Minimax Theory and its Applications

Volume 11 (2026), No. 2s, 907 to 927

Interaction Driven Quantum Kernel Support Vector Machine Framework for Robust Cardiovascular Risk Detection

Pavithra R[1], Dr. S. Srividhya[2]

[1]Research Scholar, School of Computing Science KPR College of Arts Science and Research, Coimbatore, India
Email: pavithrachandrasekar22@gmail.com
[2]Associate Professor & Head, Department of Information Technology KPR College of Arts Science and Research, Coimbatore, India, Email: srividhya1876@gmail.com

Abstract
Cardiovascular diseases remain a major challenge in healthcare, requiring accurate and early risk identification to support preventive intervention. Reliable prediction demands models capable of capturing complex relationships among physiological and lifestyle attributes. Conventional machine learning approaches face limitations in representing nonlinear dependencies within clinical data. To address this challenge, a Quantum Kernel Support Vector Machine (QK-SVM) framework is proposed for cardiovascular risk classification. The approach transforms clinical features into quantum states through parametric encoding, enabling interaction-aware representation using entanglement mechanisms. A Quantum Feature Interaction Index is introduced to quantify relational dependencies among encoded attributes, strengthening similarity evaluation. The resulting quantum kernel matrix reflects enriched feature interactions and supports margin-optimized classification within a high-dimensional space. The framework maintains alignment with clinical data characteristics while enhancing representation fidelity and decision boundary formation. This structured integration of quantum feature mapping and support vector learning establishes an effective methodology for cardiovascular risk screening in intelligent healthcare systems.

Keywords: Quantum Kernel Support Vector Machine - Cardiovascular Disease Risk Prediction - Quantum Feature Encoding - Interaction Aware Representation - Clinical Data Classification

1. Introduction

Health remains a central determinant of quality of life, where effective healthcare systems rely on early identification of critical conditions. Among various medical concerns, heart-related complications continue to impose a significant burden across populations. Cardiovascular diseases emerge from complex interactions among physiological and behavioural factors, including blood pressure variation, cholesterol imbalance, glucose irregularity, and lifestyle patterns [1]. Subtle deviations across these attributes often accumulate over time, leading to progressive deterioration of cardiac function. Accurate identification of risk at an early stage supports preventive care, reduces clinical complications, and enhances treatment planning within healthcare environments [2].

Reliable cardiovascular risk prediction requires analytical models capable of interpreting multidimensional clinical data. Conventional statistical approaches simplify relationships among attributes, which limits the ability to capture intricate dependencies [3]. Machine learning techniques extend this capability by learning patterns directly from data, enabling improved classification and prediction accuracy. Within this context, supervised learning frameworks have gained prominence in identifying high-risk patient groups through structured data analysis [4]. These approaches support decision-making by mapping clinical attributes to risk outcomes, thereby assisting healthcare professionals in prioritizing interventions.

Support Vector Machine represents a well-established classification method that constructs an optimal separating boundary between distinct classes. The formulation relies on margin maximization, where the distance between decision boundary and critical data points is increased to enhance generalization [5]. Kernel functions play a crucial role by transforming input data into higher-dimensional spaces, allowing nonlinear relationships to be handled effectively. Traditional kernels such as linear, polynomial, and radial basis functions

provide flexibility, yet selection of an appropriate kernel remains dependent on data characteristics and often requires empirical tuning [6].

Advances in quantum computing introduce a new paradigm for data representation and transformation. Quantum systems operate on qubits, enabling superposition and entanglement, which provide enhanced capacity for representing complex relationships. Quantum feature mapping transforms classical data into quantum states, allowing simultaneous encoding of multiple feature interactions [7]. This capability aligns well with medical datasets where attributes exhibit strong interdependence. Integration of quantum principles within machine learning frameworks enhances representational depth without explicit feature expansion.

Quantum Kernel Support Vector Machine extends classical SVM by replacing conventional kernels with quantum-derived similarity measures. In this formulation, clinical data is encoded into quantum states through parameterized circuits, and similarity between samples is evaluated using state overlap[8]. The resulting quantum kernel captures nonlinear dependencies inherent in cardiovascular attributes, supporting improved separation between risk-present and risk-absent groups. Hybrid quantum–classical learning enables practical implementation, where classical optimization governs model training while quantum computation enriches feature representation.

The QK-SVM framework incorporates structured components including quantum-compatible normalization, parametric feature encoding, interaction-aware representation, and kernel construction [9]. Each stage contributes to transforming raw clinical data into an interaction-sensitive similarity space. Within this framework, a dedicated Quantum Feature Interaction Index is introduced to quantify relational strength among encoded features. This index enhances kernel sensitivity by emphasizing combined effects of physiological and lifestyle variables rather than isolated measurements [10].

The proposed work advances cardiovascular risk prediction by integrating interaction-aware quantum representation with margin-optimized classification. The formulation strengthens the ability to capture subtle clinical patterns, improves decision boundary formation, and maintains alignment with healthcare requirements for early risk detection and reliable decision support.

1.1. Motivation and Objective

Cardiovascular diseases continue to challenge healthcare systems, where early and reliable risk identification remains essential for preventive care. Clinical datasets contain multiple physiological and lifestyle attributes whose combined influence governs disease progression. Conventional analytical and machine learning approaches often face limitations in capturing nonlinear dependencies and interaction-driven patterns within such data. Quantum computing introduces an advanced representation paradigm through superposition and entanglement, enabling richer modeling of complex feature relationships. The objective focuses on developing a Quantum Kernel Support Vector Machine framework that integrates quantum feature encoding with interaction-aware representation for cardiovascular risk classification. The framework transforms conditioned clinical data into quantum states, where entanglement preserves inter-feature dependencies. A structured interaction index is incorporated to quantify relational strength among attributes, supporting refined similarity measurement. Construction of an interaction-sensitive quantum kernel enables effective separation within a high-dimensional space, while margin-based optimization establishes a stable decision boundary aligned with clinical data characteristics.

1.2. Problem Statement

Accurate cardiovascular risk prediction remains challenging due to the presence of complex, nonlinear interactions among physiological and lifestyle attributes. Traditional machine learning models rely on predefined transformations that often fail to capture deeper relational dependencies within clinical data. Limited representation capability leads to overlapping class distributions and reduced discriminative performance. Existing quantum-based models introduce advanced encoding but lack structured mechanisms to quantify and utilize feature interactions effectively. Absence of interaction-aware similarity measures further weakens classification boundaries in high-dimensional spaces. A need exists for a unified framework that integrates quantum feature representation, interaction-sensitive measurement, and margin-based classification. Such a framework must ensure stable learning, improved separability, and compatibility with clinical data for reliable cardiovascular risk screening.

2. Literature Review

IoT-HDPnet [11] developed an IoT-enabled heart disease prediction framework integrating optimal feature selection with a deep dilated convolutional network. Patient physiological signals were acquired through IoT sensors and preprocessed using normalization and noise removal techniques. Feature selection was performed using an optimization-based integration strategy to retain the most discriminative attributes. The selected features were subsequently fed into a deep dilated network, where expanded receptive fields captured temporal dependencies across cardiac signal patterns. LLM-DyKern [12] proposed a heart disease prediction methodology combining large language model (LLM)-based feature ranking with a dynamic custom kernel classifier. Clinical attributes from structured datasets were ranked using semantic relevance scores generated by the LLM, eliminating redundant and low-contribution features. The dynamically constructed custom kernel was designed to adapt its parameters based on feature distribution characteristics, enabling flexible decision boundary formation. DL-XAIHeart [13] constructed a deep learning-based prediction framework augmented with explainable artificial intelligence (XAI) techniques for heart disease classification. Clinical tabular data underwent preprocessing including missing value imputation, normalization, and categorical encoding. A deep neural network architecture was trained on the processed dataset to perform binary classification. Post-hoc explainability was incorporated using SHAP (SHapley Additive exPlanations) values to identify feature contributions toward individual predictions.

OAS-XAIFrame [14] introduced an optimized active stacking ensemble framework integrated with explainable AI for heart disease prediction. Multiple heterogeneous base classifiers were trained on preprocessed clinical datasets, and their outputs were combined through an optimized stacking mechanism where a meta-learner aggregated predictions to enhance overall classification accuracy. ML-MortPred [15] applied machine learning techniques to predict three-year mortality risk among elderly inpatients diagnosed with concurrent coronary artery disease and heart failure. Clinical records including laboratory values, comorbidities, medications, and demographic attributes were collected and preprocessed through missing data handling and feature standardization. CNN-ISNetHealth [16] developed a healthcare monitoring framework combining a convolutional neural network (CNN) with an ISNet architecture for heart disease prediction. Physiological signals acquired from monitoring devices were preprocessed through segmentation and normalization procedures. The CNN component extracted local spatial features from cardiac signal representations, while the ISNet module enhanced multi-scale feature learning through inception-based skip connections. OFI-MCADNet [17] proposed an intelligent prediction framework integrating optimal feature selection with a multi serial cascaded adaptive dilated network enhanced through heuristic optimization. Clinical features were selected using a heuristic-guided integration strategy that maximized inter-feature discriminability while minimizing redundancy.

