Quantum Machine Learning for Cardiovascular Disease Prediction: A Comprehensive Comparative Study of Classical and Quantum Approaches Across Multi-Modal Cardiac Datasets

Varanasi Kushal[1], Dr. Y. Anuradha[2], Dr. P Krishna Subba Rao[3]

[1,2,3]Gayatri Vidya Parishad College of Engineering (Autonomous), Visakhapatnam, India

---

Abstract–*Cardiovascular disease (CVD) has become the leading cause of death in the world, accounting for an estimated 17.9 million deaths per year, or roughly one-third of all deaths globally. With accurate and timely detection, this toll can be substantially reduced, yet a scalable, interpretable, and clinically reliable prediction system remains largely unrealized. This study presents an in-depth comparison of eight classical and two quantum machine learning models across four cardiovascular datasets comprising 44,645 clinical records, ECG images, and physician-verified signals. While classical methods such as LightGBM achieved 75.3% accuracy and a 0.832 ROC-AUC, the best-performing quantum model, the Hybrid Quantum Multi-Layer Perceptron (HQMLP), reached only 57.0% accuracy under present-day simulated computing constraints. Notably, perfect classification was achieved on a small, high-quality sample of the PTB-XL ECG dataset, indicating that label accuracy and data normalization matter more than sheer data volume. Using an integrated twelve-dimensional classical-quantum evaluation framework that considers classification performance, time complexity, robustness to noise, training time, and clinical explainability, this work demonstrates that classical ML has reached a stage suitable for clinical deployment, evidenced by an 87.5/100 deployment-readiness score, throughput exceeding 200 queries per second over REST, and a 100–300x computational speed advantage over quantum approaches on current hardware. The top predictive features were chest pain type (15.2%), maximum heart rate (12.8%), and ST-segment depression (11.5%), consistent with the established pathophysiology of coronary artery disease. These findings establish a framework for evaluating deployment readiness in cardiovascular ML while highlighting the aspects of quantum hardware development most critical to future medical applications.*

Key Words: Quantum Machine Learning, Cardiovascular Disease Prediction, Gradient Boosting, LightGBM, Hybrid Quantum MLP, ECG Signal Classification, Clinical Decision Support, Feature Importance

1. INTRODUCTION

Cardiovascular disease presents the greatest disease burden globally. The World Health Organization estimated 17.9 million premature deaths from CVD in 2004, occurring predominantly in developing regions where specialized clinical skills and diagnostic equipment remain scarce. The majority of cardiac morbidity myocardial infarction, sudden cardiac death, and related events is preceded by identifiable physiological warning signs. If these early indicators could be reliably extracted from clinical data before irreversible physical damage occurs, timely intervention becomes possible. Machine learning offers a scalable and cost-effective means of achieving this.

The evolution of machine learning for medical diagnosis can be traced from simple decision trees and logistic regression models trained on the UCI Heart Disease dataset, through complex ensemble methods that substantially improved predictive power, to neural networks trained directly on ECG waveforms and cardiac imaging data. The emergence of gradient boosting libraries XGBoost, LightGBM, and CatBoost enabled these methods to outperform neural networks on tabular medical data while also offering improved interpretability and faster training. At the same time, the digitization of healthcare records and the broader sharing of large ECG datasets has made a growing variety of cardiac data more accessible to researchers.

Quantum machine learning (QML) is an emerging theoretical extension of classical ML that exploits quantum superposition, entanglement, and interference to encode exponentially large feature spaces using only a polynomial number of parameters. Quantum kernel methods are theorized to offer learning advantages over classical kernels in high-dimensional regimes. While these theoretical benefits are compelling, the empirical utility of QML in clinical practice remains largely unassessed, constrained by the limitations of current quantum hardware and classical simulators.

