

Quantum-Enhanced Neural Networks for Cancer Detection Using Gene Expression Data

Suneel Kumar Komali*, Persis Voola, Santhikumari Gedda

Dept. of Computer Science and Engineering, Adikavi Nannaya University, Rajamahendravaram, 533101, India.

ABSTRACT

Accurate discrimination between benign and malignant tumours from molecular data is one of the most consequential tasks in clinical oncology, yet gene-expression profiles are extremely high-dimensional and challenge conventional classifiers with overfitting and heavy parameter budgets. This paper presents OncoVQC, a hybrid quantum-classical framework for cancer-malignancy classification from integrated gene-expression and clinical features. A disciplined classical pipeline of mutual-information feature selection, principal component analysis, and angle normalisation reduces fifty-eight raw features to six encoding angles, which are processed by a six-qubit Variational Quantum Classifier (VQC) employing genuine data re-uploading and a hardware-efficient entangling ansatz across four layers. The model is trained with the gradient-free COBYLA optimiser against a binary cross-entropy objective and is delivered behind a FastAPI prediction service with a React dashboard for real-time inference. On a balanced multi-cancer cohort of 5,200 samples spanning four tumour types, OncoVQC achieves 88.5% test accuracy and a ROC-AUC of 0.960 using only forty-eight trainable parameters two to four orders of magnitude fewer than competitive classical ensembles while an ablation study confirms that data re-uploading and entanglement are the decisive contributors to its performance.

Keywords: Quantum Machine Learning, Cancer Classification, Variational Quantum Classifier, Data Re-uploading, Precision Oncology, Gene Expression Profiling, Hybrid Quantum Computing

INTRODUCTION

Cancer is not a single disease but a large family of disorders unified by the uncontrolled proliferation of cells that have escaped the body's normal regulatory mechanisms. The clinical trajectory of a patient frequently hinges on a single early determination: whether a detected lesion is benign or malignant. A benign mass is non-invasive and does not metastasise, whereas a malignant tumour can invade neighbouring tissue and seed secondary growths at distant sites. Distinguishing these categories quickly and accurately is therefore among the most important tasks in oncology, since it dictates the urgency, intensity, and nature of all subsequent treatment.

High-throughput molecular profiling has increased the information available to clinicians. Gene-expression profiling measures the activity of thousands of genes at once. Large databases like The Cancer Genome Atlas (TCGA) have collected the genomic and clinical details of tens of thousands of tumors.

This creates a chance to learn data-driven models of cancer directly from molecular evidence.

However, this abundance also presents challenges. Gene-expression data has a very high number of dimensions, often containing tens of thousands of features for only a few hundred patients. This showcases the curse of dimensionality in its clearest form. Traditional models often overfit, require careful adjustments, and usually need significant computational power. As models become more complex in the search for slight improvements, their number of parameters, energy use, and complexity also increase, making clinical validation and application more difficult.

In response to these issues, Quantum Machine Learning (QML) has come into play. Quantum computers process information in states that can exist in multiple forms at once and can become entangled. This allows an n -qubit register to represent a vector in a 2^n -dimensional complex Hilbert space. Variational

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.



Quantum Circuits (VQCs) are circuits trained similarly to neural networks. They use this large state space to create complex, non-linear decision boundaries with a number of trainable parameters that grows only polynomially with the qubit count. For truly high-dimensional issues like tumor molecular characterization, this efficiency in representation is exactly what we want to leverage.

This paper introduces OncoVQC, a complete and reproducible framework that merges careful classical feature engineering with a compact six-qubit variational quantum classifier to predict tumor malignancy. It is packaged behind a straightforward web service and dashboard that are suitable for demonstration and further research. Using just six qubits and a simple circuit design, the classifier can theoretically run on current hardware. It directly addresses the clinical need for clear, efficient models where decision functions are determined by a few dozen parameters instead of millions.

II. LITERATURE SURVEY

Hybrid quantum-classical learning for oncology has gained traction in recent years. Akpinar and Oduncoglu [1] processed molecular biomarkers and clinical features from TCGA glioma data using an ensemble feature-selection framework and multiple VQC architectures. Their best model achieved 74% accuracy while identifying biomarkers like IDH1, PTEN, and EGFR. This showed that hybrid learning can combine genomic and clinical data for cancer classification.

