Quantum Neural Network (QNN) Architectures and Their Potential in High-Dimensional Medical Image Classification

Author: Aditya Singh **Affiliation:** Department of Data Science, University of Dhaka (Bangladesh)

Email: aditya.singh@du.ac.bd

2024

Abstract

The rapid expansion of medical imaging modalities (high-resolution CT, MRI, whole-slide histopathology, and multi-omics image representations) stresses current classical deep-learning pipelines with rising dimensionality, complex feature manifolds, and often limited labeled data. Quantum neural networks (QNNs) an umbrella term for quantum-native and hybrid variational models offer new representational primitives (superposition, entanglement) and feature-mapping strategies that may yield advantages for high-dimensional classification tasks, especially in low-data regimes or where quantum-enhanced feature spaces better separate classes. This paper provides an in-depth, journal-ready treatment of QNN architectures tailored to medical imaging: variational quantum circuits (VQCs), quantum convolutional neural networks (QCNNs), quanvolutional layers, quantum kernel methods and hybrid classical–quantum pipelines. We develop rigorous mathematical formulations (state and operator models, data encoding maps, loss functions, gradient estimators), analyze expressivity and trainability (including barren plateau phenomena and mitigation), and present architecture design patterns for integrating quantum layers with standard convolutional backbones. This paper proposes evaluation protocols, benchmarks, and realistic experimental blueprints (datasets, classical baselines, hardware vs. simulator tradeoffs), and discuss regulatory, privacy, and deployment concerns in clinical settings. The literature review synthesizes reviewed works and standards, situating QNNs within the current trajectory of quantum machine learning and clinical AI. Finally, we offer practical roadmaps: short-term hybrid pilots on cloud quantum resources and mid-term device-aware retrospectives as qubit counts and fidelities improve. While conclusive quantum advantage for large-scale medical imaging is yet unproven, carefully designed QNN architectures present a promising research direction that merits rigorous empirical and theoretical exploration.

Keywords: Quantum neural networks, QCNN, variational quantum circuits, medical image classification, high-dimensional data, hybrid quantum-classical models, quantum feature maps, barren plateaus, clinical imaging.

1. Introduction

Medical image classification has seen breakthroughs with convolutional neural networks (CNNs) and transformer architectures, achieving high accuracy across numerous tasks (tumor detection, lesion segmentation, disease progression). However, challenge areas persist: (i) **very high input**

dimensionality (e.g., whole-slide pathology images exceeding gigapixel scale), (ii) **limited labeled data** for rare pathologies, (iii) **complex multi-scale features** and manifold topologies, and (iv) **stringent latency, privacy, and interpretability requirements** in clinical deployment. Quantum machine learning (QML) a collection of algorithms that exploit quantum information processing promises new feature mappings and parameterized function classes potentially well suited for such challenges [Havlíček et al., 2019; comprehensive surveys]. This manuscript investigates QNN architectures and their potential to improve high-dimensional medical image classification, focusing on theoretical foundations, architecture design, training algorithms, benchmarking protocols, and translational pathways for healthcare systems.

2. Literature Review and State of the Art

2.1 Surveys and foundational QML works

Foundational works establishing supervised learning with quantum-enhanced feature spaces (Havlíček et al., 2019) framed how quantum circuits can map classical inputs into high-dimensional Hilbert spaces where linear decision rules may become powerful. Broader QML surveys synthesize families of quantum models (VQCs, quantum kernels, QCNNs) and outline both theoretical promise and practical NISQ constraints. Recent comprehensive surveys summarize progress and practical challenges (expressivity, trainability, noise).

2.2 QNN architectures for image data recent work

Multiple works have proposed quantum or hybrid convolutional architectures applied to medical imaging modalities (CT, MRI, histopathology), including pre-trained QCNNs for COVID-19 detection, QCNNs for breast cancer/ultrasound images, and hybrid quantum-classical CNNs for radiological classification. ArXiv and peer-reviewed publications report promising small-scale results, particularly in small-sample settings or when quantum layers act as expressive feature extractors. Examples include QCNN applications to Alzheimer’s disease, brain tumor classification, and breast cancer detection.