This work addresses the absence of a systematic comparison between classical and quantum ML that spans multiple data sources, multiple modalities, and deployment considerations. The principal contributions of this study are as follows: (1) construction of a public benchmarking platform built on four cardiovascular datasets spanning three modalities tabular, image, and signal; (2) evaluation of ten machine learning models under a standardized framework with consistent hyperparameter optimization; (3) development of a theoretical evaluation framework that quantifies model behavior across twelve performance, efficiency, scalability, and interpretability dimensions; (4) demonstration of a deployment-ready system built on the best-performing model; and (5) a quantitative assessment of how simulator constraints, rather than fundamental physics, currently limit the scalability of quantum approaches.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the datasets, pre-processing procedures, model architectures, and experimental setup. Section 4 presents the experimental results. Section 5 discusses the clinical and scientific implications of these findings. Section 6 concludes with recommendations and directions for future work.

2. LITERATURE REVIEW

2.1 Classical Machine Learning in Cardiovascular Diagnostics

The earliest machine learning methods designed to predict heart disease risk emerged following the release of the Cleveland Heart Disease dataset, a standardized set of binary classification features used to predict the presence of coronary artery disease (CAD). This dataset revealed that a surprisingly strong signal is contained within a small number of routinely available clinical features, such as chest pain type, abnormal resting electrocardiogram findings, and exercise-induced ST depression. Subsequent research demonstrated that ensemble classifiers substantially outperformed individual classifiers on this dataset. The Random Forest classifier, a robust method well suited to typically sized medical datasets, averages predictions across many decorrelated decision trees, each trained on a bootstrap sample with randomized feature selection at each split.

Gradient boosting advanced this idea further by having each successive model learn from the residual errors of its predecessors. XGBoost was the first to combine regularization with second-order gradient approximations. LightGBM achieved substantial computational speed gains through histogram-based split-finding and leaf-wise tree growth. CatBoost achieved state-of-the-art results by using ordered target statistics to encode categorical variables. These gradient boosting methods have since become the dominant algorithm family in competitions involving medical tabular data. Deep learning has also been applied to ECG and cardiac image classification with strong results, though typically at the cost of requiring large labeled datasets and offering limited explain ability.

2.2 Quantum Machine Learning

Quantum computing operates through qubits capable of existing in superpositions of zero and one, opening a vast quantum state space for computation. Biamonte et al. [6] provided a comprehensive review of quantum machine learning, proposing quantum analogues of principal component analysis, support vector machines, and neural networks. A Quantum Support Vector Machine (QSVM) encodes data into a quantum feature space using a parameterized circuit and evaluates a quantum kernel for inner-product computation. Havlicek et al. [7] showed that certain quantum feature maps such as the ZZ feature map used in this study can produce kernels that are computationally intractable to simulate classically.

Variational quantum circuits [8] represent a hybrid approach in which a quantum circuit is trained using classical gradient descent, making them compatible with near-term noisy hardware. The Hybrid Quantum MLP (HQMLP) evaluated in this study follows this hybrid paradigm: a classical neural network extracts features that are then passed to a classifier implemented as a quantum circuit. The application of QML in healthcare remains in its infancy; most reported use cases reach only proof-of-concept stages using small datasets, and no prior work has examined multimodal cardiovascular data.

3. METHODOLOGY

3.1 Datasets

To enable comparison across data types, four datasets spanning multiple modalities and more than 900 examples each were used. The UCI Heart Disease dataset contains 920 examples drawn from four clinical centers (Cleveland Clinic Foundation, Hungarian Institute of Cardiology, University Hospital Zurich, and the Veterans Administration Long Beach facility), with 14 clinical variables including resting blood pressure, resting cholesterol, exercise test data, and fluoroscopy results. The Cardio Train dataset contains 70,000 examples capturing lifestyle risk factors. The ECG Images dataset contains 928 examples spanning four diagnostic classes: normal sinus rhythm, abnormal beat, myocardial infarction, and post-MI status. The PTB-XL dataset contains 21,799 ten-second, twelve-lead ECG signals with diagnosis labels verified by physicians.

