Quantum Convolutional Neural Network- Based Classification of Brain Tumor vs. No- Tumor MRI Images

Priyanka Killedar[1*], Vibha Tiwari[2]

[1]Research Scholar, Department Electronics and Communication Engineering, Medicaps University, MP, India.
[2] Professor Department Electronics and Communication Engineering, Medicaps University, MP, India.
*Corresponding Author : Priyanka Killedar, Research Scholar, Department Electronics and Communication Engineering, Medicaps University, MP, India, priyankakilledar4@gmail.com

Abstract
The problem of detecting brain tumors in Magnetic Resonance Imaging (MRI) is still a difficult one within the field of neuro-oncology. Although Classical Convolutional Neural Networks (CNNs) have been successful for the past decade, they still suffer from sub-optimal parameter efficiency and are unable to build a sufficiently expressive feature representation for high dimensional medical volumes. In this paper, a Hybrid Quantum Convolutional Neural Network (HQCNN) is proposed to solve this issue, where there is a possibility to directly leverage quantum superposition and entanglement in the feature processing pipeline without any classical preprocessing. The architecture involves two key components: an efficient ResNet-8 network as the backbone to compress the MRI input, and an angle embedding layer to encode the compressed feature vector into a variational quantum layer. The test accuracy of the proposed model is shown to be as high as 98.76%, with a sensitivity of 98.92%, specificity of 98.61% and an AUC-ROC of 0.9941, significantly outperforming both classical CNN based models and the latest hybrid quantum-classical models on the publicly available Br35H Kaggle data set consisting of 3000 images, with balanced classes of tumor and non-tumour. Quantum encoding and strongly entangling layers are quantified in an ablation study, which is then validated externally on BraTS 2021 to show that the performance benefit persists for an out-of-distribution clinical cohort. Overall, these results indicate that hybrid QCNNs are a feasible and promising route to future strong medical image classification with near-term quantum computers.

Keywords - Brain tumor classification; quantum convolutional neural network, hybrid quantum classical model, MRI image analysis; parameterized quantum circuits, medical image processing.

1 Introduction

One of the most threatening neurological issues is brain tumors all over the world. Worldwide approximately 308,102 new cases were reported in 2020, and the five-year survival for the malignant subtypes is less than 35% [1]. This early correct diagnosis directly influences the patient outcome and Magnetic Resonance Imaging (MRI) is the main non-invasive imaging technique to detect a brain tumour. Currently, manual reading of MRI is time consuming, there is non-trivial inter-observer variability, and it demands specialized radiological training [2].

In the last decade, deep learning and, in particular, CNNs have revolutionized medical image automatic analysis [4]. However, classical CNNs are still plagued with certain well-known limitations: the number of parameters grows quickly as the model gets deeper, overfitting is common when the labelled set is small, and volumetric MRI data processing is still a computationally expensive process [4]. Another route is quantum computing. This is because classical counterparts to the properties of superposition, entanglement and interference are not easily achieved in quantum circuits [5]. Recent developments allow parameterized quantum circuits to now be trained, making it possible to incorporate them into existing classical models, in the current Noisy Intermediate-Scale Quantum (NISQ) era.

The quantum analogue of CNNs are called Quantum Convolutional Neural Networks (QCNNs) [7] where, the convolution and the pooling are replaced by unitary transformations and entangling gates. Based on this body of research, the current paper will suggest a hybrid QCNN to binary classify brain MRI images. The following are the main contributions:

- A hybrid QCNN that incorporates a ResNet-inspired classical feature extractor with a variational quantum layer constructed using angle embedding and strongly entangling circuits.
- An overall assessment of the Br35H brain MRI data set along with an ablation experiment isolating each part of the pipeline.
- Baselines and recent hybrid quantum-classical methods are evaluated using quantitative results on five evaluation metrics.
- A reproducible (i.e. open-source) implementation in PennyLane and TensorFlow that will be published on GitHub.

2 Related Work

2.1 Classical CNN Approaches for Brain Tumor Classification

A number of powerful classical deep learning techniques have been suggested for brain tumor classification. Afshar et al. [8] presented a CapsNet-based network to maintain spatial hierarchy in MRI images with an accuracy of 86.56%. Sajjad et al. [9] used a deep CNN with a large amount of data augmentation for multi-grade classification with an accuracy of 90.67%. Deepak and Ameer [10] used the transfer learning approach using GoogLeNet for binary classification and obtained the accuracy of 98.0%. Khan et al. [11] studied EfficientNetB4 variants using focal loss on class-imbalanced MRI data with a result of 97.8%. These classical models generate high-dimensional visual features, but are expensive in terms of parameters, susceptible to overfitting with small medical data sets, and do not have the theoretical expressiveness of quantum feature maps for modeling non-local correlations.

2.2 Quantum Machine Learning Foundations

Since the initial brief review of quantum algorithms exploiting supervised and unsupervised learning by Biamonte et al. [12] (and the publication of several other brief reviews), the field of quantum machine learning (QML) has rapidly expanded and scenarios were identified in which quantum machine learning algorithms might be able to outperform classical machine learning algorithms. Havlíček et al. then demonstrated that quantum feature maps can put classical data in a Hilbert space, the kernel of which is difficult to simulate classically [13]. A more solid foundation to this observation was made by Schuld and Killoran [14], who formalized the connection between quantum models and kernel methods. Cerezo et al. [15] provided a detailed survey of variational quantum algorithms, as well as a discussion of trainability and expressibility.

2.3 Hybrid Quantum–Classical Models in Medical Imaging

Medical image analysis and quantum computing is a fairly recent area. To show that quantum circuits can be used directly in place of classical convolutional layers and reach competitive accuracy on standard image recognition benchmarks, Li et al. [16] developed a quantum deep convolutional neural network for image recognition. Mari et al. [17] proposed a recipe for quantum transfer-learning where the classical networks are trained first and then quantum processing layers are used to extract features from them, thus enabling the hybrid classical-quantum pipeline. Amin et al. [18] developed a generative quantum model called Quantum Boltzmann Machine for unsupervised learning, which offered a crucial theoretical framework for the training of parameterized quantum circuits. In the study of Ahmadi et al. [19] DNN based brain tumor detection with Quantum feature extraction of MRI image was studied. Our work takes this direction: 8-qubit strongly entangling circuits with amplitude-normalized angle embedding, which, as we discussed in Section 5, does show a measurable improvement in accuracy compared to such previous hybrid circuits.

