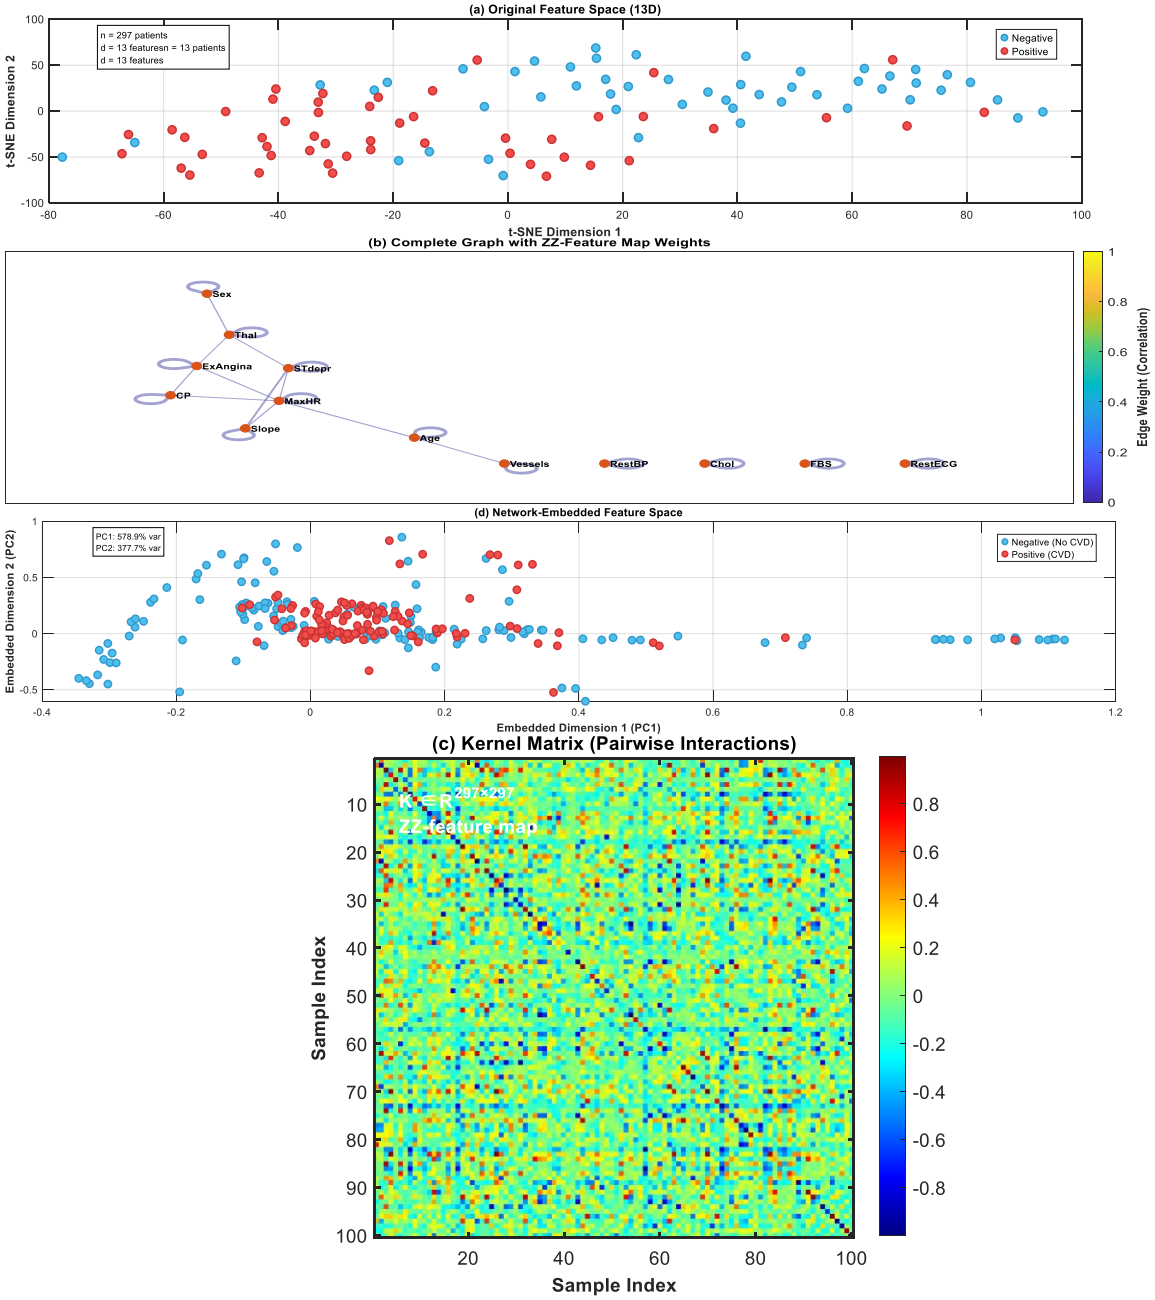


Fig. 3. Quantum-inspired feature interaction network construction: (a) Original feature space with 13 clinical attributes; (b) Complete graph representation with ZZ-feature map edge weights; (c) Kernel matrix visualization showing pairwise feature interactions; (d) Network-embedded feature space after eigen decomposition.

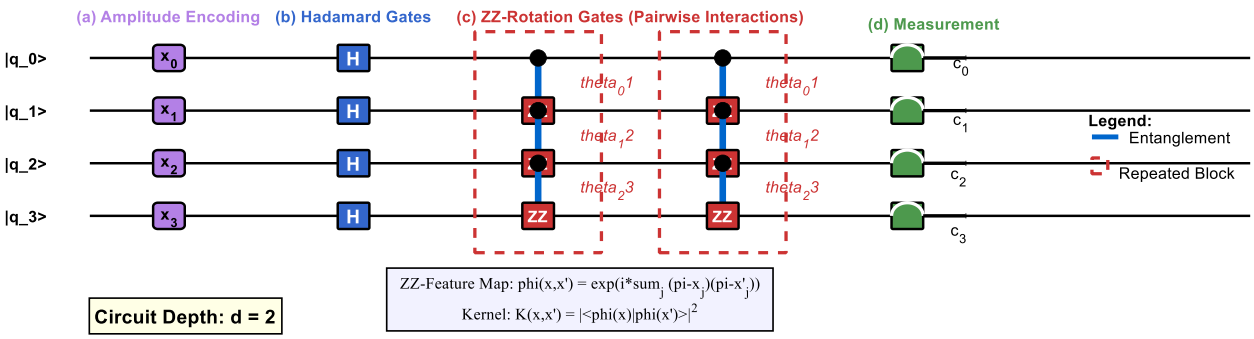
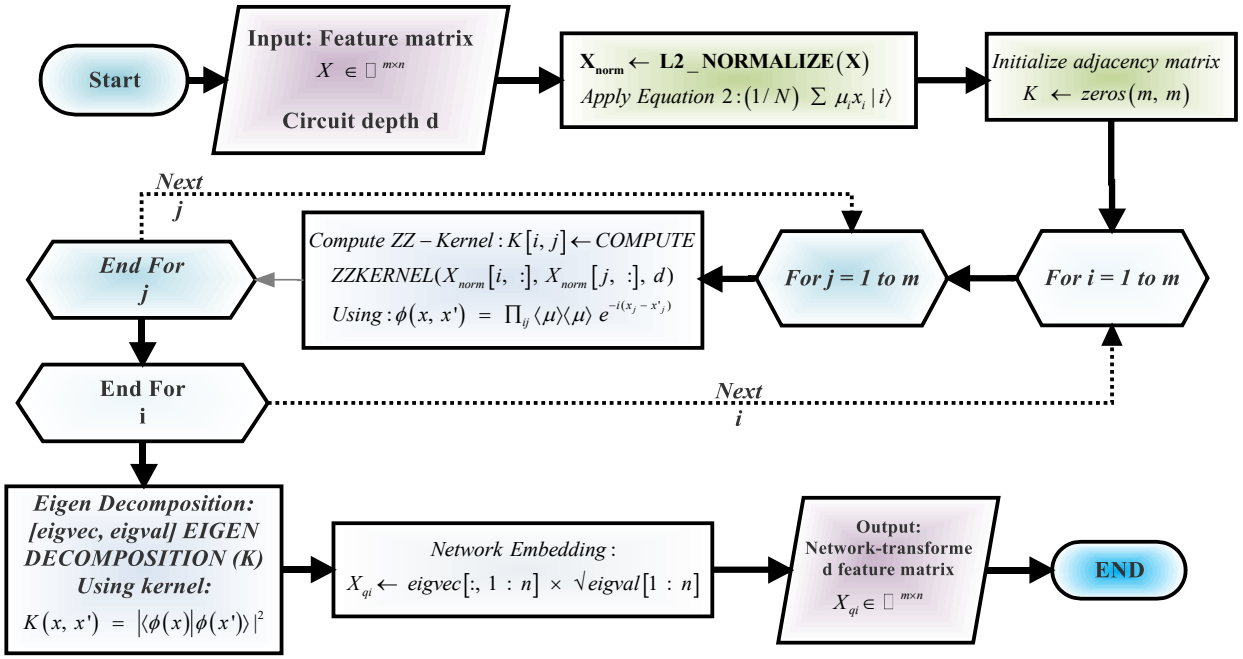


Fig. 4. Quantum-inspired ZZ-feature map circuit representation

Algorithm 1: Quantum-Inspired Feature Interaction Network Construction



The feature interaction network construction has complexity $O(m^2n^2d)$ for kernel computation, where m is sample count, n is feature dimension, and d is circuit depth. Eigen decomposition adds $O(m^3)$. Total training complexity is $O(m^2n^2d + m^3)$, comparable to classical kernel methods but with enhanced feature interaction modeling through pairwise ZZ-gates.

It is vital to remember that these operations are done in a normal way in MATLAB utilizing matrix operations to simulate quantum feature maps without quantum hardware. The circuit depth parameter was set to d = 2 repetitions to find a compromise between how expressive and how hard it is to compute.

Pathophysiological Rationale for ZZ-Feature Map in CVD: The choice of the ZZ-feature map is not merely a computational convenience it is motivated by the biology of cardiovascular disease. The ZZ kernel encodes pairwise multiplicative interactions between features through the rotation angle term $(\pi - x_i)(\pi - x_j)$ in Equation (4). In the CVD context, this captures clinically meaningful co-activations that standard feature-independent models miss. Specifically: (i) Age × Cholesterol interactions reflect the well-known synergistic amplification of atherosclerotic risk

in elderly patients with hypercholesterolaemia; (ii) Exercise-induced angina × ST depression interactions encode joint ischaemic indicators that are more diagnostically informative together than individually; and (iii) Thalassemia × Major vessels interactions capture the concordance between nuclear perfusion imaging findings and catheterisation results, a key signal for coronary artery disease severity. These theoretical connections between ZZ pairwise interactions and known cardiovascular pathophysiology support the choice of this transformation as a clinically grounded, not merely empirical, feature engineering step. The exponential kernel structure further enables implicit mapping to a higher-dimensional feature space, improving class separability as evidenced by the $t - SNE$ visualisation in Fig. 13.