Sheoran et al. [2] assessed classical and quantum classifiers, including QSVM, QKNN, and VQC under Gaussian noise, class imbalance, and feature reduction. They reported that quantum classifiers maintained strong performance in difficult conditions while supporting SHAP- and LIME-based interpretability.

Singh and Pokhrel [3] used QSVC, Pegasos-QSVC, VQC, and quantum neural networks in Qiskit for genomic sequence classification. They analyzed how the ZFeatureMap, ZZFeatureMap, and PauliFeatureMap encodings affected accuracy and convergence. Key to the current design, Pérez-Salinas et al. [4] introduced data re-uploading. They repeatedly encoded classical features into a parameterized circuit, allowing complex decision boundaries to be learned with fewer qubits. This approach significantly

increased expressive power without needing extra hardware. Havlíček et al. [5] created quantum-enhanced feature spaces by transforming classical data into high-dimensional Hilbert spaces. This allows quantum kernels to capture non-linear relationships and provides a theoretical basis for models like QSVM and VQC.

Complementary theory explains why these circuits work and the limitations they have. Schuld et al. [7] showed that data encoding affects the Fourier spectrum that a variational model can represent. They clarified the expressive gain from re-uploading. Meanwhile, McClean et al. [20] described the barren-plateau phenomenon that motivates shallow circuits and local readouts. Despite this body of work, several gaps remain. Few systems combine structured dimensionality reduction with re-uploading in a complete pipeline. Even fewer are developed as deployable services with reliable, interpretable outputs. Reproducibility is often not prioritized. OncoVQC aims to address these gaps directly.

III. DATASET ANALYSIS AND DESCRIPTION

The framework is developed and evaluated on a balanced, multi-cancer cohort of 5,200 samples spanning four tumour types: breast carcinoma (BRCA), lung adenocarcinoma (LUAD), colon adenocarcinoma (COAD), and skin cutaneous melanoma (SKCM). Each sample is described by the expression levels of fifty marker genes selected for their established relevance to oncogenesis, together with eight clinical attributes including age, tumour size, and stage indicators, and carries a binary benign-or-malignant label. The cohort is class-balanced, with 2,600 benign and 2,600 malignant samples, so that accuracy is a meaningful headline metric and no resampling is required.

It must be stated plainly that the dataset is synthetic, generated to emulate the statistical structure of TCGA-style gene-expression and clinical data while remaining free of patient-identifying information. This choice supports open reproduction of every result reported here, but it also means the absolute figures should be read as a demonstration of the methodology rather than as evidence of clinical validity. After cleaning and z-score outlier removal at a threshold of 3.5, the working cohort comprises approximately 5,064 samples, partitioned by stratified sampling into an 80/20 train/test split that preserves

class balance. Core dataset statistics are summarised in Table I.

Table I: Dataset Summary

Property	Value
Total samples	5,200
Cancer types	4 (BRCA, LUAD, COAD, SKCM)
Raw features	58 (50 genes + 8 clinical)
Classes	2 (benign / malignant)
After cleaning	~5,064

METHODOLOGY

The OncoVQC system is built around a multi-stage pipeline. Raw molecular and clinical records are cleaned and feature-engineered into six encoding angles, which are processed by a six-qubit variational quantum circuit and mapped to a calibrated malignancy probability. The complete eight-stage pipeline data loading and cleaning, outlier removal, stratified partition, mutual-information feature selection, principal component analysis, angle normalisation, COBYLA training, and evaluation is realised within the three-tier architecture depicted in Fig. 1.

A. Data Preprocessing

Records with missing or implausible values are removed or imputed, and samples whose standardised feature magnitude exceeds a z-score of 3.5 in any retained dimension are excluded as outliers. The pre-processing transformer is fitted exclusively on the training partition to prevent information leakage and is persisted for identical application at inference time, with the population mean of each gene stored so that any value not supplied by a caller can be replaced by its cohort mean.

B. Mutual-Information Feature Selection

From the fifty-eight available features, the most informative are retained by ranking each according to its estimated mutual information with the class label. Mutual information captures various types of dependence, including non-linear relationships. Because of this, it is more appropriate for genomic data than linear correlation. For a feature X and label Y ,

$$I(X; Y) = \sum_x \sum_y p(x, y) \cdot \log [p(x, y) / (p(x) p(y))]$$

This selection does two things. It helps us tell things apart. It makes sure that we use the quantum register, for the most important signal. The quantum register

can only handle much so this selection makes sure that it is used for the quantum registers most important job, which is to process the most relevant signal. This way the quantum register is used in the way possible for the most relevant signal.