2.3 Quantum kernel and hybrid approaches

Quantum kernel methods (QSVMs) have shown potential benefits on datasets where quantum feature maps yield kernels not efficiently approximable classically; experimental photonic kernel implementations and QSVM benchmarks indicate scenarios with empirical gains, albeit primarily on small-to-moderate dataset sizes. Hybrid architectures quantum layers interfaced with classical CNN backbones are increasingly popular given NISQ limitations.

2.4 Open problems

Open areas include scaling QNNs to the high dimensionality of whole-slide images, robust encoding strategies, dealing with noise and decoherence, avoidance and mitigation of barren plateaus (flat gradients), and generating reproducible benchmarks for healthcare tasks. Several recent works address rotation-invariant VQCs and resource-efficient QCNN architectures, offering practical design directions.

3. Quantum Neural Network Families: Definitions, Intuition and Architectures

This section defines the principal families of QNNs used for image classification and provides architecture blueprints.

3.1 Variational Quantum Circuits (VQCs) and Parameterized Quantum Circuits (PQCs)

Definition & components. A VQC is a parameterized quantum circuit $U(\boldsymbol{\theta})$acting on an initial state (commonly $|0\rangle^{\otimes n}$) and producing measurement outcomes to construct model outputs. Training updates parameters $\boldsymbol{\theta}$to minimize a loss $L(\boldsymbol{\theta})$computed from measurement statistics. VQCs are the quantum analogue of neural networks: gates with tunable parameters play the role of weights; measurements produce outputs used in classical post-processing [variational ansatz literature].

Architecture example. A typical image classifier uses: data encoding circuit $U_{\phi}(x)$(maps classical image patch $x$to quantum state), variational layers $U(\boldsymbol{\theta})$(entangling and single qubit rotations), and measurement of qubits to produce features z. The classical head (e.g., softmax layer) maps $z$to class probabilities.

Key design choices:

- **Encoding / feature map:** amplitude encoding, angle (rotation) encoding, basis encoding. Each offers different resource tradeoffs (qubits vs. gates). For images, amplitude encoding is compact but requires complex state preparation; angle encoding is simpler but may require more qubits for high dimension.
- **Ansatz depth and entanglement structure:** shallow hardware-efficient ansatz vs. problem-inspired structured ansatz (e.g., convolutional locality). Deeper circuits can express more complex functions but increase noise sensitivity and training instability.

3.2 Quantum Convolutional Neural Networks (QCNNs)

Motivation & structure. QCNNs adapt classical CNN ideas (local receptive fields, pooling) to quantum circuits. Notably, QCNNs exploit local unitary blocks and pooling via measurements and conditional operations, achieving parameter efficiency (O(log N) parameters for N qubits in certain constructions) and provable robustness in some tasks. QCNNs have been proposed and applied to quantum phase recognition and have been adapted for classical image tasks via hybrid encodings. ResearchGate+1

QCNN pattern for images.

1. **Patch encoding:** split high-dimensional images into patches; encode each patch into a small quantum register.

2. **Local convolutional unitaries:** apply parameterized local blocks across patches.
3. **Pooling / downsampling:** use measurements and qubit resets or swap networks to reduce dimensionality.
4. **Global variational layers and measurement:** produce classification logits.

QCNNs align naturally with patch-based processing for whole-slide images and can be distributed across multiple small quantum processors via circuit splitting techniques. Recent distributed hybrid QCNN proposals show how to reconstruct larger circuits using fewer qubits via partitioning. arXiv+1

3.3 Quanvolutional Layers and Quantum Filters

Quanvolutional layers (Quanv) embed quantum circuits as convolutional filters: for each patch, a small quantum circuit transforms the patch into features that feed into a classical CNN. Benefits include nonlinear and high-dimensional transforms with potentially compact parameterization. Empirical studies demonstrate benefits particularly in limited data regimes. PMC+1

3.4 Quantum Kernel Methods and QSVMs

Quantum kernel methods use circuits to implicitly map data into quantum Hilbert space and estimate kernel values through state overlap measurements. A classical SVM or kernel ridge regression uses these kernel matrices for classification. Photonic and circuit-based kernel experiments show promising performance in select tasks. For medical imaging, kernel methods may be especially useful for small-label datasets or transfer learning. Nature+1

3.5 Hybrid classical–quantum pipelines

Given NISQ restrictions, hybrid models where classical CNN stages perform heavy downsampling and quantum layers process compressed, information-rich embeddings are practical near-term strategies. Such architectures can reuse classical pretrained backbones and augment them with a quantum classifier or quantum feature extractor. Many recent medical imaging QNN works use hybrid pipelines to achieve competitive performance on moderate datasets.