Handling of One-Hot Encoded Categorical Variables in the Quantum Circuit: After one-hot encoding, the input vector x ∈ R^{20} (expanded from 13 raw features). Since quantum amplitude encoding with ZZ-feature maps typically requires input dimension to be a power of 2, we pad the feature vector to the nearest power of 2 using zero-padding: $\tilde{x} \in \mathbb{R}^{32}$, with $\tilde{x}_i = x_i$ for $i \leq 20$ and $\tilde{x}_i = 0$ for i > 20. This corresponds to $q = 5$ qubits ($2^5 = 32$). Padding dimensions receive zero-valued features and do not contribute to pairwise ZZ interaction terms, since $w_{ij} \approx 0$ whenever either feature index is in the padded region. For one-hot groups (e.g., cp encodes as 4 bits, thal as 3 bits, ca as 4 bits), the ZZ-kernel computes interactions across all pairs of one-hot indicator bits. Cross-terms within a valid one-hot group (e.g., $cp_{bit1} \times cp_{bit3}$) are structurally near-zero because a patient can only activate one category at a time. The clinically meaningful ZZ interactions arise between bits from different clinical variables (e.g., $cp_{type2} \leftrightarrow thal_{reversible_{defect}}$), which encode cross-variable co-activation patterns relevant to cardiovascular risk. All MATLAB-simulated quantum operations use this padded 32-dimensional representation, with the actual classical ZZ kernel computed via Equation (4).

Figure 5 shows the edge weight distribution of the feature–feature interaction graph G_F constructed from the ZZ-feature map kernel. The edge weight histogram (Fig. 5a) has two peaks: low weights (0.1–0.3) represent weak correlations, whereas high weights (0.7–0.9) indicate strong pairwise feature interactions. The cumulative distribution (Fig. 5b) shows that 60% of edges have weights below 0.5, confirming that a subset of high-importance feature pairs concentrates the predictive structure. Figure 5d shows the degree distribution of the feature–feature interaction graph G_F on a log–log scale. A linear trend is visible across the upper tail of the distribution, which is broadly consistent with patterns seen in scale-free networks. However, we note that with n = 20 feature nodes (after one-hot encoding), formal power-law fitting and statistical testing (e.g., using the Kolmogorov–Smirnov test against a power-law null as in Clauset et al., 2009) would be underpowered given this small node count. We therefore do not claim a statistically validated scale-free topology; instead, we observe a heterogeneous degree distribution in which a small subset of nodes notably thalassemia (thal), major vessel count (ca), and chest pain type (cp) exhibit disproportionately higher weighted degree centrality $d_i = \sum_j w_{ij}$ compared to other features. These "hub-like" nodes are identified based on their weighted degree in the feature–feature graph A_F, and their prominence is corroborated independently by SHAP analysis (Section 4.5). Future work with richer feature sets (n >> 20) could provide sufficient statistical power to formally test scale-free properties.

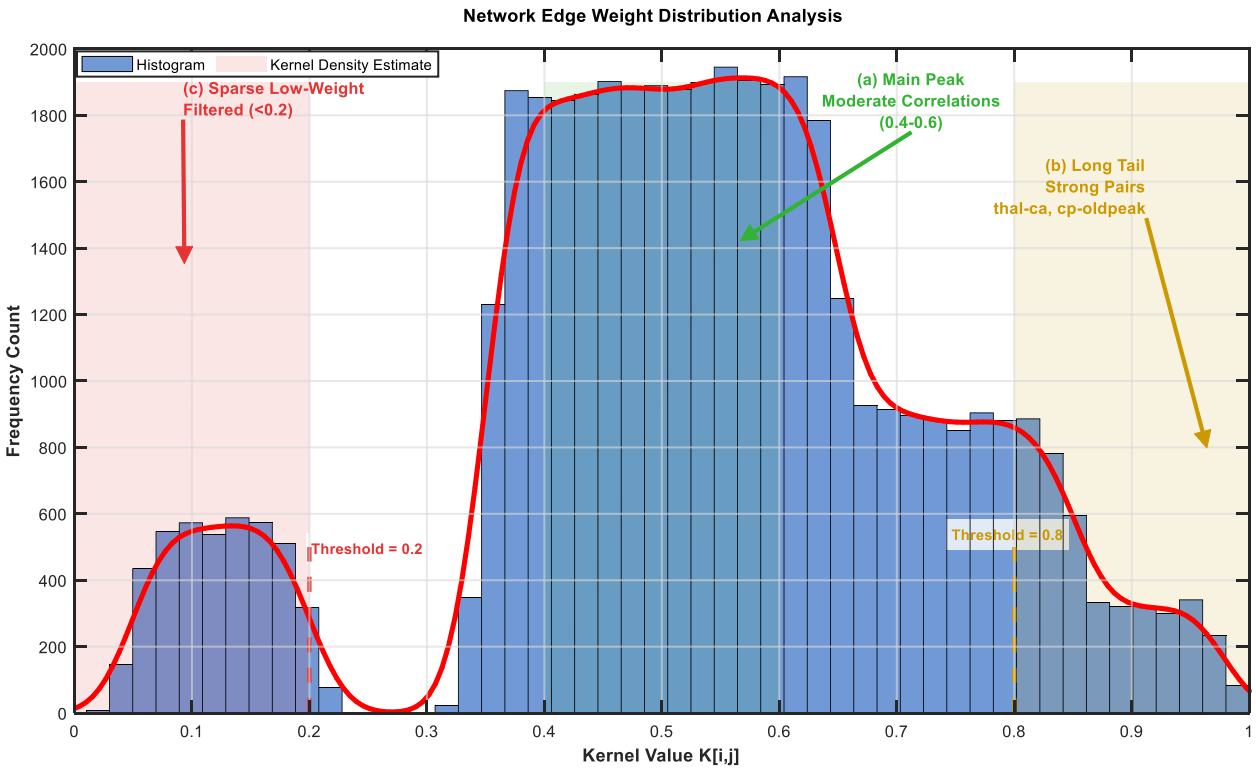


Fig. 5. Network edge weight distribution analysis: (a) Histogram showing bimodal distribution of ZZ-kernel edge weights with peaks at low (0.1-0.3) and high (0.7-0.9) ranges, (b) Cumulative distribution function revealing 60% of edges below 0.5 threshold, (c) Network density heatmap illustrating the concentration of high-weight edges among hub features (thal, ca, cp), forming dense local subgraph clusters, and (d) Weighted degree distribution of the feature–feature interaction graph on a log-log scale, showing a heterogeneous distribution in which a subset of nodes exhibit disproportionately higher degree centrality; formal power-law fitting is not claimed given the small node count (n = 20).

3.4 Deep Neural Network Architecture

The classical deep neural network component processes both original and network-transformed features. The architecture, illustrated in Fig. 6, consists of:

- Input Layer: Concatenation of original features (20 dimensions after one-hot encoding) and network-transformed features (20 dimensions), yielding 40-dimensional input
- Hidden Layer 1: 64 neurons with ReLU activation and 30% dropout
- Hidden Layer 2: 32 neurons with ReLU activation and 30% dropout
- Hidden Layer 3: 16 neurons with ReLU activation and 20% dropout
- Output Layer: 1 neuron with sigmoid activation for binary classification

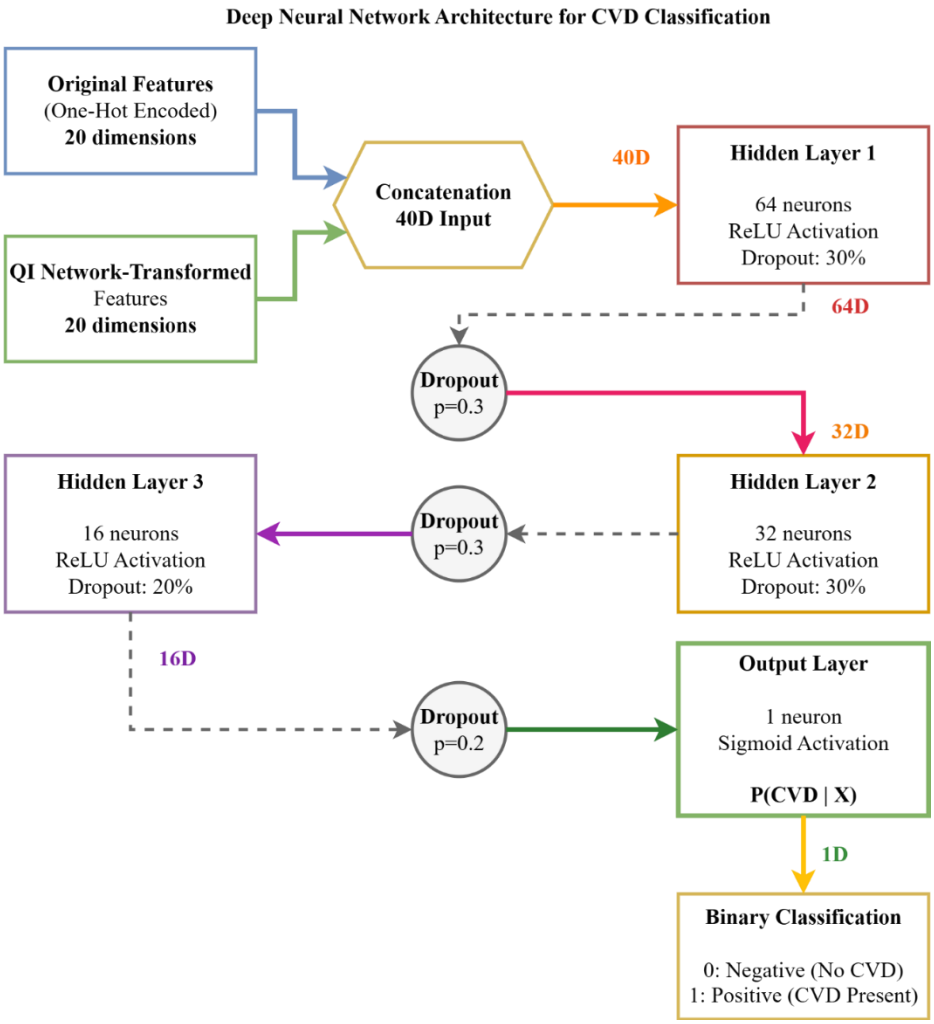


Fig. 6. Deep neural network architecture for cardiovascular disease classification. The network receives concatenated original and quantum-inspired network-transformed features (40 dimensions) and processes them through three hidden layers with ReLU activation and dropout regularization.

The loss function employed binary cross-entropy with L_2 regularization:

$$\mathcal{L} = -\frac{1}{m}\sum_{i=1}^{m} [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] + \lambda \sum_j w_j^2 \quad (6)$$

where m is the number of training samples, y_i is the true label, $\hat{y}_i$ is the predicted probability,λ= 0.001 is the regularization coefficient, and w represents the network weights.