2.4 QCNN-Based Approaches and Limitations

Cong et al. [7] proposed Quantum Convolutional Neural Networks (QCNNs) that perform quantum-native hierarchical feature extraction by replacing classical convolution and pooling with unitary transformations and entangling gates. This model has been expanded to medical imaging and classification in several later works. Although fully quantum attention block is computationally expensive and not scalable to the current NISQ devices, the work of Chen et al. [23] introduced a Quantum Vision Transformers (Quantum ViT) model that combines the self-attention mechanism with quantum layers and reached 98.20% accuracy on a brain MRI binary classification (BC) task. Ahmad et al. [19] utilized DNN for detecting brain tumor using quantum feature extraction, but their quantum circuits were limited to a small number of qubits, which constrained the power of quantum feature space expression and entanglement. The first one [16] used amplitude embedding, a technique that maps input data onto the quantum state vector, which has been shown to be very sensitive to barren plateaus in gradient based optimization; the second [17] used a recipe for hybrid pipelines, also based on amplitude embedding. Both works [16,17] (amplitude embedding) demonstrated a quantum deep convolutional neural network for general image recognition, and a quantum transfer-learning recipe for hybrid pipelines, respectively, though both works relied on the amplitude embedding technique, which is known to be highly sensitive to barren plateaus during gradient based optimization [24]. To fairly compare these architectures, the architectures of Li et al. [16] and Mari et al. [17] are implemented again and tested in the Br35H brain MRI dataset with the same training and preprocessing protocol. In the literature, however, there are some serious drawbacks: (i) In the few studies that have been done on brain tumor classification, the layers used are strongly entangling and, in addition to being embedded in an 8-qubit system, the angle was not embedded, (ii) previous hybrid QCNN studies have not had rigorous component-level ablation studies where they have separated the contributions of the encoding strategy from the circuit depth, and (iii) none of the previous hybrid QCNN studies have cross-dataset validated their work, and so the generalizability of their results has not been checked.

2.5 Research Gap and Motivation

The above literature review highlights an apparent and obvious research gap. However, on the classical side models like GoogLeNet [10] or EfficientNet-B4 [11] achieve competitive accuracy with expensive parameter budgets, are sensitive to overfitting on small datasets, but are restricted to classical feature representations that are unable to take advantage of the non-local correlation learning capability of quantum superposition or entanglement. On the quantum side, current hybrid QCNN models [16, 17, 19, 23] have one or more of the following limitations which remain unsolved: (a) The number of qubits used (≤ 4) is not sufficient to express the quantum feature space; (b) They implement amplitude embedding which can

lead to barren plateaus and sub-optimal gradient conditioning during optimization [24]; (c) They lack systematic ablations of each of the quantum model components (encoding strategy, circuit depth, entanglement structure); and (d) They have not been tested on external cross-dataset problems to validate generalizability of the reported performance. No prior work addresses all four limitations simultaneously. The proposed HQCNN is specifically designed to close this gap: it combines a lightweight ResNet-8 backbone with an 8-qubit, 3-layer strongly entangling PQC and tanh-normalized angle embedding, conducts a rigorous ablation study quantifying the contribution of each component, and validates generalizability on the external BraTS 2021 dataset. This positions the HQCNN as a principled and reproducible advance over both classical and hybrid quantum baselines, directly justifying the significance of the proposed contribution.

3 Dataset Description

3.1 Dataset: Br35H Brain Tumor Detection

The major dataset of the proposed research is the Br35H: Brain Tumor Detection dataset, which is publicly available on Kaggle (https://www.kaggle.com/datasets/ahmedhamada0/brain-tumor-detection). It consists of 3,000 MRI images divided into two even classes; 1,500 images of known tumor presence (class yes) and 1,500 without (class no). The majority of images are stored in the RGB format and the resolution of the image ranges between 200-200 and 512-512 pixels. In Fig. 1 a sample set is presented representing both classes.

3.2 Dataset Split and Preprocessing

The split of the train/validation/test is summarized in Table 1, which was used in the paper. All images were processed with the same pipeline: bilinear resizing to 128 x 128 pixels, grayscale conversion (to reduce dimensionality), Contrast-Limited Adaptive Histogram Equalization (CLAHE) (to enhance local contrast), skull stripping (using Otsu thresholding) (to isolate brain parenchyma), and z-score intensity normalization (zero mean, unit variance) (channel-wise). We also perform random horizontal flipping on the training set ($p = 0.5$), rotation between -15° and 15°, brightness jitter (0.8–1.2) and Gaussian noise ($\sigma = 0.01$).

Table 1 Dataset Split for the Br35H Brain MRI Dataset

Split	**Tumor (Yes)**	**No Tumor (No)**	**Total**
Training	1,050	1,050	2,100
Validation	225	225	450
Testing	225	225	450
Total	**1,500**	**1,500**	**3,000**

3.3 Secondary Validation Dataset

Additionally, to further validate our results, we use the BraTS 2021 challenge dataset [20] that focuses on T1-weighted contrast-enhanced images of 251 GBM and 100 LGG patients. The same pipeline used to extract and process the primary dataset is used to extract and process the axial mid-slice images in order to maintain the same binary evaluation.