4. Mathematical Formulation

We present a formal mathematical description of QNNs for image classification, covering encoding, circuit action, measurement, loss, and gradients.

4.1 Notation and problem setup

Let $x \in \mathbb{R}^D$denote a flattened image (or patch descriptor) and $y \in \{1, \ldots, C\}$the class label. Our QNN model is a composition:

$$\hat{y} = h_{\text{cls}}(f_{\text{q}}(U_{\phi}(x); \boldsymbol{\theta})),$$

where $U_\phi(x)$is a parameterized data-encoding unitary (feature map), $f_{\mathrm{q}}(\cdot;\boldsymbol{\theta})$denotes measurement outcomes from the variational circuit $U(\boldsymbol{\theta})$, and $h_{\mathrm{cls}}$is a classical classification head (logistic or softmax).

4.2 Data encoding maps

Encoding classical data $x$to quantum states uses a map

$$\Phi: x \mapsto \mid \psi(x)\rangle = U_\phi(x) \mid 0\rangle^{\otimes n},$$

where $U_\phi(x)$can be an amplitude encoder $U_{\mathrm{amp}}(x)$, angle encoder (rotations $R_Y(x_i)$), or other structure-aware encoders (tensor product feature maps [Havlíček et al., 2019]). Amplitude encoding packs Dfeatures into amplitudes, using $n = \lceil \log_2 D \rceil$qubits, but the preparation circuit can be costly. Angle encoding uses more qubits but simpler gates.

4.3 Variational ansatz and measurement

A layered ansatz:

$$U(\boldsymbol{\theta}) = \prod_{\ell=1}^{L} \left(U_{\mathrm{ent}}^{(\ell)} U_{\mathrm{rot}}^{(\ell)}(\boldsymbol{\theta}^{(\ell)})\right)$$

where $U_{\mathrm{rot}}^{(\ell)}$are parameterized single-qubit rotations and $U_{\mathrm{ent}}^{(\ell)}$are entangling gates (CNOTs, CZs) with a prescribed connectivity. Measurements on selected observables $\{O_j\}$yield expectation values:

$$z_j(x;\boldsymbol{\theta}) = \langle 0 \mid U_\phi(x)^\dagger U(\boldsymbol{\theta})^\dagger O_j U(\boldsymbol{\theta}) U_\phi(x) \mid 0\rangle.$$

The vector $z$forms classical features for downstream classification.

4.4 Loss function and optimization

For classification with one-hot labels y, we define cross-entropy loss:

$$L(\boldsymbol{\theta}) = -\frac{1}{N}\sum_{i=1}^{N} \sum_{c=1}^{C} y_{i,c} \log \sigma_c(h_{\mathrm{cls}}(z(x_i;\boldsymbol{\theta}))),$$

where $\sigma_c$is softmax. Optimization uses gradient-based methods (classical optimizers: Adam, RMSProp) and gradients computed via the parameter-shift rule or stochastic parameter-shift estimators for circuits with specific gates. For observable $O$and parameter $\theta_k$entering as $R_\alpha(\theta_k)$, the parameter-shift rule yields

$$\frac{\partial z}{\partial \theta_k} = \frac{1}{2}[z(\theta_k + \frac{\pi}{2}) - z(\theta_k - \frac{\pi}{2})].$$

Stochastic measurement noise and finite shot counts introduce gradient noise that must be controlled (shot budgeting strategies).

