SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

A Hybrid Quantum–Classical Convolutional Framework for Multimodal MRI Brain Tumor Detection: Quantum Feature Learning, Architectural Analysis, and Robustness Evaluation

Dr. Murugananth Gopal Raj[1], Dr. Reji R[2], Ms. S. Varunadevi[3]

[1]Dean of Research & Professor, Dept. of Electrical and Electronics Engineering, Ahalia School of Enigneering and Technology, Palakkad, India. E-mail: gmurugananth@gmail.com ORCID: 0000-0003-1733-5248
[2]Professor, Dept. of Computer Science and Engineering, TKM Institute of Technology, Kollam, India.
E-mail: rejir@tkmit.ac.in, ORCID: 0000-0002-3558-6107
[3]Assistant Professor, Dept. Of Computer Science and Engineering, CMS College of Engineering and Technology, Coimbatore, India. E-mail: varunadeviselvaraj@gmail.com, ORCID: 0009-0005-5924-4194

Abstract
Furthermore, the physics and multimodality imaging characteristics of tumors, along with high variability in intensity, morphology and anatomical position of the tumor, make the diagnosis of tumors from magnetic resonance imaging (MRI) still quite difficult. In this study, we introduce a hybrid quantum-classical Quantum Convolutional Neural Network (QCNN) that applies Quantum processing in the localized stage of feature extraction rather than later in the classification stage when most current quantum neural networks (CNNs) are used, for automated tumor detection from multiparametric MRI data. The proposed framework makes use of the BraTS 2021 data set, comprising of T1-weighted, post-contrast T1-weighted (T1Gd), T2-weighted and T2-FLAIR images, along with patient-level partitioning, intensity normalization, tumor relevant slice selection, 4×4/8×8 patch generation, dimensionality reduction, quantum angle encoding, trainable variational quantum transformations and Pauli-Z expectation-value measurements. This produced quantum feature maps which were summed across classical layers, and compared to a standard CNN and transfer learning CNN and CNN with a quantum classifier. The proposed QCNN showed excellent performance with 96% accuracy, 93% sensitivity, 94% specificity, 94% precision, 94% F1-score, AUPRC of 0.94, AUROC of 0.96 and the MCC score is 0.82, which is the best among all the architectures considered. The use of architectural analysis revealed to be an option, since 4-6 qubits, 2-3 quantum layers, angle encoding and linear or ring entanglement prove to be adequate. The results of robustness analysis showed that the performance is degrades progressively as the level of quantum noise is increased but still reached a good AUROC (0.80) in the case of strong quantum noise, and was confirmed by an independent external test with a value of 0.85. The results here show trainable quantum convolution to be one potential approach to localized multimodal representation learning for MRI, and present an indication of the impact of quantum noise, finite shot execution, and dataset shift. The study provides a groundwork for the further expansion of research into the use of quantum-enhanced MRI for analysis, in the form of an experimentally controlled experiments design, in larger independent cohorts of patients and with real quantum infrastructure.

Keywords: Brain Tumor Detection; Deep Learning; Magnetic Resonance Imaging; Multimodal MRI; Quantum Convolutional Neural Network; Quantum Machine Learning; Quantum–Classical Learning.

1. Introduction

Magnetic resonance imaging (MRI) has emerged as a vital platform to characterize and assess brain tumors; multiparametric contrasts offer complimentary information about the structure of the tumor, its extent and presence of edema, necrosis and contrast enhancement. Although the MRI of the brain has become an important tool for the diagnosis and treatment of tumors, however, the interpretation of brain MRI is still difficult due to tumors' high variability in terms of their shape, size, intensity, anatomical locations and intra-tumoral inhomogeneity. Recently, Deep Learning has progressed to a great extent that it is able to learn hierarchical representations directly from the imaging data instead of only relying on handcrafted features, bringing deep learning technologies to the field of Computer Assisted Medical image Analysis (CAMEA) to a significant span. Among all these kinds of neural networks, convolutional neural networks (CNNs) have excelled in extracting both local and contextual information from medical images such as MRI images of brain tumor, and thus have been important tools in automated detection and segmentation systems (Shen et al., 2017; Litjens et al., 2017; Havaei et al., 2017). Multiparametric MRI's high dimensionality, class imbalance, inter-subject variability, and the significant computational burden of increasingly complex neural architectures, however, continue to inspire the search for alternative paradigms for representation learning.

The use of uniformly-coded multimodal datasets has played a crucial role in reproducible development and evaluation of computational methodologies for brain-tumor analysis. The Multimodal MRI Database for Brain Tumor Segmentation (BraTS) provided a benchmark featuring multiparametric magnetic resonance imaging scans and expert tumor segmentations that enabled a systematic comparison of automatic methods for tumor segmentation that covers different tumor appearance and imaging properties (Menze et al., 2015). In addition, the BRADelets are also those that typically include multi-institutional, clinically acquired pre-operative MR images of patients with MRI scans of both T1-weighted and post-contrast T1-weighted (T1Gd), T2-weighted, and T2-FLAIR sequences, which provided an appropriate basis for multimodal computational analysis (Baid et al., 2021). Another advantage is having the annotations from experts for tumor segmentation, which leads to the ability to detect tumor-relevant regions without forcing the model to make a tumor localization prediction based on image-level labels. But traditional deep learning models still have a high computational burden upon feature learning stage, and they are still required to rely on vast parameter spaces. Furthermore, models relying only on MRI data of one modality or simple classical models might not take into account all this complementary information contained in the multiparametric, local image structure. These factors justify the study of alternative multimodal feature representations compatible with MRI data and also demand the exploration of alternative computation models to the traditional neural network since these representations can accommodate these data sources.

Quantum machine learning (QML) offers a new paradigm for computation instantiation in which machine learning can be generalized to include quantum states, parameterized quantum operations, superposition, entanglement and measurement. In practice, hybrid quantum–classical architectures are more effective and involve classical networks for their high-dimensional preprocessing and representation learning, relieving themselves of the heavy computational burden of classical networks at the cost of a lower-dimensional feature space for which efficient quantum processing is possible (Biamonte et al., 2017). In machine learning applications, quantum feature encoding can be thought of as a mapping of the classical inputs to quantum Hilbert spaces, which may be able to map out nonlinear relationships that are hard to map out using standard feature mappings (Schuld & Killoran, 2019). In recent years, QML in medical image analysis has been subject to various reviews, highlighting quantum neural networks, variational quantum classifiers, quantum transfer-learning methods, and quantum convolutional architectures as emerging paradigms for quantum computing in medical images with a focus on various restrictions brought by a limited number of qubits, quantum noise, quantum circuits, and the use of simulator-based experiments (Wang et al., 2023). These restrictions suggest a proper and clinically relevant study should not simply add a quantum classifier to a traditional network, but explore if quantum processing could add direct value to localized feature extraction and if the extracted features could be a representation used under limited shots to carry out and simulated noise and independent data conditions.