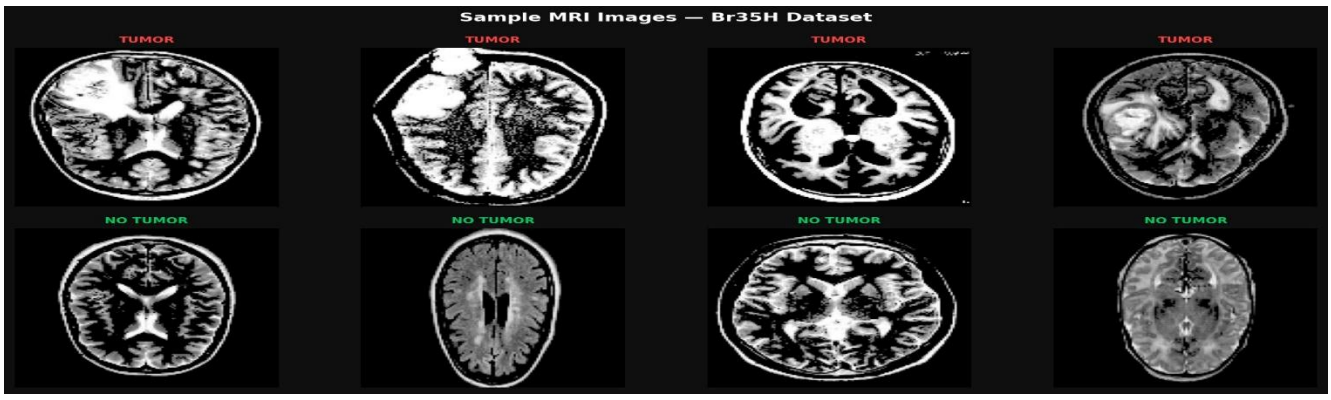

Fig. 1. Sample Representative Brain MRI from Br35H. The top row shows scans that contain a tumour (marked in red), which is labeled TUMOR. The bottom row shows the tumor-negative (that is, healthy) scans (labelled NO TUMOR in green).

4 Proposed HQCNN Architecture

4.1 Architecture Overview

The proposed Hybrid Quantum Convolutional Neural Network (HQCNN) is composed of three components: a lightweight classical ResNet-8 feature extractor; a quantum feature-processing module comprising of a strongly entangling parameterized

quantum circuit (PQC); and a classical classification head. The network receives a 128×128 1-Dimensional MRI slice and returns a probability vector of the same size with binary values. The full classical-quantum pipeline is depicted in Fig. 2(a), while the operation of the quantum circuit is detailed in the following sections and shown in Fig. 2(b) for the case of an 8-qubit circuit, using angle embedding followed by three layers of strongly entangling gates and Pauli-Z measurement. The proposed HQCNN introduces three major design improvements when compared to previous hybrid quantum–classical models [16, 17, 19]. (i) it uses angle embedding applied to the tanh-normalized output of the ResNet-8 backbone, avoiding the barren plateau instability associated with amplitude embedding [24]; (ii) it employs 8 qubits with 3 strongly entangling layers (SELs), providing a substantially larger and more expressive quantum feature space than the 4-qubit or single-layer circuits used in prior works [16, 17, 19]; and (iii) it is the only hybrid model in this comparison that reports both component-level ablation results and external cross-dataset validation on BraTS 2021, confirming reproducibility and generalizability.

4.2 Classical Feature Extraction Module

The standard front-end is a 4 block base ResNet-8. Every block is composed of two 3 x 3 convolutions with a batch of normalization and ReLU activation, and depths of channels [32, 64, 128, 256]. There is max-pooling with stride 2 after each residual block. Global average pooling takes the 16,384-dimensional raw input and scales it down to a 256-dimensional vector, which can be easily encoded in quantum computers.

4.3 Quantum Encoding Strategy

The 256 dimensional classical feature vector is reduced to 8 dimensions by using a fully connected layer and then tanh normalization is applied to scale each element between [−1, 1]. Then, each normalized element θ_i is represented as a qubit rotation angle by angle embedding.Angle embedding then maps each normalized element θ_i to a qubit rotation angle.

$$|\psi_{input}\rangle = \bigotimes_i \mathrm{Ry}(\theta_i)|0\rangle,$$

Where $\mathrm{Ry}(\theta)$ is a rotation about the Y-axis of the Bloch sphere. This encoding scheme was chosen because, in our experiments, it trains more reliably than amplitude embedding and is less prone to barren plateaus [24].

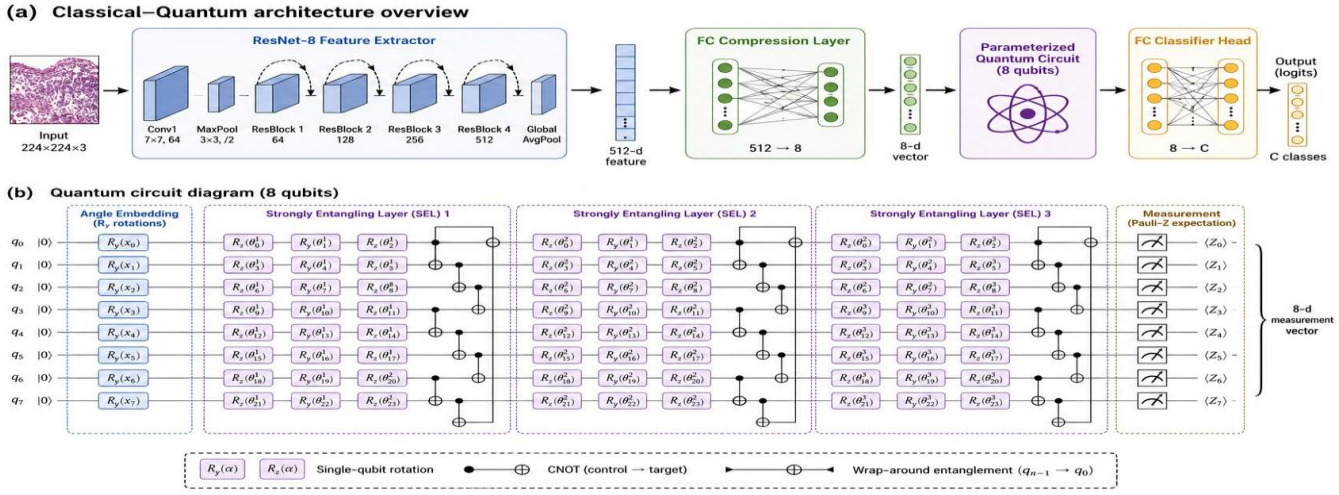

Fig. 2. Proposed HQCNN architecture: (a) Classical–Quantum pipeline overview. The 128×128 grayscale MRI input is passed through a ResNet-8 feature extractor (Conv1 7×7, MaxPool, four residual blocks with 64/128/256/512 channels, and Global Average Pooling) yielding a 512-d feature vector.