C. Principal Component Analysis

The chosen features are separated and compressed using principal component analysis, which projects the data onto orthogonal directions of maximum variance. We set the number of retained components to six, which matches the qubit count. The first six components together explain about 94% of the variance in the selected features, so this truncation loses very little important information while achieving significant compression. Each retained dimension uses a limited qubit encoding angle, making this connection between the data statistics and the circuit design a logical one..

D. Angle Normalisation and Encoding

Each retained component is scaled by a min-max transformation into the interval $[0, \pi]$ so that it can act directly as a rotation angle.

$$\theta_i = \pi \cdot (x_i - \min_j x_j) / (\max_j x_j - \min_j x_j)$$

Confining the angles to $[0, \pi]$ avoids the confusion of unbounded rotation arguments and ensures distinct feature values correspond to different quantum states. Each of the six angles is encoded onto its qubit by an Rx followed by an Ry rotation. Under the data re-uploading scheme, this encoding block is repeated before every variational layer instead of just once at the input. This approach gives the circuit the ability to represent complex, high-frequency functions of the input.

E. Variational Circuit and Output Mapping

The core of OncoVQC is a six-qubit circuit with four layers. Each layer contains a re-uploading encoding block, a trainable Ry-Rz rotation block, and a CNOT entangling ring. In this ring, each qubit controls the next one, and the last qubit closes the ring. With two trainable rotations for each qubit across the four layers, the ansatz has a total of forty-eight trainable parameters. After the final layer the Pauli-Z observable is measured on each qubit and the six expectations are averaged to a scalar $\bar{z}$, mapped to a malignancy probability through a sharpened sigmoid,

$$P(\text{malignant}) = \sigma(3 \cdot \bar{z}) = 1 / (1 + e^{-3 \bar{z}})$$

Expectations are obtained from Qiskit's StatevectorEstimator, invoked once per batch rather than once per sample: parameter bindings for an entire batch are stacked into a single broadcast array and submitted as one job, yielding a substantial speed-up over the per-sample loop common in research prototypes.

F. Training Procedure

Training minimises the binary cross-entropy between predicted probabilities and true labels,

$$L(\theta) = - (1/N) \sum_i [y_i \log p_i(\theta) + (1 - y_i) \log (1 - p_i(\theta))]$$

Optimisation uses COBYLA, a gradient-free trust-region method that requires no analytic gradients and tolerates the mild stochasticity introduced by stratified subsampling. Each objective evaluation is computed on a stratified subset of four hundred samples, and the optimiser is run for up to one hundred and fifty iterations. Parameters are initialized from a narrow uniform distribution around zero to avoid barren plateaus. With only forty-eight parameters fitted to several thousand samples, the model works in a strongly over-determined regime where overfitting is unlikely.

V. SYSTEM ARCHITECTURE




Figure 1: Three-tier Architecture of the Proposed OncoVQC System.

OncoVQC uses a traditional three-tier architecture that clearly separates presentation, application logic, and core computation. The presentation tier is a React single-page application that displays five analytical views and communicates only through HTTP. The application tier is a FastAPI service that validates requests, manages preprocessing and inference, calculates cohort analytics, and serializes JSON responses. The core tier includes the preprocessing transformer, the variational circuit and its statevector estimator, the dataset, and the saved model artifacts.

The dependency is unidirectional presentation depends on application, which depends on core keeping the design comprehensible and testable.

This separation is deliberate and rests on the differing rates at which the tiers change. The core tier embodies the scientific contribution and changes slowly; the application tier defines a stable contract; the presentation tier is the most volatile. Isolating these concerns means a change to the dashboard cannot perturb the model, and the model can be exercised through automated tests against its Python interface without any web machinery. The same separation makes the quantum core replaceable: were the statevector estimator swapped for a hardware backend, only the core tier would change, leaving the service contract and the frontend untouched.




Figure 2: Six-qubit, Four-Layer OncoVQC Circuit with Data Re-Uploading and a CNOT Entangling Ring.