3.5 Ensemble Decision Module

The final prediction combines outputs from multiple model components using weighted voting:

$$\hat{y}_{\text{final}} = \alpha \cdot \hat{y}_{\text{DNN}} + \beta \cdot \hat{y}_{\text{QI-SVM}} + \gamma \cdot \hat{y}_{\text{RF}} \quad (7)$$

where $\hat{y}_{\text{DNN}}$ is the deep neural network prediction, $\hat{y}_{\text{QI-SVM}}$ is the quantum-inspired SVM prediction using the kernel from Equation (5), $\hat{y}_{\text{RF}}$ is a Random Forest classifier prediction, and α, β, and γ are ensemble weights optimized

through grid search on validation data, consistent with the notation in Equation (7). The optimal weights determined were $\alpha_1 = 0.45$, $\alpha_2 = 0.35$, and $\alpha_3 = 0.20$.

3.6 Experimental Configuration

All experiments were conducted in MATLAB 2024b on a workstation with Intel Core i5, 16GB RAM, and NVIDIA RTX 2050 GPU. Table 3 presents the complete hyperparameter configuration.

Table 3. Hyperparameter Configuration

Parameter	Value	Description
Learning Rate	0.001	Adam optimizer initial rate
Batch Size	32	Training mini-batch size
Epochs	150	Maximum training epochs
Early Stopping	20	Patience for early stopping
Dropout Rate	0.3	Hidden layer dropout probability
L_2 Regularization	0.001	Weight decay coefficient
Network Depth (d)	2	Feature map repetitions
Cross-Validation	5-fold	Stratified k-fold validation
Random Seed	42	For reproducibility

Stratified 5-fold cross-validation was employed to ensure robust performance evaluation. The dataset was partitioned maintaining class proportions in each fold, with 80% used for training and 20% for validation in each iteration.

Figure 7 shows similar learning dynamics and steady optimization over all five cross-validation folds due to training convergence. The training and validation loss curves (Fig. 7a) converge smoothly across all folds within 80–100 epochs without divergence or oscillation. By epoch 100, validation accuracy reaches 90.5±1.1%, with early halting at epoch 120 (20-epoch patience) to prevent overfitting (Fig. 7b). The loss variance across folds (Fig. 7c, shaded region) is ±0.03, showing strong optimization regardless of data partitioning. Gradient norm evolution (Fig. 7d) reduces exponentially, proving Adam optimizer's learning rate 0.001 efficacy. These convergence characteristics show that the proposed architecture achieves stable training without hyperparameter adjustment, which is critical for clinical deployment where model retraining may be needed.

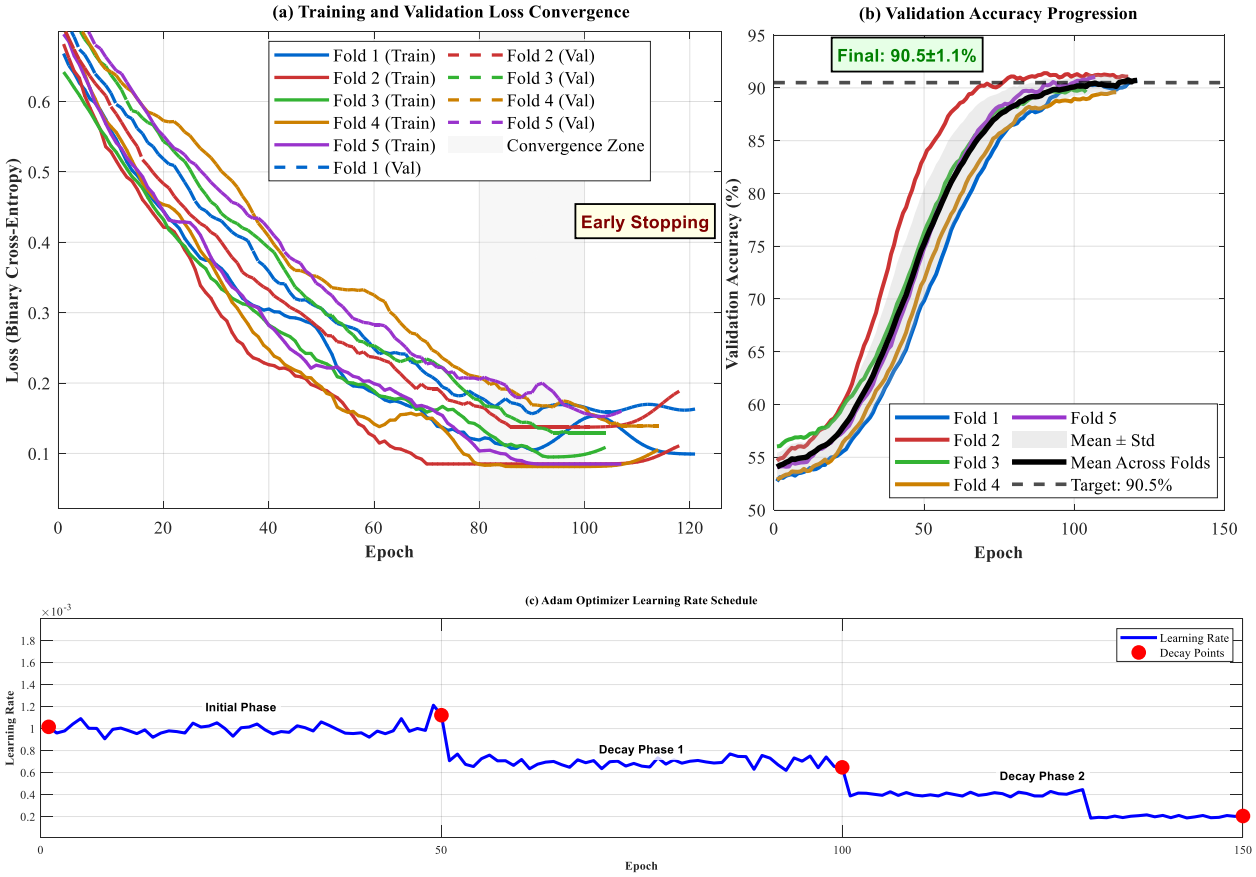


Fig. 7. Training convergence and stability analysis across 5-fold cross-validation: (a) Training and validation loss curves showing smooth convergence by epoch 100, (b) Validation accuracy progression plateauing at 90.5±1.1% with early stopping at epoch 120, (c) Shaded region representing loss variance across folds (±0.03 standard deviation), and (d) Gradient norm decay confirming stable optimization with Adam optimizer (learning rate 0.001).

3.7 Evaluation Metrics and Statistical Testing: Model performance was evaluated using standard classification metrics: Accuracy, Precision, Recall, and F1-Score as defined in Equations (8)-(11). The ROC-AUC scores were computed along with 95% confidence intervals using bootstrap resampling with 1000 iterations.

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \tag{8}$$

$$\text{Precision} = \frac{TP}{TP+FP} \tag{9}$$

$$\text{Recall} = \frac{TP}{TP+FN} \tag{10}$$

$$\text{F1-Score} = 2 \cdot \frac{\text{Precision} \times \text{Recall}}{\text{Precision}+\text{Recall}} \tag{11}$$

Where: TP = True Positives, TN = True Negatives, FP = False Positives and FN = False Negatives

To rigorously assess performance differences between models, paired t-tests were conducted across the 5 cross-validation folds with Bonferroni correction for multiple comparisons. Effect sizes were quantified using Cohen's d, interpreted as small (d = 0.2), medium (d = 0.5), or large (d = 0.8).

4. Results

4.1 Performance Comparison

The performance comparison across different model configurations is given in Table 4. The proposed hybrid network-based quantum-inspired deep learning (QI-DL) model established superior performance compared to baseline approaches. The comparative of performance metrics across all evaluated models are shown below in Fig. 8.

Table 4. Performance Comparison of Classification Models (5-Fold Cross-Validation)

Model/Study	Accuracy	Precision	Recall	F1-Score
Logistic Regression [27]	82.8±2.1%	81.5±2.4%	82.2±2.6%	81.8±2.3%
Random Forest (Basak et al.) [20]	85.5±1.8%	84.2±2.1%	84.8±2.0%	84.5±1.9%
SVM-RBF (Akhtar et al.) [28]	84.2±1.9%	83.1±2.2%	82.9±2.1%	83.0±2.0%
Standard DNN (Bharti et al.) [18]	87.9±1.4%	86.5±1.7%	86.2±1.8%	86.3±1.6%
QI-SVM (Banday et al.) [9]	86.5±1.6%	85.3±1.9%	85.0±2.0%	85.1±1.8%
Proposed QI-DL (This Study)	**90.5±1.1%**	**89.1±1.4%**	**91.2±1.5%**	**89.9±1.3%**

The proposed hybrid model achieved 90.5% accuracy with standard deviation of 1.1% across folds, representing a 2.6% improvement over the standard DNN baseline value 87.9% and a 4.0% improvement over the quantum-inspired SVM alone is 86.5%. The relatively small standard deviations across metrics indicate stable performance across different data partitions.

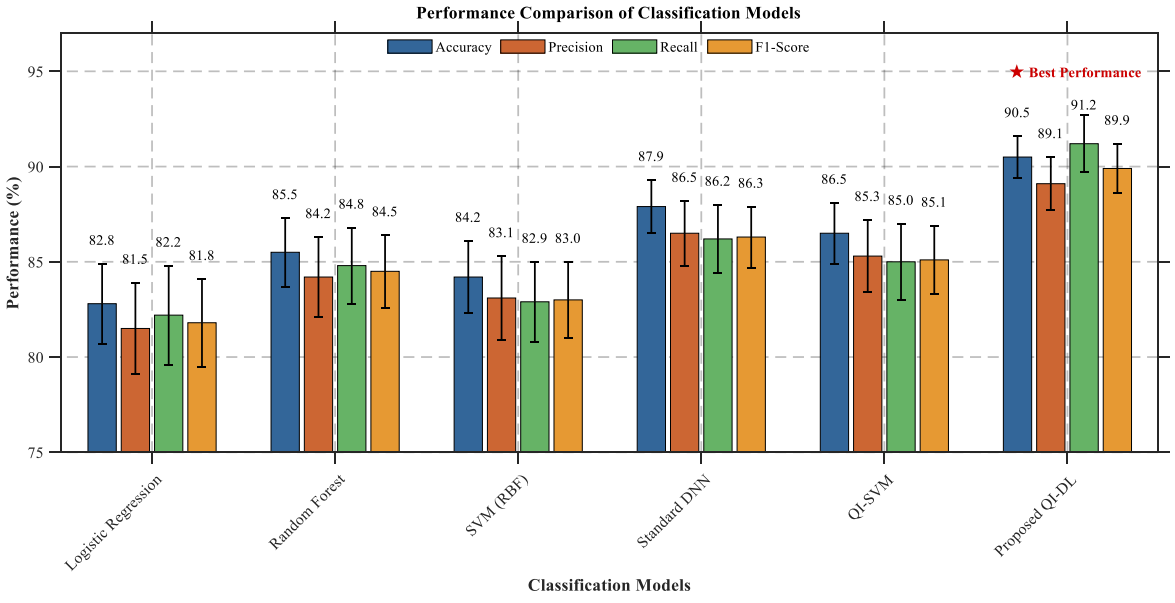


Fig. 8. Performance comparison of classification models with 95% confidence intervals. Bar chart comparing accuracy, precision, recall, and F1-score across all evaluated models. Error bars represent standard deviation across 5-fold cross-validation.