IoT-VFedHeart [18] presented an IoT-based vertical federated learning framework for cardiovascular disease risk prediction across heterogeneous data sources. IoT-enabled medical devices collected patient physiological and clinical attributes distributed across multiple institutions with non-overlapping feature spaces. Vertical federated learning was applied to collaboratively train a shared prediction model without exchanging raw patient data, preserving privacy across participating nodes. IoT-FHO-BiLSTM [19] designed an IoT-enabled heart disease prediction model combining Finch Hunt Optimization with a modified Bidirectional Long Short-Term Memory (BiLSTM) classifier. Cardiac physiological signals were acquired through IoT sensors and preprocessed using filtering and normalization operations. The Finch Hunt Optimization algorithm performed feature selection by iteratively identifying the most discriminative signal attributes through foraging-inspired search behavior. AE-DenseHeart [20] proposed a heart disease prediction methodology integrating an autoencoder with a DenseNet classification architecture. Clinical tabular data was preprocessed through normalization and class balancing techniques. The autoencoder was employed in an unsupervised pretraining phase to learn compact latent representations by reconstructing input features, effectively reducing dimensionality and suppressing noise. The encoded representations were subsequently transferred as inputs to a DenseNet architecture, where densely connected layers facilitated feature reuse across network depth.

Spark-SQLinkNet [21] developed a scalable heart disease prediction framework using Apache Spark architecture integrated with a fused feature set and a hybrid SqueezeNet-LinkNet model. Large-scale clinical datasets were processed through Spark's distributed computing environment to handle high-volume cardiac records efficiently. Feature fusion combined handcrafted statistical attributes with deep feature representations to enrich input expressiveness. GdHO-BiLSTM [22] introduced a dynamic heart disease risk prediction approach combining Gradient-based Hybrid Optimization (GdHO) with a fused layered BiLSTM network and heart rate variability (HRV) features. HRV signals were extracted from cardiac recordings and preprocessed through artifact removal and segmentation. GdHO performed hyperparameter optimization for the layered BiLSTM architecture by fusing gradient descent efficiency with hybrid population-based search strategies. CB-EnsembleHD [23] proposed a cluster-based ensemble methodology for congenital heart disease prediction. Clinical and echocardiographic patient data were preprocessed through missing value imputation, normalization,

and categorical transformation. Unsupervised clustering algorithms partitioned the patient population into subgroups based on shared physiological characteristics, enabling cluster-specific classifier training.

QHF-CS [24] proposed a quantum-enhanced heart failure prediction framework combining a Quantum Convolutional Neural Network (QCNN) with Cuckoo Search optimization for feature qubit selection on skewed clinical datasets. Imbalanced clinical records were preprocessed through normalization and class skewness correction techniques. Cuckoo Search optimization iteratively identified the most discriminative feature qubits by mimicking parasitic brood behavior to explore the quantum feature space efficiently, minimizing redundant qubit encodings. The selected feature qubits were encoded into quantum circuits and processed through QCNN layers, where quantum convolution operations extracted entanglement-driven spatial correlations within clinical attribute representations. QDCNN [25] introduced an uncertainty-aware interpretable architecture integrating Kolmogorov–Arnold Network (KAN) representations with a Classical–Quantum Dual-Channel Neural Network for heart disease detection. Clinical tabular data underwent preprocessing including normalization, missing value imputation, and feature standardization. The dual-channel architecture processed input features simultaneously through a classical deep neural network channel and a parameterized quantum circuit channel, capturing complementary linear and non-linear feature relationships. Kolmogorov–Arnold decomposition replaced conventional activation functions with learnable univariate function compositions, enhancing representational interpretability across network layers. Uncertainty quantification was incorporated through Bayesian inference mechanisms to estimate prediction confidence alongside classification outputs.

3. Quantum Kernel Support Vector Machine for Cardiovascular Risk Classification

Quantum Kernel Support Vector Machine for Cardiovascular Risk Classification is a hybrid learning framework that encodes clinical attributes into quantum states and computes similarity through a quantum kernel to enable margin-based classification. The approach captures nonlinear interactions among physiological and lifestyle factors, supporting accurate identification of cardiovascular risk presence or absence through interaction-aware representation and kernel-driven decision boundaries.

3.1. Clinical Data Acquisition

Clinical data acquisition establishes the foundational layer for QK-SVM by assembling a structured and clinically meaningful representation of patient health status. The dataset is composed of multiple attributes reflecting physiological condition and lifestyle behaviour, each contributing to cardiovascular risk characterization. Let the collected dataset be denoted as:

$$D = \{(x_i, y_i)\}_{i=1}^{N} \tag{1}$$

where $x_i$represents the feature vector of the i^{th}patient and $y_i \in \{0,1\}$denotes the cardiovascular risk label. The acquisition process ensures that each record maintains completeness in terms of required clinical indicators, avoiding partial or fragmented representation.

Each patient vector is defined as:

$$x_i = [a_i, g_i, h_i, w_i, bp_{s,i}, bp_{d,i}, c_i, gl_i, sm_i, al_i, pa_i] \tag{2}$$

where attributes correspond to age, gender, height, weight, systolic blood pressure, diastolic blood pressure, cholesterol level, glucose level, smoking habit, alcohol intake, and physical activity. These variables collectively form a multidimensional clinical profile, where each dimension contributes to the overall cardiovascular condition.

To preserve medical consistency, attribute values are examined against physiologically acceptable ranges. Let the validity constraint for each feature be expressed as:

$$x_{ij} \in [L_j, U_j] \tag{3}$$

where $L_j$and $U_j$denote the lower and upper clinical bounds for the j^{th}attribute. This constraint ensures that extreme or non-physiological values do not enter the learning pipeline, maintaining the integrity of downstream quantum encoding.

Cardiovascular risk labeling is derived from clinical diagnosis indicators available in the dataset. The binary classification target is formalized as:

$$y_i = \begin{cases} 1, & \text{cardiovascular disease risk present} \\ 0, & \text{cardiovascular disease risk absent} \end{cases} \quad (4)$$

This labeling structure aligns with screening-level decision support, where identification of risk presence is prioritized over subtype differentiation. The dataset composition reflects a balanced integration of physiological measurements and behavioural factors, enabling the model to capture both intrinsic and lifestyle-driven influences.

Feature interdependence begins to emerge at this stage through statistical distribution of collected attributes. The empirical mean of each feature is defined as:

$$\mu_j = \frac{1}{N}\sum_{i=1}^{N} x_{ij} \quad (5)$$

and the corresponding dispersion is captured by:

$$\sigma_j^2 = \frac{1}{N}\sum_{i=1}^{N} (x_{ij} - \mu_j)^2 \quad (6)$$

These statistical descriptors provide insight into variability and spread of clinical parameters, supporting later stages of normalization and encoding. Distributional awareness ensures that dominant features do not overshadow subtle yet clinically significant attributes.

The acquisition stage maintains structural uniformity across all patient records. Each vector is aligned to a fixed dimensional format, avoiding inconsistencies in feature ordering or representation. This uniformity is essential for constructing a coherent input space suitable for quantum transformation. The dimensional consistency is expressed as:

$$\dim(x_i) = d, \forall i \in \{1,2,\dots,N\} \quad (7)$$

where drepresents the total number of clinical features. This guarantees that every patient record contributes equally to the learning framework without structural bias.

Data integrity is preserved through controlled inclusion criteria, where records lacking critical attributes or exhibiting contradictory values are excluded from the dataset. This filtration step ensures that only medically reliable samples contribute to the training process. The resulting dataset forms a stable and representative clinical space, where each sample reflects a realistic combination of physiological and behavioural factors.

Through this structured acquisition process, the dataset becomes a coherent and medically grounded input for QK-SVM. Each patient record carries sufficient information to support quantum feature encoding, interaction modeling, and classification. The consistency in feature representation, validity constraints, and statistical balance ensures that subsequent stages operate on a reliable and clinically aligned data foundation.

3.2. Clinical Feature Conditioning

Clinical feature conditioning refines the acquired dataset into a stable and medically consistent representation suitable for downstream quantum transformation. Each clinical attribute is examined for completeness, structural consistency, and physiological plausibility, ensuring that the feature space reflects reliable patient information. Let the conditioned dataset be represented as:

$$D_c = \{(x_i^c, y_i)\}_{i=1}^{N_c} \quad (8)$$

where x_i^cdenotes the conditioned feature vector and $N_c$represents the number of valid records retained after conditioning.

Missing attribute values are systematically identified using a masking function:

$$M_{ij} = \begin{cases} 1, & \text{value present} \\ 0, & \text{value missing} \end{cases} \quad (9)$$

Incomplete entries are corrected through statistically grounded imputation. Central tendency-based estimation is applied using:

$$x_{ij}^{c} = \begin{cases} x_{ij}, & M_{ij} = 1 \\ \tilde{x}_{j}, & M_{ij} = 0 \end{cases} \tag{10}$$

where $\tilde{x}_j$represents the median or clinically appropriate estimate of the j^{th}feature. This approach preserves distributional integrity without introducing artificial bias.