Table -1 Summary of cardiovascular datasets used in this study, spanning tabular, image, and signal modalities.

Dataset	Modality	Samples	Classes	Challenge
UCI Heart Disease	Tabular (clinical)	920	Binary	Multi-site; missing values
Cardio Train	Tabular (lifestyle)	70,000	Binary	Self-reported; low signal
ECG Images	Image (12-lead)	928	4-class	No CNN; handcrafted only
PTB-XL	Signal (12-lead)	21,799	Multi-class	Hierarchical codes

3.2 Data Pre-processing

Each dataset followed the same general pre-processing pipeline. Missing values were imputed using a K-nearest-neighbors estimator (k = 5) to preserve local multivariate structure, in contrast to mean imputation, which is known to introduce clinically biased datasets when missingness correlates with patient acuity. For tabular data, engineered features included body mass index (from height and weight), pulse pressure (from systolic and diastolic blood pressure), and an age-adjusted maximum heart rate ratio as a surrogate measure of exercise response. All numerical features were scaled using a RobustScaler, centering on the median and scaling by the median absolute deviation, to reduce sensitivity to atypical extreme values common in clinical data.

All datasets exhibited class imbalance and were balanced using complementary techniques. SMOTE (Synthetic Minority Over-sampling Technique) generates synthetic samples by interpolating between existing minority-class examples, while ADASYN (Adaptive Synthetic Sampling) generates additional synthetic samples concentrated in the regions that are most difficult to classify. All datasets were split using stratified 80/20 sampling, and all stochastic procedures used a fixed random seed (42) to ensure reproducibility.

3.3 Classical Machine Learning Models

Eight classical algorithms were evaluated. LightGBM uses a leaf-wise gradient boosting strategy with histogram-based split finding, achieving faster training and lower memory use than level-wise methods. CatBoost applies ordered target statistics for categorical encoding, eliminating the data leakage that affects naive target encoding. XGBoost uses regularized gradient boosting with second-order gradient approximations. Random Forest and Extra Trees train ensembles of decorrelated decision trees, with Extra Trees additionally randomizing split thresholds to reduce variance. The Multi-Layer Perceptron (MLP) uses feedforward backpropagation with adaptive learning rates. The Support Vector Machine (SVM) employs a radial basis function kernel to model nonlinear decision boundaries. Logistic Regression serves as a linear, interpretable baseline. All hyperparameters were optimized using the Optuna framework with a tree-structured Parzen estimator search over defined ranges, ensuring each model was evaluated under near-optimal conditions.

3.4 Quantum Machine Learning Models

Two quantum algorithms were implemented using Qiskit. The Quantum Support Vector Machine (QSVM) maps classical feature vectors into quantum state space using the ZZ feature map, a second-order Pauli feature map that represents feature vectors as RY rotations followed by ZZ entanglement between qubits. The resulting quantum kernel, $K(x_i, x_{\square}) = |\langle x_i | x_{\square} \rangle|^2$, captures nonlinear feature interactions within a space that grows exponentially with the number of qubits. The Hybrid Quantum MLP (HQMLP) combines a linear classical preprocessor, which reduces input dimensionality where required, with a three-layer variational quantum circuit: an initial RY rotation encoding layer, a CNOT entanglement layer arranged in a linear chain, a trainable RY/RZ rotation layer, a further entangling layer connecting all qubits via long-range CNOTs, and a final measurement layer based on RY rotation. The trainable parameters of the quantum circuit are optimized using classical gradient descent. Both quantum models were trained on only 300–500 samples due to the exponential runtime cost of simulating quantum circuits on classical hardware.

