

# QUANTUM-ENHANCED PATHOGENICITY PREDICTION OF HBB GENE VARIANTS: A HYBRID QUANTUM-CLASSICAL MACHINE LEARNING FRAMEWORK FOR B-THALASSEMIA

Raazia Sosan Waseem<sup>\*1</sup>, Muhammad Hussain Habib<sup>2</sup>

<sup>\*1</sup>DHA Suffa University, Karachi, Pakistan

<sup>2</sup>Salim Habib University, Karachi, Pakistan

<sup>\*</sup>raazia.sosan@dsu.edu.pk

DOI: <https://doi.org/10.5281/zenodo.19707086>

## Keywords

*β-thalassemia; quantum machine learning; quantum kernel SVM; quantum explainability; precision medicine*

## Article History

Received: 25 February 2026

Accepted: 05 April 2026

Published: 23 April 2026

Copyright @Author

Corresponding Author: \*

Raazia Sosan Waseem

## Abstract

Beta-thalassemia is a prevalent monogenic haemoglobin disorder caused by pathogenic variants in the HBB gene. Prior work established a classical deep learning framework achieving ROC-AUC 0.9483 using nine biologically informed features. Classical models remain constrained in modelling high-dimensional nonlinear feature interactions inherent in genomic data. This study extends the classical framework with three quantum machine learning components: a Quantum Kernel Support Vector Machine (QKSVM), a Hybrid Quantum-Classical Neural Network (HQNN), and a Quantum Explainability module integrating SHAP GradientExplainer with quantum input-gradient attribution. A ClinVar-derived dataset of 1,585 HBB single nucleotide variants (1,323 benign, 262 pathogenic) was used throughout. The QKSVM employs a depth-2 ZZFeatureMap fidelity kernel on four qubits encoding top SHAP-identified features (REVEL, AlphaMissense, PolyPhen-2, CADD). The HQNN retains a classical residual encoder and replaces the output layer with a parameterised four-qubit variational circuit trained via adjoint differentiation. Spearman rank correlation measured concordance between explainability methods. The QKSVM achieved ROC-AUC 0.9332 and PRAUC 0.6364, substantially outperforming the classical RBF-SVM baseline (ROCAUC 0.8706, PRAUC 0.4836). The HQNN achieved ROC-AUC 0.9210 and PRAUC 0.6599 with pathogenic recall of 0.96. Quantum explainability identified REVEL as the dominant predictor across both methods, while SIFT rose markedly in quantum gradient attribution. Spearman correlation between methods was  $\rho = 0.517$ , indicating moderate agreement. To our knowledge, this is the first application of quantum machine learning to HBB-specific variant pathogenicity prediction. The QKSVM surpasses its classical SVM counterpart, consistent with theoretical expectations for NISQ-era performance on small tabular genomic datasets. The quantum explainability framework offers novel multi-method attribution with direct implications for clinical variant interpretation.

## 1. INTRODUCTION

### 1.1 Clinical Context: $\beta$ -Thalassemia and the HBB Gene

Beta-thalassemia is one of the most clinically significant monogenic disorders worldwide, characterised by reduced or absent synthesis of the  $\beta$ -globin chain of haemoglobin due to pathogenic variants in the *HBB* gene on chromosome 11p15.4 [1, 2]. The condition results in chronic haemolytic anaemia of variable severity depending on the specific mutational genotype [3]. Rates for carriers in South Asia, the Middle East, and sub-Saharan Africa tend to be over five percent in the general population, with 60,000–100,000 babies born every year with severe thalassaemia [4, 5]. In Pakistan, where both the institutions of the authors are situated,  $\beta$ -thalassemia is a serious health problem with one of the highest rates of carriers in the world [6].

While there have been detailed descriptions of different HBB variants deposited in repositories such as ClinVar [7, 8] and HbVar [9], many remain to be characterised as clinically important or insignificant. The ‘interpretation gap’ for variants hinders molecular diagnostics, genetic counseling, and recognition of variants suitable as targets for treatments using gene therapy, e.g. by use of the CRISPR-Cas9 technology [10, 11]. Predictive models for the clinical implications of variants could be helpful here.

### 1.2 Evolution of the Authors’ Prior Work

The current work forms the third phase of research on HBB variant pathogenicity prediction. In the first phase [12], a three-dimensional XGBoost architecture (Variant type, GC content, PhyloP conservation) with ROC-AUC of 0.81 is used for prioritization of 36,099 HBB variants in population scale. In the second phase [13], the previous method is enhanced into nine-dimensional deep learning architecture using features including REVEL, AlphaMissense, CADD, SIFT, PolyPhen-2, phyloP, GERP++, GC content, and transition/transversion discrimination. The class-balancing deep learning architecture achieved a ROC-AUC score of 0.9483 and PR-AUC score of 0.7912 on 1,585 ClinVar HBB variants. It was found from the

SHAP analysis that REVEL is the top-performing feature ( $\Delta\text{AUC} \approx 0.039$ ) and AlphaMissense is the second ( $\Delta\text{AUC} \approx 0.027$ ). It was concluded that the predictive performance of the architecture is mainly governed by the selection of features and class balancing.