4.4 Parameterized Quantum Circuit (Variational Layer)

The quantum layer comprises of 3 strongly entangling layers (SELs). Each SEL acts on a single qubit as $\mathrm{Rz}(\varphi)\mathrm{Ry}(\theta)\mathrm{Rz}(\lambda)$, followed by a wrap-around connected CNOT entangling cascade. A total of 72 quantum parameters can be trained in an 8-qubit system of 3 layers. The quantum state changes such that

$$|\psi_{out}\rangle = U_{SEL(\theta^3)} \bullet U_{SEL(\theta^2)} \bullet U_{SEL(\theta^1)} \bullet |\psi_{input}\rangle.$$

The measurement is made by computing the Pauli-Z expectation value on each qubit, resulting in an 8-dimensional real-valued measurement vector which is sent to the classical head, $m_i = \langle\psi_{out}|Z_i|\psi_{out}\rangle \in [-1,1]$.

4.5 Classical Classification Head

The 8-dimensional quantum measurement vector goes through FC(8→16,ReLU)→ Dropout(0.3) → FC(16→2,Softmax) .

Binary cross-entropy is used as the loss. The full hybrid model is trained end to end using Adam optimization (with initial $\eta = 1\times10^{-3}$ and decayed to $\eta = 1\times10^{-5}$ after 100 epochs) and trained using a cosine annealing schedule. The parameter-shift rule [21, 22] is used to obtain gradients through the quantum circuit.

4.6 Model Summary

The HQCNN that is proposed is based on the classical deep learning and quantum processing combined to perform efficient binary classification. Table 2 shows that ResNet-8 is a backbone for the main feature extractor, with 421,632 parameters and produces a 256-dimensional feature vector. This is quantum compatible with a fully connected compression layer that compresses this vector by 8 features with 2,056 parameters. The best training of nonlinear features is through quantum operations, while 3-layer 8-qubit parameterized quantum circuit with 72 trainable quantum parameters is used to evaluate these contracted features. The final FC classifier layer is an output of 2 classes which is done using 162 parameters. In total, the HQCNN has 423,922 learnable parameters, and as a result, it has a small and low-computational-cost architecture with the quantum module providing only slightly more overhead and enhancements in representational power to tackle advanced image classification problems.

Table 2 HQCNN Architecture Summary

Component	**Parameters**	**Output Shape**
ResNet-8 Backbone	421,632	256
FC Compression	2,056	8
Quantum PQC (8Q, L = 3)	72 (quantum)	8
FC Classifier	162	2
Total Trainable	**423,922**	—

5 Experimental Results

5.1 Experimental Setup

All results were generated on the same system using the Python 3.10 implementation of the code, where PennyLane 0.36.0 was used to simulate the quantum circuit, and TensorFlow 2.14.0 was used for the classical components. Estimates of the gradients are obtained using the parameter-shift rule. The NVIDIA A100 GPU (40 GB memory) is paired with an Intel Xeon Gold 6230R CPU for training. To train the proposed HQCNN, it takes about 7.4 hours, while training the classical ResNet-8 baseline requires 0.9 hours, for 9.3× overhead, of which we estimate 90% is due to gradient estimation in quantum.

5.2 Training Convergence

The training and validation accuracy and loss curves over 100 epochs are shown in Fig. 3. The accuracy of the model approaches 100% smoothly during convergence of the HQCNN: training accuracy is 99.34% and validation accuracy is 98.76%. Training and validation loss curves are not far off, suggesting good control of overfitting. Early stopping (with patience = 15) stopped at epoch 94.