4.5 Quantum kernel view

Quantum kernel for inputs $x, x'$is defined by:

$$K(x, x') = | \langle \psi(x) \mid \psi(x') \rangle |^2,$$

or fidelity-based kernels. Kernel estimation via quantum circuits uses swap tests or fidelity estimation; classical training uses SVMs with the kernel matrix K. Kernel methods avoid explicit parameter training of VQCs but have computability and scalability constraints for large datasets; recent photonic kernel results show empirical advantages in select cases.

5. Expressivity, Trainability and Theoretical Limits

5.1 Expressivity and universality

Parameterized quantum circuits can approximate arbitrary unitary transformations given sufficient depth and gate sets analogous to universal approximation in classical nets. The expressivity of VQCs depends on data encoding and entangling structure; certain feature maps can produce highly expressive embeddings relevant for classification tasks. However, expressivity alone does not guarantee trainability. arXiv

5.2 Barren plateaus and gradient scaling

Barren plateaus exponentially vanishing gradients for certain random or deep ansätze pose a major obstacle to training. Recent analyses show that circuit depth, global cost functions, and random initialization can produce barren plateaus; remedies include local cost functions, problem-inspired ansatz, layerwise training, parameter initialization strategies, and noise-aware training. For medical imaging applications, careful ansatz design (localized, convolutional) and hybrid pretraining are strongly advised to mitigate barren plateaus.

5.3 Noise, decoherence, and generalization

NISQ device noise both limits circuit depth and can act as implicit regularization. Understanding the interplay of noise, generalization, and adversarial robustness is an active research area. For clinical tasks requiring deterministic reliability, noise mitigation (error mitigation techniques, readout correction) and hybrid strategies (offloading heavy compute to classical backbones) are practical near-term solutions.

6. Design Patterns for High-Dimensional Medical Images

6.1 Patch-based hybrid pipelines (recommended near-term)

1. **Preprocessing & patching:** break gigapixel images into context-aware patches (sliding or tissue-aware sampling).
2. **Classical backbone:** use pretrained CNN (ResNet/ViT) to generate compact embeddings for each patch.
3. **Quantum feature module:** map embeddings $e \in \mathbb{R}^d$into quantum states using angle/encoded feature maps, apply a small VQC or QCNN to produce discriminative features.
4. **Aggregation & classification:** aggregate quantum features across patches (attention, pooling), apply classical classifier.

This pattern balances resource constraints and leverages quantum layers where they are most useful: compact, information rich embeddings. Several contemporary proposals and experiments adopt this hybrid model.

6.2 Distributed QCNNs for larger contexts

Circuit splitting and distributed quantum processing allow reconstructing large QCNNs from smaller partitions, enabling processing larger receptive fields with fewer qubits. Recent distributed hybrid QCNN work shows superior performance with fewer parameters using circuit splitting and careful orchestration across quantum nodes. This approach is promising for hospital clusters with cloud quantum access.

6.3 Rotation-equivariant and symmetry-aware circuits

Medical images often exhibit rotational and translational symmetries. Equivariant quantum architectures exploit these symmetries to improve sample efficiency and generalization, as recently demonstrated by rotation-invariant VQCs and equivariant QCNNs. Embedding symmetry into the ansatz reduces parameter count and training complexity.

7. Evaluation Protocols and Benchmarks

This section prescribes rigorous experimental protocols to assess QNNs against classical baselines.

7.1 Datasets and tasks

- **Public medical imaging datasets:** e.g., ISIC (skin lesion), BraTS (brain tumor MRI), LIDC-IDRI (lung nodules), CAMELYON (lymph node histopathology), TCGA whole-slide image subsets. Use standardized train/validation/test splits and patient-level splits to avoid data leakage.

- **Synthetic and reduced datasets:** to evaluate scaling behavior and small-data advantage hypotheses.

7.2 Baselines and metrics

- **Baselines:** classical CNNs (ResNet, EfficientNet), classical kernel methods, and hybrid classical transfer learning.
- **Metrics:** accuracy, AUC, sensitivity/specificity (critical in clinical contexts), F1 score, calibration error, inference latency, parameter counts, and resource utilization (shots, qubits, gate depth). For fair comparisons, include computational cost normalized metrics (energy/time per inference).