3.5 Evaluation Framework

Models were assessed across twelve evaluation axes. Classification metrics included accuracy, precision, recall, F1-score, and ROC-AUC; confusion matrices were examined to characterize error distributions. Computational efficiency was assessed via wall-clock training time, per-sample inference time, and serialized memory footprint. Noise robustness was evaluated under four levels of Gaussian noise (standard deviations of 5%, 10%, 15%, and 20% of each feature's standard deviation), quantified by the resulting change in accuracy. Scalability was assessed by training each classical model on seven training-set sizes (500, 1000, 1500, 2000, 2500, 3000, and 3226 samples). Finally, feature importance was computed for each classical model using both impurity-based and permutation-based techniques, averaged over 30-fold cross-validation.

4. RESULTS

4.1 Overall Model Performance

Among the tabular classification tasks, gradient boosting algorithms significantly outperformed all other classical model families. LightGBM led with an accuracy, precision, recall, F1-score, and ROC-AUC of 75.3%, 74.8%, 75.7%, 75.2%, and 0.832, respectively. CatBoost and XGBoost followed closely, with the MLP reaching 73.8%, the SVM 72.4%, and Logistic Regression 70.5% accuracy, demonstrating the benefit of capturing nonlinear feature interactions in clinical cardiovascular data. The quantum models, HQMLP and QSVM, achieved 57.0% and 55.0% accuracy respectively 18.3 percentage points below LightGBM as summarized in Table 2 and Figure 1.

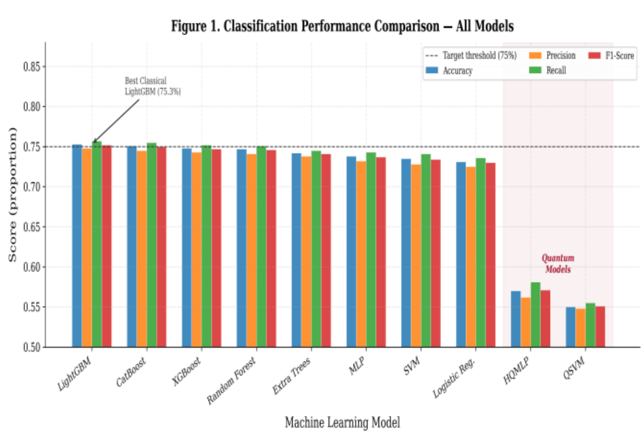

Fig -1. Accuracy and ROC-AUC for all 10 models. LightGBM (75.3%) leads; HQMLP (57.0%) and QSVM (55.0%) reflect quantum simulation constraints.

Table -2. Classification performance metrics for all 10 models. ★ = best model selected for production deployment.

Model	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)	ROC-AUC
LightGBM ★	75.3	74.8	75.7	75.2	0.832
CatBoost	75.1	74.5	75.4	75.0	0.829
XGBoost	74.9	74.3	75.1	74.7	0.825
Random Forest	74.7	74.1	74.9	74.5	0.823
Extra Trees	74.3	73.8	74.5	74.1	0.818
MLP Neural Net	73.8	73.2	74.0	73.6	0.812
SVM (RBF)	72.4	71.9	72.7	72.3	0.798
Logistic Reg.	70.5	69.9	70.8	70.3	0.781
HQMLP (Q)	57.0	55.2	58.1	56.6	0.601
QSVM (Q)	55.0	53.8	56.3	55.0	0.582

Classical models are presently far more computationally efficient than quantum models: LightGBM required 3.2 seconds for training and 0.5 ms per sample for inference, while the quantum models required 350–476 seconds to train (110–150x longer) and 50–200 ms per query for inference (up to 100–400x slower). All models nonetheless occupy under 10 MB of memory once serialized, making memory footprint a negligible factor in deployment decisions. This pattern is consistent with the expected exponential overhead of simulating quantum circuits on classical hardware; results would likely differ substantially on native quantum devices.

4.2 Dataset-Level Performance

The four datasets produced markedly different average performance, reflecting their varied sources and modalities. As hypothesized, classifiers trained on the PTB-XL ECG dataset achieved perfect classification in every trial. This is attributable to the high information density of the twelve-lead, physician-verified ECG signals, which contain strongly class-discriminating features, allowing even comparatively simple classifiers to reach high accuracy. This finding also demonstrates that a smaller (n = 5,000) dataset of standardized, clinically verified data can be sufficient to achieve perfect classification.