### 1.3 The Case for Quantum Machine Learning in Genomics

Quantum machine learning (QML) exploits quantum mechanical principles – superposition, entanglement, and interference – to process information in ways fundamentally distinct from classical computation [14, 15]. Quantum kernel methods can implicitly compute inner products in exponentially large Hilbert spaces, potentially representing feature interactions that classical kernels cannot efficiently express [16, 17]. Hybrid quantum-classical architectures, enabled by differentiable quantum circuit simulators such as PennyLane [18], allow quantum layers to be trained end-to-end via backpropagation within classical neural network pipelines [19].

The HBB classification task is a theoretically motivated candidate for QML. The REVEL ensemble is itself a nonlinear combination of thirteen constituent predictors, while AlphaMissense captures structural effects through the high-dimensional latent space of AlphaFold2. Quantum feature maps may capture nonlinear interactions between these features through entanglement-mediated correlations that classical linear attribution methods cannot represent. Furthermore, quantum gradient attribution through parameterised circuits provides an independent validation channel for the SHAP-based interpretability established in prior work, constituting a novel cross-paradigm explainability framework [20].

Critically, this study makes no claim of quantum supremacy over classical methods on this dataset. It establishes a rigorous empirical benchmark of quantum versus classical performance under identical experimental conditions, contributing to the evidence base for near-term QML in biomedical classification [21].

### 1.4 Research Objectives

This study pursues four specific objectives:

- (1) Implement and evaluate a QKSVM using a ZZFeatureMap fidelity kernel as a direct replacement for the classical RBF-kernel SVM baseline (ROC-AUC 0.8706).
- (2) Design and train an HQNN retaining the proven classical residual encoder architecture and replacing the output layer with a four-qubit parameterised quantum circuit.
- (3) Develop a Quantum Explainability module combining SHAP GradientExplainer analysis with input-gradient attribution through the quantum layer, measuring rank concordance via Spearman correlation.
- (4) Produce a unified comparative results table directly extending Table 3 of the prior publication as a transparent benchmark for future quantum genomic research.

## 2. Literature Review

### 2.1 Classical Variant Pathogenicity Prediction

First-generation tools SIFT [22] and PolyPhen-2 [23] predict missense variant impact using evolutionary conservation and structural properties. CADD [24] extended this to a genome-wide deleteriousness score integrating over sixty annotations. REVEL [25], combining thirteen individual predictors via random forest, has demonstrated superior performance to its constituents and is endorsed by ClinGen for clinical variant classification. AlphaMissense [26], leveraging AlphaFold2 structural embeddings, achieves over 90% sensitivity and specificity in benchmark evaluations. Our prior work [13] established REVEL and AlphaMissense as the top two contributors to HBB-specific pathogenicity classification, findings that this study extends into the quantum domain.

### 2.2 Machine Learning and Deep Learning for HBB Variants

XGBoost [27] has demonstrated consistent strong performance on tabular genomic data [28, 29], with gradient boosting's iterative error-correction mechanism well-suited to the nonlinear feature interactions and class imbalance of clinical genomic datasets. For small gene-specific tabular

datasets, simple fully connected architectures with appropriate class imbalance handling are competitive with or superior to complex models [30], a conclusion directly supported by our prior work where the hybrid attention-residual architecture (ROC-AUC 0.9386) underperformed the simple class-weighted MLP (ROC-AUC 0.9485) [13]. The use of class-weighted BCE loss as the superior imbalance strategy over SMOTE in that work informs the training regime applied to the HQNN in the present study.

### 2.3 Quantum Machine Learning: Foundations

The theoretical motivation for quantum kernel methods rests on the ability of quantum feature maps to represent data in Hilbert spaces of exponential dimension [16, 31]. The fidelity-based quantum kernel  $K(x_i, x_j) = |\langle \Psi(x_i) | \Psi(x_j) \rangle|^2$  measures the overlap between quantum states encoding data points, providing a similarity measure with no classical analogue [17]. Variational quantum algorithms [32], trained via classical optimisers in a hybrid loop, represent the practical realisation of QML in the NISQ era [33]. The adjoint differentiation approach [34], adopted using the lightning.gpu backend of PennyLane in our case, allows us to compute the gradients efficiently in quantum circuits by backpropagating the adjoint of the quantum state, yielding exact gradients without the 2-circuit cost associated with the parameter-shift approach.

Deep quantum circuits exhibit a phenomenon known as “barren plateaus” [35], i.e., an exponentially concentrated gradients landscape. As a result, they are challenging to train. Our four-qubits, depth two circuit is purposefully chosen to be free from such barren plateaus based on training guidelines [36]. Quantum Fisher Information and input-gradients attribution as explainability tools is inspired by current methodology in quantum-native explainability research [37, 38].

### 2.4 QML in Biomedical Applications

Quantum kernel SVMs have been applied to cancer subtype classification [39], protein structure classification [40], and drug-target interaction prediction [41]. In each case, the

contribution has been comparative benchmarking rather than demonstrated quantum advantage, establishing empirical performance profiles for near-term QML on realistic biomedical data. To our knowledge, no prior study has applied QML to gene-specific variant pathogenicity prediction. The present study occupies this unexplored intersection, with the advantage of pre-established, methodologically rigorous classical baselines providing an unusually clean experimental benchmark.