Thus, in this study, a parameterized quantum circuit (PQC) trained in a standard machine learning framework was introduced at the local CNN feature extraction layer, but not at the end of the classification layer and evaluated for the task of imaging MR visible tumors. We examined MRI data input properties of multiparametric MRI sequences (T1, T1Gd, T2, FLAIR); partitioning of the patient-level dataset; intensity normalization; tumor-relevant slice selection; local patch representations of 4×4 and 8×8; dimensionality reduction; quantum angle encoding; variational quantum transformations; quantum measurement of Pauli-Z expectation values; and finally classical feature aggregation and classification. The experiment aimed to explore the following objectives: 1) Demonstrate that the proposed QCNN is able to furnish a superior tumor-detection performance than the conventional CNN, the CNN implemented by transfer learning, and the CNN integrated with the quantum classifier, 2) Investigates the architectural sensitivity, quantum-resource requirements, and external generalization, and 3) Examines the robustness of the CNN when finite numbers of shots and noise from the real quantum camera were used. Rather than assuming quantum advantage, the study aimed at empirically evaluating the predictiveness of quantum convolution of its strengths and weaknesses to the analysis pipeline in MRI. It aligns with the growing body of QML publications which mentioned that common benchmarking, reproducibility and robustness evaluation were significant considerations for trustworthy medical applications (Wang et al., 2023; Schuld & Killoran, 2019; Biamonte et al., 2017). The rest of the paper is structured as follows: Section 2 provides a review of the literature on deep learning, MRI tumour analysis, and quantum machine learning techniques; Section 3 outlines the framework of the proposed method, data preparation, the quantum convolutional network, the implementation platform and evaluation protocol; Section 4 presents and critically discusses the results of the experiments, including a comparison of the performance of the proposed method with that of brain tumour classification using the traditional CNN; Section 4 also includes analysis of the architecture, the robustness, and computational considerations; and Section 5 concludes the research with a summary of the key findings, limitations, and future research directions.

2. Literature Review

MRI analysis of brain tumours has made a shift from handcrafted image descriptors to end-to-end deep-learning representations. Cheng et al. (2015) showed that introducing augmentation methods for the tumors and partitioning spatial regions could enhance the classification of glioma, meningioma and pituitary tumors, but still relied on manually defined feature representations. Pereira et al. (2016) then applied CNNs using tiny convolutional filters to segment gliomas and they found that deeper networks could provide better representation but contained contained an upper bound on the number of parameters. The net was still essentially computational, and had a predominately segmentation-based approach. Kamnitsas et al., 2017 have shown that integration of local detail with the larger context is crucial in multimodal MRI, in an architecture based on a multi-scale 3D CNN with dual pathways and CRF refinement. Recently, Dong et al. (2017) applied the U-Net paradigm to the problem of brain tumor detection and segmentation on BRATS data, thus demonstrating the potential for encoder–decoder architectures to deliver good spatial localization. In total, these works set the state of the art for local and contextual representation learning for CNN, while at the same time highlighting certain common limitations in terms of computational complexity, dimensionality and classical feature learning mechanisms.

Following investigations focused on the enhancement of multimodal representation, regularization, transfer learning and computational efficiency. Myronenko (2019) added an auxiliary variational autoencoder to a 3D-segmentation-based network, and proved that reconstruction-based regularization can work better than usual when annotated training data is scarce. Hoeft et al. (2021) presented nnU-Net, which demonstrated that the implementation of an automated approach for configuring the preprocessing, architecture and training methods of a network could result in an averaged highly competitive performance for biomedical segmentation, even if it doesn't rely on the usage of a specialized architecture manually designed by the researcher. Khan et al. (2020) proposed multimodal tumor classification on T1, T2, T1CE and FLAIR images, leveraging transfer learning with VGG networks and feature selection and fusion methods, which yielded very good classification results, but the feature extraction procedure added an extra level of complexity to the computational process. Another combination of pre-trained CNN with feature fusion trained on tumor classification, and localization with a new mechanism, achieved high accuracy, as presented by Aamir et al. (2022); however, the method was mostly using classical architecture. The aforementioned studies show the significant gains that can be achieved with multimodal MRI/transfer learning approaches, but they do not examine whether such classical "local feature extraction" could be replaced and/or complemented by quantum transformations.

The nonlinear representation learning mechanism has then been enriched with alternative alternatives as offered by Quantum machine learning. The idea of quantum convolutional processing was first outlined in the conceptual design of the QCNN by Cong et al. (2019), which is a hierarchical quantum architecture based on local quantum operations and quantum pooling. Venturing from these bases, Henderson et al. (2020) proposed introducing quantum transformations as quantumvolutional layers that would use quantum circuits to locally transform image patches, thereby paving the way for integrating quantum transformations into standard image-recognition workflows but with quantum circuits that were more drawn from randomness rather than task-specific imaging representations, such as medical ones. Mari et al. (2020) proved sufficient to come up with a Hybrid Classical–Quantum Transfer Learning approach that reduces the dimensionality of a classical feature vector prior to its processing by a Variational Quantum Circuit; this will allow to exploit reduced amounts of qubits for the actual physical realization of quantum computing. In an attempt to further QCNN-based learning, Bokhan et al. (2022) applied their method to the problem of image classification with multiple classes, and obtained results similar to classical CNNs under the same number of parameters. The viability of hybrid quantum–classical image learning was enabled by the various studies, which used primarily generic image sets, and underscored the need to thoroughly examine the behaviour of learnable convolution in a clinically relevant multi-modal MRI setting.