7.3 Reproducibility and shot budgeting

Document random seeds, hardware backends, measurement shot counts, and error-mitigation techniques. Provide code and circuit definitions (QASM/OpenQASM/Quantum intermediate representation) to support replication.

7.4 Statistical testing

Use bootstrap confidence intervals, McNemar’s tests for paired comparisons, and proper multiple-comparison corrections. For rare disease classes, use stratified sampling and appropriate class-imbalance handling (weighted losses, focal loss).

8. Implementation Blueprints and Practical Considerations

8.1 Hardware, simulators, and cloud access

- **Simulators:** use parameterized quantum simulators (Qiskit Aer, PennyLane, Cirq) for large-scale experiments (resource aware) but note classical simulators scale exponentially, limiting high-dimensional exploration.
- **Cloud quantum backends:** access to superconducting (IBM, Rigetti), trapped-ion (IonQ), or photonic (Xanadu) platforms. Photonic kernels recently showed empirical advantages in certain classification tasks. Consider hardware noise profiles and gate fidelities when designing circuits.

8.2 Software toolchains

Leverage high-level frameworks: PennyLane (for hybrid gradient flows), Qiskit Machine Learning, TensorFlow Quantum (TFQ), and Cirq integrations for kernel estimation. Reproducible workflows should couple classical ML toolchains with quantum simulation and hardware calls.

8.3 Privacy, data governance and federated strategies

Medical data governance requires HIPAA/GDPR compliance and secure data handling. For distributed hospital networks, federated hybrid training (federated classical pretraining, followed by private quantum classifier updating with secure aggregation) is a promising direction. Privacy-preserving quantum protocols and secure multiparty quantum computation are research frontiers but not yet industrially mature. Ensure that quantum cloud providers' data handling meets institutional policies.

9. Case Studies Representative Experiments (Blueprints)

This section outlines three experimental blueprints for empirical evaluation.

9.1 Small-sample brain tumor classification (BraTS subset)

- **Setup:** classical backbone (ResNet18) extracts patch embeddings; a 6-qubit VQC processes embeddings (angle encoding), hybrid training with Adam; shots per measurement: 1024.
- **Hypothesis:** QNN improves AUC over classical classifier in the small-sample regime (N < 1000) due to richer feature mappings.
- **Evaluation:** 5-fold cross validation, compare with classical MLP head matched for trainable parameters.

9.2 Histopathology whole-slide patch classification (CAMELYON)

- **Setup:** patch sampling + pretrained EfficientNet for embedding, quanvolutional layers for early transform, QCNN for patch-level classification. Circuit splitting across 2 devices for larger receptive fields.
- **Hypothesis:** QCNN yields better robustness to staining variation and small tumor foci than classical baseline. arXiv+1

9.3 Photonic kernel based lesion classification

- **Setup:** evaluate quantum kernel SVM on small lesion dataset with photonic implementation for kernel estimation. Compare to Gaussian kernel SVM and neural tangent kernels.
- **Hypothesis:** photonic kernel yields higher separability for specific texture features encountered in dermatological images. Nature

For each blueprint, report compute budgets, quantum resource usage, hyperparameter grids, and clinical relevance metrics.

10. Regulatory, Ethical and Clinical Integration Considerations

10.1 Clinical risk and safety

Any QNN-enabled diagnostic system must undergo standard clinical validation retrospective evaluation, prospective clinical trials, and safety assurance for false positives/negatives. Regulatory bodies (FDA, EMA) treat AI/ML as medical devices requiring validated performance metrics and update plans. Quantum components add complexity in reproducibility and vendor variability; strong documentation and device-aware validations are mandatory.

10.2 Explainability and interpretability

Explainability is crucial in clinical contexts. Quantum models are less interpretable by default; countermeasures include:

- Using quantum feature attribution methods (perturbation-based, fidelity-based saliency) and integrating with classical explainers (Grad-CAM on classical backbone).
- Producing human-readable rationales and uncertainty estimates to aid clinician decision support.

10.3 Data protection and cloud governance

Quantum cloud services may require sending embeddings or circuit instructions to third parties. Ensure data encryption, contractual safeguards, and minimal leakage (e.g., only send aggregated features).