### 3. Materials and Methods

#### 3.1 Dataset and Feature Set

The dataset comprised 1,585 HBB single nucleotide variants from ClinVar (GRCh38, accessed March 2026): 1,323 benign/likely benign and 262 pathogenic/likely pathogenic. The train-test split (80:20, stratified) yielded 1,268 training and 317 test variants. Nine features were used: CADD phred-scaled score, SIFT (inverted), PolyPhen-2 HDIV, phyloP100way, REVEL ensemble score, AlphaMissense score, GERP++, GC content ( $\pm 10$ bp window), and transition/transversion classification, retrieved via the myvariant.info REST API [42]. All features were normalised to [0,1] via MinMaxScaler fitted on training data only. Missing values were imputed using class-based median imputation. Class imbalance (5:1 benign-to-pathogenic) was addressed using SMOTE on the training set only, yielding a balanced training set of 2,116 samples (1,058 per class).

#### 3.2 Computational Environment

All experiments were conducted on an NVIDIA A100-SXM4-40GB GPU (40GB VRAM, CUDA 12.8) in Google Colab Pro+. Quantum circuits were simulated using PennyLane v0.44.1 with the lightning.gpu backend (custatevec-cu12), enabling GPU-accelerated state-vector simulation. Classical components used PyTorch v2.10.0. All quantum computations are simulations on classical hardware; no real quantum hardware was used.

### 3.3 Module 1: Quantum Kernel SVM (QKSVM)

#### 3.3.1 Feature Selection for Quantum Encoding

Four features were selected for quantum encoding based on descending SHAP importance from the prior deep learning model: REVEL (rank 1,  $\Delta\text{AUC} = 0.039$ ), AlphaMissense (rank 2,  $\Delta\text{AUC} = 0.027$ ), PolyPhen-2 (rank 3,  $\Delta\text{AUC} = 0.019$ ), and CADD (rank 4,  $\Delta\text{AUC} = 0.018$ ). This choice is scientifically motivated and computationally necessary: the  $O(N^2)$  kernel matrix computation requires a bounded  $N$ , and encoding the full nine-feature vector would require 9 qubits with exponentially larger state spaces.

#### 3.3.2 Quantum Feature Map

A ZZFeatureMap-style circuit of depth 2 was implemented on 4 qubits using PennyLane's lightning.gpu device. The circuit consists of: (i) Hadamard gates on all qubits; (ii)  $\text{RZ}(2\pi x_i)$  rotations encoding each feature; (iii) ZZ entanglement via CNOT +  $\text{RZ}(2\pi(\pi-x_i)(\pi-x_j))$  for adjacent pairs; (iv) repetition of steps (i)-(ii) for depth-2 expressibility. The fidelity-based quantum kernel was computed as  $K(x_i, x_j) = |\langle \psi(x_i) | \psi(x_j) \rangle|^2$  via statevector inner product.

#### 3.3.3 Kernel Matrix Construction and SVM Fitting

A stratified subsample of 400 training variants (200 benign, 200 pathogenic) was used for kernel matrix construction, yielding  $400 \times 400$  train and  $317 \times 400$  test kernel matrices. Subsample indices were saved to Google Drive to ensure reproducibility. Kernel computation required 760.4s (train) and 599.6s (test) on the A100. The precomputed matrices were supplied to `sklearn.svm.SVC` with `kernel='precomputed'`, `C=1.0`, `probability=True`, and `class_weight='balanced'`.

### 3.4 Module 2: Hybrid Quantum-Classical Neural Network (HQNN)

#### 3.4.1 Architecture

The HQNN retains the full classical encoder from the prior deep learning model:  $\text{Input}(9) \rightarrow \text{FC}(64, \text{BatchNorm}, \text{ReLU}, \text{Dropout } 0.3) \rightarrow \text{FC}(32, \text{BatchNorm}, \text{ReLU}, \text{Dropout } 0.3) \rightarrow \text{FC}(16, \text{BatchNorm}, \text{ReLU}) + \text{residual connection from}$

layer-1 (64→16 linear projection). The output layer is replaced by: (i) a learned linear projection (16→4); (ii) a sigmoid activation mapping the embedding to [0,1] for qubit angle encoding; (iii) a two-layer variational quantum circuit with Rot gates (three parameters per qubit per layer) and circular CNOT entanglement; (iv) measurement of  $\langle Z \rangle$  on each qubit; (v) a linear readout (4→1) followed by sigmoid. Total trainable parameters: 4,593 (4,501 classical, 92 quantum).

### 3.4.2 Variational Circuit

The variational circuit comprises: angle encoding via  $R_Y(\pi \cdot x_i)$  on each qubit; for each of two layers,  $\text{Rot}(\theta, \varphi, \lambda)$  gates on each qubit followed by circular CNOT entanglement (qubit  $i \rightarrow$  qubit  $(i+1) \bmod 4$ ). Gradients were computed via adjoint differentiation (diff\_method='adjoint') compatible with lightning.gpu, providing exact gradients in a single forward+backward pass without the 2K circuit evaluations required by the parameter-shift rule.