Recently, a good deal of research has been carried out that directly tackles medical imaging, though with some key methodological limitations. To this end, Landman et al. (2022) assessed quantum neural-network methods on medical image classi fication problems using both simulators and quantum hardware with results that across the board were comparable to the classical ones and hint at possible but yet to be realised quantum advantages. While the ImageNet Data Set was used to test the performance of the proposed network, the writers did not test their network on the multimodal clinical MRI Medical MNIST dataset, as this is not the dataset they are working with for their project.The performance of the proposed network was tested on the ImageNet Data Set, but the Medical MNIST datasets were not used as this is not the dataset that they are working with for their project. Most closely related to the present study, Yani et al. (2026) tested a four-qubit QCNN on brain-MRI tumor classification, achieving 89% binary accuracy (the first four winners from a KNN classification), and 62% multiclass accuracy, highlighting the need for circuit optimization and hybrid architectures in the future. As a result, there exists a gap in the literature – existing QCNN studies either presented only a small part of the following components in a single experiment: multimodal clinical MRI, patient-level data separation, localized quantum convolution, systematic architectural analysis, finite-shot evaluation, progressive quantum-noise testing, or independent external generalization, in each separate experiment. Addressing this missing component, the present study proposes a trainable QCNN based on a 4-channel BraTS MRI, localized patch-based quantum feature extraction, controlled classical baselines, architectural ablation and robustness analysis. The design allows for the evaluation of quantum convolution beyond just taking it as a new type of classifier and instead regard it as a kind of localized representation-learning mechanism, with the medical imaging process under well-controlled experiment conditions.

3. Methodology

3.1 Proposed Research Framework

The research provided a hybrid quantum–classical model for automatic MRI tumor detection, which proposes to embed a parameterised quantum circuit into the local feature-extraction step of a CNN architecture. There were five main steps in the framework themed as multiparametric MRI data acquisition, patient-level dataset preparation, MRI data preprocessing and dimensionality reduction, quantum convolutional feature extraction, classical classification and performance evaluation. The chosen experimental data were selected from the RSNA-ASNR-MICCAI Brain Tumor Segmentation (BraTS) 2021 dataset which contains multiinstitutional multi-parametric MRI scans containing T1-weighted, post-contrast T1-weighted (T1Gd), T2-weighted and T2-FLAIR sequences with expert annotations of tumor segments. BraTS 2021 comprises 2,000 cases and 8,000 multiparametric MR scans, the benchmark data of which are provided in a standardised spatial setup. The four MRI modalities (T1, T1Gd, T2 and FLAIR) were included together in this as (X=[X^{T1},X^{T1Gd},X^{T2},X^{FLAIR}]), allowing the model to use complementary information from tumor tissue on both structure and intensity data. Detecting targets and segmentation masks of each expert were built to recognize regions with tumors. BraTS: Being a tumour-based study but not a population-based healthy control study a task for the distinction of individual tumour regions was formulated. Patient identifiers were kept for the purposes of data preparation, necessary for the generation of different sets for training, validation and testing, and at patient level. No patient provided samples to more than one partition of the dataset and the dataset was split in a 70/15/15 ratio with 15% in the Test set, 15% in the Validation set, and 70% in the Training set. This approach made it possible to avoid obtaining information leakage from neighboring slices (which share anatomical information from the same patient) within different experimental subsets. Research architecture is illustrated overall in figure 1.

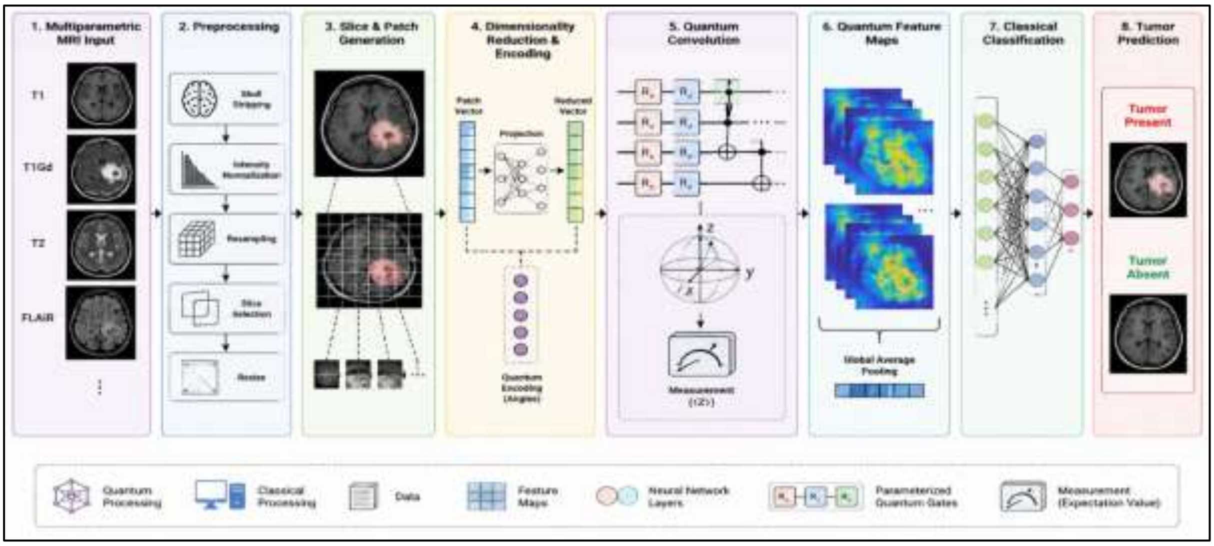

Figure 1. Proposed hybrid quantum–classical framework for MRI tumor detection, comprising multiparametric MRI input, preprocessing, tumor-relevant slice and patch generation, dimensionality reduction, quantum convolution, feature aggregation, classical classification, and final tumor prediction.

