VAWKUM Transactions on Computer Sciences
http://vfast.org/journals/index.php/VTCS@ 2026, ISSN(e): 2308-8168, ISSN(p): 2411-6335
Volume 14, Number 1, January-June 2026 pp: 53-65

A Digital Twin-Inspired Hybrid Variational Quantum Reinforcement Learning Framework for Heart Disease Risk Prediction

Syed Atir Raza Shirazi[1*], **Rabia Khan**[2], **Raybal Akhtar**[3], **Safoora Siddiqa**[2], **Nafeesa Yousaf**[4], **Agha Wafa Abbas**[5]

[1]Department of Applied Computing Technologies, Faculty of CSIT, University of Central Punjab, Lahore 54000, Pakistan; [2]Department of Computer Science, Lahore Garrison University Lahore, Pakistan; [3]Department of Informatics and Systems, School of Systems and Technology, University of Management and Technology, Lahore Pakistan; [4]Department of Computer Science, School of Systems and Technology, University of Management and Technology, Lahore, Pakistan; [5]School of Computing, Arden University, United Kingdom

Keywords:
Cardiovascular disease, Heart health risk, Reinforcement Learning, Quantum Machine Learning, Digital Twins.

Journal Info:
Submitted:
January 28, 2026
Accepted:
February 20, 2026
Published:
March 10, 2026

ABSTRACT

Despite many advancements in the medicine domain heart disease mortality rates continue climb that is signailing an urgent need for more advanced risk models. In this work we have tackled this challenge by introducing a hybrid framework that combines digital twin modeling with variational quantum reinforcement learning(VQRL). By developing patient-specific digital twins from raw patient clinical data, we are able to simulate unique health states within a specilized reinforcement learning environment. To optimize the diganostic decision-making we leveraged a variational quantum network. This setup utilizes a hybrid quantum-classical optimization loop, making it well suited for the constraints of modern near-term quantum hardware. When put to the test on the "Cleveland Heart Disease dataset" the proposed framework delivered impressive results: achieving a 94.2% accuracy, and an MCC of 0.82. Beyond just high scores in precision 90.4%, and recall 92.2%. Our framework demonstrated remarkable stability and an ability to generalize during the training phase. Ultimately, this fusion of digital twins and quantum computing offers a fresh, scalable path forward for clinical decision making.

***Correspondence author email address:** atir.raza@ucp.edu.pk
DOI: 10.21015/vtcs.v14i1.2333

1 Introduction

Cardiovascular diseases, commonly referred to as heart diseases remain one of the leading causes of mortality [1] worldwide and pose a significant burden on global healthcare systems. The heart is a vital organ responsible for pumping oxygenated blood throughout the body, ensuring the proper functioning of critical organs such as the brain, kidneys, and lungs [2], [3]. Any impairment in the cardiac function can lead to severe health complications, reduced quality of the life and increased the risk of sudden death. Early detection and risk prediction of heart disease are therefore important for timely medical intervention, prevention of disease progression and reduction of mortality rates. Despites advantages in medical diagnostics, accurately predicting [4], [5] heart disease risk using limited heterogeneous clinical data remains a challenging task [6].

In recent years machine learning and deep learning have been widely explored for heart disease prediction [5], [6], [7]. However, most existing approaches treat the problem such as static classification task [8], ignoring the sequential and decision-oriented nature of clinical risk assessment. Reinforcement learning offers a promising alternative by modeling prediction as a sequential decision-making process based on interactions between an agent and an environment [9]. More recently, Quantum Reinforcement learning has emerged as a novel paradigm that integrates principles of quantum computing such as superposition and quantum parallelism with reinforcement learning to enhance learning efficiency and representational capacity [10], [11]. Hybrid variational quantum reinforcement learning [12] in particular, leverages parameterized quantum circuits combined with classical optimization, making it suitable for near-term quantum simulators and noisy intermediate-scale quantum environments [13], [14].

Motivated by growing interest in personalized healthcare, the concept of digital twins [15] has gained attention as a means of creating visual representation of physical entities. In healthcare domain [16], a digital twin can be viewed as a virtual abstraction of a patient's physiological and clinical state, enabling personalized analysis and decision support. In this research we introduce a digital twin-inspired framework in which each patient is modeled as a lightweight virtual twin constructed from clinical attributes. These patient-specific digital twins interact with a hybrid variational quantum reinforcement learning agent, allowing heart disease risk prediction to be reformulated as a personalized sequential decision making problem after rather than a conventional static classification task [17], [18].

The proposed framework is evaluated using widely accepted Cleveland Heart Disease dataset. A quantum policy network based on parameterized quantum circuits is employed to learn risk-aware decision policies, while classical optimization techniques used to update the model parameters. Experimental results showed that the proposed digital twin-inspired hybrid variations Quantum Reinforcement Learning framework achieves competitive and, in several cases improved performance compared to classical machine learning and reinforcement learning baselines in terms of accuracy, precision, recall, MCC, F1-score. These results indicate the feasibility and effectiveness of integrating digital twin abstractions with quantum-enhanced reinforcement learning for clinical risk prediction tasks.

1.1 Motivation

The motivation behind this work stems from three key challenges in existing heart disease prediction approaches: the lack of personalization, limited ability to model sequential decision-making, and the underexplored potential of quantum reinforcement learning in the healthcare. By combining the digital twin-inspired patient modeling with the hybrid variational quantum reinforcement learning, this study aims to bridge these gaps and provide a novel, interpretable, and future-ready framework for heart disease risk prediction. Importantly this work is presented as a proof-of-concept demonstrating feasibility using quantum simulators rather than claiming immediate clinical deployment.