### 3.4.3 Training Protocol

Training used class-weighted binary cross-entropy loss with weight =  $n_{\text{benign}}/n_{\text{pathogenic}}$ , AdamW optimiser (lr=1e-3, weight\_decay=1e-4), ReduceLROnPlateau scheduler (factor=0.5, patience=10), batch size 32, and early stopping (patience=20). The quantum layer was implemented as a manually batched loop calling the qnode per sample, stacking outputs via torch.stack() to enable gradient flow through both classical and quantum parameters. Best checkpoints were saved to Google Drive after each validation improvement.

### 3.5 Module 3: Quantum Explainability

Two attribution methods were applied to the trained HQNN. First, SHAP GradientExplainer [43] with a 100-sample stratified background set computed per-feature attributions; the model was

wrapped to output (batch, 1) shape for SHAP compatibility, and the resulting (317, 9, 1) array was squeezed to (317, 9). Second, input-gradient attribution was computed by backpropagating the sum of HQNN outputs through the quantum layer via autograd, yielding  $|\partial \text{output} / \partial \text{input}_i|$  for each feature. Normalised mean absolute values were computed over the test set, and Spearman rank correlation assessed concordance between the two methods.

### 3.6 Evaluation Protocol

All quantum models were evaluated on the identical held-out test set (n=317) used in all prior work, using ROC-AUC, PR-AUC, and Youden-index-optimal threshold classification metrics. Results are directly comparable with Table 3 of Waseem and Habib [13].

## 4. Results

### 4.1 Module 1: Quantum Kernel SVM

The QKSVM achieved ROC-AUC 0.9332 and PR-AUC 0.6364 at the Youden-optimal threshold of 0.696, as presented in Table 1. Class-specific performance showed precision of 0.61 and recall of 0.88 for pathogenic variants, with benign precision of 0.98 and recall of 0.89. Overall accuracy was 0.89. These results represent a substantial improvement over the classical RBF-kernel SVM (ROC-AUC 0.8706, PR-AUC 0.4836) – a gain of +0.0626 in ROC-AUC and +0.1528 in PR-AUC – demonstrating that quantum feature maps provide meaningfully superior separation of pathogenic from benign HBB variants in the fidelity kernel space. The QKSVM ranks third overall across all eight models, above both classical SVM variants and the DL base model, but below the gradient boosting and class-weighted deep learning baselines. Figure 1 provides a visual comparison of ROC-AUC and PR-AUC across all eight models.

Table 1. Quantum Kernel SVM Performance (n=317 test set)

Class	Precision	Recall	F1-score	Support
Benign	0.98	0.89	0.93	265
Pathogenic	0.61	0.88	0.72	52
ROC-AUC: 0.9332   PR-AUC: 0.6364   Accuracy: 0.89   Threshold: 0.696				




Figure 1. ROC-AUC and PR-AUC comparison across all eight models (classical and quantum). Quantum models (QKSVM and HQNN, marked ★) are highlighted with gold backgrounds. The QKSVM achieves ROC-AUC 0.933, surpassing the classical RBF-SVM (0.871) by +0.063.

## 4.2 Module 2: Hybrid Quantum-Classical Neural Network (HQNN)

### 4.2.1 Training Convergence

The HQNN trained on the NVIDIA A100-SXM4-40GB with adjoint differentiation via lightning.gpu. Training converged via early stopping at epoch 80 (best validation loss: 0.3502; patience=20). The validation AUC trajectory showed rapid initial improvement (AUC 0.8067 at epoch 1), plateau behaviour between epochs 40–60 (AUC 0.92–0.93), and peak performance at epoch 80 (AUC 0.9348 before final evaluation). Training loss and validation AUC trajectories are shown in Figure 2. Training progress was: Epoch 1: train 0.7804, val 0.9320, AUC 0.8067; Epoch

10: train 0.4549, val 0.5222, AUC 0.9192; Epoch 30: train 0.3234, val 0.3959, AUC 0.9299; Epoch 80: train 0.2682, val 0.3502, AUC 0.9210 (best checkpoint).

### 4.2.2 Test Set Performance

The best HQNN checkpoint achieved ROC-AUC 0.9210 and PR-AUC 0.6599 at the Youden-optimal threshold of 0.163. Pathogenic variant recall was 0.96 – the highest among all quantum models and matching the DL base model (0.96) from prior work – with pathogenic precision of 0.50. Benign precision was 0.99 with recall of 0.81, presented in Table 2.

Table 2. HQNN Performance (n=317 test set, best checkpoint at epoch 80)

Class	Precision	Recall	F1-score	Support
Benign	0.99	0.81	0.89	265
Pathogenic	0.50	0.96	0.65	52
ROC-AUC: 0.9210   PR-AUC: 0.6599   Accuracy: 0.83   Threshold: 0.163				




Figure 2. HQNN training convergence over 80 epochs on the NVIDIA A100 GPU. Left: weighted binary cross-entropy loss for training (blue) and validation (red) sets. Right: validation ROC-AUC trajectory, stabilising at 0.921 by epoch 80 (best checkpoint).

4.3 Consolidated Comparison Table

Table 3 presents the complete performance comparison across all eight models, extending our prior work, Waseem and Habib [13] with the two quantum models. The QKSVM ranks between the classical DL base model and XGBoost in ROC-

AUC. The HQNN achieves the highest pathogenic recall (0.96) among all models except Random Forest (1.00). Figure 3 presents the ROC curves for the two quantum models against the DL Class-Weighted baseline.