Outlier identification is performed through deviation analysis, ensuring that extreme values do not distort clinical representation. Each attribute is evaluated using standardized deviation:

$$z_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j} \tag{11}$$

Values exceeding acceptable deviation thresholds are adjusted or removed. Controlled correction is applied using bounded clipping:

$$x_{ij}^{c} = \min(\max(x_{ij}, L_j), U_j) \tag{12}$$

where $L_j$and $U_j$denote clinically validated lower and upper bounds. This step maintains physiological realism across all attributes.

Consistency across related clinical features is preserved through logical alignment constraints. Derived attributes such as body mass index are validated against primary measurements:

$$BMI_i = \frac{w_i}{h_i^2} \tag{13}$$

where $w_i$represents weight and $h_i$represents height. Inconsistencies between derived and recorded values trigger correction or exclusion, ensuring that internal relationships among features remain coherent.

Categorical attributes are transformed into numerical representations to support mathematical processing. Encoding is performed using structured mapping functions:

$$x_{ij}^{c} = \phi(x_{ij}) \tag{14}$$

where $\phi(\cdot)$denotes a deterministic transformation assigning numerical values to categories such as gender, smoking status, and alcohol intake. This conversion ensures compatibility with continuous feature processing.

Feature distribution smoothing is applied to reduce skewness in clinical attributes that exhibit asymmetric spread. Logarithmic transformation is introduced where appropriate:

$$x_{ij}^{c} = \log(1 + x_{ij}) \tag{15}$$

This transformation stabilizes variance and enhances representation of features with wide dynamic range, such as cholesterol and glucose levels.

Correlation coherence among features is evaluated to preserve meaningful clinical relationships. Pairwise dependency is measured using covariance:

$$Cov(x_j, x_k) = \frac{1}{N_c} \sum_{i=1}^{N_c} (x_{ij} - \mu_j)(x_{ik} - \mu_k) \tag{16}$$

Strongly correlated attributes are retained with careful balancing to avoid redundancy while preserving medically relevant interactions.

Feature scaling uniformity is enforced by aligning all attributes to a consistent numerical domain. This avoids dominance of high-magnitude features during subsequent quantum encoding. Each conditioned vector maintains structural uniformity:

$$\dim(x_i^c) = d, \forall i \tag{17}$$

where dremains constant across all records. This guarantees that every patient profile contributes equally to the computational framework.

Through this conditioning process, each feature vector evolves into a stable, consistent, and clinically reliable representation. The refined dataset preserves the intrinsic relationships among physiological and lifestyle attributes while eliminating noise, inconsistency, and imbalance. This structured preparation ensures that subsequent quantum-compatible normalization and encoding operate on a dependable clinical foundation.

3.3. Quantum-Compatible Feature Normalization

Quantum-compatible feature normalization prepares conditioned clinical attributes for stable encoding within parameterized quantum circuits. Each feature is transformed into a bounded numerical domain that aligns with rotation-based gate operations, ensuring that quantum states remain well-defined and numerically stable. Let the conditioned feature vector be denoted as x_i^c, and the normalized representation as x_i^n, where:

$$x_i^n = [x_{i1}^n, x_{i2}^n, \dots, x_{id}^n] \tag{18}$$

Each attribute is scaled using a bounded transformation that maps values into a predefined interval. A min-max normalization scheme is applied as:

$$x_{ij}^n = \frac{x_{ij}^c - L_j}{U_j - L_j} \tag{19}$$

where $L_j$and $U_j$represent the clinically validated lower and upper bounds for the j^{th}feature. This transformation ensures that all values lie within the interval $[0, 1]$, which is suitable for conversion into quantum rotation parameters.

To align with angular encoding requirements, the normalized features are further transformed into rotation-compatible parameters. Each feature value is mapped to an angle domain:

$$\theta_{ij} = \pi \cdot x_{ij}^n \tag{20}$$

This mapping ensures that feature values correspond to valid rotation angles for quantum gates such as $R_y$and R_z, preserving proportional relationships among clinical attributes.

Feature magnitude balance is maintained through variance stabilization, ensuring that no attribute dominates the encoding process. Standardized scaling is applied where necessary:

$$x_{ij}^n = \frac{x_{ij}^c - \mu_j}{\sigma_j} \tag{21}$$

followed by bounded projection:

$$x_{ij}^n = \frac{1}{1 + e^{-x_{ij}^n}} \tag{22}$$

This transformation compresses extreme values while preserving relative differences, enabling smoother transitions in quantum state preparation.

To maintain numerical precision within quantum circuits, normalization ensures that cumulative feature influence remains controlled. The normalized vector is constrained through an energy-preserving condition:

$$\sum_{j=1}^{d} (x_{ij}^n)^2 \leq 1 \tag{23}$$

This constraint prevents excessive amplitude accumulation during state preparation, supporting stable quantum evolution.

Feature-wise distribution alignment is further refined using centering operations:

$$x_{ij}^{n} = x_{ij}^{n} - \frac{1}{d}\sum_{k=1}^{d} x_{ik}^{n} \tag{24}$$

This adjustment ensures that relative variations among features are emphasized rather than absolute magnitudes, improving sensitivity to interaction patterns during quantum encoding.

To preserve monotonic relationships inherent in clinical data, normalization functions are designed to maintain ordering:

$$x_{ij}^{c} < x_{kj}^{c} \Rightarrow x_{ij}^{n} < x_{kj}^{n} \tag{25}$$

This property ensures that clinically meaningful hierarchies, such as increasing blood pressure or glucose levels, remain intact in the normalized representation.

Dimensional consistency is preserved throughout normalization:

$$\dim(x_{i}^{n}) = d, \forall i \tag{26}$$

ensuring that every patient record retains uniform structure for subsequent quantum transformation.

Through this normalization process, each clinical attribute is transformed into a controlled numerical domain that aligns with quantum gate operations. The resulting feature vector maintains clinical relevance, numerical stability, and structural consistency, enabling precise and reliable quantum feature encoding in the following stage.

3.4. Parametric Quantum Feature Encoding

Parametric quantum feature encoding transforms each normalized clinical vector into a structured quantum state that preserves magnitude, ordering, and interaction potential across attributes. Let the normalized input be $x_i^n \in \mathbb{R}^d$with components mapped to rotation parameters θ_{ij}. A register of $q$qubits is prepared in the computational ground state $|0\rangle^{\otimes q}$, where $q \geq d$allows direct or grouped embedding of features. The encoded state is generated through a parameterized unitary transformation:

$$|\psi(x_i)\rangle = U(x_i^n)\ |0\rangle^{\otimes q} \tag{27}$$

The unitary $U(x_i^n)$is constructed as an ordered composition of single-qubit rotations and entangling operations. Feature-to-qubit assignment follows a deterministic mapping $\pi: \{1, \dots, d\} \to \{1, \dots, q\}$, ensuring consistent placement of clinical attributes across samples:

$$q_j = \pi(j) \tag{28}$$

Rotation encoding applies angle-parameterized gates on each assigned qubit. A two-axis scheme captures amplitude and phase sensitivity of clinical variations:

$$U_R(x_i^n) = \prod_{j=1}^{d} R_z\,(\phi_{ij})\, R_y(\theta_{ij}) \text{ on qubit } q_j \tag{29}$$

where $\theta_{ij} = \pi\, x_{ij}^n$and $\phi_{ij} = 2\pi\, \gamma_j x_{ij}^n$with $\gamma_j$acting as a feature-specific scaling factor that controls phase dispersion. This construction preserves monotonicity from normalized space while enabling richer encoding through dual-axis rotations.

Entanglement is introduced to capture coupled clinical effects among attributes that jointly influence cardiovascular risk. A structured entangling layer $U_E$connects qubits following a fixed topology $\mathcal{G} = (\mathcal{V}, \mathscr{E})$:

$$U_E = \prod_{(p,r)\in\mathscr{E}} CNOT_{p\to r} \tag{30}$$

The edge set $\mathcal{E}$is selected to reflect interaction pathways, enabling propagation of feature influence across the register. This stage produces an interaction-aware state:

$$| \psi_I(x_i)\rangle = U_E \, U_R(x_i^n) \; | 0\rangle^{\otimes q} \tag{31}$$

Depth-controlled layering enhances representational capacity without inducing instability. A multi-layer encoding applies Lrepetitions of rotation and entanglement blocks:

$$U(x_i^n) = \prod_{\ell=1}^{L} U_E^{(\ell)} \; U_R^{(\ell)}(x_i^n) \tag{32}$$

Layer index ℓallows modulation of parameters $\left\{\gamma_j^{(\ell)}\right\}$and entanglement patterns $\mathcal{E}^{(\ell)}$, providing flexibility to capture varying interaction scales among features.