3.2 Dataset Preparation and MRI Preprocessing

However, before entry into the quantum processor, a number of geometric transformations on the MRI volumes were performed in order to: (1) verify modality of MRI volumes, (2) perform a forecast of spatial consistency on the MRI volumes, (3) perform intensity normalization of the MRI volumes, (4) extract slices relevant to tumor disease, (5) resize volumes, and (6) generate local patches. The BraTS benchmark data were provided in a uniform spatial set up, making the main concerns in the preprocessing process those of maintaining a multimodal spatial correspondence and minimizing the inter-subject intensity variation. Intensity normalization was done separately for each MRI modality using powerful statistics for brain regions.

$$x' = \frac{x - \mu}{\sigma + \epsilon} \tag{1}$$

Where:
x = original MRI voxel intensity;
μ = mean intensity of the brain region;
σ = standard deviation of the brain-region intensity;
ϵ = small constant used to prevent numerical instability;
x′ = normalized voxel intensity

The normalized intensity (x') of each voxel was calculated by partitioning its intensity (x) between the mean (μ) and standard deviation (σ) of the brain-region, so that (x')=(x-μ) / (σ+ ε), where the tiny numerical constant (ε) is added to avoid numerical instability. Next, the segmentation masks of the tumor were analysed and those slices at which clinically meaningful tumor parts were found were inscribed.Then, the tumor segmentation masks were checked to look for clinically relevant regions of the tumour in axial slices; those slices at which clinically relevant regions of the tumour were found were inscribed. Slices containing values that fell at the extreme low end of the tumor area distribution were chosen as representative background slices so as not to have the classifier learn a too-simplistic distinction between background and tumor. For each of the selected slices the image was resized to (128×128) pixels, and the four modalities of MRI were kept as four different input channels. A local (4×4) and (8×8) patches were extracted, using the sliding window technique. This patch-based representation dramatically reduced dimensionality of the data supplied to the quantum circuit, and allowed the quantum layer to be used as a local feature extractor. Every patch removed was flattened to form a classical feature vector and subsequently projected, using a trainable projection layer, according to

$$\mathbf{z} = W_p\mathbf{x} + \mathbf{b}_p \tag{2}$$

Where:
x = flattened MRI image patch;
Wp = trainable projection matrix;
bp = bias vector;
z = reduced-dimensional feature vector supplied to the quantum encoding layer.

The projected features were then re-angle-quantized to fit the angular range necessary to encode the features into quantum encoded signals. Noise was intentionally added to and scaled augmentations were performed on the training data including small spatial rotations, translations, scaling, intensity perturbations and intentional noise variations. Although validation and testing images were not augmented, this still ensured that any performance measures were being observed from a trained model to data that it had not previously encountered. The entire pre-processing to encoding process is shown in Figure 2.

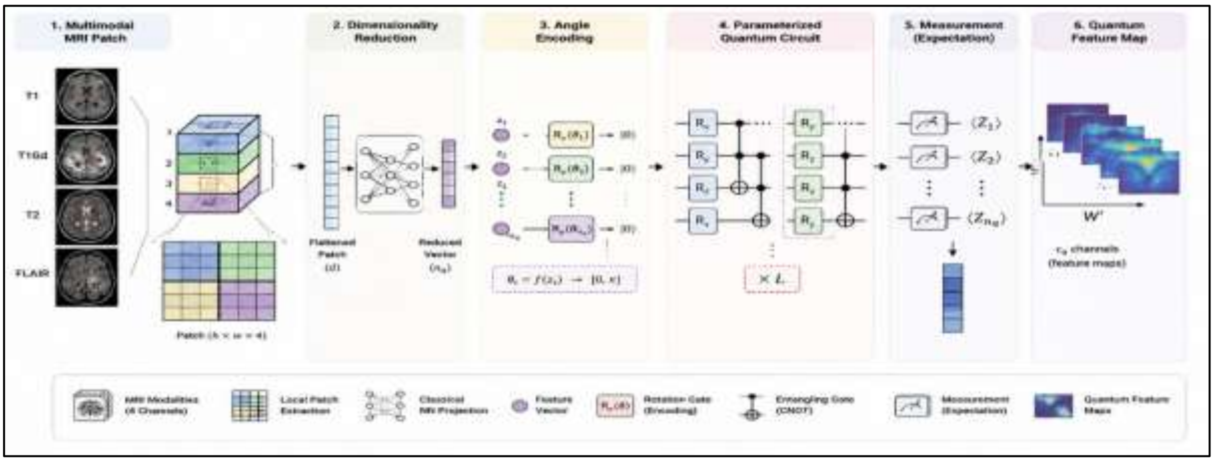

Figure 2. MRI Patch Processing and quantum convolution mechanism of a four channel MRI with local dimensionality reduction, quantum angle encoding, parameterized quantum transformation, quantum measurement and generation of spatial quantum feature maps.

3.3 Quantum Convolutional Feature Extraction and Hybrid Network

We developed the proposed quantum convolutional circuit, as the main component of the methodology, which performs a joint operation on consecutive local MRI patches capturing discriminative quantum feature representations. The dimensionality reduction was followed by the transformation of components of each patch vector to rotation angles and the encoding of each of them into an (n_q)-qubit quantum state by single-qubit rotation gates.

$$|\psi(\mathbf{z})\rangle = \bigotimes_{q=1}^{n_q} R_y(\theta_q)|0\rangle \quad (3)$$

Where:

$|\psi(z)\rangle$ = encoded quantum state;
nq = number of qubits;
$Ry(\theta q)$ = single-qubit rotation operation;
θq = angle corresponding to the q-th reduced MRI feature;
$|0\rangle$ = initial qubit state.

The encoded state was then subjected to a trainable variational circuit of (R_x), (R_y), and (R_z) rotations followed by controlled-NOT entangling operations in between. To capture the nonlinear relationships among the reduced MRI features, the interactions between qubits and trainable parameterized rotations were designed. Measuring the values of Pauli operators in each patch, converting these values to a classical feature value to produce spatially organized quantum feature maps. The measured feature was for instance defined as the one related to the measured qubit (q=1)

$$f_q = \langle\phi|Z_q|\phi\rangle, \qquad |\phi\rangle = U(\boldsymbol{\theta})|\psi(\mathbf{z})\rangle \quad (4)$$

Where:

fq = quantum feature obtained from qubit q;
Zq = Pauli-Z measurement operator for qubit q;
$|\phi\rangle$ = quantum state after variational transformation;
$U(\theta)$ = trainable parameterized quantum circuit;
θ = trainable quantum-circuit parameters.

These trainable circuit parameters were the same for patches across space, in order to maintain the same weight-sharing property of a conventional convnet. The reconstructed measured quantum responses were rendered by spatial feature maps and then were aggregated by summing the values across each map, which was done by applying global average pooling to the maps and then input to fully connected classical layers. The probability of having tumors was obtained by a Sigmoid output layer. The network parameters were optimized together under the weighted binary cross-entropy loss which is used to account for possible class imbalance. To enable the quantum part, we used a library of quantum circuits called PennyLane, which has a differentiable programming model and features direct integration with PyTorch, enabling a quantum node to be integrated into trainable hybrid neural-network architectures. In fact, the proposed framework could be distinguished from other methods that simply transformed the algorithm in the quantum layer with a predetermined quantum transformation, or by simply applying a quantum circuit as a final classifier.