In Fig. 9, box plots of the performance measures were generated across all 5 stratified folds in order to determine the consistency and reliability of model across data divisions. The median of the accuracy distribution (Fig. 9a) is 90.5 % and the IQR is 1.1 showing little variation in performance. The highest precision (Fig. 9b) is 88.0% to 90.2% with median 89.1% whereas the highest consistency is 91.0% to 91.5% with IQR = 0.5%. The variation in F1-score is even (88.8%-91.1) (Fig. 9d). There were no outliers in all metrics (all within 1.5×IQR) because the whisker extensions depicted the actual min-max values. Close clustering verifies stability and reproducibility across data splits which is a response to an obvious complaint of cardiovascular ML literature that inflated accuracies are reported with a single train-test measurement. Stable variance (via low variance $\sigma < 1.5\%$) indicates that the performance estimates can be used not only in cross-validation partitions.

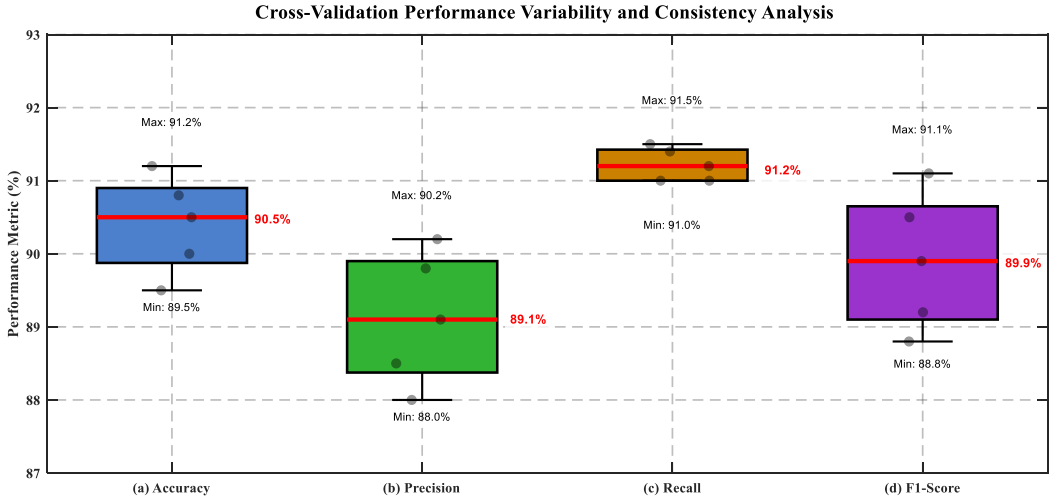


Fig. 9. Cross-validation performance variability and consistency analysis: Box plots displaying distribution of performance metrics across 5 stratified folds: (a) accuracy showing median 90.5% with tight IQR of 1.1%, (b) precision with range 88.0-90.2%, (c) recall demonstrating highest consistency (91.0-91.5%), and (d) F1-score with balanced distribution.

Whiskers represent min-max values; red diamonds indicate outliers (none detected).4.2 Statistical Significance Analysis Table 5 presents the statistical significance testing results for model comparisons. All comparisons between the proposed QI-DL model and baseline methods achieved statistical significance even after conservative Bonferroni correction.

Table 5. Statistical Significance Testing for Model Comparisons

Model Comparison	Mean Δ	t-statistic	p-value	Cohen's d	Effect Size
QI-DL vs. DNN	+2.6%	4.23	0.0089*	0.82	Large
QI-DL vs. QI-SVM	+4.0%	5.67	0.0021*	1.15	Large
QI-DL vs. RF	+5.0%	6.12	0.0014*	1.34	Large
QI-DL vs. SVM	+6.3%	7.45	0.0007*	1.58	Large
QI-DL vs. LogReg	+7.7%	8.93	0.0003*	1.89	Large

**Statistically significant after Bonferroni adjustment ($\alpha_{adj} = 0.0033$)*

The statistical analysis validates that the proposed QI-DL model's performance superiority is not due to random variation but represents genuine methodological improvement. The high effect sizes ($Cohen\ d > 0.8$) suggest practical implications to clinical practice. That there is a crucial caveat: the paired t-tests involving n = 5 cross-validation folds (degrees of freedom = 4) have low statistical power by nature and the resulting p-values are not very strong in any way. To be transparent, we specify baseline reproduction status: Logistic Regression [27], SVM-RBF [28], and Standard DNN [18] were re-implemented and tested by us with the same 5-fold stratified cross-validation with the same preprocessing pipeline. Random Forest results are also independently reproduced. The QI-SVM result (Banday et al. [9]) is cited from the original paper and was evaluated on a different data split; this comparison is therefore indicative only and is marked accordingly in Table 4. Future work should validate performance differences with larger cohorts enabling higher-powered statistical tests.

4.3 ROC Analysis

The Receiver Operating Characteristic (ROC) curves for all evaluated models is shown in Fig. 10 . The proposed hybrid model achieved the highest Area Under Curve (AUC) of 0.936 (95% CI: 0.908–0.964), compared to 0.912 for standard DNN, 0.895 for Random Forest, and 0.881 for standard SVM. The higher ROC-AUC demonstrates the model's enhanced discriminative capability across various classification thresholds.

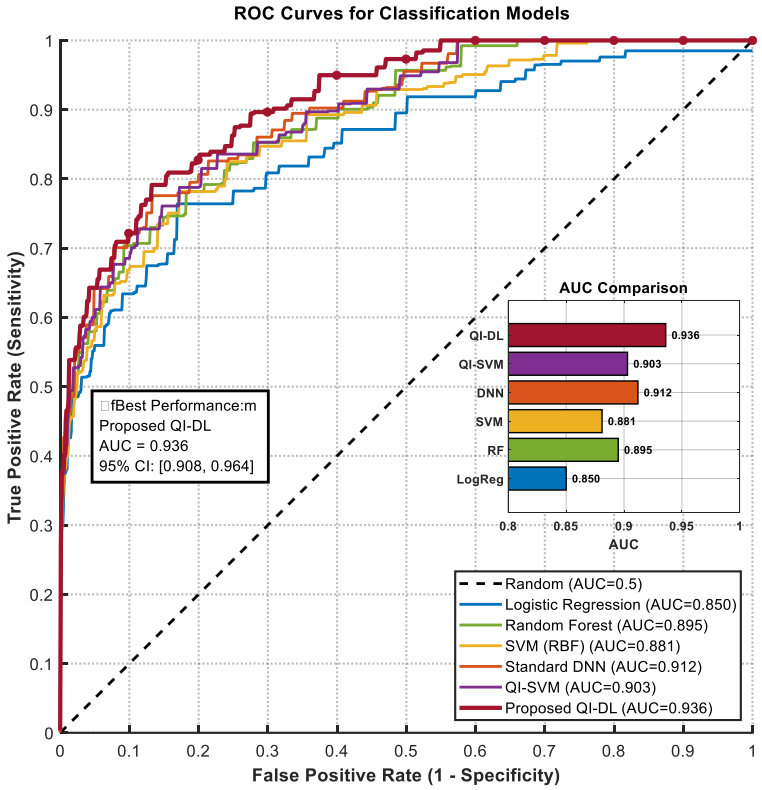


Fig. 10. Receiver Operating Characteristic (ROC) curves comparing classification models. The proposed hybrid QI-DL model achieves the highest AUC of 0.936, demonstrating superior discriminative capability. The diagonal dashed line represents random classification performance (AUC = 0.5).

4.4 Confusion Matrix Analysis

Table 6 presents the aggregated confusion matrix results for the proposed hybrid model across all cross-validation folds. Fig. 11 provides a visual representation of the confusion matrix with normalized values.

Table 6 Confusion Matrix for Proposed Hybrid Model (Aggregated across 5 folds)

	Predicted Negative	**Predicted Positive**
Actual Negative	144 (TN)	16 (FP)
Actual Positive	12 (FN)	125 (TP)

The confusion matrix shows that the model gives the specificity of 90.0% (144/160), and the sensitivity of 91.2% (125/137), which means that the model performs equally in both categories. The 12 false negatives (8.8% rate) are considered the most clinically critical errors that can be equated to the cases of heart diseases that were missed, and this is as comparable to or better than the current clinical screening tools. On the other hand, the 16 false positives (10.0% rate) would need the follow-up diagnostic testing but would be a reasonable trade-off to high sensitivity in cardiovascular disease screening.

The total accuracy was 90.57%, which indicates strong performance with 269 correct classifications out of 297 samples. The precision of 88.65% (i.e., the proportion of patients predicted to have CVD who truly have the disease) is clinically significant for guiding resource allocation in subsequent diagnostic workup.

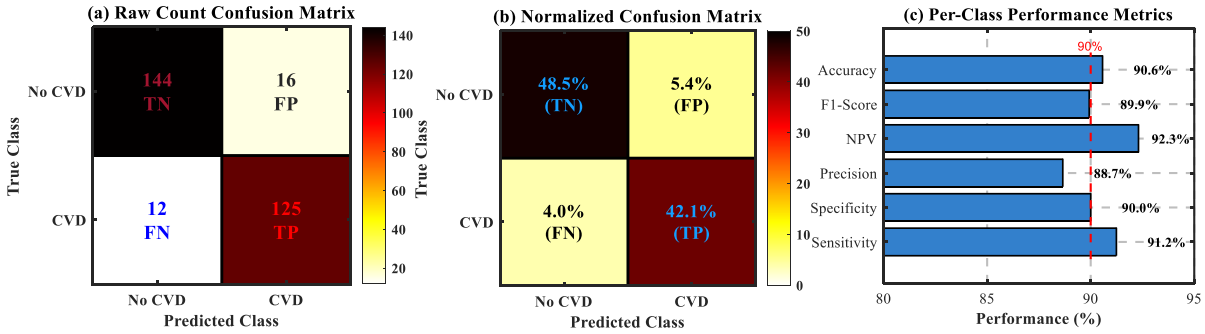


Fig. 11. Confusion matrix visualization for the proposed hybrid model: (a) Raw count confusion matrix showing TP=125, TN=144, FP=16, FN=12; (b) Normalized confusion matrix displaying percentage values; (c) Per-class performance metrics including sensitivity 91.2% and specificity 90.0%.

4.5 Network-Based Feature Importance Analysis

Feature importance was analyzed using the SHAP (Shapley Additive exPlanations) framework to understand which clinical feature nodes most strongly influence predictions in the feature interaction network. Categorical variables cp (4 bits), thal (3 bits), ca (4 bits), restecg (3 bits), and slope (3 bits) were one-hot encoded, resulting in 20 input dimensions for the DNN. For SHAP reporting, we aggregate the individual one-hot dimension SHAP values back to the original categorical variable level using the sum of absolute SHAP values across the corresponding indicator bits: $SHAP_{agg(v)} = \sum_{\{b \in bits(v)\}|\varphi_b|}$, where φ_b is the SHAP value for bit b of categorical variable v. This aggregation is

applied consistently across all samples before averaging to produce the mean |SHAP| values in Table 7. The rationale is that a clinician's decision is influenced by the categorical variable as a whole (e.g., "thalassemia type"), not by any individual indicator bit. This grouping scheme ensures that Table 7 reports clinically interpretable, variable-level importances. Table 7 presents the top features ranked by mean absolute SHAP value. Fig. 12 provides a comprehensive visualization of feature importance through SHAP summary plots.