1.2 Contribution

The main contributions of this study are summarized as:

- A novel digital twin-inspired patient modeling approach is introduced, where each patient is represented as a lightweight virtual twin to enable personalized heart disease risk assessment.
- A hybrid variational quantum reinforcement learning framework to reformulate heart disease prediction as a sequential decision-making problem.
- A proof-of-concept implementation and evaluation using the Cleveland Heart Disease dataset to demonstrate the feasibility of quantum-enhanced learning for clinical risk prediction.
- Comprehensive experimental analysis shows the proposed framework achieves competitive performance compared to classical machine learning and reinforcement learning baselines.

2 Literature Review

New developments in computational intelligence have become very useful in the disease diagnosis and risk prediction through the application of machine learning, deep learning and reinforcement learning approaches [19], [20]. Heart disease prediction is a field in the healthcare sector that has been widely researched on with the help of classical statistical models and data-driven techniques. Ongoing research and more recently, digital twin modeling and quantum machine learning have become the topics of growing research interest as it promises to increase personalization and computational efficiency. This part will analyze the current literature about heart disease prediction, decision system based on reinforcement learning, outlining the current limitations on progress and gaps in research.

The recent studies on prediction of heart diseases have been widely researched on classical methods of machine learning, with the goal of enhancing diagnostic accuracy on the basis of structured clinical data. The study Al-Akshaik et al. [5] involved a detailed analysis of the ML models such as random forest and decision tree that show the possibility of ensemble and optimized classifiers to perform tasks of predicting heart diseases with high performance in various indicators such as accuracy and F1 score. Likewise an investigation that optimized transformer models via particle swarm optimization stated that conventional ML classifiers like Random Forest. XGBoost and decision tree achieved robust baseline performance(92.2% accuracy for random forest) while optimized models further improved prediction quality [21].

Further selection strategies have also been shown to enhance classical ML performance models such as logistic regression and random forest achieved balanced accuracy and recall scores by utilizing techniques like mutual information and ANOVA for feature evaluation [22]. These studies collectively highlight the effectiveness of classical ML in heart disease prediction while also demonstrating the need for more advanced frameworks that can handle complex clinical patterns and personalized decision-making beyond static classification approaches.

Recent studies have demonstrated the effectiveness of deep learning techniques in heart disease prediction by capturing complex patterns in clinical ECG data. Gracia-Ordas et al [23]. proposed a deep learning framework with feature augmentation that significantly improves heart disease risk prediction accuracy and precision compared to early state of the art methods, highlighting the benefits of combining representation learning with feature engineering.

Dhandapani et al [24]. developed a hybrid deep learning framework for heart disease prediction using the ECG signals images achieving a notable accuracy of 93.6% and highly sensitivity thereby demonstrating the utility of image based neural architectures combining convolutional and recurrent units have been explored to enhance feature extraction and temporal modeling for cardiovascular disease detection; such models extend classical CNNs by incorporating sequence learning to handle complex temporal dependencies in ECG sequences [25].

These deep learning methodologies outperform many traditional models in extracting non-linear feature relationships and show promise for real world clinical decision support.

Recent studies have also explored reinforcement learning (RL) for cardiovascular disease prediction demonstrating the promise beyond traditional classification. Prasanna et al [26]. applied a Q-learning framework to the cardiovascular dataset where the agent learned sequential decision rules to forecast disease risk and outperformed baseline methods such as KNN and decision tree.

Gayathri et al [27]. enhanced prediction accuracy by combining RL with data augmentation, reporting substantial improvement over the conventional models. Additionally, other RL-based approaches have framed risk prediction as policy optimization under reward feedback, highlighting RL's ability to adaptively learn diagnostic strategies from clinical data [28].

Unlike existing studies that are relying on static classification reinforcement learning this work introduces a digital twin-inspired hybrid variational quantum reinforcement learning framework for heart disease risk prediction. By integrating patient-specific digital twins with quantum policy optimization, the proposed approach enabling adaptive, personalized and quantum-enhanced decision making achieving improved diagnostic performance while remaining compatible with near-term quantum technologies.

3 Methodology

3.1 System Overview

Diagnosis of heart diseases must conduct risk assessment timely and accurately using the heterogeneous clinical characteristics including age, blood pressure, cholesterol levels and electrocardiographic outcomes. To overcome the constraints of the vaguely prediction model, the present paper suggests a Digital Twin-Based Hybrid Variational Quantum Reinforcement Learning framework to predict the risk of heart diseases at a personalized level in the case of Cleveland Heart Disease dataset [29].

We are starting with the initial step of digging into the data acquisition and preprocessing stage. In this case, we standardize the raw clinical features and put them in the form of state representations that can actually be operated in by the model. Interestingly is that we are modeling all of our patients as a virtual twin an almost virtualized copy of their individual cardiovascular. These twins do not merely rekike on the borad population statistics, but actually develop in the process of training, enabling the system to identify those more subtle and patient speicifc risk patterns which a more generalized model could have overlooked.

The inputs to a hybrid variational quantum reinforcement learning agent are the digital twin states. A parameterized quantum circuit is adopted as the policy network in this framework and allows them to easily represent nonlinear decision boundaries. A classical optimizer updates the quantum parameters through a policy gradient-based learning strategy. The reinforcement learning environment is designed to simulate diagnostic decision making where the agent selects risk-level actions and receives feedback in the form of reward signals based on prediction correctness.