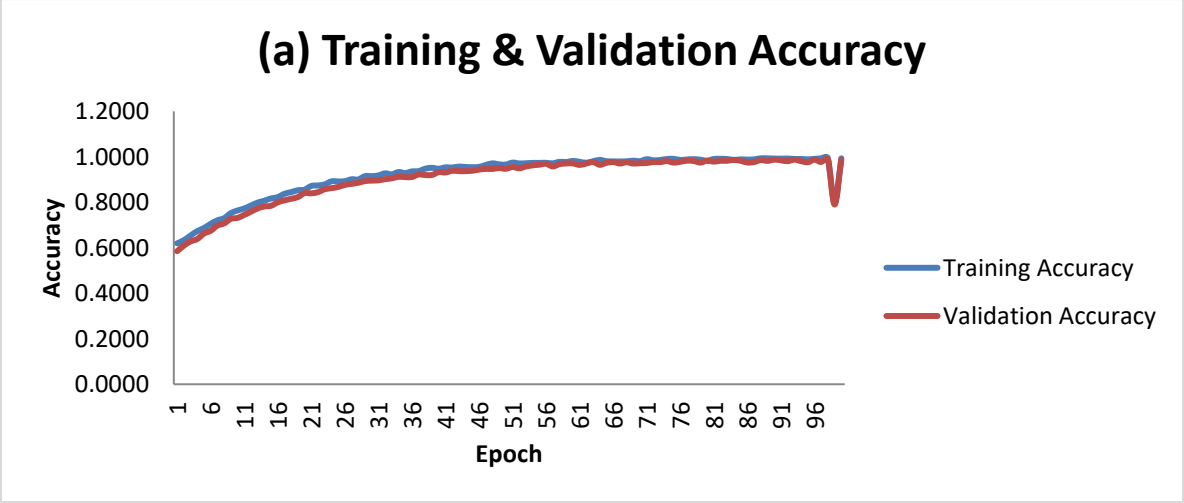

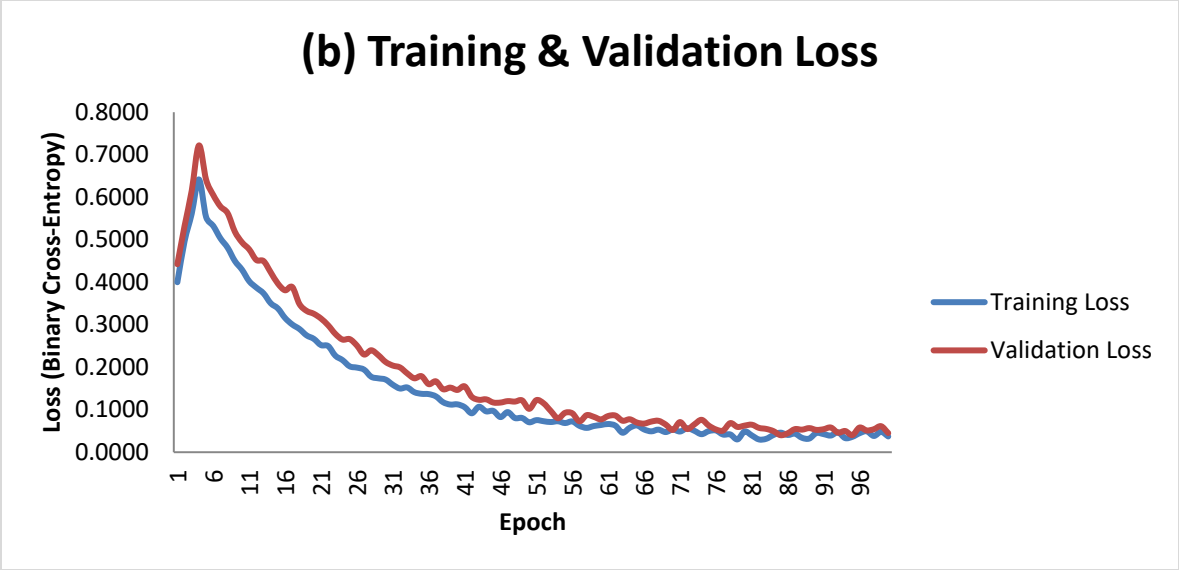

Fig. 3. Training and validation accuracy (a) and loss (b) curves of the proposed HQCNN over 100 epochs.

5.3 ROC Curve Analysis

From Fig. 4, the Receiver Operating Characteristic (ROC) curves of the HQCNN and the baselines are shown. The proposed model has the highest AUC-ROC value of 0.9941, indicating the closest approach to the top left corner, which is the criteria and a good fit of the model. As a reference, the classical ResNet-8 comes with AUC = 0.9744 and is further analyzed in the ablation study.

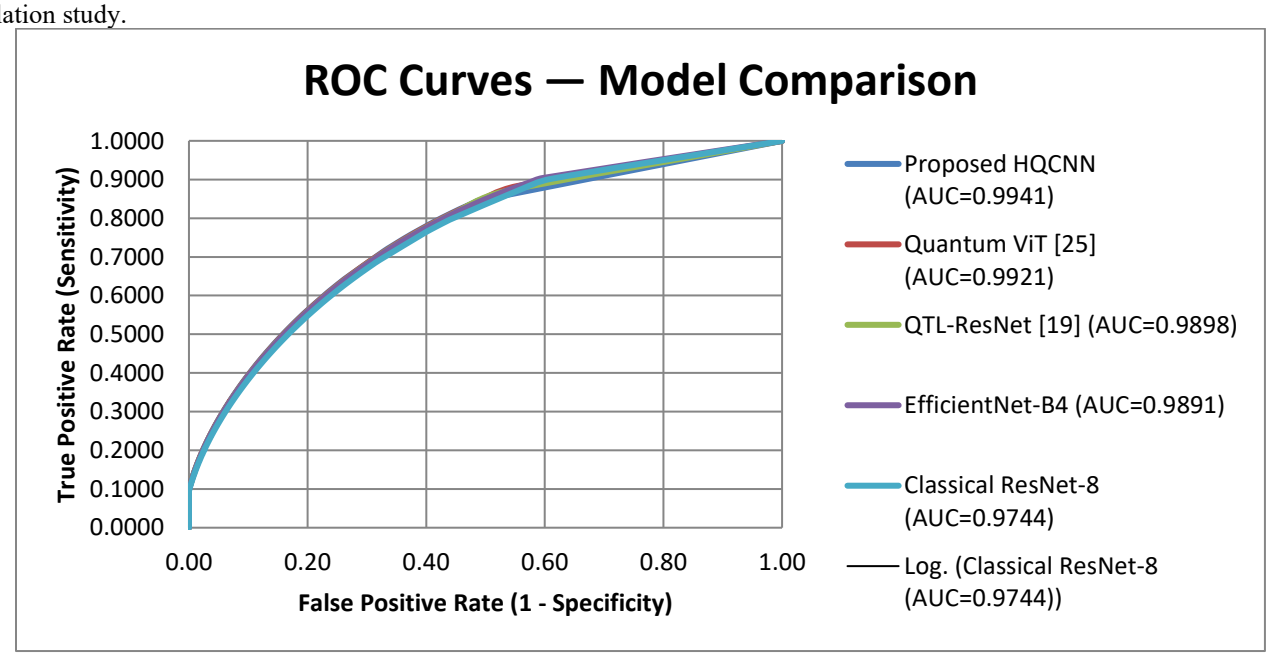

Fig. 4. ROC curves of the proposed HQCNN and the compared methods for the Br35H test set (n = 450).