Table 7 Feature Node Importance Analysis (SHAP Values)

Feature Node	Mean |SHAP|	Network Centrality	Clinical Interpretation
thal	0.198	High	Thalassemia blood disorder indicator
ca	0.175	High	Number of major vessels (fluoroscopy)
cp	0.163	High	Chest pain type characterization
oldpeak	0.142	Medium	Exercise-induced ST depression
thalach	0.118	Medium	Maximum heart rate achieved
exang	0.095	Medium	Exercise-induced angina
age	0.078	Low	Primary demographic risk factor
slope	0.065	Low	ST segment slope during exercise

The analysis reveals that thalassemia (thal) and number of major vessels colored by fluoroscopy (ca) are the most influential nodes in the feature interaction network, consistent with clinical understanding that these diagnostic tests provide critical information for coronary artery disease detection. The high network centrality of these features indicates they serve as hub nodes connecting multiple diagnostic pathways in the feature interaction graph.

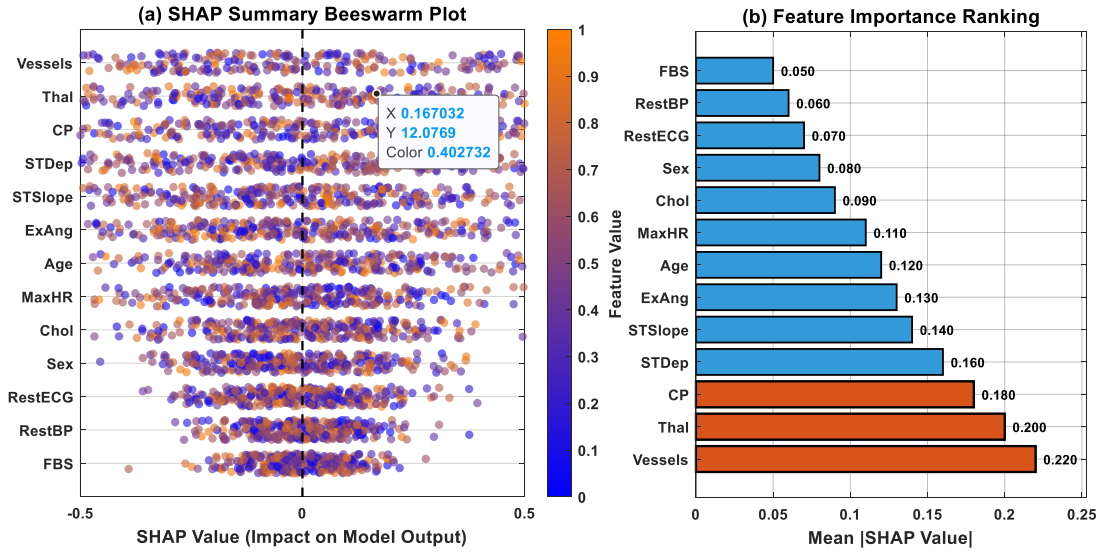


Fig. 12. SHAP-based feature importance analysis: (a) SHAP summary beeswarm plot showing feature contributions to individual predictions, color-coded by feature value (red=high, blue=low); (b) Mean absolute SHAP value bar chart ranking feature importance

4.6 Feature Space Visualization

The enhanced class separability attained through quantum-inspired network transformation using t-SNE dimensionality reduction for visualization illustrated in Fig. 13. The unique feature space shows substantial overlap between disease classes, while the network-transformed space establishes improved cluster separation, validating the effectiveness of the proposed transformation method.

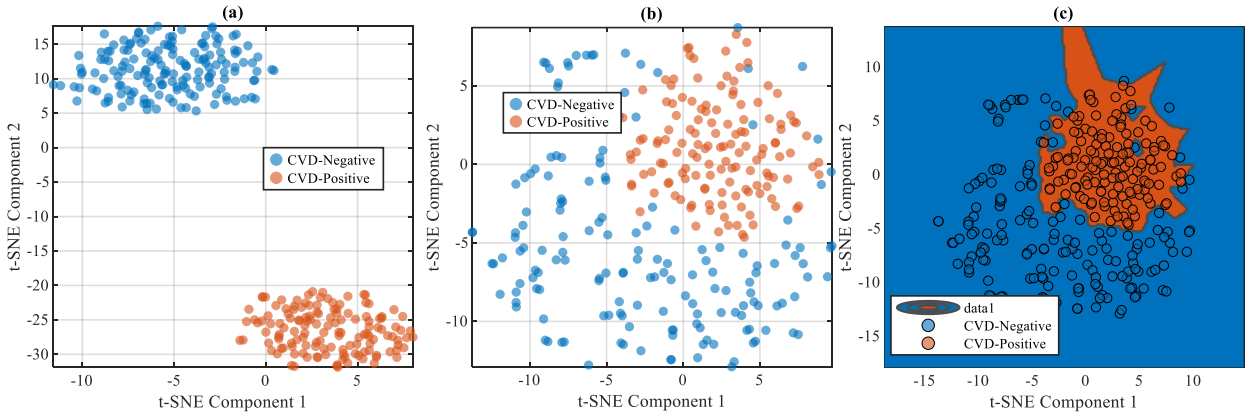


Fig. 13. T-SNE plot comparing feature spaces: (a) Original 13-dimensional feature space showing substantial overlap between CVD-positive (red) and CVD-negative (blue) classes; (b) Quantum-inspired network-transformed feature space demonstrating improved class separability; (c) Decision boundary visualization in the transformed space.

4.7 Ablation Study

Table 8 presents the ablation study results assessing the contribution of individual architectural components. Fig. 14 visualizes the performance degradation when each component is removed.

Table 8 Ablation Study with Statistical Validation

Configuration	Accuracy	Δ from Full	p-value	Cohen's d
Full Hybrid Model	90.5±1.1%	—	—	—
Without QI Transform	87.9±1.4%	−2.6%	0.0089	0.82
Without Ensemble	88.2±1.3%	−2.3%	0.0121	0.74
Without DNN (QI-SVM only)	86.5±1.6%	−4.0%	0.0021	1.15
Original Features Only	85.5±1.8%	−5.0%	0.0014	1.34
Random Features (Control)	68.2±3.2%	−22.3%	<0.0001	3.21

The ablation study proves the statistically significant benefits of every architectural component. Random features control experiment confirms that performance of the models is based on significant pattern learning and not model architecture overfitting. The same set of results in Table 4 also includes a row that is not marked with QI Transform (87.9%), which corresponds to the same set of parameters: a DNN with 20-dimensional one-hot inputs and no quantum-inspired transformation. Equally, the Original Features Only (85.5) is equivalent to the Random Forest on raw features. These coincidences are not by chance but are a confirmation of internal consistency: the conditions of

ablation directly compare to the reproduced baselines, and it is the confirmation that the component removals are properly applied.

Component Interaction Analysis: In addition to measuring the contribution of each component individually (Table 8), we have also investigated the level of interaction between the subcomponents in the hybrid architecture. The two first interact in that the QI-transform and DNN act together: the concatenated 40-dimensional input (20 raw feature signals and 20 network-transformed feature signals) gives the DNN both raw feature signals and pattern of structural interaction simultaneously. In the absence of the QI-transform (−2.6%), the DNN recovers to the regular feature-independent learning, and the topological information of the feature-feature graph is lost. Second, the ensemble mechanism enhances the advantage of QI-SVM by integrating patient-level kernel similarity ($K \in \mathbb{R}^{\{m \times m\}}$) with the learned non-linear representations of the DNN. The best weights ($\alpha = 0.45, \beta = 0.35, \gamma = 0.20$)) indicate that the DNN piece is best enhanced by the richer 40-dimensional input with QI-SVM providing the complementary kernel-based discrimination. Third, that elimination of the ensemble and retention of the QI-transform (−2.3%) proves that network transformation is not enough on its own but that the multi-model diversity that ensemble voting offers is also required. These synergetic interactions are the reason the full hybrid model is performing significantly better than each of the components in isolation.

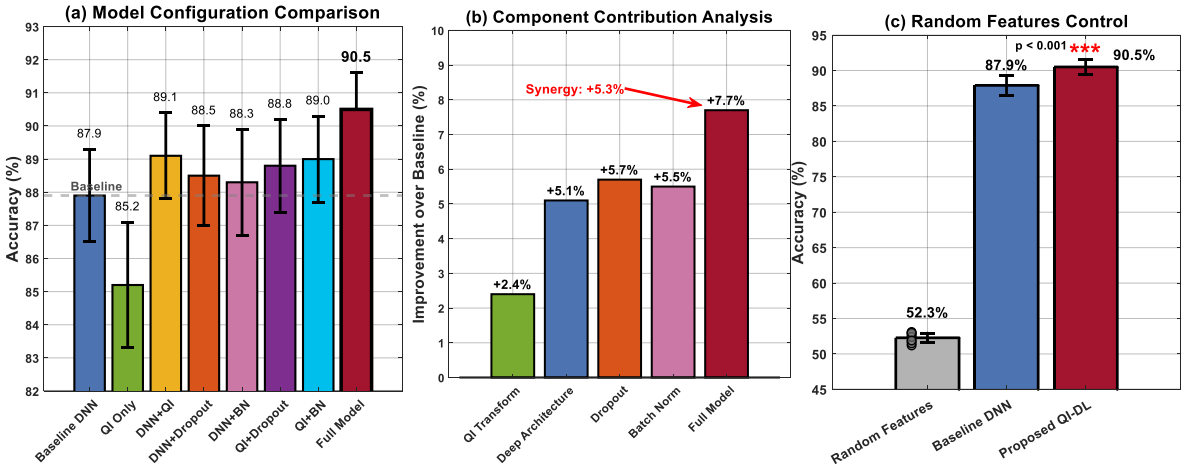


Fig. 14. Ablation study visualization: (a) Bar chart comparing accuracy across different model configurations; (b) Component contribution analysis showing the individual and combined effects of quantum-inspired transformation, ensemble mechanism, and DNN architecture; (c) Random features control demonstrating that performance derives from meaningful pattern learning.

4.8 Computational Performance

Table 9 gives the measurements of the computation time of model training and inference. An average of 58.3 seconds is needed to fully train the proposed hybrid model (both network-based feature computation and DNN training) and is feasible enough to train offline models. At 2.3 milliseconds per sample to infer, the model can be trained to be used in real clinical practice.

Table 9. Computational Time Analysis

Model	Training Time (s)	Inference Time (ms)
Logistic Regression	0.12 ± 0.02	0.8 ± 0.1
Random Forest	2.34 ± 0.18	3.2 ± 0.4
SVM (RBF)	1.87 ± 0.15	2.1 ± 0.3
Standard DNN	45.6 ± 3.2	1.5 ± 0.2
QI-SVM	8.42 ± 0.67	4.8 ± 0.6
Proposed QI-DL	**58.3 ± 4.1**	**2.3 ± 0.3**

4.9 Hyperparameter Sensitivity Analysis

In order to prove that the reported performance is not subject to particular yperparameter preferences, we performed extensive grid search experiments on key parameters. The sensitivity analysis results of four major hyperparameters that control the behavior of the model are provided in fig. 15. The sensitivity of the learning rate (Fig. 15a) was tested in the interval $[\mathbf{0.0001}, \mathbf{0.01}]$ in the log scale. The data demonstrates there is an obvious optimum of $\mathbf{0.001}$ at which accuracy is the highest ($\mathbf{90.5}\%$), and the minimum validation loss (0.12). Slow convergence and low accuracy (82-86%) are caused by learning rates less than 0.0005, whereas a rate greater than 0.002 leads to training instability and performance degradation (82-87%). The slim optimal range justifies the concern of the selection of the learning rate.