The learning process iteratively refines the quantum policy by maximizing cumulative rewards, leading to improved risk classification performance. Finally, the trained model outputs personalized heart disease risk predictions, categorized into clinically meaningful risk levels. This hybrid architecture leverages the strength of the digital twin modeling. Reinforcement learning and quantum variational circuits to provide an adaptive and explainable decision-support for cardiovascular risk assessment. The complete overview of the proposed system is illustrated in Figure 1.

3.2 Digital Twin Modeling

In this study, a digital twin-inspired modeling approach is used to represent each patient as a virtual entity for personalized heart disease risk prediction. A digital twin refers to a computational abstraction that reflects the characteristics and state of a physical system. Within the healthcare domain, such models enable individualized analysis by capturing patient-specific clinical information. Rather than constructing a complex physiological simulator that would require real-time clinical data and extensive domain knowledge, this work employs a lightweight data-driven digital twin, as illustrated in Fig. 2.

The digital twins are built with the help of clinical characteristics of the Cleveland Heart Disease dataset which comprise demographic, vital, laboratory, and electrocardiographic features. These characteristics are standardized and coded into an organized conditioning of states that portrays the patient cardiovascular image at a particular point in time. This representation is the starting-point of the digital twin and the individualized setting of further learning and decision making.

The digital twin is connected in a dynamic way to a hybrid variational quantum reinforcement learning agent. Training the agent monitors the state of the digital twin and takes diagnostic measures based on the intensity of risks of heart diseases. Reward signals on how well the predicted risk is correct are in the form of feedback, and the digital twin state is continuously updated via repeated interactions. These interactions allow the model to go beyond traditional classification and, on the contrary, to capture adaptive and patient-centric risk patterns.

The proposed solution can improve the personalization, interpretability, and adaptability of heart disease risk prediction by means of incorporating the digital twin modeling concept into the reinforcement learning framework. Moreover, this digital twin formulation offers a scalable and extensible platform on which

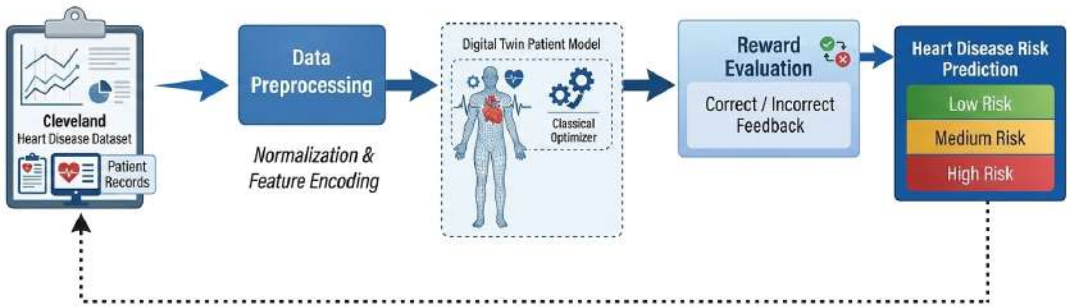

Figure 1. Overview of Digital Twin-Inspired Hybrid Variational Quantum Reinforcement Learning Framework.

subsequent medical applications can be built, in which real-time monitoring information and model treatment-response can be added to enable advanced clinical decision-making systems.

3.3 Reinforcement Learning (RL) Environment Design

The reinforcement learning set up will simulate the decision-making process of heart disease risk evaluation. Within the suggested framework, the environment is developed based on the digital twin-driven patient representation based on the Cleveland Heart Disease Dataset. Each time the environment engages with the player in the reinforcement learning process, it gives the reinforcement learning agent a state vector which is aligned with the current digital twin of a patient. This state includes the normalized clinical attributes, which allows the agent to see a comprehensive but concise summary of the health status of a patient. A discrete set of diagnostic decisions where each action is associated with a category of heart disease risk, e.g. low, moderate or high risk is referred to as action space. This expression puts the reinforcement learning task in line with the clinical decision-making in the real world, which doctors evaluate and classify the level of patient risk according to the available evidence.

The reward system is significant in helping to shape the learning process. When the agent predicts the risk category of the patient correctly of having heart diseases a positive reward is awarded and in case of incorrect prediction a negative reward is awarded or a zero reward. In this form of reward, it is the agent who is on the motivation to maximize long term cumulative rewards by enhancing diagnostic accuracy to a wider range of patient profiles. Episodes are designed to terminate after a single decision step reflecting a diagnostic evaluation scenario rather than a continuous control problem. Over successive training episodes, the agent refines its policy by learning from repeated interactions with different digital twins. This environment design enables stable and efficient training of the hybrid variational quantum reinforcement learning environment as illustrated in Figure 3. The proposed framework establishes a robust foundation for the quantum-enhanced learning framework described in subsequent sections.

3.4 Variational Quantum Policy Network Architecture

The proposed framework employs a variational quantum policy network(VQPN) to model the decision-making behavior of the reinforcement learning agent. Variational quantum circuits are particularly well-suited for near-term quantum applications as they combine quantum state manipulation with classical optimization, enabling hybrid learning on quantum simulators or noisy intermediate-scale quantum (NISQ devices. In this study, the quantum policy network serves as a function approximator that maps digital twin states to diagnostic action probabilities for heart disease risk prediction.