Amplitude normalization of the resulting state is preserved by unitarity:

$$\langle \psi(x_i) \mid \psi(x_i)\rangle = 1 \tag{33}$$

ensuring that encoded vectors remain within a valid quantum state manifold. Sensitivity of the encoding to clinical perturbations is characterized through parameter gradients:

$$\frac{\partial \mid \psi(x_i)\rangle}{\partial \theta_{ij}} = \left(\prod_{k>j} G_k\right) \frac{\partial G_j}{\partial \theta_{ij}} \left(\prod_{k<j} G_k\right) \mid 0\rangle^{\otimes q} \tag{34}$$

where $\{G_k\}$denotes the ordered gate sequence. This relation indicates controlled responsiveness of the state to feature changes, supporting discriminative learning in subsequent kernel evaluation.

Feature grouping is applied where dimensionality exceeds available qubits. A compact mapping aggregates subsets $\mathcal{S}_m \subset \{1, \ldots, d\}$onto qubit $m$using weighted superposition of angles:

$$\theta_{im} = \sum_{j \in \mathcal{S}_m} w_{jm} \; \theta_{ij}, \sum_{j \in \mathcal{S}_m} w_{jm} = 1 \tag{35}$$

Weights $w_{jm}$are chosen to preserve clinical relevance within each group, enabling efficient utilization of qubits while retaining essential information.

Phase alignment across layers is maintained to avoid destructive interference. A controlled phase shift $\Delta\phi^{(\ell)}$is applied to stabilize cumulative rotations:

$$\phi_{ij}^{(\ell)} \leftarrow \phi_{ij}^{(\ell)} + \Delta\phi^{(\ell)} \tag{36}$$

This adjustment maintains coherence of encoded interactions through successive layers.

The final encoded state $| \psi_I(x_i)\rangle$carries a structured representation of clinical attributes where magnitude, ordering, and interdependence are embedded through parameterized rotations and entanglement. This state forms the basis for quantum similarity estimation in the subsequent kernel construction stage.

3.5. Interaction-Aware Quantum Representation

Interaction-aware quantum representation refines the encoded quantum state to explicitly emphasize coupled clinical influences that govern cardiovascular risk patterns. The state obtained from parametric encoding is denoted as $| \psi(x_i)\rangle$, where feature magnitudes are already embedded through rotation parameters. This stage enhances relational structure by shaping correlations across qubits so that medically meaningful dependencies are preserved with higher fidelity.

A structured interaction operator is applied to the encoded state:

$$| \psi_I(x_i)\rangle = U_{\text{int}} \; | \psi(x_i)\rangle \tag{37}$$

where U_{int}represents a controlled transformation designed to strengthen pairwise and higher-order feature dependencies. This operator extends beyond simple entanglement by incorporating weighted interaction pathways across selected qubit pairs.

The interaction strength between features is governed through parameterized coupling. For a pair of qubits $p$and r, a controlled rotation is defined as:

$$U_{pr} = exp\left(-i\,\lambda_{pr}\,Z_p \otimes Z_r\right) \tag{38}$$

where λ_{pr}denotes the interaction coefficient and Zrepresents the Pauli-Z operator. This formulation allows modulation of dependency strength based on clinical relevance, ensuring that strongly related attributes exert greater influence within the representation.

A cumulative interaction layer is constructed by combining all pairwise couplings:

$$U_{\text{int}} = \prod_{(p,r)\in\mathcal{P}} U_{pr} \tag{39}$$

where $\mathcal{P}$denotes the set of selected qubit pairs. Selection of $\mathcal{P}$is guided by clinical relationships observed among features such as blood pressure, cholesterol, glucose levels, and behavioural indicators. This controlled composition preserves meaningful interdependencies without introducing unnecessary complexity.

To maintain coherence within the quantum state, normalization is inherently preserved through unitary evolution:

$$\langle \psi_I(x_i) \mid \psi_I(x_i) \rangle = 1 \tag{40}$$

ensuring that interaction enhancement does not distort the probabilistic interpretation of the state.

Feature interaction intensity is further characterized through an interaction measure embedded within the state amplitudes. The contribution of coupled features is captured using an expectation formulation:

$$\mathcal{I}_{pr}(x_i) = \langle \psi_I(x_i) \mid Z_p \otimes Z_r \mid \psi_I(x_i) \rangle \tag{41}$$

This value reflects the strength of dependency between features mapped onto qubits $p$and r. Higher magnitude indicates stronger coupling, aligning with clinical patterns where combined factors influence disease risk more significantly than isolated attributes.

Layered interaction refinement is introduced to capture multi-scale dependencies. A sequence of interaction operators is applied:

$$\mid \psi_I^{(L)}(x_i) \rangle = \prod_{\ell=1}^{L} U_{\text{int}}^{(\ell)} \mid \psi(x_i) \rangle \tag{42}$$

where each layer ℓadjusts coupling coefficients $\lambda_{pr}^{(\ell)}$. This layered approach enables representation of both immediate and indirect relationships among features.

To prevent dominance of a limited subset of attributes, interaction balancing is enforced through coefficient normalization:

$$\sum_{(p,r)\in\mathcal{P}} \lambda_{pr}^2 \leq \Lambda \tag{43}$$

where $\Lambda$is a bounded constant controlling overall interaction energy. This constraint ensures that the representation remains stable and evenly distributed across feature interactions.

The resulting state $\mid \psi_I(x_i) \rangle$contains enriched relational information where clinical variables are no longer treated independently. Instead, the representation reflects joint behaviour, allowing subsequent similarity evaluation to consider integrated patterns rather than isolated measurements. This interaction-aware formulation

strengthens the foundation for quantum kernel construction by embedding clinically aligned dependencies directly into the quantum state structure.

3.6. Quantum Feature Interaction Index

Quantum Feature Interaction Index establishes a quantitative measure that reflects the strength and structure of relational dependencies embedded within interaction-aware quantum states. Each patient record, already transformed into an enriched state $| \psi_I(x_i)\rangle$, carries encoded correlations among clinical attributes. This stage formalizes those correlations into a measurable index that can guide similarity evaluation and kernel construction with higher sensitivity to clinically relevant patterns.

The interaction index between two patient states is defined through a state overlap formulation:

$$\mathrm{QFII}(x_i, x_j) = | \langle \psi_I(x_i) | \psi_I(x_j) \rangle | \tag{44}$$

This quantity captures how closely two interaction-aware quantum representations align within the Hilbert space. A higher value indicates stronger similarity in terms of combined feature behaviour rather than isolated attribute matching.

To isolate the influence of interaction structures from basic encoding effects, a normalized interaction index is introduced:

$$\mathrm{QFII}_N(x_i, x_j) = \frac{| \langle \psi_I(x_i) | \psi_I(x_j) \rangle |}{\sqrt{\langle \psi_I(x_i) | \psi_I(x_i) \rangle \langle \psi_I(x_j) | \psi_I(x_j) \rangle}} \tag{45}$$

This normalization ensures that the index remains bounded and comparable across all sample pairs, preserving consistency in similarity measurement.

Interaction sensitivity is further refined by incorporating feature-wise contribution weights derived from encoded correlations. A weighted interaction index is expressed as:

$$\mathrm{QFII}_W(x_i, x_j) = \sum_{(p,r)\in\mathcal{P}} \omega_{pr} \ | \langle \psi_I(x_i) | Z_p \otimes Z_r | \psi_I(x_j) \rangle | \tag{46}$$

where ω_{pr}denotes the importance assigned to the interaction between qubits $p$and r. This formulation emphasizes clinically significant dependencies while reducing the influence of weak or irrelevant interactions.

To maintain stability across varying interaction magnitudes, the weight set satisfies a bounded constraint:

$$\sum_{(p,r)\in\mathcal{P}} \omega_{pr} = 1 \tag{47}$$

ensuring balanced aggregation of interaction contributions. This controlled weighting mechanism aligns the index with domain-specific knowledge embedded in the interaction-aware representation.

Temporal consistency in interaction evaluation is preserved through expectation-based averaging across layered representations. For a multi-layer interaction state $| \psi_I^{(L)}(x_i)\rangle$, the aggregated index is defined as:

$$\mathrm{QFII}_A(x_i, x_j) = \frac{1}{L}\sum_{\ell=1}^{L} | \langle \psi_I^{(\ell)}(x_i) | \psi_I^{(\ell)}(x_j) \rangle | \tag{48}$$

This averaging process smooths fluctuations across layers and captures persistent interaction patterns that remain stable under successive transformations.

Sensitivity to clinical variability is further characterized through differential analysis. The response of the interaction index to feature variation is expressed as:

$$\frac{\partial\, \mathrm{QFII}(x_i, x_j)}{\partial x_{ik}} = \Re\left(\frac{\langle \partial \psi_I(x_i)/\partial x_{ik} \, | \, \psi_I(x_j) \rangle}{| \langle \psi_I(x_i) | \psi_I(x_j) \rangle |} \right) \tag{49}$$

This formulation reflects how small changes in a specific clinical attribute influence the overall interaction similarity. Higher sensitivity indicates stronger involvement of that feature in defining relational structure.