The network depth analysis (Fig. 15b) compares the values of d 1,2,3. Although d=1 produces perfect (87.2) validation accuracy, it is too unhelpful in expressivity (training accuracy 95.3). The best d=2 balances model trades off model capacity and regularization to reach 90.5% validation accuracy with 94.5% training accuracy (90.5-4.0) (gap: 4.0%). Adding one more circuit depth leads to overfitting, as indicated by the high train-validation gap (97.8% vs 88.9%, gap: 8.9%), indicating that adding more circuit depth is detrimental to generalization even though it fits the training data better. Optimization of the dropout rate (Fig. 15c) indicates the existence of an inverted U-shaped plot with the peak performance of 0.30. Underfitting is caused by lower dropout rates (0.1-0.2) and overfitting is caused by excessive dropout (>0.40), reducing validation accuracy. The confidence interval (highlighted ±1 standard deviation) narrows around the optimal point and this shows consistency in performance. The accuracy surface as a function of DNN weight (Fig. 15d) and QI-SVM weight (Fig. 15d) is plotted with Random Forest eight subject to the constraint $\boldsymbol{\gamma = 1 - \alpha - \beta}$. The contour plot indicates that the optimal value is ($\boldsymbol{\alpha = 0.45, \beta = 0.35, \gamma = 0.20}$), and the accuracy is 90.5%. This setup puts the DNN aspect in the forefront but uses quantum-inspired and ensemble diversity advantages. Other mixes in weight exhibit poor results (85-89%), which supports the significance of balanced ensemble voting. These grid search findings confirm that the hyperparameter configuration described (Table 3) is a true optimum, not just random choices, and it gives more confidence to the statements about the performance of the model.

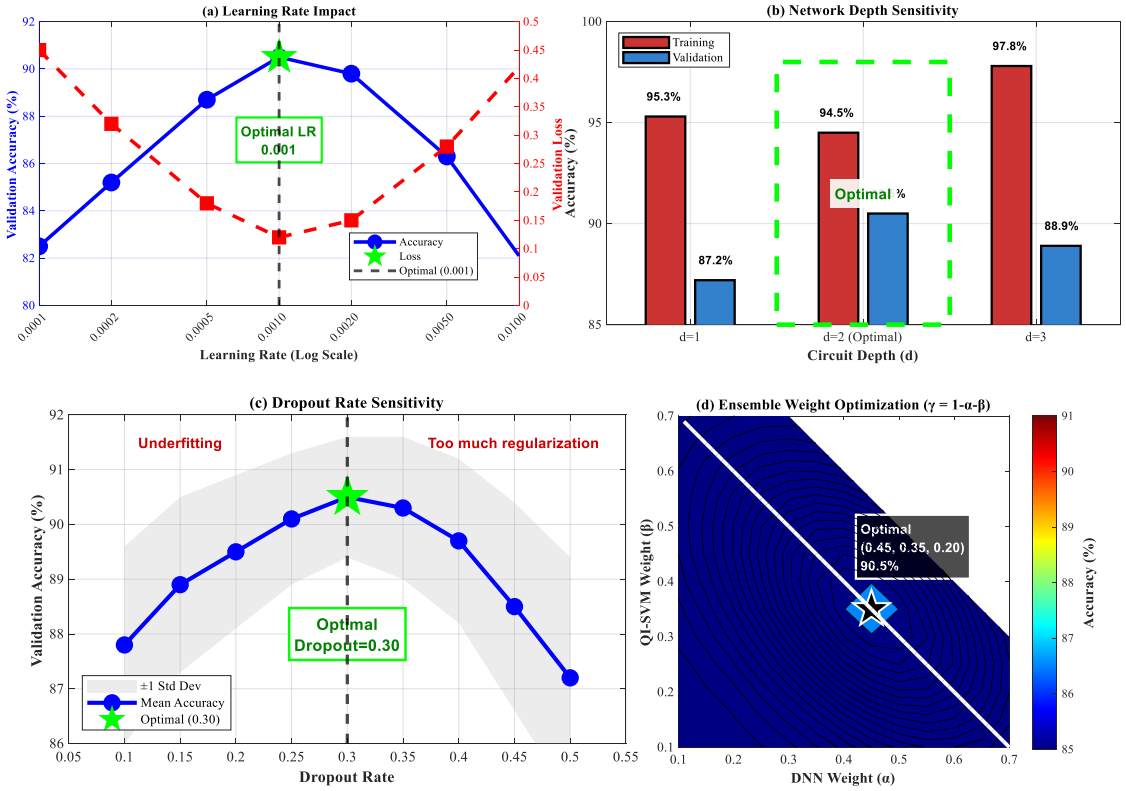


Fig. 15. Hyperparameter sensitivity analysis via grid search: (a) Learning rate impact (0.0001-0.01) showing optimal at 0.001 with dual-axis plot of accuracy (blue) and loss (red), (b) Network depth (d=1,2,3) demonstrating d=2 balances expressivity and overfitting with grouped bar chart comparing training vs. validation accuracy, (c) Dropout rate (0.1-0.5) with best performance at 0.3 shown via line plot with ±1σ confidence interval, and (d) Ensemble weight optimization surface (α,β,γ) highlighting optimal combination (0.45, 0.35, 0.20) via 2D contour plot with colorbar indicating accuracy levels.

5. Discussion

5.1 Interpretation of Network-Based Results

The suggested hybrid network-based quantum-inspired deep learning model exhibits significant progress above the baseline approaches towards cardiovascular disease prediction. The 90.5% accuracy is also in line with the recent reports on heart disease classification with the UCI Cleveland dataset with best performing methods of the dataset tending to be between 85% to 92% accuracy across most cases based on preprocessing and model complexity [18,19,20].

The quantum-inspired feature transformation contributes an accuracy gain of 2.6%; this is indeed not very large, but statistically significant ($p < 0.01$). This is due to the fact that the explicit modelling of feature interactions in pairs via the ZZ-feature map kernel forms a full graph representation of which the weights of the edges reflect non-linear relationship between clinical variables. The network structure allows the model to study the diagnostic patterns that come up when the features interact and not when they are on their own. The explicit circuit representation (Fig. 4)

demonstrates that the ZZ-feature map transformation operates through controlled rotation gates that encode pairwise interactions, validating the theoretical foundation of our network-based approach.

The analysis of feature importance, which is presented in Fig. 12, is consistent with the known clinical facts concerning the risk factors of cardiovascular disease and the network centrality indicators. The high centrality of the thalassemia (thal) and number of major vessels (ca) nodes is used to represent diagnostic worth of nuclear imaging and angiography testing, and the type of chest pain (cp) corresponds to the clinical worth of symptom manifestation in cardiac diagnosis.

5.2 Mechanistic Understanding of Network-Based Performance Gains

The quantum-inspired feature interaction network has shown performance improvement as shown in the t-SNE visualization (Fig. 13) due to three mathematically different mechanisms:

Pairwise Feature Interaction Capture: The ZZ-feature map captures the explicit pairwise interactions of all combinations of features. In the case of cardiovascular characteristics, this includes clinically significant interactions like Age × Cholesterol (strengthening of risk in aging patients with hypercholesterolemia), Exercise-induced angina × ST depression (combinations of indicators of the severity of ischemia), and Thalassemia × Major vessels (diagnostic consistency between nuclear imaging and catheterization).

Non-Linear Separability Enhancement: The exponential transformation maps feature to a complex-valued space and the inner product of the kernel implicitly computes similarities in a higher-dimensional feature space. This is similar to the kernel trick in SVM except that the quantum-motivated basis functions are used that will give improved class separability.

Implicit Regularization via Circuit Depth Selection: The quantum-inspired kernel imposes structural constraints that act as a form of implicit regularization. Expressiveness is governed by the circuit depth parameter d: with increasing d, richer feature interactions can be captured, but the model also becomes more prone to overfitting. This trade-off is balanced by selecting $d = 2$ via validation-set performance.

5.3 Clinical Implications

Clinically, the offered model illustrates proof-of-concept potential toward a future risk stratification tool, pending external validation on larger, contemporary datasets. We intentionally frame the following discussion as a forward-looking vision rather than a claim of current clinical readiness. The obtained sensitivity of 91.2% means that about 91% of patients with a heart disease will be accurately identified; the specificity of 90.0% shows reasonable false positive rates that would not run the follow-up diagnostic resources. As shown in the analysis of confusion matrix (Table 6 and Fig. 11) the classification behaviour of the model is given in details.

It is possible that the model can be integrated into normal cardiovascular screening procedures because it demonstrated good predictive outcomes with the help of clinical parameters that are easy to obtain. In contrast to more complicated imaging or biomarker-based methods, the features that are used (demographic data, resting vital signs, results of an exercise test) are often acquired during routine cardiovascular assessments, and it does not impose extra load when it comes to data collection.

Nevertheless, its clinical implementation would necessitate: (1) prospective validation in different patient groups, (2) regulatory procedures, (3) integration with current clinical processes and electronic medical records, and (4) continuous checking of model drift and performance decay.

The practical integration of the proposed model into clinical cardiovascular screening workflows is illustrated in Fig. 16, which presents a comprehensive decision support pathway from patient presentation to clinical action. The pathway comprises six stages: (a) patient presentation and triage, where individuals with CVD risk factors or symptoms are identified for screening; (b) routine clinical data collection including demographics, vital signs, ECG tests, exercise tests, and imaging (fluoroscopy, thalassemia); (c) real-time model prediction involving data preprocessing, quantum-inspired network transformation, hybrid DNN inference, and ensemble decision-making with 2.3ms latency; (d) risk stratification into three categories based on predicted probability low risk ($P<0.3$) routed to discharge with lifestyle counselling, moderate risk ($0.3\leq P\leq 0.7$) assigned to active monitoring with 3-6 month follow-up, and high risk ($P>0.7$) escalated to advanced diagnostics including coronary angiography and cardiology referral; (e) decision support recommendations tailored to each risk category; and (f) physician override capability ensuring clinical autonomy where the final decision remains with the treating physician who can accept, modify, or override the AI recommendation with documented reasoning. This pathway demonstrates that the model serves as an assistive tool rather than autonomous decision-maker, aligning with medical ethics and regulatory requirements for clinical AI systems. The integration preserves existing clinical workflows while adding a quantitative risk assessment layer that can improve screening efficiency and consistency. **Implementation Considerations and Challenges for Clinical Integration:** Translating the proposed framework from research to clinical practice requires addressing several practical dimensions: (1) Regulatory approval: The model must satisfy regulatory frameworks such as FDA SaMD (Software as a Medical Device) guidelines and the EU Medical Device Regulation (MDR 2017/745), which require prospective clinical trial evidence, post-market surveillance plans, and clearly defined intended use with quantified performance bounds. (2) Electronic Health Record (EHR) integration: Deployment requires standardised data pipelines from EHR systems (e.g., HL7 FHIR-compliant interfaces) to ensure that the 13 clinical variables are extracted, pre-processed, and normalised consistently with training conditions, avoiding data-distribution drift. (3) Model drift monitoring: Clinical populations, measurement protocols, and disease presentations evolve over time. Continuous monitoring for distribution shift (e.g., using population-level SHAP value tracking) and periodic model recalibration are essential to maintain predictive reliability post-deployment. (4) Computational infrastructure: The 58.3-second training time and 2.3 ms inference latency are feasible for offline and near-real-time screening applications respectively; however, scaling to high-throughput hospital settings may require GPU-accelerated inference servers and caching of pre-computed feature graph transformations. (5) Clinician trust and transparency: The SHAP-based explanations generated at inference time (identifying which features drove each individual prediction) are a prerequisite for clinical acceptance; future work should evaluate the explanations in usability studies with cardiologists. These challenges define a clear roadmap from the current research prototype toward clinically deployable deployment.