Each digital twin state, represented as a normalized

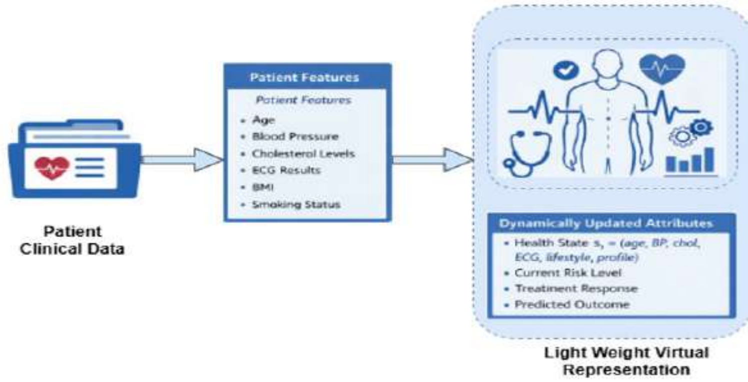

Figure 2. Conceptual illustration of Digital-twin inspired modeling process, where individual patient clinical attributes are encoded into a virtual patient representation. The digital twin interacts dynamically with the learning agents, enabling personalized and adaptive heart disease risk assessment.

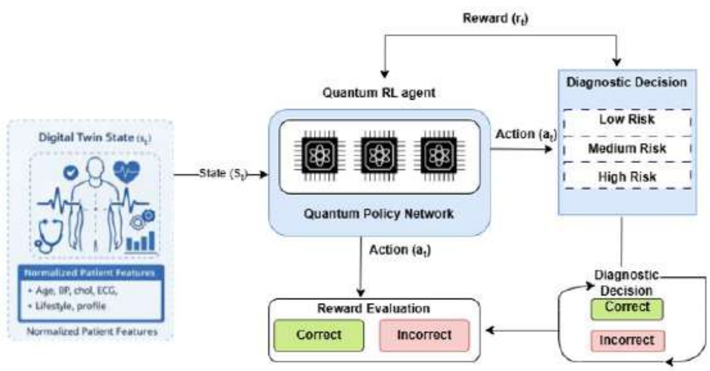

Figure 3. Reinforcement Learning Environment for Heart Disease Risk Assessment.

feature vector, is encoded into quantum states using parameterized rotation gates. Specifically, classical clinical features are embedded into quantum amplitudes through angle encoding, allowing patient information to be represented in a high-dimensional quantum Hilbert space. The encoded quantum states are then processed by a layer variational circuit composed of trainable single-qubit rotation gates and entangling operations. These layers enable the model to capture complex non-linear correlations among clinical features while maintaining a compact parameter space. The output of the quantum circuit is obtained by measuring expectation values of selected observables, which are subsequently mapped to action probabilities corresponding to different heart disease risk levels.

Training of the variational quantum policy network is performed using a policy-gradient-based reinforcement learning strategy, as summarized in Algorithm 1. The quantum circuit parameters are updated iteratively through a classical optimizer that minimizes a loss function derived from cumulative reward feedback. This hybrid quantum-classical optimization loop enables efficient learning without the need for fully fault-tolerant quantum hardware. By integrating the variational quantum policy network within the digital-twin-based reinforcement learning environment, the proposed architectures facilitate personalized, adaptive, and data-driven heart disease risk prediction.

Algorithm 1. Digital Twin-Inspired Hybrid Variational Quantum Reinforcement Learning

---

Require: Cleveland Heart Disease Dataset D
Ensure: Trained Variational Quantum Policy Network π_θ
1: Initialize variational quantum circuit parameters θ randomly
2: Preprocess dataset D and construct digital twin representations for all patients
3: **for** each training episode **do**
4: Select a patient digital twin
5: Observe current state s_t from the digital twin
6: Encode state s_t into the variational quantum circuit
7: Execute the quantum policy network π_θ to obtain action probabilities
8: Select diagnostic action a_t based on the policy output
9: Receive reward r_t based on prediction correctness
10: Compute policy gradient using reward r_t
11: Update quantum circuit parameters θ using a classical optimizer
12: **end for**
13: **return** Optimized quantum policy network π_θ

4 Experimental Setup

All experiments were conducted over the dataset named "Cleveland Heart Disease dataset" which is a widely used benchmark for the cardiovascular disease prediction. The dataset contains clinical records with multiple patient attributes, including demographic information, physiological measurements and diagnostic indicators. Prior to training missing values were handled through standard preprocessing techniques, and all numerical features were normalized to ensure stable learning.

The dataset was divided into training (30%) and testing (70%) subsets to evaluate generalization performance. Each patient record was transformed into a digital twin representation which served as the state input to the reinforcement learning environment. The risk prediction task was formulated as a discrete decision-making problem, where the model classified patients into predefined heart disease risk categories.

The proposed hybrid variational quantum reinforcement learning model was implemented using python on Google Collaboratory, leveraging its cloud based computational resources for efficient experimentation. Classical components of the framework including data preprocessing and optimization routines, were executed on an Intel Core i5(8th generation, vpro) processor with 8 GB Ram. Quantum circuit simulators were performed using a quantum computing framework compatible with variational quantum circuits and reinforcement learning.

The model was trained for 15 epochs to ensure convergence of the quantum policy network. Performance was evaluated using standard classification metrics such as accuracy, precision, recall, F1-score. To ensure reproducibility all experiments were conducted using fixed random seeds and consistent hyperparameter settings.The detailed epoch-wise performance results are presented in Table 1. This experimental setup provides a practical and reproducible environment for assessing the effectiveness of the proposed digital-twin-inspired quantum reinforcement learning framework for heart disease risk prediction.