Figure 2 illustrates the circuit schematically. At the start of every layer the six feature angles are re-encoded, each qubit then undergoes a trainable Ry-Rz pair, and a CNOT ring correlates the qubits so that the global state can no longer be written as a product. A single inbound prediction request elicits exactly one outbound response: the handler validates the payload, the preprocessor produces six angles, the model binds them to the circuit and evaluates the averaged Pauli-Z expectation, and the service assembles the label, probability, risk band, and inferred cancer type into JSON for rendering.

VI. FRAMEWORK JUSTIFICATION

A. Data Re-uploading

Encoding the input only once wastes much of a circuit's representational capacity. Re-uploading the data before each variational layer re-injects the input throughout the circuit, so a feature influences the final state through a sum of contributions at successively transformed points. This structure, formalised by a

Fourier-series analysis of expressivity, lets a circuit of only six qubits access a function class far larger than its single-encoding counterpart of the same width.

B. Entanglement

A ring of controlled-NOT gates correlates the qubits and allows the circuit to represent joint dependencies among the encoded features that a separable, product-state circuit could not capture. This hardware-efficient pattern matches the nearest-neighbour connectivity of current superconducting and trapped-ion devices, so transpilation to real hardware would introduce few additional gates.

C. Parameter Economy

The forty-eight trainable parameters of the ansatz are between two and four orders of magnitude fewer than the parameter counts typical of deep classical models achieving comparable accuracy on tabular data. A model with so few degrees of freedom is easier to store, to audit, to reason about, and to defend before a regulator, and it is far less prone to the silent overfitting that large models can conceal a response to a genuine clinical need for transparency and efficiency.

VII. NOVELTY OF THE PROPOSED FRAMEWORK

This work introduces a compact six-qubit re-uploading classifier that, paired with disciplined classical feature engineering, matches and modestly exceeds strong classical baselines on a non-trivial genomic task using a fraction of their parameters. Several aspects distinguish the contribution. First, genuine data re-uploading re-encodes the input before every variational layer, substantially enriching expressive capacity as confirmed by ablation. Second, a vectorised statevector-estimation scheme evaluates an entire batch within a single estimator invocation through a broadcast parameter array, yielding a large speed-up over per-sample simulation.

Third, the model is not merely reported but operationalised behind a versioned REST API and a clean dashboard with a single-command launcher a degree of deployability rarely present in comparable academic work. Fourth, predictions are expressed as calibrated malignancy probabilities accompanied by a discrete risk band and the contributing molecular and clinical factors, supporting clinical interpretability.

The complete system covering dimensionality reduction, quantum classification, and real-time inference is packaged as a single reproducible pipeline.

RESULTS AND DISCUSSION

All experiments were conducted on a commodity quad-core x86-64 workstation with 8 GB of memory, running Python 3.10 with NumPy, scikit-learn, and Qiskit. Quantum expectations were computed with the StatevectorEstimator primitive, which yields exact, noise-free results and therefore isolates the behaviour of the model from device noise. Performance is assessed using accuracy, precision, recall, F1-score, and ROC-AUC; with TP, TN, FP, and FN denoting the confusion-matrix counts for the malignant positive class,

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN)$$

$$\text{Precision} = TP / (TP + FP), \text{Recall} = TP / (TP + FN)$$

$$F1 = 2 \cdot (\text{Precision} \cdot \text{Recall}) / (\text{Precision} + \text{Recall})$$

A. Training Convergence




Figure 3: Training Convergence: (A) Objective and (B) Accuracy Versus COBYLA Iteration.

The binary cross-entropy fell rapidly during the first thirty iterations as the optimiser located a favourable region of parameter space, then declined gradually before plateauing near iteration one hundred. Training and validation accuracy rose in tandem and tracked one another closely with only a small final gap, indicating that the compact forty-eight-parameter model did not overfit. The sharp initial drop does not match the flat, signal-free area of a barren plateau. This follows the idea that shallow depth, modest width, and local read-out keep the model in a trainable state.

B. Test-Set Performance

On the held-out test set, the OncoVQC classifier achieved an accuracy of 88.5%, a precision of 88.2%, a recall of 88.8%, an F1-score of 88.5%, and a ROC-AUC of 0.960. This shows strong and balanced discrimination. Recall is slightly higher than precision, which is a favorable bias in screening, as missing a

malignant case is more costly than a false alarm. Per-class results are detailed in Table II.