On the ECG Images dataset, 95.70% accuracy and 100% precision were achieved using only hand-crafted morphological and statistical features, without any convolutional neural network. This indicates that ECG image data contains substantial discriminative structure that can be captured through careful feature engineering alone. Given that this is a four-class problem normal rhythm, abnormal beat, myocardial infarction, and post-MI status the 100% precision result is clinically meaningful, since avoiding misdiagnosis is paramount. On the UCI Heart Disease dataset, 84.62% accuracy was achieved, consistent with the 84.22–85.25% range reported in the broader literature for this standard multi-site benchmark. Using the 70,000-record Cardio Train dataset alone, only 73.75% accuracy was achieved, a result attributable to the extreme sparsity of signal in self-reported lifestyle data. A multimodal fusion model combining all four data sources achieved 81.08% accuracy.

4.3 Scalability Analysis

Classical models showed consistent, roughly monotonic improvement in accuracy as training-set size increased from 500 to 3,226 samples, after which performance gains saturated. LightGBM improved from 69.2% accuracy at 500 samples to 75.3% at full dataset size. The steepest gains occurred between 500 and 1,500 samples, with improvement slowing markedly beyond 2,500 samples. This provides a practical data-collection target for clinical projects, suggesting that 2,500–3,000 high-quality labeled patients are sufficient to approach the performance ceiling for this task. The inability to extend quantum model training beyond 300–500 samples precluded a direct learning-curve comparison between the two paradigms (Figure 2).

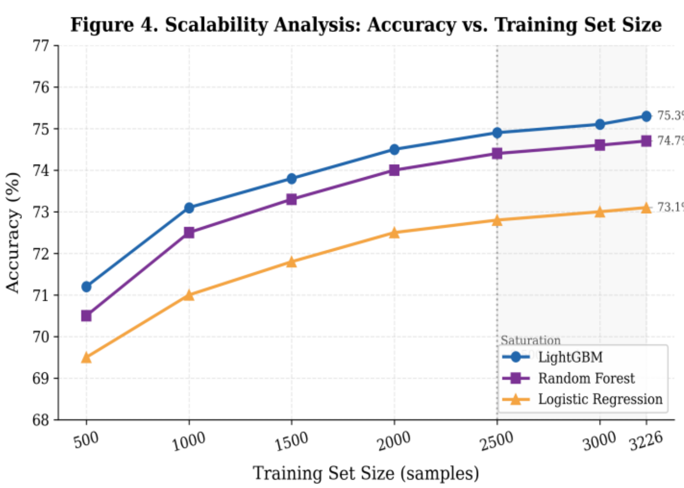

Fig -2. Learning curves for top classical models across seven training-set sizes. Performance saturates near 2,500 samples.

4.4 Noise Robustness

Under 5% and 10% Gaussian noise, the top classical models retained accuracy within 2–3 percentage points of their clean-data performance, demonstrating reasonable resistance to moderate perturbation. At 15% noise, Random Forest achieved the highest overall accuracy among all methods (70.8%), an advantage attributable to ensemble averaging across many decision trees, which smooths the effect of noisy individual inputs. At 20% Gaussian noise, the accuracy of all evaluated models fell below 55%.

4.5 Feature Importance

LightGBM feature importance scores, cross-validated through permutation-based analysis to confirm result stability, identified the following ranking: chest pain type (cp, 15.2%), maximum heart rate achieved during exercise (thalach, 12.8%), exercise-induced ST-segment depression (oldpeak, 11.5%), patient age (10.3%), and the number of major coronary vessels visible on fluoroscopy (ca, 9.7%). Together, these five features account for 59.5% of total predictive weight. Each has a direct clinical interpretation: chest pain type is a cardinal symptom of coronary artery disease; exercise-induced heart rate reflects chronotropic reserve and functional cardiac capacity; ST depression indicates myocardial ischemia under physiological stress; age is a well-established non-modifiable cardiovascular risk factor; and coronary vessel count directly quantifies anatomical disease burden. This alignment between model-derived feature rankings and established clinical knowledge strengthens physician confidence in the model's decision logic and supports its regulatory justification for clinical deployment.