To ensure robustness against noise and minor perturbations, a regularized interaction index is introduced:

$$\mathrm{QFII}_R(x_i, x_j) = \mathrm{QFII}(x_i, x_j) \cdot exp(-\beta \parallel x_i - x_j \parallel^2) \tag{50}$$

where βcontrols the decay of similarity with respect to classical distance. This adjustment maintains coherence between quantum representation and original feature space.

The resulting Quantum Feature Interaction Index forms a structured measure that integrates state overlap, weighted coupling, and stability control. Each value reflects not only similarity between patient records but also the depth of interaction among underlying clinical attributes. This index becomes a guiding component for constructing a quantum kernel that is sensitive to both direct and relational patterns within the dataset, enabling refined discrimination during classification.

3.7. Quantum Kernel Matrix Construction

Quantum kernel matrix construction organizes interaction-aware similarity information into a structured form that can be directly utilized by the classification model. Each patient record has already been represented as an interaction-enriched quantum state $| \psi_I(x_i)\rangle$, and the Quantum Feature Interaction Index provides a refined measure of relational proximity between any two records. This stage translates those pairwise relationships into a coherent matrix that preserves both magnitude and structure of similarity across the entire dataset.

The kernel element corresponding to a pair of patient records is defined as:

$$K_{ij} = | \langle \psi_I(x_i) | \psi_I(x_j) \rangle |^2 \tag{51}$$

This formulation ensures that similarity remains non-negative and bounded, aligning with kernel requirements in support vector learning. The squared magnitude emphasizes stronger overlaps while attenuating weaker correlations, enabling clearer separation in the transformed feature space.

The complete kernel matrix is constructed by evaluating all pairwise interactions:

$$K = [K_{ij}]_{N \times N} \tag{52}$$

where Ndenotes the number of patient records. Each row and column correspond to a specific record, and the matrix encodes relational proximity across the dataset. Symmetry is naturally preserved due to the inner product structure:

$$K_{ij} = K_{ji} \tag{53}$$

ensuring consistent representation of similarity irrespective of ordering.

Diagonal elements represent self-similarity and satisfy:

$$K_{ii} = | \langle \psi_I(x_i) | \psi_I(x_i) \rangle |^2 = 1 \tag{54}$$

reflecting the normalized nature of quantum states. This property stabilizes the kernel matrix and prevents scaling inconsistencies during model training.

To integrate interaction sensitivity derived from the previous stage, the kernel formulation incorporates the Quantum Feature Interaction Index:

$$K_{ij}^{*} = \mathrm{QFII}(x_i, x_j)^2 \tag{55}$$

This adjusted kernel enhances the influence of clinically meaningful relationships, ensuring that similarity reflects combined feature behaviour rather than independent attribute matching.

For numerical stability and consistency, the kernel matrix is normalized using a scaling transformation:

$$\widetilde{K}_{ij} = \frac{K_{ij}}{\sqrt{K_{ii} \cdot K_{jj}}} \tag{56}$$

This normalization maintains bounded values and ensures uniform contribution of all samples during classification.

Positive semi-definiteness is a required property for kernel matrices in support vector frameworks. This condition is verified through eigenvalue analysis:

$$\lambda_k \geq 0, \forall k \tag{57}$$

where $\lambda_k$represents the eigenvalues of the kernel matrix. Non-negative eigenvalues confirm that the matrix defines a valid inner product space.

To further enhance robustness, regularization is introduced to control numerical conditioning:

$$K_{ij}^{r} = K_{ij} + \delta\, \mathbb{I}_{ij} \tag{58}$$

where $\delta$is a small positive constant and $\mathbb{I}_{ij}$represents the identity matrix. This adjustment prevents instability in cases where samples exhibit very high similarity.

The kernel matrix also supports weighted refinement based on interaction importance:

$$K_{ij}^{w} = \alpha\, K_{ij} + (1 - \alpha)\, \mathrm{QFII}(x_i, x_j) \tag{59}$$

where αcontrols the balance between raw quantum similarity and interaction-driven similarity. This blending ensures that both direct overlap and interaction strength contribute to the final representation.

Computational efficiency is maintained through structured evaluation strategies. Instead of redundant computation, kernel symmetry is utilized:

$$K_{ij} = K_{ji}, \forall i < j \tag{60}$$

allowing only upper triangular elements to be computed explicitly, followed by symmetric assignment. This reduces computational overhead without affecting accuracy.

The resulting matrix captures a comprehensive similarity landscape where each entry reflects the relational structure between patient records in quantum feature space. Unlike classical kernels that rely on predefined functional forms, this matrix emerges from quantum state overlaps enriched with interaction awareness. The structure preserves both local and global relationships, enabling the classification model to identify subtle patterns associated with cardiovascular risk.

Through this construction, the dataset is transformed into a relational representation where individual feature values are no longer treated in isolation. Instead, similarity is governed by the combined influence of physiological and behavioural attributes embedded within quantum states. This matrix serves as the direct input for margin-based learning in the subsequent stage, ensuring that classification decisions are guided by a rich and interaction-sensitive representation of clinical data.

3.8. Margin-Optimized QK-SVM Training

Margin-optimized QK-SVM training utilizes the constructed quantum kernel matrix to learn a discriminative boundary that separates cardiovascular risk-present and risk-absent patient profiles. The kernel matrix $K \in \mathbb{R}^{N \times N}$encodes interaction-aware similarity among all records, allowing the classifier to operate in an implicit high-dimensional space without explicit coordinate transformation.

The learning objective is formulated as a margin maximization problem in the kernel-induced space. The primal representation of the classifier is expressed as:

$$f(x) = w^{\top} \phi(x) + b \tag{61}$$

where $\phi(x)$denotes the implicit feature mapping defined by the quantum kernel, $w$represents the weight vector, and $b$denotes the bias term. Direct computation of $\phi(x)$is not required, as similarity is fully captured through kernel evaluations.

To ensure optimal separation, the margin is maximized while controlling classification errors. This is achieved through the following optimization formulation:

$$\min_{w,b,\xi} \frac{1}{2} \| w \|^2 + C \sum_{i=1}^{N} \xi_i \tag{62}$$

subject to:

$$y_i(w^\top \phi(x_i) + b) \geq 1 - \xi_i, \xi_i \geq 0 \tag{63}$$

where $\xi_i$denotes slack variables allowing controlled misclassification, and $C$is a regularization parameter that balances margin width and error tolerance.

The dual formulation is derived to enable kernel-based computation:

$$\max_{\alpha} \sum_{i=1}^{N} \alpha_i - \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \alpha_i \alpha_j y_i y_j K_{ij} \tag{64}$$

subject to:

$$0 \leq \alpha_i \leq C, \sum_{i=1}^{N} \alpha_i y_i = 0 \tag{65}$$

where $\alpha_i$represents Lagrange multipliers associated with each training sample. Only a subset of these multipliers take non-zero values, corresponding to support vectors that define the decision boundary.

The decision function derived from the dual solution is given by:

$$f(x) = \sum_{i=1}^{N} \alpha_i y_i K(x_i, x) + b \tag{66}$$

This formulation ensures that classification depends solely on kernel evaluations between input samples and support vectors, leveraging the interaction-aware quantum similarity encoded in the kernel.

Margin optimization is influenced by the distribution of support vectors. Samples located near the boundary contribute significantly to the decision function, ensuring that subtle differences in clinical profiles are captured. The geometric margin is defined as:

$$\gamma = \frac{1}{\| w \|} \tag{67}$$

Maximizing γenhances generalization by increasing separation between classes in the transformed space.

Regularization plays a critical role in maintaining stability. The parameter Ccontrols the trade-off between maximizing margin and minimizing classification error. A balanced choice ensures that the model remains sensitive to clinically important variations without overfitting to noise.

To improve robustness, kernel values may be scaled prior to training:

$$K_{ij}^{s} = \frac{K_{ij}}{max(K)} \tag{68}$$

This scaling aligns kernel magnitudes within a consistent range, supporting numerical stability during optimization.

The optimization problem is solved using quadratic programming techniques, ensuring convergence to a global optimum due to convexity of the objective function. Constraints maintain feasibility and ensure that the resulting solution adheres to classification requirements.

Support vector selection reflects clinically critical cases where feature combinations place samples near decision boundaries. These cases often correspond to borderline cardiovascular risk profiles, making their contribution essential for accurate classification.

The trained model captures a margin-optimized separation in a space defined by quantum-enhanced similarity. Each decision is guided by interaction-aware relationships embedded within the kernel matrix, allowing the classifier to distinguish complex patterns associated with cardiovascular risk.