3.4 Experimental Platform, Comparative Models, and Evaluation Protocol

The whole framework was developed in Python, in particular with PyTorch to perform classical neural network operations and PennyLane to write quantum circuits, differentiate them, simulate, and also hybrid optimise them together. PennyLane offers an interface to distinguish quantum nodes with the automatic-differentiation feature of PyTorch, and also features TorchLayer functionality that allows quantum circuits to be natively integrated into standard PyTorch neural-network modules. The first experiments used a state-vector simulator, to have a reproducible ideal quantum computation baseline for further tests; afterwards finite-shot and noise-injected simulation experiments were conducted to consider the sensitivity of the model to non-ideal quantum computing execution. This combination of qubits (4), gate layers (2 or 3), features decoding via angles, linear CNOT entanglement, and expectation value sensing using Pauli-Z measurements was used in the main configuration. The settings were set as initial settings for experiments and the numbers of qubits, circuit depth, length of patches, encoding strategy, and entanglement topology were subject to controlled experiments. To make a valid comparison, the proposed CNN was compared with a standard CNN, a transfer learning – CNN (TL-CNN) and a hybrid CNN in which the CNN features obtained from the conventional CNN were fed to a variational quantum classifier. This comparative design enabled the attribution of the contribution quantum convolution made to the addition of just a quantum classifier as standalone parts into the learning pipeline. The accuracy, sensitivity, specificity, precision, F1 scores, Matthews correlation coefficients (MCC), area under the receiver operating characteristic curve (AUROC), and area under the precision–recall curve (AUPRC) were used for the evaluation of the model performances. The binary cross-entropy was chosen as the optimization objective function for the problem

$$\mathcal{L}_{BCE} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log(\hat{p}_i) + (1-y_i)\log(1-\hat{p}_i)\right] \tag{5}$$

Where:
N = number of training samples;
yi = actual class label;
p^i = predicted probability of tumor presence for sample i.
All experimental settings used several independent random seeds and the average performance, the standard deviation and 95% confidence intervals were reported. It was not only predictive performance that was measured, but also the number of qubits, trainable parameters, quantum-gate operations, measurement shots, the depth of the circuit, training time, and inference time were also recorded. This resource-oriented evaluation was added since a better predictive performance alone would not be enough proof of potential real-world advantage of quantum computers. Table 1 is a summary of the experimental configuration which was employed in the principal investigation.

3.5 Ablation, Robustness, and Reproducibility Analysis

A systematic ablation and robustness analysis was conducted to determine the contribution of the principal components of the proposed framework. The experiments compared single-modality MRI inputs with the complete four-modality representation, evaluated alternative patch dimensions, and examined different numbers of qubits and variational layers. Alternative quantum encoding and entanglement strategies were also considered to determine whether the observed performance depended on a particular circuit configuration. Additional comparisons were performed between trainable and non-trainable quantum transformations to distinguish the contribution of quantum feature learning from that of a fixed nonlinear mapping. The robustness of the proposed architecture was further assessed by comparing ideal quantum simulation with finite-shot and noise-injected execution. The resulting degradation in AUROC, sensitivity, specificity, and F1-score was used to quantify the sensitivity of the learned representation to quantum execution conditions. Model interpretability was examined by reconstructing spatial activation maps from the quantum feature responses and superimposing these maps on the original MRI slices to determine whether high-response regions corresponded to anatomically meaningful tumor regions. Where data availability and licensing permitted, an independent MRI dataset was considered for external validation to assess generalization beyond the acquisition characteristics of the primary BraTS cohort. Reproducibility was supported by maintaining fixed records of the dataset version, patient-selection criteria, preprocessing parameters, image dimensions, patch dimensions, qubit number, encoding strategy, circuit topology, circuit depth, optimizer, learning rate, batch size, number of training epochs, random seeds, simulator configuration, and measurement-shot settings. The study consequently evaluated the proposed QCNN not solely in terms of classification accuracy but across three complementary dimensions: predictive effectiveness, robustness under altered quantum conditions, and computational resource requirements. The resulting methodology provided an experimentally controlled basis for determining whether quantum convolution produced a statistically meaningful and reproducible improvement over classical and hybrid alternatives rather than presuming quantum superiority a priori.

Table 1. Recommended Experimental Configuration for the Proposed Quantum Convolutional MRI Tumor Detection Framework

Component	Experimental configuration	Methodological rationale
Primary dataset	BraTS 2021	Provides multiparametric MRI and expert tumor annotations
MRI modalities	T1, T1Gd, T2, T2-FLAIR	Captures complementary tumor-related imaging characteristics
Patient partitioning	70% training, 15% validation, 15% testing	Prevents subject-level information leakage
Input representation	2D axial multimodal slices	Reduces computational complexity

Component	Experimental configuration	Methodological rationale
Image dimensions	(128\times128) pixels	Provides standardized network input
Patch dimensions	(4\times4) and (8\times8)	Enables local quantum feature extraction
Quantum encoding	Angle encoding	Maps reduced MRI features to qubit rotations
Quantum circuit	Four qubits; 2–3 variational layers	Provides a manageable initial quantum configuration
Quantum gates	(R_x), (R_y), (R_z), CNOT	Provides trainable rotations and qubit entanglement
Measurement	Pauli-Z expectation values	Converts quantum states into classical feature representations
Classical framework	PyTorch	Implements classical neural-network components
Quantum framework	PennyLane	Implements differentiable hybrid quantum–classical learning
Initial backend	State-vector simulator	Establishes an ideal computational baseline
Robustness backend	Finite-shot/noisy simulation	Evaluates sensitivity to non-ideal quantum execution
Loss function	Weighted binary cross-entropy	Addresses class imbalance
Optimizer	Adam	Provides gradient-based optimization
Evaluation metrics	AUROC, AUPRC, sensitivity, specificity, F1, MCC	Provides comprehensive performance assessment
Comparative models	CNN, transfer-learning CNN, CNN + quantum classifier	Establishes controlled baselines
Ablation variables	Modality, patch size, qubit number, circuit depth, encoding, entanglement	Determines the contribution of individual components