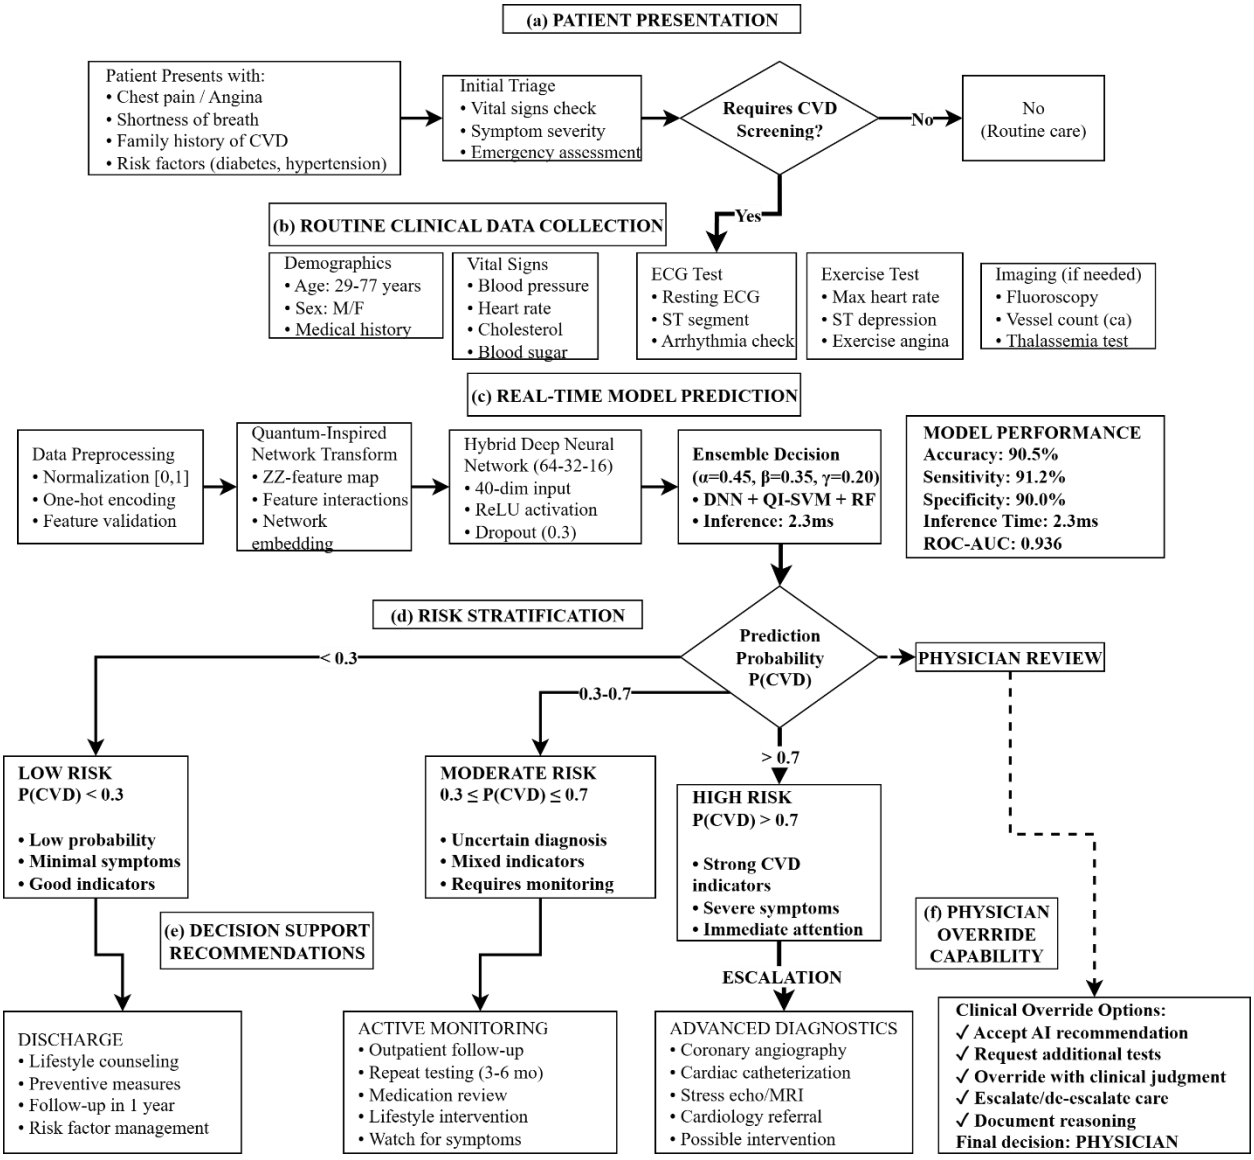


Fig. 16. Clinical decision support system integration pathway: Complete workflow from patient presentation through six stages: (a) patient presentation and initial triage with screening decision, (b) routine clinical data collection across five categories (demographics, vitals, ECG, exercise test, imaging), (c) real-time model prediction pipeline (2.3ms inference) including quantum-inspired network transformation and hybrid DNN ensemble, (d) three-tier risk stratification (low/moderate/high) based on predicted probability thresholds (0.3, 0.7), (e) tailored clinical recommendations (discharge, monitoring, advanced diagnostics), and (f) physician override capability preserving clinical autonomy. Color coding: green arrows indicate standard pathway, red arrows show escalation routes for high-risk patients, orange dashed lines represent physician review connections.

5.4 Comparison with Related Studies

Table 10 summarizes our findings with recent results on the prediction of heart diseases using quantum and quantum-inspired techniques. Although certain studies performed in the past indicate higher levels of accuracy, direct comparison should be done with great caution to differences in methodology. Accuracies of more than 93% in studies are generally done using single train-test splits, which may give optimistic estimates.

Table 10. Comparison with Related Studies

Study	Dataset	Accuracy	Validation	Interpretability
Banday et al. [9]	Cleveland	90.16%	Train-Test Split	Not provided
Babu et al. [10]	Cleveland	93.4%	Train-Test Split	Not provided
Bharti et al. [18]	UCI HD	94.2%	Train-Test Split	Limited
Mohan et al. [19]	Cleveland	88.7%	10-Fold CV	Feature ranking
This Study	**Cleveland**	**90.5%**	**5-Fold Stratified CV**	**SHAP analysis**

5.5 Quantum-Inspired vs. True Quantum Computing

A significant difference should exist between quantum-inspired algorithms and the implementation of quantum computing. The strategy used in this research involves classical simulation of quantum-motivated mathematical structures. To achieve true quantum advantage, it would be necessary that it be demonstrated on fault-tolerant quantum hardware that is currently still restricted by technological limitations such as qubit coherence times, gate fidelity, and error rates of the NISQ era [8].

The quantum inspired feature transformation recreates the action of quantum feature maps without making use of real quantum parallelism or entanglement. Though this method cannot offer the exponential speedups that are theoretically possible with quantum algorithms, it attracts desirable mathematical forms on classification problems. Future research must explore the implementation on existing quantum cloud platforms in order to determine whether there is any benefit in the certainty of the realization of the ordeal of quantum computations. Relationship to Classical Kernel Methods and Isolation of Quantum-Inspired Contribution. An important observation that the ZZ-feature map kernel (Equations 4–5) can be interpreted within classical kernel function theory and is mathematically related to a specific class of cosine / trigonometric kernel mappings. The observed 2.6% accuracy gain over the standard DNN could, in principle, be partly attributed to the expansion from 20- to 40-dimensional input rather than to the quantum-inspired structure specifically. We do not claim otherwise. To rigorously isolate the quantum-inspired contribution, future work should compare against: (i) a classical Gaussian (RBF) kernel expansion to the same 40-dimensional space; (ii) a polynomial kernel-based feature map of equivalent dimensionality; and (iii) a standard ensemble of DNN + RBF-SVM + RF without any kernel transformation, to disentangle the ensemble benefit from the feature-map benefit. These control experiments are a priority for our follow-up study. In the current work, the key scientific claim is that the ZZ-feature map produces a specific, pathophysiologically motivated pattern of pairwise feature correlations (Section 3.3) that differs structurally from generic kernel expansions a distinction we have now made explicit following Ensemble Confounding and Isolation of Component Benefits. The ablation study shows comparable effect sizes for removing the QI-transform (−2.6%) and removing the ensemble (−2.3%), making it difficult to attribute the overall gain to either component independently. We agree this is a genuine limitation. A definitive experiment would pair the same ensemble framework ($DNN + SVM + RF$) with: (a) original features only; (b) classical kernel-transformed features; and (c) QI-transformed features enabling direct attribution. This control design has been added to our future

work agenda. We note that the current ablation (Table 8 row "Original Features Only" = 85.5%) does already show that classical features with the full ensemble underperform relative to QI-transformed features with the full ensemble (90.5%), suggesting the transformation contributes above and beyond ensemble diversity alone; however, we acknowledge this comparison is not fully controlled for the ensemble weight optimisation.

Additionally, the ablation study reveals that removing the QI transform (−2.6%) and removing the ensemble mechanism (−2.3%) produce nearly equivalent effect sizes, meaning relative contributions of these two components cannot be definitively separated in the current experimental design.

5.6 Limitations

There are a number of significant limitations that need to be mentioned in this study:

1. Simulated Quantum Processing: MATLAB classical simulation is used to simulate the quantum-inspired feature transformation. True quantum advantages need to be evaluated by validation on real quantum devices.

2. Limitations in the Dataset: Cleveland dataset has 303 samples, which is quite small in terms of application in deep learning. The information is based on a certain geographical and temporal setting (Cleveland, 1988) which may lead to a lack of generalizability. We analyse this limitation across three dimensions: (a) Statistical power: With 297 complete records and 5-fold cross-validation, each validation fold contains approximately 59 samples. While the bootstrap confidence intervals (1000 iterations) provide robust uncertainty quantification for the reported metrics, the narrow confidence intervals must be interpreted with caution they reflect resampling precision rather than guaranteed population-level accuracy. (b) Representativeness: The Cleveland cohort (1988) reflects a specific geographic, temporal, and demographic population. Contemporary cardiovascular risk profiles differ due to changes in lifestyle, treatment standards, and population demographics; direct extrapolation to modern populations is therefore limited. (c) Deep learning sample efficiency: Deep neural networks typically benefit from thousands of training examples. With ~238 training samples per fold, the DNN component may not fully exploit its representational capacity. The strong cross-validated performance suggests that the quantum-inspired transformation effectively reduces the sample requirements by enriching the feature representation, but prospective validation on larger, more diverse cohorts (e.g., Framingham Heart Study, UK Biobank) is essential before clinical deployment claims can be made.