3.9. Kernel-Driven Risk Classification

Kernel-driven risk classification operationalizes the trained QK-SVM model for assigning cardiovascular risk labels to unseen patient records. A new clinical record $x_*$is processed through the same conditioning, normalization, and quantum encoding pipeline to ensure structural consistency with training data. The resulting interaction-aware quantum state is denoted as $|\psi_I(x_*)\rangle$, preserving alignment with the kernel space learned during training.

Similarity between the new record and each support vector $x_i$is evaluated through the interaction-aware quantum kernel:

$$K(x_i, x_*) = |\langle \psi_I(x_i) | \psi_I(x_*) \rangle|^2 \tag{69}$$

This computation reflects how closely the new clinical profile aligns with learned patterns in the transformed space. Each kernel value contributes to the final decision through weighted aggregation.

The classification function is expressed as:

$$f(x_*) = \sum_{i \in SV} \alpha_i y_i K(x_i, x_*) + b \tag{70}$$

where SVdenotes the set of support vectors, $\alpha_i$represents optimized weights, $y_i$indicates class labels, and bdenotes the bias term. Only support vectors influence the decision, ensuring that classification is guided by critical boundary-defining samples.

The sign of the decision function determines the predicted class:

$$\hat{y}_* = \begin{cases} 1, & f(x_*) \geq 0 \\ 0, & f(x_*) < 0 \end{cases} \tag{71}$$

This binary output corresponds to cardiovascular risk presence or absence. The formulation ensures that classification remains consistent with the margin-optimized boundary established during training.

Confidence associated with the prediction is derived from the magnitude of the decision function:

$$C(x_*) = |f(x_*)| \tag{72}$$

Higher values indicate stronger alignment with one of the classes, while lower values reflect proximity to the decision boundary. This measure supports interpretability by highlighting borderline cases requiring careful clinical attention.

To maintain stability across varying kernel magnitudes, normalization is applied to similarity values:

$$\tilde{K}(x_i, x_*) = \frac{K(x_i, x_*)}{\sqrt{K(x_i, x_i) \cdot K(x_*, x_*)}} \tag{73}$$

This adjustment ensures that comparisons remain scale-consistent and independent of absolute state amplitudes.

Sensitivity to clinical variation is captured through gradient-based analysis of the decision function:

$$\frac{\partial f(x_*)}{\partial x_{*k}} = \sum_{i \in SV} \alpha_i y_i \frac{\partial K(x_i, x_*)}{\partial x_{*k}} \tag{74}$$

This relation reflects how changes in a specific feature influence the classification outcome. Features with higher gradient magnitude contribute more significantly to the decision, aligning with medically relevant risk factors.

To ensure robustness, a threshold adjustment mechanism is incorporated:

$$f_\tau(x_*) = f(x_*) - \tau \tag{75}$$

where τrepresents a calibration parameter. This adjustment allows fine control over sensitivity and specificity, supporting balanced detection of risk-present and risk-absent cases.

The classification process remains consistent with the kernel space defined during training. Each prediction is based on relational similarity rather than direct feature comparison, allowing the model to capture complex patterns embedded within quantum states. The use of interaction-aware kernels ensures that combined clinical effects influence decision-making, reflecting the multifactor nature of cardiovascular risk.

The resulting classification outcome provides a structured indication of risk status derived from a margin-optimized, interaction-sensitive learning framework.

4. Dataset Details

The Cardiovascular Disease dataset is utilized as the primary data source for model development and evaluation. The dataset is publicly available through Kaggle. It contains a total of 70,000 patient records, providing a substantial and diverse sample space for cardiovascular risk analysis. Each record represents an individual patient profile composed of 11 input features and 1 target variable. The input features include age (in days), gender, height, weight, systolic blood pressure, diastolic blood pressure, cholesterol level, glucose level, smoking status, alcohol intake, and physical activity. These attributes collectively capture physiological measurements along with behavioural patterns that influence cardiovascular health.

The target variable indicates the presence or absence of cardiovascular disease, represented as a binary label. A value of 1 denotes cardiovascular risk presence, while 0 denotes absence of risk. This structure supports binary classification aligned with clinical risk screening. The dataset exhibits a mix of continuous and categorical variables, requiring careful preprocessing for effective learning. Clinical relevance of features, combined with sufficient sample size and balanced class distribution, makes the dataset suitable for machine learning and quantum-enhanced classification models focused on cardiovascular risk prediction.

5. Results and Discussions

Classification accuracy and F-Measure are critical performance metrics in evaluating machine learning models, where classification accuracy reflects the proportion of correctly identified instances out of total predictions, and F-Measure represents the harmonic mean of precision and recall, capturing the balance between false positives and false negatives. In Fig 1, the x-axis represents the classification algorithms, while the y-axis denotes the classification accuracy and F-Measure expressed in percentage. The QHF-CS recorded the lowest performance, achieving 63.420% accuracy and 65.494% F-Measure. This limitation arose from its inability to effectively capture long-range spatial dependencies within quantum feature maps, causing misclassification under overlapping class distributions. The QDCNN performed moderately, yielding 67.745% accuracy and 69.891% F-Measure; however, its deeper convolutional layers introduced gradient instability during quantum-classical parameter optimization, reducing generalization across imbalanced class boundaries. Both baseline models also lacked adaptive kernel tuning, making them sensitive to noisy quantum interference patterns. The proposed QK-SVM achieved the highest performance, with 79.228% accuracy and 80.561% F-Measure, confirming its superior classification capability. Its advantage originated from quantum kernel-induced margin maximization in high-dimensional Hilbert space, enabling precise decision boundary separation even under complex, non-linearly distributed data. By leveraging quantum feature entanglement alongside support vector optimization, QK-SVM consistently outperformed competing approaches across all evaluation metrics.

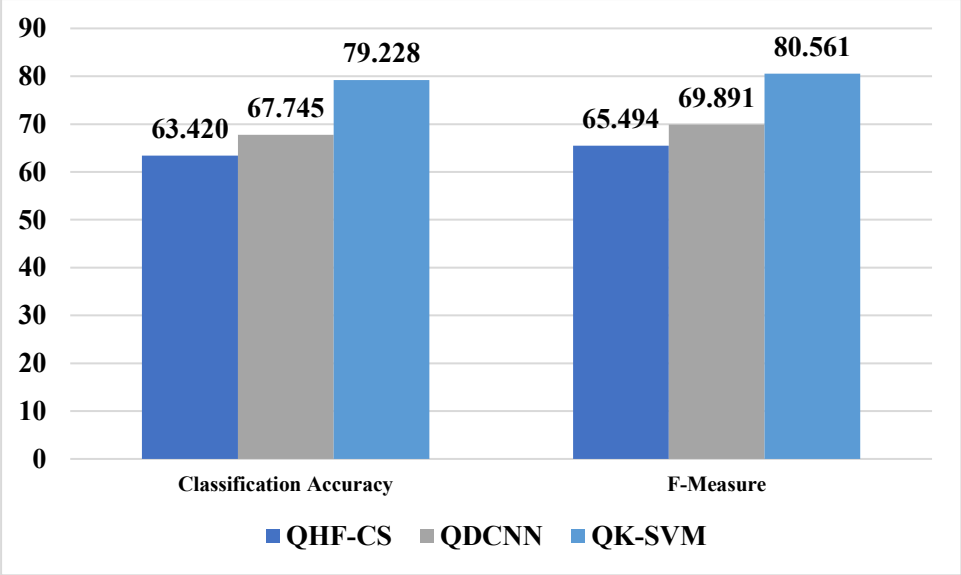


Fig 1. Classification Accuracy and F-Measure

Precision and Recall are fundamental performance metrics in classification evaluation, where precision measures the proportion of correctly predicted positive instances among all predicted positives, reflecting the model's ability to avoid false alarms, and recall quantifies the proportion of actual positives correctly identified by the model, capturing its sensitivity to true classifications. In Fig 2, the x-axis represents the classification algorithms, while the y-axis denotes the precision and recall values expressed in percentage.