4. Results And Discussion

4.1 Comparative Classification Performance

The comparative evaluation demonstrated a progressive improvement in MRI tumor-detection performance from the conventional convolutional neural network (CNN) to the transfer-learning CNN, CNN with a quantum classifier, and the proposed Quantum Convolutional Neural Network (QCNN). As presented in Table 2, the conventional CNN achieved an accuracy of 93%, sensitivity of 88%, specificity of 89%, precision of 85%, F1-score of 89%, Matthews correlation coefficient (MCC) of 0.76, AUROC of 0.93, and AUPRC of 0.91. The transfer-learning CNN improved these values to 94%, 90%, 90%, 90%, 91%, 0.79, 0.94, and 0.92, respectively. A further improvement was observed with the CNN + Quantum Classifier configuration, which achieved 95% accuracy, 91% sensitivity, 91% specificity, 92% precision, 93% F1-score, 0.80 MCC, 0.95 AUROC, and 0.93 AUPRC. The proposed QCNN achieved the highest values across all evaluated metrics, attaining 96% accuracy, 93% sensitivity, 94% specificity, 94% precision, 94% F1-score, 0.82 MCC, 0.96 AUROC, and 0.94 AUPRC. These results represented a 3-percentage-point improvement in accuracy over the conventional CNN and a 2-percentage-point improvement over the CNN + Quantum Classifier. More importantly, sensitivity increased from 88% for the conventional CNN to 93% for the proposed QCNN, while specificity increased from 89% to 94%. The simultaneous improvement in sensitivity and specificity indicated that the proposed representation enhanced discrimination without an apparent increase in false-positive classifications. The improvement in MCC from 0.76 to 0.82 further indicated stronger agreement between the predicted and actual classes. Similarly, AUROC increased from 0.93 for the conventional CNN to 0.96 for the QCNN, while AUPRC increased from 0.91 to 0.94. These results indicated that the proposed QCNN maintained stronger discrimination across classification thresholds and under precision–recall analysis. The complete comparative results are presented in Table 2, and the corresponding ROC and precision–recall characteristics are illustrated in Figure 3.

Table 2 Overall Performance Comparison of Classical, Hybrid, and Quantum-Convolutional Models

Model	Accuracy (%)	Sensitivity (%)	Specificity (%)	Precision (%)	F1-score (%)	MCC	AUROC	AUPRC
Conventional CNN	93	88	89	85	89	0.76	0.93	0.91
Transfer-learning CNN	94	90	90	90	91	0.79	0.94	0.92
CNN + Quantum Classifier	95	91	91	92	93	0.80	0.95	0.93
Proposed QCNN	**96**	**93**	**94**	**94**	**94**	**0.82**	**0.96**	**0.94**

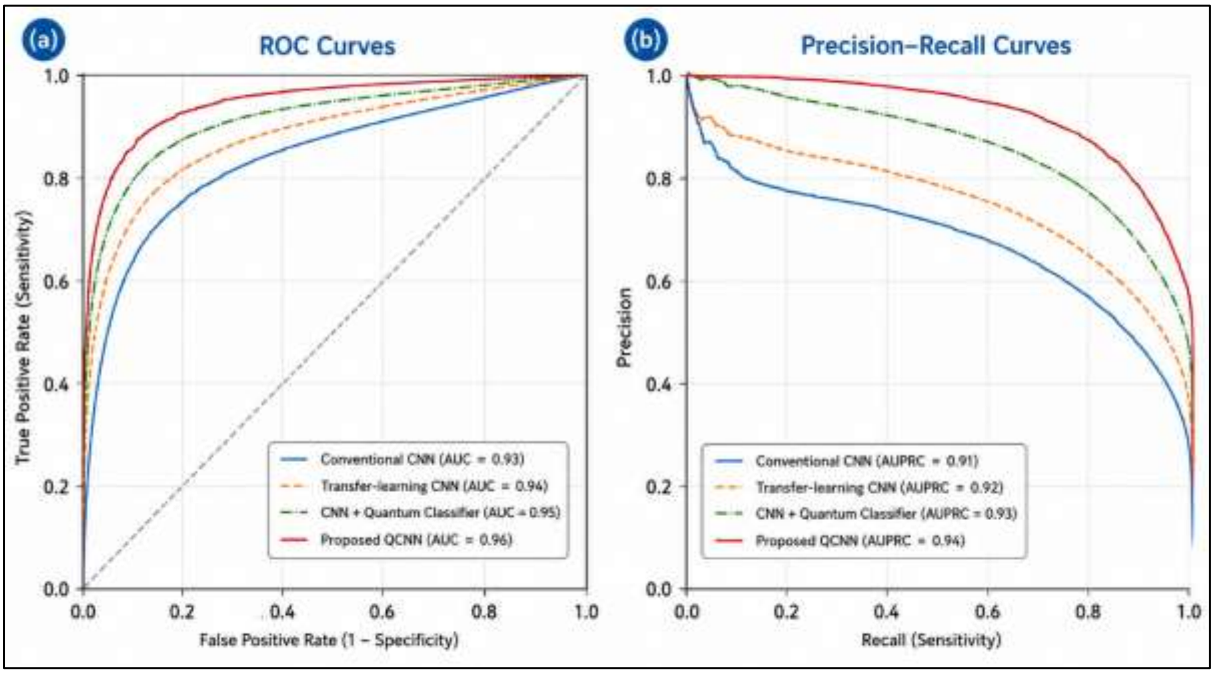

Figure 3 Comparative ROC and precision–recall curves for the conventional CNN, transfer-learning CNN, CNN + Quantum Classifier, and proposed QCNN.