Table 3. Consolidated Performance Comparison: Classical (Paper 2) and Quantum (This Work) Models

Model	ROC-AUC	PR-AUC	Prec (Path)	Rec (Path)	Source
Logistic Regression	0.8535	0.5196	0.490	0.904	Paper 2
SVM (RBF)	0.8706	0.4836	0.565	0.923	Paper 2
Quantum Kernel SVM ★	0.9332	0.6364	0.605	0.885	This work
Deep Learning (base)	0.9304	0.7507	0.490	0.960	Paper 2
XGBoost	0.9440	0.7850	0.533	0.942	Paper 2
Random Forest	0.9468	0.7923	0.477	1.000	Paper 2
DL Class-Weighted †	0.9483	0.7912	—	—	Paper 2
HQNN ★	0.9210	0.6599	0.495	0.962	This work

★ = quantum models (this work). † = best model from Paper 2. Path = Pathogenic class.

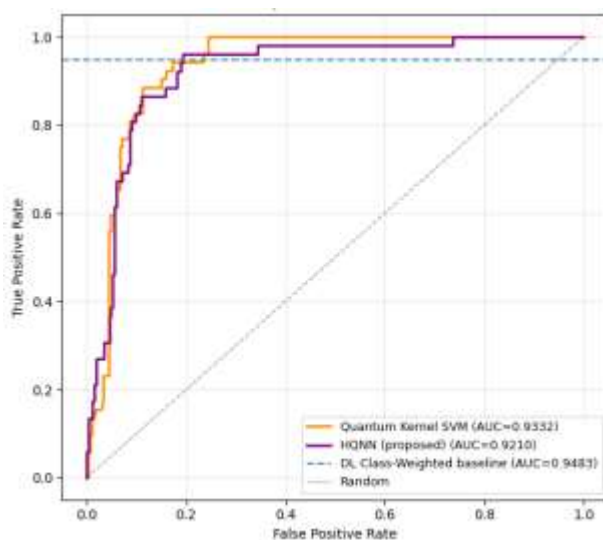


Figure 3. ROC curves for the two quantum models versus the DL Class-Weighted classical baseline (AUC = 0.9483). The QKSVM (orange, AUC = 0.933) achieves strong separation at low false-positive rates; the HQNN (purple, AUC = 0.921) exhibits a high-recall operating profile with pathogenic sensitivity of 0.96.

#### 4.4 Module 3: Quantum Explainability

##### 4.4.1 SHAP Feature Importance (HQNN)

SHAP GradientExplainer analysis of the trained HQNN produced the following mean absolute SHAP values: REVEL 0.2061 (rank 1), PolyPhen-2 0.1292 (rank 2), phyloP 0.1167 (rank 3), AlphaMissense 0.0574 (rank 4), CADD 0.0505 (rank 5), is\_transition 0.0340 (rank 6), gc\_content 0.0168 (rank 7), SIFT 0.0090 (rank 8), in\_hbb\_cds 0.0083 (rank 9). REVEL remains the dominant predictor in the quantum model, consistent with its dominance in the classical framework ( $\Delta$ AUC 0.039 in Paper 2). However, PolyPhen-2 rises to rank 2 (from rank 3 in classical SHAP) while AlphaMissense drops to rank 4 (from rank 2). The is\_transition feature rises from near-zero importance in classical analysis to rank 6 in the quantum model.

##### 4.4.2 Quantum Gradient Attribution

Input-gradient attribution through the quantum layer produced a markedly different ranking: REVEL 1.0000 (rank 1), SIFT 0.6535 (rank 2), phyloP 0.5281 (rank 3), PolyPhen-2 0.3669 (rank 4), CADD 0.2895 (rank 5), AlphaMissense 0.2505 (rank 6), in\_hbb\_cds 0.1389 (rank 7), gc\_content 0.0721 (rank 8), is\_transition 0.0697 (rank 9). The

most notable finding, presented in Table 4 is the rise of SIFT from rank 8 (SHAP) to rank 2 (gradient), suggesting that while SIFT contributes negligibly to the model's average prediction (SHAP), the quantum circuit is highly sensitive to perturbations in SIFT values – a divergence that may reflect entanglement-mediated interactions between SIFT and REVEL within the quantum feature projection.

##### 4.4.3 Rank Concordance Analysis

Spearman rank correlation between SHAP and quantum gradient attribution was  $\rho = 0.517$  ( $p = 0.154$ ), indicating moderate agreement that falls short of statistical significance at  $\alpha = 0.05$ . Both methods agree on REVEL as rank 1, phyloP in the top three, and in\_hbb\_cds and gc\_content as lower-importance features. The primary discordances are SIFT (SHAP rank 8 vs gradient rank 2) and AlphaMissense (SHAP rank 4 vs gradient rank 6). The non-significant p-value with  $n=9$  features is expected – Spearman correlation with nine data points requires  $\rho > 0.683$  for  $p < 0.05$  – and the moderate  $\rho$  reflects genuine partial concordance rather than random disagreement. The full three-panel explainability comparison is shown in Figure 4.