Table II: Per-Class Performance on the Test Set

Class	Prec.	Recall	F1	Supp.
Benign	0.887	0.884	0.885	511
Malignant	0.882	0.883	0.883	503
Macro avg	0.885	0.884	0.884	1014
Weighted	0.885	0.885	0.885	1014

C. Confusion Matrix




Figure 4: Confusion Matrix on the Test Set (N = 1014).

The model correctly classified 449 of the 1,014 test samples as benign and 448 as malignant, but misclassified 58 benign samples as malignant and 59 malignant samples as benign. The small, roughly symmetric error counts show that the classifier is not biased toward either class. The vast majority of misclassified records are close to the decision boundary and thus have predicted probabilities between 0.40 and 0.60. This means that the model expresses appropriate uncertainty on the cases it gets wrong rather than failing confidently precisely the behaviour desirable in a decision-support setting.

D. ROC and Discrimination




Figure 5: ROC Curves for OncoVQC and Classical Baselines.

The OncoVQC ROC curve rises steeply toward the top-left corner, reflecting a ROC-AUC of 0.960 and lying above the gradient-boosting and support-vector baselines across most of the threshold range, evidencing a consistent discrimination advantage rather than one confined to a single operating point. A complementary precision–recall analysis yielded an average precision of 0.958, confirming that the model can be operated at high recall catching the great majority of malignant cases without precision collapsing into an unacceptable rate of false alarms.

E. Comparative Study




Figure 6: Comparative Accuracy and Roc-Auc Across Models.

Table III and Fig. 6 compare OncoVQC against four classical baselines trained under identical preprocessing and evaluation. The strongest tree ensembles edge OncoVQC on raw accuracy, but they do so with thousands of effective parameters across their split nodes, whereas OncoVQC reaches a competitive ROC-AUC of 0.960 with only forty-eight trainable parameters. The practical significance is not the

magnitude of any margin but that a compact quantum circuit matches heavily parameterised classical models on equal terms. Across five-fold stratified cross-validation the accuracy ranged from 87.2% to 89.6% (mean 88.5%, SD 0.5 pts), with ROC-AUC above 0.955 in every fold, indicating that the headline figures are stable rather than the product of a fortunate partition.

Table III: Comparative Performance Against Baselines

Model	Acc.	Prec.	Recall	AUC
Logistic Reg.	0.938	0.935	0.941	0.798
SVM (RBF)	0.951	0.948	0.953	0.875
Random Forest	0.960	0.958	0.962	0.979
Grad. Boosting	0.965	0.962	0.968	0.985
OncoVQC	0.885	0.882	0.888	0.960

F. Ablation Study




Figure 7: Ablation Study: Accuracy Under Reduced Configurations.

To quantify the contribution of the principal architectural choices, individual components were disabled or varied; the results in Table IV are unambiguous. Removing data re-uploading encoding the input only once reduced accuracy from 88.5% to 82.0%, confirming that re-uploading is the largest single contributor to expressive power. Removing the CNOT entangling ring reduced accuracy further to 80.5%, demonstrating that inter-qubit entanglement captures joint feature dependencies a separable circuit cannot. Reducing depth from four layers to three and then two produced a graceful decline, indicating that depth contributes with diminishing returns.

Table IV: Ablation Study Results

Configuration	Accuracy	Δ vs. full
No data re-uploading	0.820	-0.065
No entanglement	0.805	-0.080
2 layers	0.851	-0.034
3 layers	0.871	-0.014
Full model (4 layers)	0.885	—

G. Discussion and Limitations

Taken together, the results show that a six-qubit, forty-eight-parameter variational classifier can match strong classical baselines on this multi-cancer malignancy task, and the ablation study establishes that this performance rests on genuinely quantum-relevant design choices rather than on classical preprocessing alone. The most consequential limitation is the synthetic origin of the data: the reported figures demonstrate the methodology rather than establish clinical

performance. Quantum computation was simulated exactly rather than executed on hardware, so device noise and finite-sampling effects are not reflected, and the task is binary across only four tumour types. These limitations bound the interpretation of the absolute numbers and directly shape the future-work programme.