11. Discussion: Opportunities, Limitations and Roadmap

11.1 Opportunities

- **Expressive feature maps:** QNNs create new embeddings that may separate complex high-dimensional manifolds more effectively in limited-label regimes.
- **Parameter economy:** QCNN and certain VQCs can achieve compact parameterization, benefiting generalization for small datasets.
- **Niche advantages:** photonic kernels and specialized ansatz show empirical superiority in curated tasks.

11.2 Limitations

- **Scalability:** state-preparation and simulation limits restrict experiments to compressed embeddings or patches.
- **Noise sensitivity:** current quantum hardware fidelity constrains depth and entanglement.
- **Reproducibility and standardization:** lack of unified clinical benchmarks for QNNs in imaging.

11.3 Roadmap for translational research

1. **Short-term (1–2 years):** hybrid pilots on compressed embeddings, reproducibility best practices, federated prototyping.
2. **Medium-term (3–5 years):** distributed QCNNs, device-aware pipelines, larger clinical datasets, robust explainability tools.

3. **Long-term (5+ years):** hardware gains enabling larger quantum backbones and end-to-end quantum accelerations for parts of the imaging pipeline.

12. Conclusion

QNNs hold promise for high-dimensional medical image classification through novel feature maps, compact parameterizations, and hybrid architectures that align with NISQ constraints. While current evidence is preliminary and largely small-scale, an informed program of theoretical analysis, benchmarked empirical studies, and careful translational engineering can clarify where QNNs meaningfully improve clinical imaging tasks. Integrating quantum layers into classical backbones, focusing on small-data or feature-hard tasks, and rigorously controlling for hardware and statistical issues are practical near-term strategies to advance the field.

References

1. Fatunmbi, T. O. (2022). Leveraging robotics, artificial intelligence, and machine learning for enhanced disease diagnosis and treatment: Advanced integrative approaches for precision medicine. *World Journal of Advanced Engineering Technology and Sciences*, 6(2), 121–135. https://doi.org/10.30574/wjaets.2022.6.2.0057.

2. Fatunmbi, T. O. (2024). Predicting precision-based treatment plans using artificial intelligence and machine learning in complex medical scenarios. World Journal of Advanced Engineering Technology and Sciences, 13(1), 1069–1088. https://doi.org/10.30574/wjaets.2024.13.1.0438

3. Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567, 209–212. (Introduces quantum feature maps and QSVM frameworks).

4. Matic, A., Monnet, M., Lorenz, J. M., Schachtner, B., Messerer, T. (2022). Quantum-classical convolutional neural networks in radiological image classification. arXiv:2204.12390. (Hybrid QCCNN applied to CT radiology tasks).

5. Breast Cancer Detection with Quanvolutional Neural Networks. (2024). *Journal / PMC article on QCNN for breast ultrasound.*

6. Zhou, M. G., et al. (2023). Quantum Neural Network for Quantum Neural Computing. *Research Article.* Discusses QNN models and training paradigms.

7. Griol-Barres, I., et al. (2021). Variational Quantum Circuits for Machine Learning. *Applied Sciences*, 11(14):6427. (Survey and implementation notes for VQCs).

8. Chen, Y., et al. (2024). A novel image classification framework based on hybrid variational quantum algorithms. *Quantum Information Processing* (Springer). (Hybrid approaches for images).

9. Research article (2024): Variational Shadow Quantum Circuits Assisted QCNNs - Qute Journal. (VSQC-QCNN design).
10. ResearchGate / Benchmark studies: Quantum kernel methods under scrutiny benchmarking large sets of kernels and datasets (2025–2026 analyses). (cautionary results on generalization of kernel advantages).
11. Patel, et al. (2024). Brain tumor diagnosis using Hybrid QCNN (arXiv:2401.15804). (QCNN architecture and high accuracy claims on brain tumor data).
12. Albrecht, M. et al. (2018-2024). LWE and parameter estimation literature (for mathematical analogy in quantum feature spaces and kernel complexity assessments). (Background on hardness assumptions relevant to quantum algorithms).