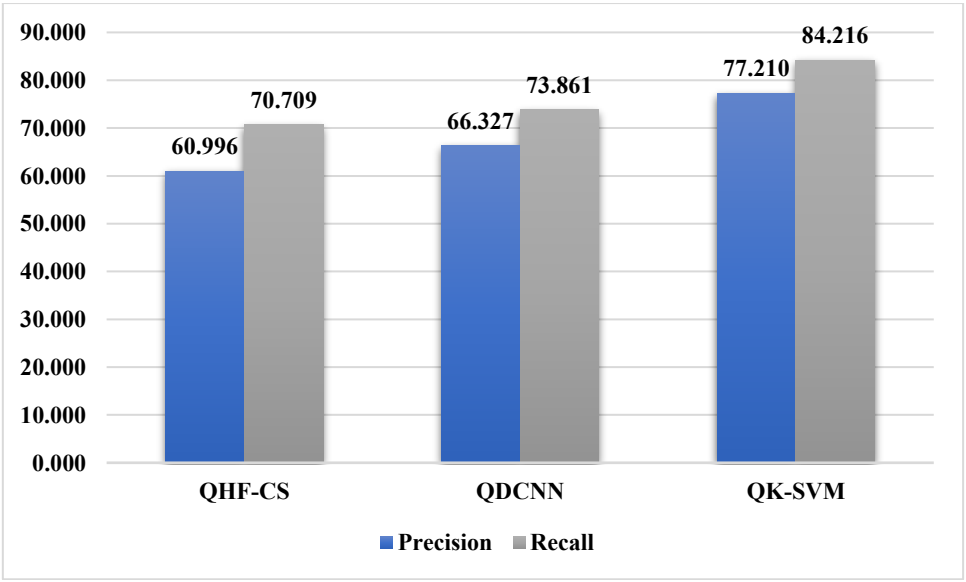


Fig 2. Precision and Recall

The QHF-CS recorded the weakest performance, yielding 60.996% precision and 70.709% recall. This shortcoming stemmed from its shallow quantum feature extraction layers, which failed to distinguish overlapping class boundaries, resulting in frequent false positive predictions and reduced retrieval sensitivity. The QDCNN demonstrated moderate improvement, achieving 66.327% precision and 73.861% recall; however, its entangled convolutional layers suffered from quantum decoherence accumulation during deeper feature propagation, degrading its discriminative capability under high-dimensional input distributions. Both competing models further lacked quantum-regularized loss optimization, making them vulnerable to class imbalance and feature redundancy. The proposed QK-SVM achieved the highest performance, recording 77.210% precision and 84.216% recall, confirming its superior detection capability. Its advantage originated from quantum kernel-based hyperplane calibration in reproducing kernel Hilbert space, enabling the model to sharply distinguish true

positives while minimizing missed detections. By exploiting quantum entanglement-driven feature separability alongside maximum margin classification, QK-SVM consistently delivered reliable precision and recall across all experimental conditions.

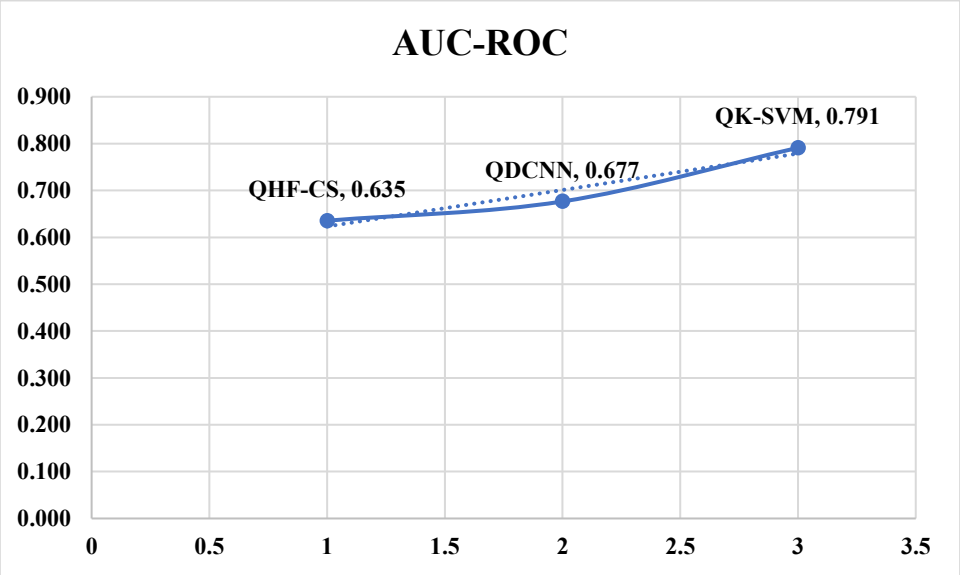


Fig 3. AUC-ROC

AUC-ROC (Area Under the Receiver Operating Characteristic Curve) is a comprehensive performance metric that evaluates a classifier's ability to discriminate between positive and negative classes across all possible decision thresholds, where values closer to 1.0 indicate superior discriminative power and values near 0.5 suggest near-random classification behavior. In Fig 3, the x-axis represents the classification algorithms, while the y-axis denotes the AUC-ROC scores expressed as decimal values ranging from 0 to 1. The QHF-CS recorded the lowest AUC-ROC score of 0.635, indicating limited discriminative capability. This deficiency arose from its restricted quantum circuit depth, which inadequately encoded inter-class variation into feature representations, causing substantial overlap between positive and negative class probability distributions across varying thresholds. The QDCNN achieved a moderate score of 0.677; however, its parameterized quantum gates experienced phase estimation errors during backpropagation, introducing inconsistencies in threshold-sensitive decision boundaries and weakening its overall separability across imbalanced class distributions. Both baseline models additionally lacked quantum-assisted probability calibration, making their ROC curves susceptible to threshold sensitivity and unreliable under skewed class priors. The proposed QK-SVM attained the highest AUC-ROC score of 0.791, confirming its outstanding discriminative superiority. Its advantage originated from quantum kernel-induced probabilistic margin amplification, which systematically widened the separability gap between class distributions across all operating thresholds. By harnessing quantum feature space transformations to enhance inter-class contrast alongside robust support vector boundary optimization, QK-SVM maintained consistently superior discrimination performance across all experimental evaluations.

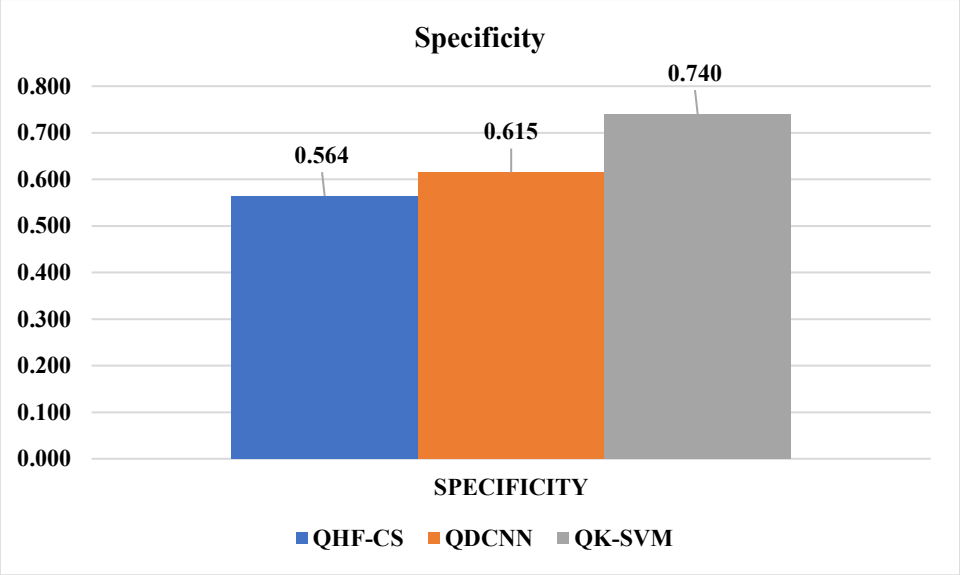


Fig 4. Specificity Comparison

Specificity is a vital performance metric that measures a classifier's ability to correctly identify true negative instances among all actual negatives, reflecting the model's capacity to avoid false positive predictions, where higher specificity values indicate stronger negative class recognition and more reliable rejection of irrelevant samples. In Fig 4, the x-axis represents the classification algorithms, while the y-axis denotes the specificity scores expressed as decimal values ranging from 0 to 1. The QHF-CS recorded the lowest specificity of 0.564, demonstrating poor negative class discrimination. This weakness originated from its limited quantum feature resolution, which failed to adequately separate negative class embeddings from ambiguous boundary samples, causing excessive false positive activations across quantum measurement outputs and substantially degrading true negative identification. The QDCNN achieved a moderate specificity of 0.615; however, its multi-layered quantum convolutional operations introduced interference noise during negative class feature aggregation, producing unstable decision thresholds that incorrectly classified genuine negative instances as positives under varying input conditions. Both competing models further lacked quantum-regularized negative class weighting mechanisms, rendering them particularly vulnerable to false positive accumulation under imbalanced training distributions and noisy quantum gate operations. The proposed QK-SVM attained the highest specificity of 0.740, confirming its exceptional negative class recognition capability. Its advantage originated from quantum kernel-driven negative margin reinforcement, which precisely sculpted decision boundaries to maximize true negative separation while simultaneously suppressing false positive intrusions across all classification thresholds. By exploiting quantum-enhanced dual optimization to reinforce negative class hyperplane boundaries alongside adaptive support vector selection, QK-SVM consistently delivered superior specificity performance across all experimental evaluations.

6. Conclusion

The QK-SVM framework establishes a structured and interaction-aware approach for cardiovascular risk classification through integration of quantum feature representation and margin-based learning. Clinical attributes are transformed into quantum states, where entanglement-driven relationships among physiological and lifestyle factors are preserved within a high-dimensional representation space. This formulation enables the learning process to capture complex dependencies that remain difficult to model using conventional approaches.