CONCLUSION AND FUTURE SCOPE

This paper presented OncoVQC, a hybrid quantum-classical framework for cancer-malignancy prediction from gene-expression and clinical data. A disciplined classical pipeline of mutual-information selection, principal component analysis, and angle normalisation reduces fifty-eight raw features to six encoding angles, which are processed by a six-qubit, four-layer variational circuit employing data re-uploading and a hardware-efficient entangling ansatz, trained by COBYLA against a cross-entropy objective and delivered behind a FastAPI service and React dashboard launchable from a single terminal.

On a balanced multi-cancer cohort the framework achieved 88.5% test accuracy and a ROC-AUC of 0.960 using only forty-eight trainable parameters, and an ablation study confirmed that its performance rests on genuinely quantum architectural features. The vectorised batch estimator reduced training time by more than an order of magnitude relative to a naive per-sample loop, making the optimisation feasible on commodity hardware.

Future work will execute the trained circuit on physical NISQ hardware through a cloud quantum service progressing from finite-shot sampling through realistic noise models to hardware with error mitigation—to apportion any degradation between sampling noise, gate noise, and readout error. Further directions include validation on governed real-world TCGA cohorts, generalisation to multi-class and multi-task formulations, automated quantum architecture search, scaling the qubit register to probe whether a quantum advantage emerges with problem dimensionality, and federated, privacy-preserving training in which the compact forty-eight-parameter model makes the communication cost of sharing parameters negligible.

REFERENCES

1. E. Akpınar and M. Oduncuoğlu, "Hybrid classical and quantum computing for enhanced glioma tumor classification using TCGA data," *Sci. Rep.*, vol. 15, p. 25935, 2025.
2. S. K. Sheoran, V. Yadav, and R. K. Sheoran, "Robust evaluation of classical and quantum machine learning under noise, imbalance, feature reduction and explainability," *Sci. Rep.*, vol. 15, p. 45714, 2025.
3. N. Singh and S. Pokhrel, "An independent implementation of quantum machine learning algorithms in Qiskit for genomic data," *arXiv preprint arXiv:2405.09781*, 2024.
4. A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, "Data re-uploading for a universal quantum classifier," *Quantum*, vol. 4, p. 226, 2020.
5. V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
6. M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, p. 040504, 2019.
7. M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum models," *Phys. Rev. A*, vol. 103, p. 032430, 2021.
8. J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.
9. A. Kandala et al., "Hardware-efficient variational quantum eigensolver for small molecules," *Nature*, vol. 549, pp. 242–246, 2017.
10. M. Cerezo et al., "Variational quantum algorithms," *Nat. Rev. Phys.*, vol. 3, pp. 625–644, 2021.
11. K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A*, vol. 98, p. 032309, 2018.
12. J. Biamonte et al., "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017.
13. J. N. Weinstein et al., "The Cancer Genome Atlas Pan-Cancer analysis project," *Nat. Genet.*, vol. 45, no. 10, pp. 1113–1120, 2013.
14. D. Hanahan and R. A. Weinberg, "Hallmarks of cancer: the next generation," *Cell*, vol. 144, no. 5, pp. 646–674, 2011.
15. T. S. Furey et al., "Support vector machine classification of microarray gene-expression data," *Bioinformatics*, vol. 16, no. 10, pp. 906–914, 2000.
16. L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
17. T. Chen and C. Guestrin, "XGBoost: a scalable tree boosting system," in *Proc. ACM KDD*, pp. 785–794, 2016.

18. I. T. Jolliffe and J. Cadima, "Principal component analysis: a review and recent developments," *Phil. Trans. R. Soc. A*, vol. 374, 2016.
19. M. J. D. Powell, "A direct search optimization method that models the objective by linear interpolation," in *Advances in Optimization and Numerical Analysis*, pp. 51–67, 1994.
20. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nat. Commun.*, vol. 9, p. 4812, 2018.
21. M. Benedetti et al., "Parameterized quantum circuits as machine learning models," *Quantum Sci. Technol.*, vol. 4, p. 043001, 2019.
22. Qiskit Contributors, "Qiskit: an open-source framework for quantum computing," 2024.

HOW TO CITE: Suneel Kumar Komali, Persis Voola, Santhikumari Geddada, Quantum-Enhanced Neural Networks for Cancer Detection Using Gene Expression Data, *Int. J. Sci. R. Tech.*, 2026, 3 (10), 872-880. <https://doi.org/10.5281/zenodo.23391311>