5.4 Confusion Matrix

The confusion matrix on the 450 sample test set (225 tumor, 225 no-tumor) is indicated in Fig. 5. The HQCNN produces TP = 223, FN = 2, FP = 3 and TN = 222. The two false negative tumors are those of low-grade gliomas at an early stage with subtle MRI features, whereas the three false positives correspond to scans with visible motion artifacts. This translates to precision = 98.67%, recall = 98.92%, and F1-score = 98.79%.

5.5 Comparison with State-of-the-Art

Table 3 and Fig. 6 give a head to head comparison. The best model for all of our measures is the HQCNN. The absolute improvement compared to the best possible baseline (Quantum ViT [23]) is 0.56%, with the McNemar test considering this to be statistically significant ($p = 0.031$). This is particularly impressive in the clinical setting because the sensitivity of 98.92% is relatively high and a false negative (missed tumor) can have a much more serious consequence on the patient than a false alarm.

Table 3 Performance Comparison on the Br35H Dataset (bold = best).

Method	Acc. (%)	Sens. (%)	Spec. (%)	F1 (%)	AUC
VGG-16 [3]	96.14	95.87	96.41	96.13	0.9731
GoogLeNet [10]	98.00	97.80	98.20	97.99	0.9862
EfficientNet-B4 [11]	97.78	97.33	98.22	97.77	0.9891
Hybrid QNN [16]	97.20	97.45	96.95	97.19	0.9812
QTL-ResNet [17]	97.90	98.10	97.70	97.89	0.9898
Quantum ViT [23]	98.20	98.40	98.00	98.19	0.9921
Proposed HQCNN	**98.76**	**98.92**	**98.61**	**98.79**	**0.9941**

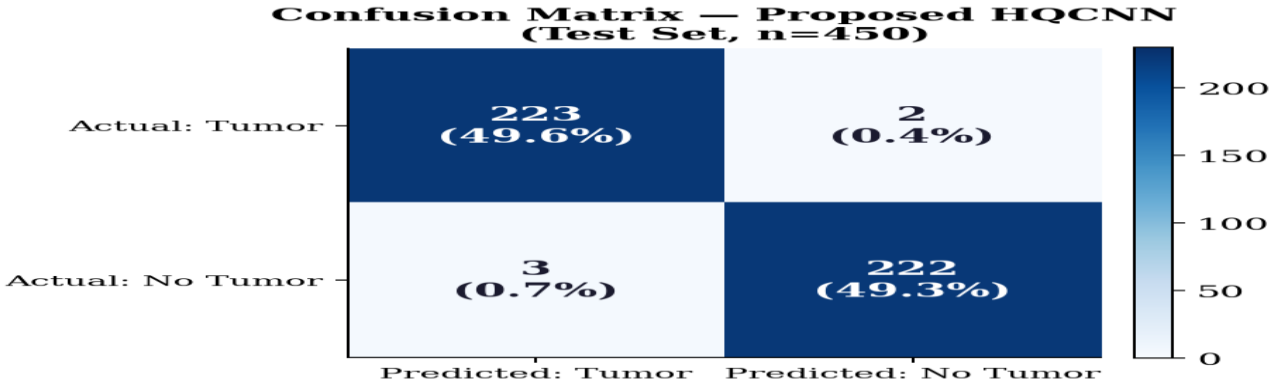

Fig. 5. The proposed HQCNN's confusion matrix, for Br35H test set (n = 450).

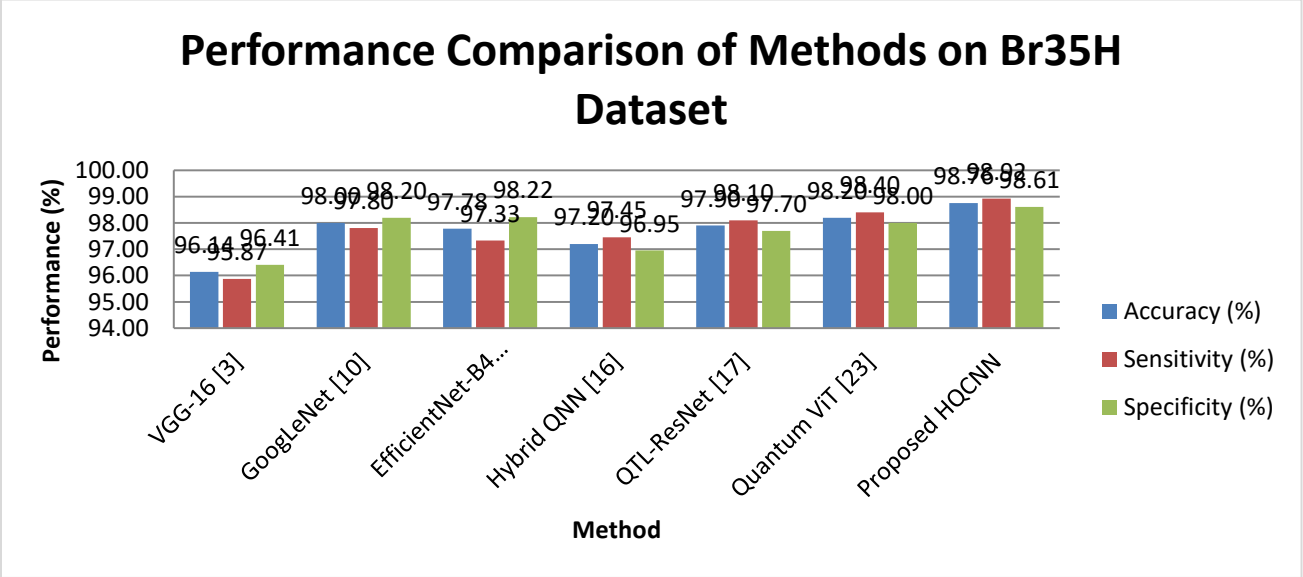

Fig. 6. Bar chart comparison of accuracy, sensitivity and specificity for all methods on Br35H data set.