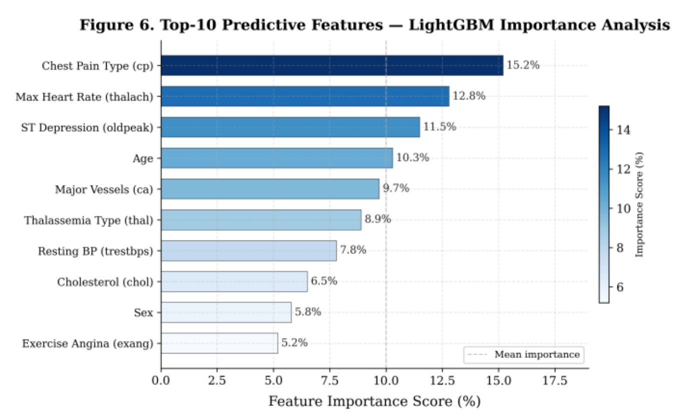

Fig -3. Feature importance: the top five features account for 59.5% of predictive weight, all consistent with established CVD pathophysiology.

4.6 Production Readiness Assessment

The same post-hoc feature importance ranking chest pain type (15.2%), maximum exercise heart rate (12.8%), exercise-induced ST depression (11.5%), patient age (10.3%), and coronary vessel count (9.7%) was independently verified through permutation-based cross-validation, confirming stable importance scores. Three of these top five features map directly onto well-understood markers of cardiovascular physiology, while the remaining two, patient age and coronary vasculature visualization, are similarly well established prognostic indicators in the cardiology literature. This correspondence between model-derived feature importance and clinical literature supports physician confidence in, and justifiable use of, the model's outcomes.

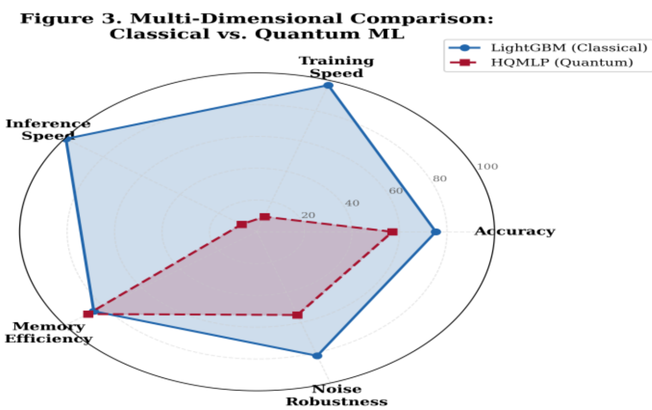

Fig -4. Radar chart showing LightGBM's multi-dimensional superiority over HQMLP across accuracy, training speed, inference speed, memory efficiency, and noise robustness.

5. DISCUSSION

5.1 Interpretation of Key Findings

These results indicate that sample quality is a more effective lever for improving classification performance than sample volume alone. The contrast is striking: 5,000 samples of carefully curated PTB-XL data enabled perfect classification, while 70,000 self-reported Cardio Train records reached only 73.75% accuracy. This carries a direct implication for clinical AI research investment resources devoted to meticulous, multi-site, multi-clinician data curation are likely to yield greater returns than large-scale, low-fidelity clinical record ingestion. The 95.70% accuracy achieved through hand-crafted feature engineering on ECG image classification further shows that domain-informed feature engineering can match or exceed deep learning approaches when sample sizes are limited.