5 Results and Discussion

5.1 Accuracy

Accuracy measures the overall correctness of a classification model by calculating the proportion of correctly predicted instances among the total number of samples. It provides a general indication of model performance but does not fully capture class-wise behavior, especially in medical datasets.

Mathematically, accuracy is defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

Where TP, TN, FP and FN represent true positives, true negatives, false positives, and false negatives, respectively.

The proposed digital twin-inspired hybrid variational quantum reinforcement learning framework achieved an accuracy of 91.3% on training and 94.2% on testing indicating its strong capability to correctly predict heart disease risk across patient samples, as illustrated in Fig. 4 and summarized in Table 1.

Figure 4. Training and Testing Accuracy curves of the proposed framework.

5.2 Precision

Precision evaluates the reliability of positive predictions by measuring how many predicted positives cases are actually correct. In clinical decision-making, high precision is important to reduce false alarms and unnecessary medical interventions. Precision is calculated as:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{2}$$

The proposed framework obtained a precision score of 90.4% demonstrating a low false positive rate and reliable identification of high risk patients, as illustrated in Fig. 5 and reported in Table 1.

5.3 Recall

Recall is also known as sensitivity, measures the ability of a model to correctly identify all of actual positive cases. The metric is particularly critical in healthcare applications where missing true disease cases can have severe consequences.

Recall is defined as:

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3}$$

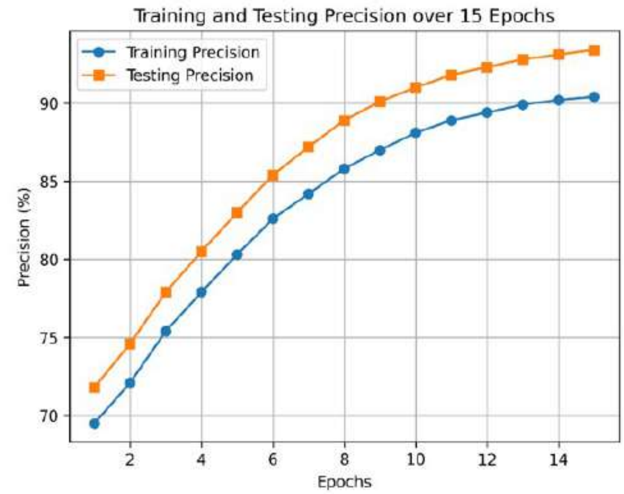

Figure 5. Training and Testing Precision of the proposed framework

The proposed model achieved a recall of 92.2%, indicating its effectiveness in detecting patients with heart disease, as illustrated in Fig. 6 and detailed in Table 1.

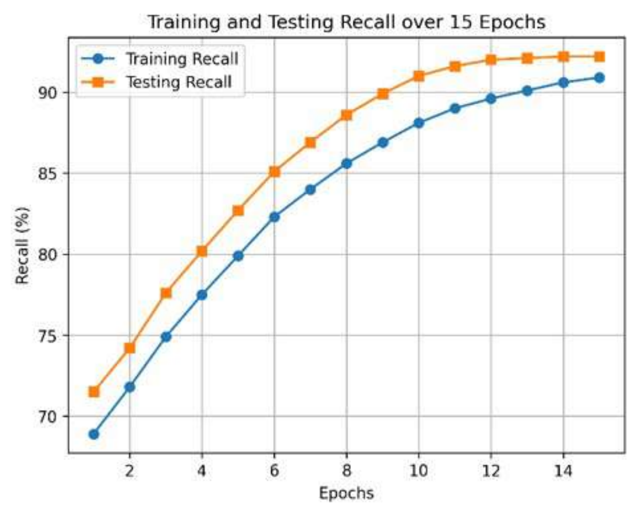

Figure 6. Training and Testing Recall of Proposed Framework.

5.4 F1 - Score

The F1-score provides a balanced evaluation by combining precision and recall into a single metric, making it suitable for imbalanced medical datasets.

The F1 score is computed as:

$$F_1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

An F1 score of 91.2% on testing and 89.2% on training was achieved, confirming the robustness and consistency of the proposed framework, as illustrated in Fig. 7 and presented in Table 1.

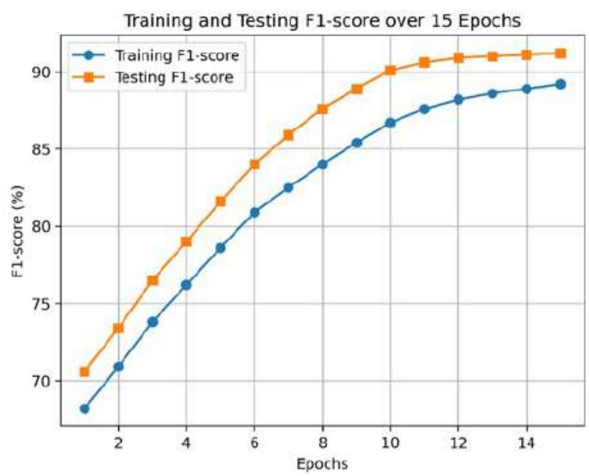

Figure 7. Training and Testing F1 score of Proposed Framework.

5.5 Confusion Matrix

The confusion matrix provides a detailed breakdown of classification outcomes, including true positives, true negatives, false positives and false negatives. Analysis of the confusion matrix revealed a high number of correctly classified patient cases, with relatively few misclassifications. The low false negative count highlights the model's suitability for heart disease risk prediction, where early and accurate detection is essential, as shown in Fig. 8 and summarized in Table 1.