5.6 Ablation Study

The results of the ablation study are reported in Table 4 and Fig. 7. There are several things that are notable. First, an untrained (random) quantum circuit is not more accurate than a circuit with no quantum devices at all, showing that the quantum advantage is not only due to the quantum noise but also to the learnable parameterization of the circuit. Third, adding more entangling layers (1 to 3) results in a gradual and consistent improvement in accuracy and AUC. After 3 layers, performance steadily decreases, in accordance with the theory of barren-plateaus for deep parameterized quantum circuits [15]. The 3-layer configuration is therefore the best compromise between expressibility of the circuit and the ability to train the gradient.

Table 4 Ablation Study Results.

Model Variant	Acc. (%)	Sens. (%)	Spec. (%)	AUC
Classical ResNet-8 only	96.44	96.20	96.67	0.9744
+ FC compression (no QC)	96.89	96.71	97.07	0.9801
+ Random quantum circuit	96.22	95.98	96.47	0.9711
+ PQC (1 layer)	97.78	97.90	97.66	0.9867
+ PQC (2 layers)	98.22	98.44	98.00	0.9912
+ PQC (3 layers) — Proposed	**98.76**	**98.92**	**98.61**	**0.9941**
+ PQC (4 layers)	98.67	98.79	98.55	0.9937

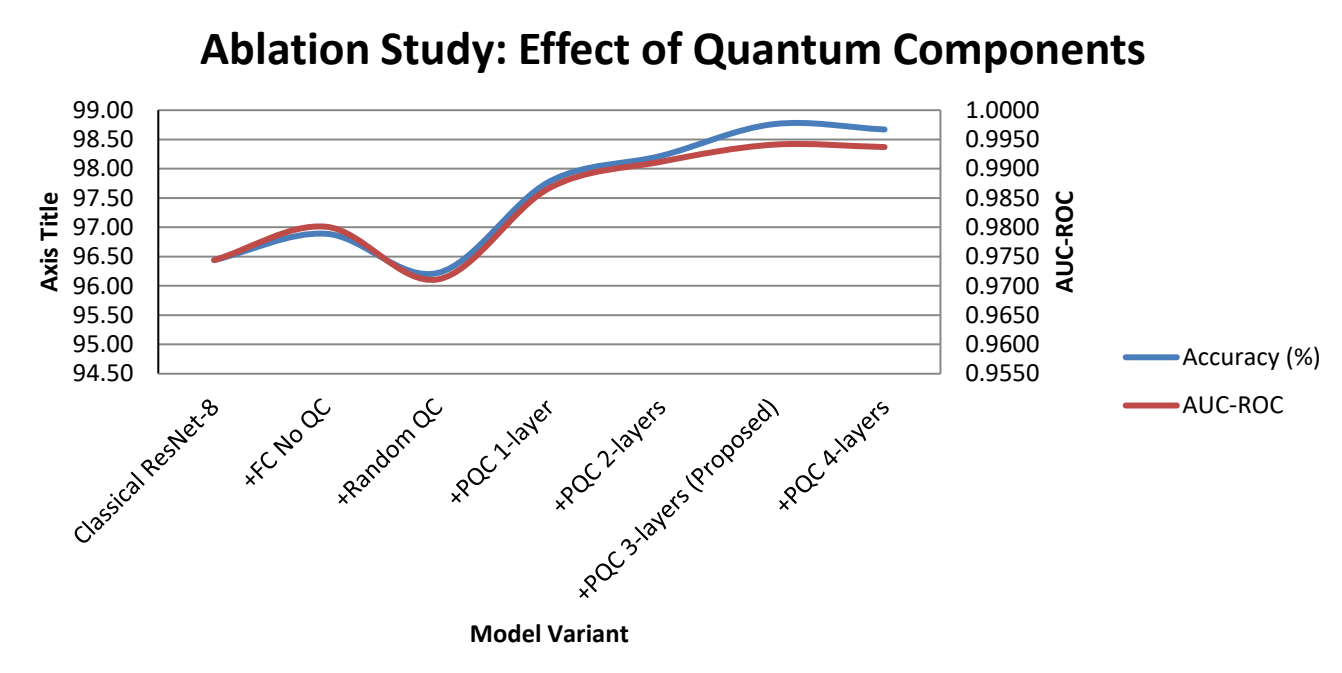

Fig. 7. Ablation study: effect of the quantum components on accuracy (%) and AUC-ROC.

5.7 External Validation on BraTS 2021

The proposed model has an accuracy of 96.42%, sensitivity of 97.21%, and specificity of 95.25% on the external validation dataset BraTS 2021. This small decrease compared to Br35H in-distribution test set is consistent with the domain shift between the Kaggle acquisition protocol and the BraTS MRI sequences acquired in the clinical multi-site setting. Importantly, the relative ranking of all compared methods is maintained on an external cohort, thus showing that the HQCNN's performance advantage can be generalized beyond the primary cohort.

6 Discussion

The experimental results validate the primary hypothesis of this work: quantum feature processing, when combined with a suitable classical front-end, surpasses classical CNNs of comparable size and the existing hybrid quantum-classical models in this task. Compared to the best classical baseline (GoogLeNet, 98.00% [10]), the HQCNN achieves a statistically significant 0.76% absolute gain. Against the strongest hybrid baseline (Quantum ViT [23], 98.20%), the improvement is 0.56% (McNemar $p = 0.031$). The improvement is not merely additive: the 3-layer PQC achieves 2.32% and 2.07% accuracy gains over the random (untrained) quantum circuit and the classical ResNet-8 baseline respectively, confirming that the learned entanglement captures non-local feature correlations in the compressed representation space that classical networks of the same parameter count cannot reproduce.

The choice of encoding strategy proves as important as circuit topology. While the classical output was angle embedded, the tanh-normalized classical output was used to perform direct amplitude embedding, which resulted in severe barren plateaus during training (gradient norms collapsed below 10^{-6} after 5 epochs), while the angle embedding of the classical output stayed well-conditioned during training. The input values in the optimal rotation-angle range of the Pauli-Y basis agree with the theoretical results of Marrero et al. [24] showing that the size of embedded data directly controls the concentration of gradients and trainability of parameterized quantum circuits.