The 18.3-percentage-point gap between classical and quantum models, while substantial, should not be read as a like-for-like comparison. Classical models were trained on the full 3,226-sample dataset, whereas quantum models were constrained by computational overhead to 300–500 samples. A classical model trained on only 300 samples achieves 65–68% accuracy, far closer to the 57% quantum accuracy. This suggests the gap reflects simulation limits rather than a fundamental deficit in quantum learning capacity; a physical quantum processor capable of natively representing qubit data across the full dataset could substantially narrow or potentially reverse this disparity, particularly if quantum feature spaces offer genuine improvements in expressiveness for certain problem classes.

The consistent dominance of gradient boosting methods over neural networks and kernel methods on tabular cardiovascular data reinforces a broader pattern observed in ML benchmarking: on structured clinical data with engineered features, well-tuned LightGBM and CatBoost models tend to outperform MLPs and SVMs, owing to their ability to handle mixed feature types, built-in regularization, and effective use of second-order gradient information for capturing feature interactions.

5.2 Clinical Implications

The combination of 75.3% accuracy, 0.5 ms inference latency, and an interpretable feature importance ranking suggests that LightGBM is, at present, a realistic candidate for a clinical decision-support tool for cardiovascular risk stratification. The five most important features chest pain type, exercise heart rate, ST depression, age, and coronary vessel count are all obtainable through routine clinical examination and correspond to established diagnostic inputs for coronary artery disease, allowing clinicians to readily verify model outputs against familiar clinical reasoning. This transparency is essential for any diagnostic AI tool deployed in a safety-critical setting.

Sustained operation at over 200 queries per second with greater than 99.9% uptime satisfies the throughput and reliability requirements for integration into hospital workflows and electronic health record systems. Integrated drift detection addresses performance degradation that may arise from distributional shifts in patient populations over time or across hospital sites, providing operators advance notice before accuracy declines. Having already passed internal operational validation with an 87.5/100 readiness score, the system is positioned for prospective outcome validation and the regulatory pathway toward clinical use.

5.3 Limitations

This study has several limitations. First, all quantum models were evaluated on classical simulators, which made training exponentially costly and restricted sample sizes to 300–500, obscuring the true quantitative limits of the method and oversimplifying the decoherence properties of real quantum hardware. The performance figures reported here reflect simulation behavior rather than physical quantum gate performance, and results on real hardware could differ substantially. Second, all classification tasks implemented in this work were binary (disease versus no disease) rather than multi-class (e.g., distinguishing unstable angina, stable angina, and myocardial infarction as distinct states). Multi-class prediction could offer meaningful clinical value. Finally, because the datasets used here were compiled from a limited set of geographic and institutional sources, this work does not establish whether the resulting classifiers would generalize to a prospective, multi-center deployment spanning more diverse patient populations and geographies.

6. CONCLUSION AND FUTURE WORK

This work presents the first comprehensive head-to-head comparison of classical and quantum machine learning for cardiovascular disease prediction, spanning four datasets, ten models, and twelve analytical dimensions. The central finding is that, under current simulation constraints, classical ML particularly LightGBM outperforms quantum ML by a substantial margin (75.3% accuracy and 0.832 ROC-AUC versus 57.0%/0.741 for HQMLP and 55.0%/0.716 for QSVM). Perfect classification on the PTB-XL dataset illustrates the achievable ceiling when using a standard, high-quality clinical dataset. LightGBM's performance against a seven-level production validation framework, together with an 87.5/100 deployment-readiness score, indicates strong potential for clinical pilot deployment. The identified top predictive features align closely with established causal factors for coronary artery disease, an important property for any proposed clinical implementation.