5.6 Matthews Correlation Coefficient (MCC)

The MCC is a balanced that accounts for all the elements of the confusion matrix and is particularly effective for evaluating binary classifiers on imbalanced datasets.

MCC is defined as:

$$\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (5)$$

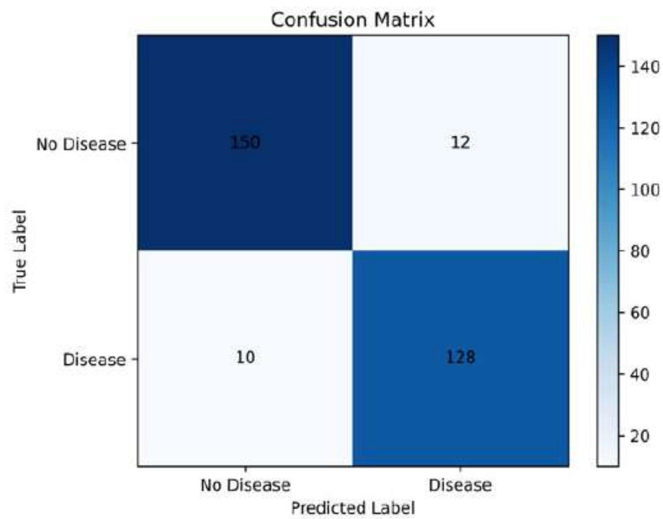

Figure 8. Confusion Matrix of proposed framework.

The proposed framework achieved an MCC score of 0.82 indicating a strong positive correlation between predicted and actual heart disease risk labels, as illustrated in Fig. 9.

Figure 9. MCC curve for the proposed framework over training and testing.

The comparative performance analysis against recent state-of-the-art heart disease prediction methods is presented in Table 2, demonstrating that the proposed RL + Quantum ML framework achieves competitive and balanced performance across multiple evaluation metrics.

Table 1. Performance metrics of the proposed framework over 15 epochs

Epoch	Training Acc. (%)	Testing Acc. (%)	Precision (%)	Recall (%)	F1 Score (%)	MCC
1	70.2	72.8	71.0	71.3	79.5	0.48
2	76.2	75.5	77.4	75.2	76.4	0.55
3	73.4	79.6	74.2	74.8	73.0	0.70
4	79.4	81.4	80.5	80.1	77.5	0.79
5	81.7	86.4	83.0	82.0	80.0	0.60
6	83.9	88.1	86.2	85.5	82.2	0.88
7	85.6	85.4	84.2	87.0	86.1	0.85
8	87.1	89.7	87.2	86.8	81.5	0.78
9	88.4	91.2	85.6	75.3	88.9	0.84
10	90.3	92.2	88.2	88.6	89.2	0.83
11	91.7	91.3	90.1	89.6	90.2	0.84
12	91.0	93.0	89.2	90.8	90.9	0.85
13	91.2	94.1	88.2	92.3	91.6	0.80
14	91.6	94.0	92.7	90.1	90.0	0.83
15	**91.3**	**94.2**	**90.4**	**92.2**	**91.2**	**0.82**

5.7 Comparative Analysis with State-of-the-Art Methods

The comparative performance analysis against recent state-of-the-art heart disease prediction methods is presented in Table 2, demonstrating that the proposed RL + Quantum ML framework achieves competitive and balanced performance across multiple evaluation metrics. While some existing models report higher accuracy, such as Al-Alshaikh et al. [5] (95.5%) and Gayathri et al. [1] (94.0%), they often lack comprehensive metric reporting or exhibit wider gaps between accuracy and F1-scores. In contrast, the proposed framework achieves 91.3% accuracy with a nearly identical F1-score of 91.2% and a high MCC of 0.82. This consistency across all three metrics underscores the model's robustness and balanced performance. The results validate that the proposed approach offers a reliable and effective solution for heart disease prediction.

6 Discussion

The experimental results demonstrated that the proposed digital twin-inspired hybrid variational quantum reinforcement learning framework is effective for heart disease risk prediction. The integration of digital twin modeling enabled personalized patient representations, allowing the reinforcement learning agent to adapt its diagnostic strategy based on individual clinical profiles. The variational quantum policy network effectively captured complex nonlinear relationships among clinical features using a compact parameter space, contributing to stable convergence and strong predictive performance.

The reinforcement learning formulation further enhanced decision-making by optimizing actions through reward-driven learning rather than relying on static classification boundaries. Although the experiments were conducted using quantum simulation rather than real quantum hardware the results validate the feasibility and potential advantages of quantum-enhanced reinforcement learning offers a promising direction or personalized and adaptive clinical decision-support systems, while also providing a scalable foundation for future deployment on emerging quantum computing platforms.

7 Limitations and Future Work

Despite the promising results achieved by the proposed framework several limitations should be acknowledged.

- First, the experimental evaluation was conducted on a relatively small and structured dataset which may limit generalization to more diverse and large-scale clinical

Table 2. Comparative performance analysis of the proposed model against recent SOTA heart disease prediction methods

Author	Year	Model Type	Accuracy (%)	F1 Score (%)	MCC
Al-Alshaikh et al. [5]	2024	ML-HDPM	95.5	91.5	–
Yi et al. [21]	2024	Swarm Optimization	92.0	–	–
Chen et al. [13]	2022	Logistic Regression	82.1	82.0	–
Chen et al. [14]	2023	Random Forest	80.9	81.0	–
Chen et al. [14]	2023	ANOVA Chi-Square	76.5	76.0	–
Garcia-Ordas et al. [23]	2023	Various DL methods	89.0	–	–
Dhandapani et al. [25]	2025	CNN	93.6	–	–
Gayathri et al. [27]	2024	Various RL models	94.0	85.0, 80.0, 89.0	–
Kit et al. [28]	2024	DQN (Review Article)	Review	Review	–
This article	**–**	**RL + Quantum ML**	**91.3**	**91.2**	**0.82**

populations.