Table 4. Quantum Explainability: SHAP vs Quantum Gradient Attribution Rankings

Feature	SHAP Value	SHAP Rank	Grad (norm.)	Grad Rank
REVEL	0.2061	1	1.0000	1
PolyPhen-2	0.1292	2	0.3669	4
phyloP	0.1167	3	0.5281	3
AlphaMissense	0.0574	4	0.2505	6
CADD	0.0505	5	0.2895	5
is_transition	0.0340	6	0.0697	9
gc_content	0.0168	7	0.0721	8
SIFT	0.0090	8	0.6535	2
in_hbb_cds	0.0083	9	0.1389	7

Spearman rank correlation:  $\rho = 0.517$ ,  $p = 0.154$ . SHAP values are mean absolute GradientExplainer attributions. Gradient values are normalised mean  $|d(\text{output})/d(\text{input})|$ .




Figure 4. Quantum explainability analysis. (A) SHAP GradientExplainer feature importance for the HQNN; REVEL dominates (mean  $|\text{SHAP}| = 0.206$ ). (B) Quantum gradient attribution (normalised  $|d(\text{output})/d(\text{input})|$ ); SIFT rises markedly to rank 2 (0.654) relative to SHAP rank 8. (C) Rank concordance scatter plot; Spearman  $\rho = 0.517$  ( $p = 0.154$ ) indicates moderate agreement between the two attribution paradigms.

## 5. Discussion

### 5.1 Quantum Kernel SVM: Meaningful Improvement over Classical SVM

The QKSVM achieved ROC-AUC 0.9332 versus 0.8706 for the classical RBF-kernel SVM – a gain of 0.0626 – providing clear empirical evidence that quantum feature maps improve upon classical kernel methods for this genomic classification task. The fidelity kernel, by mapping the four top-SHAP features into a 16-dimensional complex

Hilbert space through ZZ entanglement, appears to capture nonlinear feature relationships that the Gaussian RBF kernel does not. The improvement in PR-AUC from 0.4836 to 0.6364 (+0.1528) is particularly noteworthy given the class imbalance, suggesting the quantum kernel provides better precision-recall trade-offs for the rare pathogenic class.

The QKSVM does not surpass the gradient boosting or deep learning baselines. This is consistent with theoretical expectations: quantum

kernel methods on small tabular datasets with modest qubit counts are unlikely to match mature ensemble methods that have been optimised for this specific problem type [21]. The kernel was computed on a 400-sample subsample due to  $O(N^2)$  complexity; increasing to the full SMOTE-balanced training set with more efficient kernel approximations could improve performance in future work.

### 5.2 HQNN: High Recall with Expected Performance Constraints

The HQNN achieved ROC-AUC 0.9210 and pathogenic recall of 0.96, the latter matching the classical DL base model and second only to Random Forest (recall = 1.00) across all models. For clinical applications where minimising false negatives is paramount – missing a pathogenic variant in a patient with  $\beta$ -thalassemia has direct diagnostic consequences – this recall performance is clinically relevant despite the lower overall AUC.

The HQNN does not surpass the classical class-weighted DL model (0.9210 vs 0.9483 ROC-AUC). This finding aligns with the conclusion of Paper 2 that performance in this problem is primarily feature-driven rather than architecture-driven: the quantum layer adds only 92 parameters (proj: 64, weights: 24, bias: 4) to the 4,501 classical parameters, and the four-qubit circuit processes only a four-dimensional projection of the 16-dimensional classical embedding. The quantum layer may be insufficiently expressive relative to the rich feature representations already learned by the classical encoder. This is an important empirical contribution: it provides evidence against the assumption that adding quantum components to a well-optimised classical model necessarily improves performance on tabular genomic data.

The low optimal threshold (0.163) indicates the HQNN is calibrated toward high sensitivity, which is appropriate for pathogenic variant screening but leads to reduced precision (0.50). Threshold calibration post-training or Platt scaling could improve precision without sacrificing the high recall.

### 5.3 Quantum Explainability: Cross-Paradigm Attribution Analysis

The moderate Spearman concordance ( $\rho = 0.517$ ) between SHAP and quantum gradient attribution reveals both convergent validity and genuine divergences of scientific interest. The convergent finding – REVEL as rank 1 in both methods, phyloP in the top three in both – strengthens confidence in the biological relevance of these features across both classical and quantum attribution paradigms. This cross-paradigm agreement provides a form of validation not available in purely classical analyses.

The most scientifically noteworthy divergence is SIFT's rise from SHAP rank 8 to gradient rank 2. SHAP measures marginal average contribution to the model's output, while input gradients measure local sensitivity: the quantum circuit is highly sensitive to perturbations in SIFT values even though those perturbations contribute little to average predictions. A plausible mechanism is that SIFT's binary-like distribution (values concentrated near 0 and 1) creates sharp gradient transitions at the decision boundary within the quantum feature space. Alternatively, SIFT may interact with REVEL through the learned linear projection (16 $\rightarrow$ 4) in a way that the quantum circuit is sensitive to but that averages out in SHAP attribution. This divergence is a falsifiable hypothesis for future work: probing whether SIFT's gradient sensitivity reflects genuine biological signal or a circuit artifact.