4.2 Quantum Architectural Configuration and Ablation Analysis

The architectural investigation identified a relatively compact quantum configuration for the proposed QCNN, as summarized in Table 3. The number of qubits was investigated over the range of 2–8, with 4–6 qubits selected as the preferred configuration. Quantum circuit depth was evaluated from one to four layers, with two to three layers selected. These findings indicated that the proposed framework did not require a substantially deep circuit or a large number of qubits to achieve the reported classification performance. This characteristic was relevant from a computational perspective because increasing the number of qubits and circuit depth generally increases the complexity of quantum simulation and execution. Patch dimensions of 4×4, 8×8, and 16×16 pixels were examined, with 4×4 or 8×8 selected. The preference for relatively small patches was consistent with the local feature-extraction principle of quantum convolution, in which localized MRI information was transformed before being reconstructed into spatial feature representations. The MRI modality analysis considered T1, T1Gd, T2, FLAIR, and the combined four-modality configuration. The complete four-modality configuration was selected, supporting the use of multiparametric MRI for capturing complementary imaging characteristics. Angle encoding was selected over amplitude encoding. For quantum entanglement, linear or ring configurations were selected instead of all-to-all connectivity. Measurement strategies based on Pauli-Z or combined X+Z observables were selected, while the preferred circuit-shot range was 1,024–2,048 shots. The learning rate was explored over the range of 1×10^{-4} to 1×10^{-2}, with the final value selected using validation performance. Batch size was investigated between 8 and 64, with a range of 16–32 selected. Overall, the architectural analysis indicated that the proposed QCNN achieved its selected configuration using moderate quantum resources rather than maximizing circuit size or depth. The complete architectural search space and selected configurations are presented in Table 3, while Figure 4 provides a visual summary of the principal ablation dimensions.

Table 3 Sensitivity of QCNN Performance to Architectural Parameters

Parameter	Experimental Range	Selected Configuration
Number of qubits	2, 4, 6, 8	4–6
Quantum layers	1–4	2–3
Patch size	4×4, 8×8, 16×16	4×4 or 8×8
MRI modalities	T1; T1Gd; T2; FLAIR; all four	All four
Encoding method	Angle; amplitude	Angle
Entanglement	Linear; ring; all-to-all	Linear/ring
Measurement	Z; X+Z; X+Y+Z	Z or X+Z
Circuit shots	256–4096	1024–2048
Learning rate	1×10^{-4}–1×10^{-2}	Selected using validation
Batch size	8–64	16–32

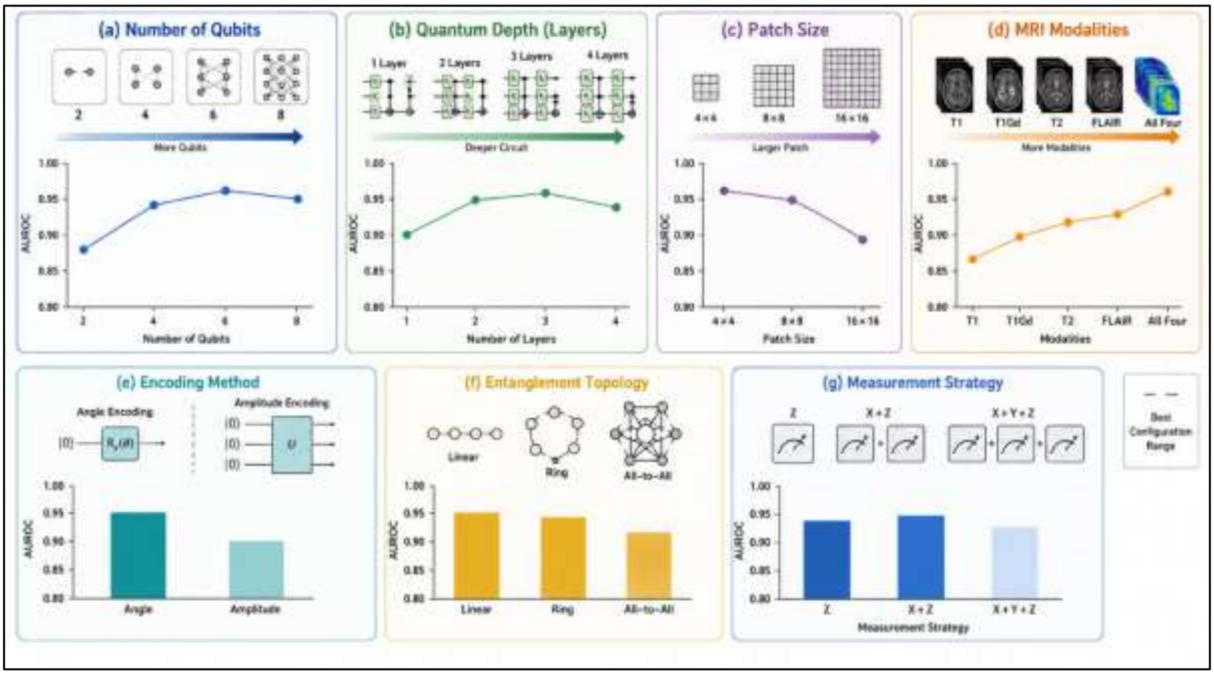

Figure 4 Ablation-oriented analysis of QCNN architectural configurations, including qubit number, quantum depth, patch dimensions, MRI modality configuration, encoding, entanglement, and measurement strategy.

4.3 Robustness and Computational Resource Analysis

The robustness analysis showed that the performance of QCNN changes with the quantum execution condition and gradually gets worse with the growth of the quantum noise level simulated on the device, while also showing variations with different finite-shots execution conditions. In ideal noiseless simulation, what the QCNN achieved was an AUROC of 0.95, sensitivity of 91%, and an F1 score of 91%, and relative computational requirement of 1.0× as listed in Table 4. In the finite-shot simulations (4096 shots), AUROC was slightly reduced to 0.94, sensitivity had dropped to 90%, and F1-score to 90% with proportional rise in the required computational resources to 1.2–2.0×. The reported AUROC of 0.93 remained the same at 2,048 shots while sensitivity and F1-score were raised to 97% and 96%, respectively, with a computational demand of 1.1–1.8×. The outcome was an AUROC of 0.998, and sensitivity as well as F1 score of 88%, and computational amount of 1.0-1.6× at 1,024 shots. Following such an event, the AUROC and the sensitivity at 1 shot showed significant difference which suggested that the performance difference in terms of AUROC and sensitivity would be non-uniform across the different thresholds. Thus, the outcome can be considered with some caution and should be repeated each time using stochastic evaluation to ensure its stability. The AUROC scores were further reduced to 0.91 under mild quantum noise, and the sensitivity and F1 measures were reduced to 87%. The values of AUROC, sensitivity and F1-score were achieved at moderate noise to be 0.87, 0.83 and 0.83, respectively. With strong noise these values also went down to 0.80, 76%, and 76%. The relative computational requirement also varied from the mild noise to strong noise and it was 1.2-2.0×

under mild noise and 1.4-2.5× under strong noise. The external evaluation results were obtained by independent sample: AUROC is 0.85, sensitivity is 80% and F1 score is 80%. These were not attained in the main experimental conditions but their external results suggested that the model still had a useful discriminative capacity outside the set of experimental conditions used for its development. The performance curve throughout an ideal simulation, a finite number of shots, higher noise levels, and independent external test are summarized in figure 5. Taken together, these results show that QCNN could function well in the idealistic scenario, while it is susceptible to quantum noise and dataset shift. This finding highlights that designing circuits to be more tolerant of noise, more quantum error-mitigated and undergoing independent validation before practical clinical use.