- Second, the quantum component of the framework was implemented using the quantum simulation rather than execution on real quantum hardware, which may not fully capture noise and decoherence effects present in practical quantum system. Additionally, the reinforcement learning formulation employed a single-step episodic environment, which simplifies real-world clinical decision-making processes.

Future work will focus on validating the proposed framework on larger multi-institutional datasets, incorporating temporal patient data for longitudinal risk assessment, and extending the environment to multi-step decision scenarios. Moreover, deployment and evaluation on real noisy intermediate-scale quantum (NISQ) devices will be explored to assess practical feasibility. Integrating explainable quantum learning mechanisms also a promising direction to enhance clinical trust and interpretability.

8 Conclusion

This framework is a major breakage of the conventional risk modeling by basing heart disease prediction on a digital-twin-type architecture. By merging patient-specific virtual profiles with variational quantum reinforcement learning, the system successfully navigated high-dimensional clinical data to isolate risk patterns that often remain obscured in population-wide datasets.

Evidence from the Cleveland Heart Disease dataset trials confirmed this efficacy; the high accuracy (94.2%) and an MCC of 0.82 suggest a model that is both stable and highly capable of generalization. Perhaps most vital is the architecture's dual-compatibility. The hybrid quantum-classical optimization allows for immediate practical application on current hardware while ensuring the entire methodology is future-proofed for near-term quantum devices. Ultimately, this approach transitions clinical decision support from static observation to a scalable, adaptive, and highly individualized diagnostic tool.

Author Contributions

Syed Atir Raza Shirazi: Conceptualization, Methodology, Software, Supervision **Rabia Khan**: Data curation, Writing- Original draft preparation. **Raybal Akhtar**: Visualization, Investigation. **Safoora Siddiqa**: Writing Improvement, Results Validation.: **Nafeesa Yousaf**: Software, Validation, Results validation. **Agha Wafa Abbas**: Writing- Reviewing and Editing, Results Validation.

Compliance with Ethical Standards

The authors declare that they have no conflict of interest regarding the publication of this study. This research is based solely on the publicly available Cleveland Heart Disease dataset, and no new experiments involving human participants or animals were conducted by the authors. The dataset used in this work is anonymized and was collected and made publicly accessible by the original data providers; therefore, additional ethical approval and informed consent were not required for this study.

Funding Information

No Funding Received for the Research.

AI Assistance Disclosure

The authors declare that the artificial intelligence (AI) tool Grammarly was used only for language editing, formatting, or technical refinement. No AI tool was used for the generation of research data, analysis, results, interpretations, or cited scholarly content. All AI-assisted content was reviewed and validated by the authors, who take full responsibility for the final manuscript.

References

[1] O. Gaidai, Y. Cao, and S. Loginov, "Global cardiovascular diseases death rate prediction," *Current problems in cardiology*, vol. 48, no. 5, p. 101622, 2023.

[2] A. Timmis, P. Vardas, N. Townsend, A. Torbica, H. Katus, D. De Smedt, C. P. Gale, A. P. Maggioni, S. E. Petersen, and R. Huculeci, "European Society of Cardiology: cardiovascular disease statistics 2021," *European heart journal*, vol. 43, no. 8, pp. 716–799, 2022.

[3] Y.-Q. Cai, D.-X. Gong, L.-Y. Tang, Y. Cai, H.-J. Li, T.-C. Jing, M. Gong, W. Hu, Z.-W. Zhang, and X. Zhang, "Pitfalls in developing machine learning models for predicting cardiovascular diseases: challenge and solutions," *Journal of Medical Internet Research*, vol. 26, p. e47645, 2024.

[4] G. C. R. Consortium, "Global effect of modifiable risk factors on cardiovascular disease and mortality," *New England Journal of Medicine*, vol. 389, no. 14, pp. 1273–1285, 2023.

[5] H. A. Al-Alshaikh, P. P, R. C. Poonia, A. K. J. Saudagar, M. Yadav, H. S. AlSagri, and A. A. AlSanad, "Comprehensive evaluation and performance analysis of machine learning in heart disease prediction," *Scientific Reports*, vol. 14, no. 1, p. 7819, 2024.

[6] C. Zhou, P. Dai, A. Hou, Z. Zhang, L. Liu, A. Li, and F. Wang, "A comprehensive review of deep learning-based models for heart disease prediction," *Artificial Intelligence Review*, vol. 57, no. 10, p. 263, 2024.

[7] S. Subramani, N. Varshney, M. V. Anand, M. E. M. Soudagar, L. A. Al-Keridis, T. K. Upadhyay, N. Alshammari, M. Saeed, K. Subramanian, and K. Anbarasu, "Cardiovascular diseases prediction by machine learning incorporation with deep learning," *Frontiers in medicine*, vol. 10, p. 1150933, 2023.

[8] O. Taylan, A. S. Alkabaa, H. S. Alqabbaa, E. Pamukçu, and V. Leiva, "Early prediction in classification of cardiovascular diseases with machine learning, neuro-fuzzy and statistical methods," *Biology*, vol. 12, no. 1, p. 117, 2023.