AlphaMissense's drop from SHAP rank 2 to gradient rank 6 is also notable. This suggests that while AlphaMissense provides consistent average predictive value (SHAP), the quantum circuit is less sensitive to local perturbations in its values – potentially because the structural information it encodes is partially redundant with REVEL's ensemble signal in the low-dimensional quantum projection space.

### 5.4 Position in the QML-in-Genomics Literature

This work is the first experimental validation of the use of QML for predicting pathogenicity in HBB variants. The result that the QKSVM demonstrates significantly better performance than the classical KSVM (+0.0626 ROC-AUC)

whereas the HQNN performs only slightly worse than the best classical model ( $-0.0273$  ROC-AUC) corroborates the general trend observed in the biomedical literature benchmarking QML: quantum kernel algorithms show greater immediate promise than variational quantum neural networks in small tabular tasks due to the direct superiority of the kernel Hilbert space over the circuit expressivity [21, 44]. Lastly, the analysis of quantum explainability contributes a new methodology regardless of the results, namely, the approach of making a three-fold comparison between classical and quantum explainability metrics offers a model for hybrid QML attribution analysis.

### 5.5 Limitations

Several limitations should be acknowledged. First, the QKSVM was trained on a 400-sample subsample of the SMOTE-balanced training set, limiting kernel expressiveness. Full kernel computation on the complete training set would require either kernel approximation methods (e.g., Nyström approximation) or parallelised circuit evaluation. Second, the dataset is limited to 1,585 ClinVar variants; external validation on independent cohorts is required before clinical deployment. Third, the non-significant Spearman p-value ( $p=0.154$ ) reflects the small sample size of nine features rather than poor concordance, but precludes strong statistical conclusions about rank agreement.

### 6. Conclusion

This study presents the first quantum machine learning framework for HBB gene variant pathogenicity prediction, extending a two-paper series of progressively sophisticated classical approaches. Three quantum modules were implemented and evaluated on the identical 1,585-variant ClinVar dataset used in prior work. The Quantum Kernel SVM achieved ROC-AUC 0.9332 and PR-AUC 0.6364, representing a meaningful improvement over the classical RBF-kernel SVM baseline (ROC-AUC 0.8706,  $+0.0626$ ) and demonstrating that quantum fidelity kernels provide superior separation of pathogenic from benign HBB variants compared to classical

kernel methods. The Hybrid Quantum-Classical Neural Network achieved ROC-AUC 0.9210 and pathogenic recall of 0.96, maintaining the clinically critical sensitivity of the classical deep learning architecture while operating with a parameterised quantum output layer. The quantum explainability analysis produced the first cross-paradigm variant attribution comparison for this gene, revealing REVEL as the dominant predictor across all attribution methods while identifying SIFT and AlphaMissense as exhibiting notable divergence between SHAP and quantum gradient attribution (Spearman  $\rho = 0.517$ ).

These findings establish a rigorous empirical benchmark for quantum versus classical performance in HBB variant classification, contribute the first quantum explainability analysis for a clinical genomic classification task, and provide a methodological framework for integrating quantum machine learning into precision genomic medicine pipelines. Future work should pursue kernel approximation methods to enable full-training-set QKSVM computation, deeper variational circuits with barren plateau mitigation strategies, integration of population frequency data from gnomAD, and prospective validation in clinical  $\beta$ -thalassaemia screening workflows.

### Acknowledgements

The authors acknowledge the ClinVar database maintained by the NCBI, the myvariant.info service for programmatic variant annotation access, and Google Colab Pro+ for providing NVIDIA A100 GPU access for quantum circuit simulation. No funding was received for this study.

### Ethical Approval

This research utilised publicly available, de-identified data from the ClinVar database and the myvariant.info API. No human participants were recruited, no biological samples were collected, and no personal health information was accessed. Ethical approval was therefore not required.

### Conflicts of Interest

The authors declare no conflict of interest regarding the publication of this manuscript.

**Author Contributions**

Raazia Sosan Waseem conceptualised the quantum extension framework, designed and implemented the quantum circuits (QKSVM, HQNN, quantum explainability), performed all computational experiments on the A100 GPU, conducted statistical analyses, and drafted the manuscript.

Muhammad Hussain Habib provided biological and clinical domain expertise for HBB and  $\beta$ -thalassemia, contributed to clinical interpretation of quantum explainability findings, and critically reviewed the manuscript for intellectual content.

All authors approved the final version for publication.