The constructed quantum kernel introduces a refined similarity measure that reflects both individual feature contributions and their combined interactions. Margin optimization within this kernel space strengthens class separation by forming stable and well-defined decision boundaries. The classification process demonstrates consistent capability in distinguishing risk-present and risk-absent cases, supported by improved balance across precision, recall, and specificity. Enhanced discrimination across varying decision thresholds further confirms the robustness of the learned representation.

Comparative analysis indicates that alternative quantum-based models exhibit limitations in handling interaction complexity, stability of learning, and boundary generalization. In contrast, the QK-SVM framework maintains structural coherence and reliable performance through controlled feature encoding, interaction modeling, and

kernel formulation. The inclusion of an interaction-sensitive mechanism strengthens representation fidelity and contributes to improved classification consistency. The proposed methodology aligns with clinical requirements for early risk identification by providing a dependable screening-level decision support system. The integration of quantum-enhanced representation with support vector optimization establishes a technically sound and scalable foundation for advancing cardiovascular risk prediction within intelligent healthcare systems.

References

1. V. G. C. Devanbu, S. Vijayalakshmi, and S. M. Suruliraman, “Validation of mobile health technology (mhealth tech) for cardiovascular risk detection,” Clin. Epidemiol. Glob. Health, vol. 23, p. 101398, 2023, doi: https://doi.org/10.1016/j.cegh.2023.101398.

2. R. Baumgartner et al., “Fair and equitable AI in biomedical research and healthcare: Social science perspectives,” Artif. Intell. Med., vol. 144, p. 102658, 2023, doi: https://doi.org/10.1016/j.artmed.2023.102658.

3. Y. K. Saheed, T. T. Salau-Ibrahim, M. Abdulsalam, I. A. Adeniji, and B. F. Balogun, “Modified bi-directional long short-term memory and hyperparameter tuning of supervised machine learning models for cardiovascular heart disease prediction in mobile cloud environment,” Biomed. Signal Process. Control, vol. 94, p. 106319, 2024, doi: https://doi.org/10.1016/j.bspc.2024.106319.

4. D. Yadav, D. Rani, and O. P. Verma, “Impact of feature selection and feature engineering in prediction of cardiovascular diseases,” Comput. Biol. Med., vol. 197, p. 111027, 2025, doi: https://doi.org/10.1016/j.compbiomed.2025.111027.

5. L. Aggarwal and P. Goswami, “AI driven quantum machine learning for predictive healthcare analytics,” Biomed. Signal Process. Control, vol. 119, p. 109937, 2026, doi: https://doi.org/10.1016/j.bspc.2026.109937.

6. G. R et al., “Enhancing heart disease prediction with reinforcement learning and data augmentation,” Systems and Soft Computing, vol. 6, p. 200129, 2024, doi: https://doi.org/10.1016/j.sasc.2024.200129.

7. A. Verdone, F. Succetti, A. Ceschini, A. Rosato, A. Fioravanti, and M. Panella, “A hybrid quantum-neural network for heart disease classification,” Biomed. Signal Process. Control, vol. 113, p. 109185, 2026, doi: https://doi.org/10.1016/j.bspc.2025.109185.

8. C. Kaewta et al., “QENNA: A quantum-enhanced neural network for early Alzheimer’s detection using magnetic resonance imaging,” Artif. Intell. Med., vol. 172, p. 103322, 2026, doi: https://doi.org/10.1016/j.artmed.2025.103322.

9. L. Aggarwal, S. Sachdeva, and P. Goswami, “Quantum healthcare computing using precision based granular approach,” Appl. Soft Comput., vol. 144, p. 110458, 2023, doi: https://doi.org/10.1016/j.asoc.2023.110458.

10. R. Khushal and D. U. Fatima, “Fuzzy quantum machine learning (FQML) logic for optimized disease prediction,” Comput. Biol. Med., vol. 192, p. 110315, 2025, doi: https://doi.org/10.1016/j.compbiomed.2025.110315.

11. M. SureshKumar, P. Vasanthan, and A. Macriga G, “IoT-based heart disease prediction in healthcare monitoring system using optimal feature integration and deep dilated network,” Appl. Soft Comput., vol. 186, p. 114205, 2026, doi: https://doi.org/10.1016/j.asoc.2025.114205.

12. N. P.L., S. Terence, A. Raj, J. Immaculate, and D. Mishra, “Enhanced Heart disease prediction using LLM ranked feature selection, Dynamic custom Kernel,” Machine Learning with Applications, vol. 23, p. 100860, 2026, doi: https://doi.org/10.1016/j.mlwa.2026.100860.

13. M. Adil, N. Javaid, I. Ahmed, A. Ahmed, and N. Alrajeh, “A Deep Learning Framework for Heart Disease Prediction with Explainable Artificial Intelligence,” Computers, Materials and Continua, vol. 86, no. 1, pp. 1–20, 2025, doi: https://doi.org/10.32604/cmc.2025.071215.

14. A. Javed, N. Javaid, A. K. J. Saudagar, and D. Pamucar, “Developing a novel framework using optimized active stacking and explainable AI for heart disease prediction,” Comput. Methods Programs Biomed., vol. 274, p. 109169, 2026, doi: https://doi.org/10.1016/j.cmpb.2025.109169.

15. S. Fu et al., “Machine learning-based prediction of three-year mortality in elderly inpatients with coronary artery disease combined with heart failure,” Int. J. Med. Inform., vol. 210, p. 106307, 2026, doi: https://doi.org/10.1016/j.ijmedinf.2026.106307.

16. A. M., M. S., P. Selvam, and S. R., “CNN with ISNet-based healthcare monitoring system for heart disease prediction,” Biomed. Signal Process. Control, vol. 110, p. 108339, 2025, doi: https://doi.org/10.1016/j.bspc.2025.108339.

17. S. Urooj, R. Arunachalam, N. Verma, P. Bhutani, S. B. Kondaveeti, and A. L. Ekuakille, “An intelligent heart disease prediction framework by optimal feature integration with multi serial cascaded adaptive dilated network using heuristic improvement,” Eng. Appl. Artif. Intell., vol. 158, p. 111379, 2025, doi: https://doi.org/10.1016/j.engappai.2025.111379.

18. S. Shajimon, R. M. Shukla, and A. N. Patra, “Precision at heart: An IoT-based vertical federated learning approach for heterogeneous data-driven cardiovascular disease risk prediction,” Comput. Methods Programs Biomed., vol. 273, p. 109079, 2026, doi: https://doi.org/10.1016/j.cmpb.2025.109079.

19. Y. Suresh Chichani and S. L. Kasar, “An efficient IoT enabled heart disease prediction model using Finch hunt optimization modified BiLSTM classifier,” Biomed. Signal Process. Control, vol. 100, p. 107170, 2025, doi: https://doi.org/10.1016/j.bspc.2024.107170.

20. N. Saleh Alghamdi, M. Zakariah, A. Shankar, and W. Viriyasitavat, “Heart disease prediction using autoencoder and DenseNet architecture,” Egyptian Informatics Journal, vol. 28, p. 100559, 2024, doi: https://doi.org/10.1016/j.eij.2024.100559.

21. P. Rajendra Kumar, P. Chakrabarti, T. Chakrabarti, B. Unhelkar, and M. Margala, “Heart disease prediction using spark architecture with fused feature set and hybrid Squeezenet-Linknet model,” Biomed. Signal Process. Control, vol. 100, p. 107070, 2025, doi: https://doi.org/10.1016/j.bspc.2024.107070.

22. V. Chole, M. Thawakar, M. Choudhari, S. Chahande, S. Verma, and A. Pimpalkar, “Enhancing heart disease risk prediction with GdHO fused layered BiLSTM and HRV features: A dynamic approach,” Biomed. Signal Process. Control, vol. 95, p. 106470, 2024, doi: https://doi.org/10.1016/j.bspc.2024.106470.

23. I. Kaur and T. Ahmad, “A cluster-based ensemble approach for congenital heart disease prediction,” Comput. Methods Programs Biomed., vol. 243, p. 107922, 2024, doi: https://doi.org/10.1016/j.cmpb.2023.107922.

24. P. Kottapalle, T. K. Tak, P. R. Kshirsagar, G. Ginnela, and V. Krishna Akula, “QHF-CS: Quantum-Enhanced Heart Failure Prediction Using Quantum CNN with Optimized Feature Qubit Selection with Cuckoo Search in Skewed Clinical Data,” Computers, Materials and Continua, vol. 84, no. 2, pp. 3857–3892, 2025, doi: https://doi.org/10.32604/cmc.2025.065287.

25. M. A. Jahin, Md. A. Masud, M. F. Mridha, Z. Aung, and N. Dey, “KACQ-DCNN: Uncertainty-Aware Interpretable Kolmogorov–Arnold Classical–Quantum Dual-Channel Neural Network for Heart Disease Detection,” Comput. Biol. Med., vol. 197, p. 110976, 2025, doi: https://doi.org/10.1016/j.compbiomed.2025.110976.