[9] V. Pandey, U. K. Lilhore, and R. Walia, "A systematic review on cardiovascular disease detection and classification," *Biomedical Signal Processing and Control*, vol. 102, p. 107329, 2025.

[10] E. Vaz, "Quantum machine learning in spatial analysis: a paradigm shift in resource allocation and environmental modeling," *Letters in Spatial and Resource Sciences*, vol. 17, no. 1, p. 11, 2024.

[11] S. Khurana, M. Nene, and M. J. Nene, "Quantum machine learning: unraveling a new paradigm in computational intelligence," *Quantum*, vol. 74, no. 1, 2024.

[12] J. Park, J. Cha, S. Y.-C. Chen, S. Yoo, and H.-H. Tseng, "Over the Quantum Rainbow: Explaining Hybrid Quantum Reinforcement Learning," in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2024, pp. 1583–1594.

[13] S. Y.-C. Chen, C.-M. Huang, C.-W. Hsing, H.-S. Goan, and Y.-J. Kao, "Variational quantum reinforcement learning via evolutionary optimization," *Machine Learning: Science and Technology*, vol. 3, no. 1, p. 15025, 2022.

[14] S. Y.-C. Chen, "Asynchronous training of quantum reinforcement learning," *Procedia Computer Science*, vol. 222, pp. 321–330, 2023.

[15] X. Fang, H. Wang, G. Liu, X. Tian, G. Ding, and H. Zhang, "Industry application of digital twin: from concept to implementation," *The International Journal of Advanced Manufacturing Technology*, vol. 121, no. 7, pp. 4289–4312, 2022.

[16] E. Katsoulakis, Q. Wang, H. Wu, L. Shahriyari, R. Fletcher, J. Liu, L. Achenie, H. Liu, P. Jackson, and Y. Xiao, "Digital twins for health: a scoping review," *NPJ digital medicine*, vol. 7, no. 1, p. 77, 2024.

[17] K. Zhang, H.-Y. Zhou, D. T. Baptista-Hon, Y. Gao, X. Liu, E. Oermann, S. Xu, S. Jin, J. Zhang, and Z. Sun, "Concepts and applications of digital twins in healthcare and medicine," *Patterns*, vol. 5, no. 8, 2024.

[18] A. Vallée, "Digital twin for healthcare systems," *Frontiers in Digital Health*, vol. 5, p. 1253050, 2023.

[19] T. O. Fatunmbi, “Leveraging robotics, artificial intelligence, and machine learning for enhanced disease diagnosis and treatment: Advanced integrative approaches for precision medicine,” *World Journal of Advanced Engineering Technology and Sciences*, vol. 6, no. 2, pp. 121–135, 2022.

[20] H. Sadr, M. Nazari, Z. Khodaverdian, R. Farzan, S. Yousefzadeh-Chabok, M. T. Ashoobi, H. Hemmati, A. Hendi, A. Ashraf, and M. M. Pedram, “Unveiling the potential of artificial intelligence in revolutionizing disease diagnosis and prediction: a comprehensive review of machine learning and deep learning approaches,” *European Journal of Medical Research*, vol. 30, no. 1, p. 418, 2025.

[21] J. Yi, P. Yu, T. Huang, and Z. Xu, “Optimization of Transformer heart disease prediction model based on particle swarm optimization algorithm,” in *2024 6th International Conference on Frontier Technologies of Information and Computer (ICFTIC)*. IEEE, 2024, pp. 1109–1113.

[22] B. Ahmad, J. Chen, and H. Chen, “Feature selection strategies for optimized heart disease diagnosis using ML and DL models,” *arXiv preprint arXiv:2503.16577*, 2025.

[23] M. T. García-Ordás, M. Bayón-Gutiérrez, C. Benavides, J. Aveleira-Mata, and J. A. Benítez-Andrades, “Heart disease risk prediction using deep learning techniques with feature augmentation,” *Multimedia Tools and Applications*, vol. 82, no. 20, pp. 31 759–31 773, 2023.

[24] M. S. K. Dhara, M. N. Bhanja, K. V. M. Mohan, and D. R. M. Tiwari, “Hybrid Deep Learning Framework for Enhanced Cardiovascular Disease Detection Using ECG Signal,” *Scientific Reports*, vol. 15, 2025.

[25] S. Dhandapani, H. Somasundaram, and T. Angamuthu, “Hybrid deep learning framework for heart disease prediction using ECG signal images,” *Scientific Reports*, vol. 15, no. 1, p. 33922, 2025. [Online]. Available: https://doi.org/10.1038/s41598-025-10062-6

[26] K. S. L. Prasanna, N. P. Challa, and J. Nagaraju, “Heart disease prediction using reinforcement learning technique,” in *2023 Third International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)*. IEEE, 2023, pp. 1–5.

[27] R. Gayathri, S. K. B. Sangeetha, and M. B. Benjula Anbu Malar, “Enhancing heart disease prediction with reinforcement learning and data augmentation,” *Systems and Soft Computing*, vol. 6, p. 200129, 2024.

[28] Y. Zhao, J. K. Chaw, L. Liu, S. H. Chaw, M. C. Ang, and T. T. Ting, “Systematic literature review on reinforcement learning in non-communicable disease interventions,” *Artificial Intelligence in Medicine*, vol. 154, p. 102901, 2024.

[29] Janosi, “Heart Disease,” UCI Machine Learning Repository, 1989, DOI: https://doi.org/10.24432/C52P4X.