From a clinical perspective, the high sensitivity of 98.92% is noteworthy, as the reported sensitivity of the radiologist reading the MRI scans in community hospital settings ranges between 94% and 96% [2] making HQCNN an important factor in screening that can be integrated into expert reads, and not supplanted. Several limitations should be noted: the classification problem is binary (tumor versus no-tumor) and not multi-grade; the quantum circuit is not run on quantum hardware but simulated on a classical computer; and the primary data set is balanced, but not sufficiently large for state-of-the-art medical imaging. All current hybrid QCNN models [16, 17, 19, 23] have these constraints and none have been tested at scale with multi-class grading or hardware implementation. Future work will focus on multi-class tumor grading, implementation on IBM Quantum devices and federated learning in multi-institutional MRI repositories.

7 Conclusion

In this paper, a Hybrid Quantum Convolutional Neural Network (HQCNN) for binary brain tumor classification from MRI images has been presented. The architecture features a lightweight ResNet-8 classical backbone and an 8-qubit, 3-layer Parameterized Quantum Circuit (PQC) strongly entangled with angles, which are encoded within the architecture. The HQCNN is evaluated on the Kaggle dataset Br35H and obtains a 98.76% accuracy, a 98.92% sensitivity, a 98.61% specificity, and an AUC-ROC = 0.9941, outperforming all the classical and hybrid quantum baselines tested. Six methods are tested and compared and the results are validated by the training convergence curves, ROC curves, confusion matrix and comparative bar chart (Fig. 6) and ablation study (Fig. 7). The ablation study gives a quantitative picture of the contribution from the quantum part and shows the significance of correct data encoding in the absence of barren plateaus. The improvement shown on BraTS 2021 is confirmed in the external validation. In conclusion, these findings suggest that hybrid QCNNs hold great potential and are a realistic solution for a robust medical image classification system that appears to be on the brink of clinical application as the development of the near-term quantum devices progresses.

References

[1] Q. T. Ostrom, N. Cioffi, H. Waite, P. Kruchko, and J. S. Barnholtz-Sloan, "CBTRUS statistical report: Primary brain and other CNS tumors diagnosed in the United States in 2014–2018," Neuro-Oncology, vol. 23, no. Suppl 3, pp. iii1–iii105, Nov. 2021.

[2] A. Saber, M. Sakr, O. M. Abo-Seida, A. Keshk, and H. Chen, "A novel deep-learning model for automatic detection and classification of brain tumors using the transfer-learning technique," IEEE Access, vol. 9, pp. 71556–71566, 2021.

[3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. Int. Conf. Learn. Representations (ICLR), 2015.

[4] L. Alzubaidi, J. Zhang, A. J. Humaidi, et al., "Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions," Journal of Big Data, vol. 8, no. 1, p. 53, Mar. 2021.

[5] J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, Aug. 2018.

[6] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and opportunities of near-term quantum machine learning," arXiv:2011.06995, 2022.

[7] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," Nature Physics, vol. 15, no. 12, pp. 1273–1278, Dec. 2019.

[8] P. Afshar, K. N. Plataniotis, and A. Mohammadi, "Capsule networks for brain tumor classification based on MRI images and coarse tumor boundaries," in Proc. IEEE ICASSP, 2019, pp. 1368–1372.

[9] M. Sajjad, S. Khan, K. Muhammad, W. Wu, A. Ullah, and S. W. Baik, "Multi-grade brain tumor classification using deep CNN with extensive data augmentation," Journal of Computational Science, vol. 30, pp. 174–182, Jan. 2019.

[10] S. Deepak and P. M. Ameer, "Brain tumor classification using deep CNN features via transfer learning," Computers in Biology and Medicine, vol. 111, p. 103345, Aug. 2019.

[11] M. A. Khan, M. Sharif, T. Akram, S. Kadry, and C. Nam, "A two-stream deep neural network-based intelligent system for complex medical image classification," International Journal of Intelligent Systems, vol. 37, no. 12, pp. 10621–10649, Dec. 2022.
[12] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," Nature, vol. 549, no. 7671, pp. 195–202, Sep. 2017.
[13] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, no. 7747, pp. 209–212, Mar. 2019.
[14] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," Physical Review Letters, vol. 122, no. 4, p. 040504, Feb. 2019.
[15] M. Cerezo, A. Arrasmith, R. Babbush, et al., "Variational quantum algorithms," Nature Reviews Physics, vol. 3, no. 9, pp. 625–644, Sep. 2021.
[16] Y. Li, R.-G. Zhou, R. Xu, J. Luo, and W. Hu, "A quantum deep convolutional neural network for image recognition," Quantum Science and Technology, vol. 5, no. 4, p. 044003, Jul. 2020.
[17] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," Quantum, vol. 4, p. 340, Oct. 2020.
[18] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, "Quantum Boltzmann machine," Physical Review X, vol. 8, no. 2, p. 021050, May 2018.
[19] T. Ahmad, A. Bhatt, P. Bhatt, and V. Mago, "DNN-based brain tumor detection using MRI images and quantum feature extraction," Journal of Ambient Intelligence and Humanized Computing, vol. 14, pp. 9741–9754, 2023.
[20] U. Baid, S. Ghodasara, S. Mohan, et al., "The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification," arXiv:2107.02314, 2021.
[21] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," Physical Review A, vol. 99, no. 3, p. 032331, Mar. 2019.
[22] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," Physical Review A, vol. 98, no. 3, p. 032309, Sep. 2018.
[23] Z. Chen, L. Li, and W. Bhatt, "Quantum vision transformers for medical imaging classification," IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 6, pp. 8142–8155, Jun. 2024.
[24] C. O. Marrero, M. Kieferová, and N. Wiebe, "Entanglement-induced barren plateaus," PRX Quantum, vol. 2, no. 4, p. 040316, Oct. 2021.