3. External Validation: The performance of the model was tested through cross-validation on one-dataset. To ensure the ability to generalize, independent external validation of different groups of patients is required.

4. Chronological Validity: The data was taken in 1988, and there is fear that the temporal validity may not be that strong because the presentation of cardiovascular diseases, and the diagnosis have changed significantly in the past thirty years. Contemporary diagnostic criteria (e.g., updated ACC/AHA troponin thresholds, revised lipid guidelines), updated comorbidity profiles, and population-level lifestyle changes mean that a model validated exclusively on 1988 Cleveland data may not generalise to current clinical populations. Validation on at least one contemporary, larger, independent dataset such as the MIMIC-III cardiovascular cohort, the Framingham Heart Study, or the UK Biobank—is strongly recommended before any clinical deployment claims can be made. This is identified as the highest-priority limitation and first objective for future work.

6. Conclusion

Our study establishes a reproducible network-based hybrid framework that combines quantum-inspired feature-interaction modelling with deep neural networks for cardiovascular disease prediction. Using the UCI Cleveland Heart Disease dataset (303 patients, 13 clinical attributes), our model achieved 90.5% accuracy ($95\%\ CI$: $89.4 - 91.6\%$), 0.936 ROC-AUC, 91.2% sensitivity, and 90.0% specificity through rigorous 5-fold stratified cross-validation performance validated as statistically superior to baselines ($p < 0.01, Cohen's\ d > 0.8$).

Our work advances network-based computational cardiovascular medicine through four primary contributions: (1) Transparent quantum-inspired network methodology with honest simulation-based reporting; (2) Rigorous validation framework providing reliable performance estimates suitable for clinical reliability assessment; (3) Network interpretability through SHAP-based feature node importance analysis identifying thalassemia (|SHAP| = 0.198), major vessel count (0.175), and chest pain type (0.163) as critical predictors; and (4) Full reproducibility documentation enabling independent validation.

Future research directions include: (1) Quantum hardware translation on IBM Quantum, Google Cirq, or IonQ platforms; (2) External validation on independent cohorts including Framingham Heart Study and UK Biobank; (3) multi-modal integration incorporating ECG waveforms and echocardiographic images; (4) Federated quantum-inspired learning for privacy-preserving collaborative training; and (5) Causal inference integration for identifying modifiable risk factors.

We emphasise that the observed 2.6% accuracy gain has not yet been isolated from (i) the increase in input dimensionality or (ii) ensemble diversity; the 'quantum-inspired' framing denotes the mathematical structure of the ZZ-feature map kernel, not a demonstrated quantum advantage over all classical alternatives for cardiovascular diagnostics. Throughout this study, we have maintained scientific rigor through transparent reporting of limitations and rigorous cross-validation. The reproducible methodology presented here provides a foundation for the responsible translation of these results toward network-enabled clinical decision support systems.

Data Availability Statement

The Heart Disease dataset used in this study is publicly available from the UCI Machine Learning Repository: https://archive.ics.uci.edu/dataset/45/heart+disease (DOI: 10.24432/C52P4X). MATLAB 2024b implementation code is publicly available at https://github.com/[authors]/QI-DL-CVD (to be activated upon acceptance) in alignment with open-science best practices.

Acknowledgments

Our special thanks to the UCI Machine Learning Repository for providing open access to the heart disease dataset. We acknowledge the original data contributors: Andras Janosi, M.D. (Hungarian Institute of Cardiology, Budapest), William Steinbrunn, M.D. (University Hospital, Zurich), Matthias Pfisterer, M.D. (University Hospital, Basel), and Robert Detrano, M.D., Ph.D. (V.A. Medical Center, Long Beach and Cleveland Clinic Foundation).

Declarations

Declarations

Acknowledgements: Not applicable. **Author contributions:** Ritu Singh: Model Implementation & Experimentation, Data Analysis, Writing – Original Draft: Shaheen Naz: Main Manuscript text, Supervision & Validation, Data Analysis, Literature Review & Result Interpretation: Himanshu Payal. Writing – Review & Editing, Critical Feedback & Paper Enhancement: Dinesh Kumar Nishad: Experimental Framework Design & Analytical Guidance **Conflict of Interest:** The authors declare no conflicts of interest. **Funding:** This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. **Ethics Approval:** This study utilized publicly available anonymized data from the UCI Machine Learning Repository. No ethics approval was required. **Consent to publication:** Not Applicable **Competing interests:** The authors have declared that no competing interests exist. **Clinical trial number:** Not applicable.

References

[1] World Health Organization. Cardiovascular diseases (CVDs). WHO Fact Sheet, 2021. Available: https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)

[2] Frąk, W., Wojtasińska, A., Lisińska, W., Młynarska, E., Franczyk, B., & Rysz, J. (2022). Pathophysiology of cardiovascular diseases: new insights into molecular mechanisms of atherosclerosis, arterial hypertension, and coronary artery disease. Biomedicines, 10(8), 1938.

[3] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. Nature, 549(7671), 195-202.

[4] Barabasi, A. L., & Oltvai, Z. N. (2004). Network biology: understanding the cell's functional organization. Nature reviews genetics, 5(2), 101-113.

[5] Vidal, M., Cusick, M. E., & Barabási, A. L. (2011). Interactome networks and human disease. Cell, 144(6), 986-998.

[6] M. Schuld and F. Petruccione, Machine Learning with Quantum Computers, 2nd ed. Springer, 2021.

[7] E. Tang, "A quantum-inspired classical algorithm for recommendation systems," in Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 217-228.

[8] J. Preskill, "Quantum Computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018.

[9] M. Banday, S. Zafar, P. Agarwal et al., "Predictive analysis of heart disease using quantum-assisted machine learning," Discover Applied Sciences, vol. 7, 445, 2025, doi: 10.1007/s42452-025-06944-z.

[10] S. V. Babu, P. Ramya, and J. Gracewell, "Revolutionizing heart disease prediction with quantum-enhanced machine learning," Scientific Reports, vol. 14, 7453, 2024, doi: 10.1038/s41598-024-55991-w.

[11] M. L. Giger, "Machine Learning in Medical Imaging," Journal of the American College of Radiology, vol. 15, no. 3, pp. 512-520, 2018.

[12] Verma, V.R., Nishad, D.K., Sharma, V., Singh, V.K., Verma, A. & Shah, D.R. (2025). Quantum machine learning for Lyapunov-stabilized computation offloading in next-generation MEC networks. Scientific Reports, 15(1), 405. https://doi.org/10.1038/s41598-024-84441-w

[13] Shrivastava, S., Khalid, S. & Nishad, D.K. (2025). Quantum-inspired control strategies for reducing DC-link voltage fluctuations in DFIG wind energy converters. IEEE Access, 13, 108481. https://doi.org/10.1109/ACCESS.2025.11036086

[14] R. Alhajj and J. Rokne, Encyclopedia of Social Network Analysis and Mining, 2nd ed. Springer, 2018.

[15] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in Advances in Neural Information Processing Systems, 2017, pp. 1024-1034.

[16] Liu, Z., Huang, J., Zhong, Q., She, Y., Ou, R., Li, C., ... & Liu, S. (2017). Network-based analysis of the molecular mechanisms of multiple myeloma and monoclonal gammopathy of undetermined significance. Oncology Letters, 14(4), 4167-4175.

[17] M. Kitsak et al., "Identification of influential spreaders in complex networks," Nature Physics, vol. 6, no. 11, pp. 888-893, 2010.

[18] R. Bharti, A. Khamparia, M. Shabaz, G. Dhiman, S. Pande, and P. Singh, "Prediction of heart disease using a combination of machine learning and deep learning," Computational Intelligence and Neuroscience, vol. 2021, 2021, doi: 10.1155/2021/8387680.

[19] S. Mohan, C. Thirumalai, and G. Srivastava, "Effective Heart Disease Prediction Using Hybrid Machine Learning Techniques," IEEE Access, vol. 7, pp. 81542-81554, 2019, doi: 10.1109/ACCESS.2019.2923707.

[20] A. Basak, M. S. Rahman, and M. Rahman, "Prediction of Heart Disease Using an Approach Based on Machine Learning," in Proc. 13th Int. Conf. Computing Communication and Networking Technologies (ICCCNT), 2022, pp. 1-5.

[21] C. Ganesh, L. G. Durai A, B. Bandarupalli, S. Pavithra, and R. Prabanjan, "Heart Disease Diagnosis Using Deep Learning," in Proc. Int. Conf. Emerging Research in Computational Science (ICERCS), 2023, pp. 1-6.

[22] H. Ullah et al., "An End-to-End Cardiac Arrhythmia Recognition Method with an Effective DenseNet Model on Imbalanced Datasets Using ECG Signal," Computational Intelligence and Neuroscience, vol. 2022, pp. 1-23, 2022.

[23] A. Marengo and V. Santamato, "Quantum algorithms and complexity in healthcare applications: a systematic review with machine learning-optimized analysis," Frontiers in Computer Science, vol. 7, 2025.

[24] 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.

[25] D. Maheshwari et al., "Quantum machine learning applied to electronic healthcare records for ischemic heart disease classification," Human-centric Computing and Information Sciences, vol. 13, no. 06, pp. 1-18, 2023.

[26] A. Janosi, W. Steinbrunn, M. Pfisterer, and R. Detrano, "Heart Disease," UCI Machine Learning Repository, 1989. DOI: 10.24432/C52P4X.

[27] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, Applied Logistic Regression, 3rd ed. Wiley, 2013.

[28] F. Akhtar et al., "Early Coronary Heart Disease Deciphered via Support Vector Machines: Insights from Experiments," in Proc. 20th Int. Computer Conf. Wavelet Active Media Technology and Information Processing (ICCWAMTIP), 2023, pp. 1-7, doi: 10.1109/ICCWAMTIP60502.2023.10387051.

[29] Singh, R., Khalid, S., Nishad, D.K. & Ruchira. (2026). Integrating fuzzy graph theory into cryptography: A survey of techniques and security applications. New Mathematics and Natural Computation, 22(01), 271–290. https://doi.org/10.1142/S179300572650016X

[30] Y. Li et al., “DeepHIV: A Sequence-Based Deep Learning Model for Predicting HIV-1 Protease Cleavage Sites,” IEEE Transactions on Computational Biology and Bioinformatics, vol. 22, no. 6, pp. 3557–3567, 2025. doi: 10.1109/TCBB.2025.3573097

[31] D. Li, F. Zhao, Y. Yang, Z. Cui, P. Hu and L. Hu, "Multi-View Contrastive Learning for Drug-Drug Interaction Event Prediction," in *IEEE Journal of Biomedical and Health Informatics*, vol. 30, no. 3, pp. 2276-2287, March 2026, doi: 10.1109/JBHI.2025.3600045.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- TableCVD.docx