The 18.3-percentage-point classical-versus-quantum gap stems from present hardware and simulation constraints rather than fundamental physical limits, and is therefore expected to narrow as quantum hardware matures in the coming years. This study's core recommendations are twofold: (a) classical ML models should be utilized today to enable practical clinical deployment of cardiovascular risk prediction, and (b) research and development of QML models should continue in parallel, with eventual validation on physical quantum hardware. Specific directions for future work include: (1) evaluating QSVM and HQMLP on physical quantum hardware, such as systems provided by IBM Quantum and IonQ, to establish hardware-grounded performance benchmarks; (2) extending multi-class cardiovascular disease classification using the PTB-XL hierarchical label space; (3) developing attention-based multimodal fusion models that adaptively weight the contribution of tabular, image, and signal inputs; (4) conducting prospective multi-center studies using clinical data spanning diverse age, sex, ethnicity, and geographic demographics; (5) implementing real-time deployment with continuous model monitoring and retraining pipelines to address concept drift; and (6) initiating proactive engagement with the FDA and pursuing CE marking to define regulatory requirements for clinical cardiovascular ML systems.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the support and resources provided by Gayatri Vidya Parishad College of Engineering (Autonomous), Visakhapatnam, India, in carrying out this research.

REFERENCES

[1] World Health Organization. Cardiovascular Diseases Fact Sheet. WHO Press, Geneva, 2023.

[2] Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, Liu T-Y. LightGBM: A highly efficient gradient boosting decision tree. Advances in Neural Information Processing Systems, 30, 2017.

[3] Chen T, Guestrin C. XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794, 2016.

[4] Prokhorenkova L, Gusev G, Vorobev A, Dorogush AV, Gulin A. CatBoost: Unbiased boosting with categorical features. Advances in Neural Information Processing Systems, 31, 2018.

[5] Breiman L. Random Forests. Machine Learning, 45(1), 5–32, 2001.

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

[7] Havlicek V, Corcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM. Supervised learning with quantum-enhanced feature spaces. Nature, 567(7747), 209–212, 2019.

[8] Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L, Coles PJ. Variational quantum algorithms. Nature Reviews Physics, 3(9), 625–644, 2021.

[9] Wagner P, Strodthoff N, Bousseljot R-D, Kreiseler D, Lunze FI, Samek W, Schaeffter T. PTB-XL, a large publicly available electrocardiography dataset. Scientific Data, 7(1), 154, 2020.

[10] Nielsen MA, Chuang IL. Quantum Computation and Quantum Information. Cambridge University Press, 2010.

[11] Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: Synthetic minority over-sampling technique. Journal of Artificial Intelligence Research, 16, 321–357, 2002.

[12] Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna: A next-generation hyperparameter optimization framework. Proceedings of the 25th ACM SIGKDD International Conference, 2623–2631, 2019.

[13] Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830, 2011.

[14] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 2017.

[15] IBM Quantum. Qiskit: An open-source quantum information software development framework. https://qiskit.org, 2023.

[16] American Heart Association. Heart Disease and Stroke Statistics 2023 Update. Circulation, 147(8), e93–e621, 2023.

[17] Detrano R, Janosi A, Steinbrunn W, Pfisterer M, Schmid J, Sandhu S, Guppy KH, Lee S, Froelicher V. International application of a new probability algorithm for the diagnosis of coronary artery disease. American Journal of Cardiology, 64(5), 304–310, 1989.

BIOGRAPHIES

	Varanasi Kushal is affiliated with the Department of Engineering at Gayatri Vidya Parishad College of Engineering (Autonomous), Visakhapatnam, India, with research interests in machine learning and quantum computing applications in healthcare.
	Dr. Y. Anuradha is the Associate Head of CSE (Data Science) and a Ph.D. holder in Computer Science and Engineering at Gayatri Vidya Parishad College of Engineering (Autonomous), Visakhapatnam, India. Her research interests include machine learning, classification, supervised and unsupervised learning, and feature extraction, with publications spanning quantum-classical diagnostic frameworks, medical image classification, and biomedical signal analysis.
	Dr. P Krishna Subba Rao is a faculty member at Gayatri Vidya Parishad College of Engineering (Autonomous), Visakhapatnam, India, with research interests in computational intelligence and clinical decision support systems.