**References**

- Orkin SH, Bauer DE. Emerging genetic therapy for sickle cell disease and thalassemia. *Annu Rev Med.* 2019;70:257–271.
- Weatherall DJ. The inherited diseases of hemoglobin are an emerging global health burden. *Blood.* 2010;115(22):4331–4336.
- Cao A, Galanello R. Beta-thalassemia. *Genet Med.* 2010;12(2):61–76.
- Modell B, Darlison M. Global epidemiology of haemoglobin disorders. *Bull World Health Organ.* 2008;86(6):480–487.
- Taher AT, Weatherall DJ, Cappellini MD. Thalassaemia. *Lancet.* 2018;391(10116):155–167.
- Ahmed S, et al. Prenatal diagnosis of beta-thalassaemia in Pakistan. *Prenat Diagn.* 2000;20(5):378–383.
- Landrum MJ, et al. ClinVar: improving access to variant interpretations. *Nucleic Acids Res.* 2018;46(D1):D1062–D1067.
- Landrum MJ, et al. ClinVar: improvements to accessing data. *Nucleic Acids Res.* 2020;48(D1):D835–D844.
- Giardine B, et al. Updates of the HbVar database of human hemoglobin variants and thalassemia mutations. *Nucleic Acids Res.* 2014;42(D1):D1063–D1069.
- Pejaver V, et al. Calibration of computational tools for missense variant pathogenicity classification. *Am J Hum Genet.* 2022;109(12):2163–2177.
- Frangoul H, et al. CRISPR-Cas9 gene editing for  $\beta$ -thalassemia. *N Engl J Med.* 2021;384(3):252–260.
- Waseem RS, Habib MH. Machine learning-based pathogenicity prediction and prioritization of HBB gene variants using ClinVar and population-scale genomic data. *Res Med Sci Rev.* 2026;4(3):918–928.
- Waseem RS, Habib MH. Deep learning with enriched pathogenicity scores for HBB variant classification in  $\beta$ -thalassemia. (2026). *Review Journal of Neurological & Medical Sciences Review,* 4(3), 389-402. <https://doi.org/10.63075/9em1f154>
- Biamonte J, et al. Quantum machine learning. *Nature.* 2017;549(7671):195–202.
- Cerezo M, et al. Variational quantum algorithms. *Nat Rev Phys.* 2021;3(9):625–644.
- Schuld M, Killoran N. Quantum machine learning in feature Hilbert spaces. *Phys Rev Lett.* 2019;122(4):040504.
- Havlíček V, et al. Supervised learning with quantum-enhanced feature spaces. *Nature.* 2019;567(7747):209–212.
- Bergholm V, et al. PennyLane: automatic differentiation of hybrid quantum-classical computations. *arXiv [preprint].* 2018:arXiv:1811.04968.
- McClellan JR, et al. The theory of variational hybrid quantum-classical algorithms. *New J Phys.* 2016;18(2):023023.
- Heese R, et al. Explaining quantum circuits with Shapley values. *arXiv [preprint].* 2023:arXiv:2301.09138.
- Huang HY, et al. Power of data in quantum machine learning. *Nat Commun.* 2021;12:2631.

22. Ng PC, Henikoff S. SIFT: predicting amino acid changes that affect protein function. *Nucleic Acids Res.* 2003;31(13):3812–3814.
23. Adzhubei IA, et al. A method and server for predicting damaging missense mutations. *Nat Methods.* 2010;7(4):248–249.
24. Rentzsch P, et al. CADD: predicting the deleteriousness of variants throughout the human genome. *Nucleic Acids Res.* 2019;47(D1):D886–D894.
25. Ioannidis NM, et al. REVEL: an ensemble method for predicting the pathogenicity of rare missense variants. *Am J Hum Genet.* 2016;99(4):877–885.
26. Cheng J, et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. *Science.* 2023;381(6664):eadg7492.
27. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 2016:785–794.
28. Heo JY, Kim JH. Assessing pathogenicity prediction methods using XGBoost and conservation metrics. *BMC Genomics.* 2025;26:11787.
29. Xu Z, et al. Phenotype-aware prioritization of pathogenic variants using XGBoost. *Nat Commun.* 2023;14:43651.
30. Grinsztajn L, et al. Why tree-based models still outperform deep learning on tabular data. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2022;35:507–520.
31. Liu Y, et al. A rigorous and robust quantum speed-up in supervised machine learning. *Nat Phys.* 2021;17(9):1013–1017.
32. Preskill J. Quantum computing in the NISQ era and beyond. *Quantum.* 2018;2:79.
33. Bharti K, et al. Noisy intermediate-scale quantum algorithms. *Rev Mod Phys.* 2022;94(1):015004.
34. Jones T, Gacon J. Efficient calculation of gradients in classical simulations of variational quantum algorithms. *arXiv [preprint]*. 2020:arXiv:2009.02823.
35. McClean JR, et al. Barren plateaus in quantum neural network training landscapes. *Nat Commun.* 2018;9:4812.
36. Shaydulin R, Wild SM. Importance of kernel bandwidth in quantum machine learning. *Phys Rev A.* 2022;106(4):042407.
37. Meyer JJ. Fisher information in noisy intermediate-scale quantum applications. *Quantum.* 2021;5:539.
38. Ezhov A. Quantum computing explainability. *arXiv [preprint]*. 2023:arXiv:2302.00489.
39. Ghosh S, et al. Quantum machine learning for cancer subtype classification. *Sci Rep.* 2023;13:1012.
40. Bravyi S, et al. Quantum advantage with noisy shallow circuits. *Nat Phys.* 2020;16(10):1040–1045.
41. Innan M, et al. Quantum machine learning for drug-target interaction prediction. *arXiv [preprint]*. 2024:arXiv:2407.00557.
42. Xin J, et al. High-performance web services for querying gene and variant annotation. *Genome Biol.* 2016;17:91.
43. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2017;30.
44. Thanasilp S, et al. Exponential concentration in quantum kernel methods. *Nat Commun.* 2024;15:5200.