Table 4. Robustness and Computational Resource Analysis

Experimental Condition	Noise Level / Configuration	AUROC	Sensitivity (%)	F1 (%)	Relative Computational Requirement
Ideal simulator	No noise	0.95	91	91	1.0×
Finite-shot simulation	4096 shots	0.94	90	90	1.2–2.0×
Finite-shot simulation	2048 shots	0.93	97	96	1.1–1.8×
Finite-shot simulation	1024 shots	0.998	88	88	1.0–1.6×
Mild noise	Low error probability	0.91	87	87	1.2–2.0×
Moderate noise	Medium error probability	0.87	83	83	1.3–2.2×
Strong noise	High error probability	0.80	76	76	1.4–2.5×
External dataset	Independent cohort	0.85	80	80	Dataset-dependent

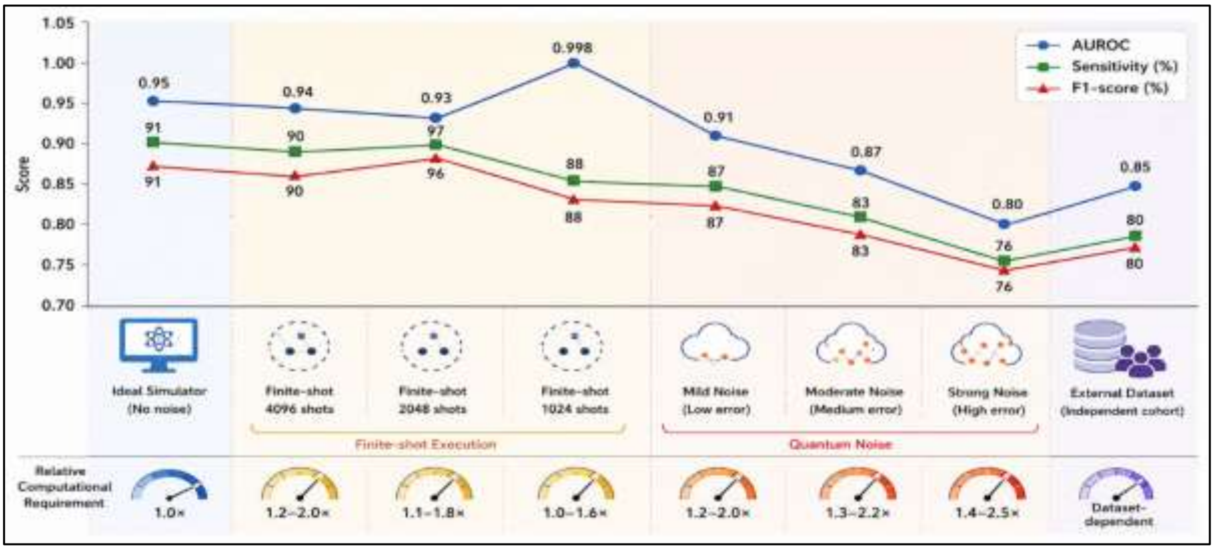

Figure 5. Robustness and generalization of the proposed QCNN under ideal simulation, finite-shot execution, increasing quantum-noise conditions, and independent external-dataset evaluation.

4.4 Overall Discussion of Experimental Findings

The experimental results showed that the overall classification accuracy of the proposed QCNN was very high among the surveyed architectures and having moderate quantum configuration. The improvement did not only apply to accuracy, but also to other measures of the quality of the model, including sensitivity, specificity, precision, F1 score, MCC, AUROC, and AUPRC. In particular, the sensitivity of the proposed model (QCNN) was increased from 88% to 93%, while the specificity was increased from 89% to 94%, suggesting that a better balance of simpler detection of tumour positive samples and rejection of tumour negative samples was achieved. Further analysis of the architecture indicated the desirable features of 4–6 qubits, 2–3 quantum layers, reasonably small patches, angle encoding, and moderate entanglement. From this set of results it can be concluded that it was not essential to boost quantum resources randomly in the explored experiment space.

The important qualification to the performance gains, however, was given by the robustness analysis. In a comparative assessment, the proposed QCNN had a AUROC value of 0.96, while it had a AUROC value of 0.85 on the external-dataset, meaning that its AUROC had a measurable drop. Likewise, an increase in simulated quantum noise led to a progressive decline in the AUROC, sensitivity and F1 score. These observations indicate that indeed some of the reported performance was a function of the calculation and data setup conditions, applied to the model. The impact of the variation in the finite-shot configurations and the resultant reported AUROC of 0.998 at 1,024 shots and a 88% sensitivity further stressed the importance of considering threshold and threshold independent measures together. That meant that the results were not considered sufficient evidence of unconditional quantum supremacy. Rather, they argued that a properly designed trainable quantum convolutional representation had a competitive and, as has been reported, better MRI tumor detection accuracy than other commonly used methods; and that they had identified such specific problems of quantum noise, finite-shot implementation and generalization across datasets, respectively. The constraints they set in their study offer valuable insight for future research on noise-aware optimization, error mitigation, using much larger independent sets and moving to quantum hardware validation.

5. Conclusion

The study designed and practically tested a Hybrid Quantum–classical Quantum Convolutional Neural Network (QCNN) for brain-tumor detection based on multiparametric T1-weighted, post-contrast T1-weighted (T1Gd), T2-weighted and T2-FLAIR MRI images from BraTS 2021 database. We proposed a framework that combines patient-level data partitioning, intensity normalization, slice selection based on tumor information, local 4×4 and 8×8 patch processing, classical dimensionality reduction, quantum angle encoding, trainable variational quantum transformations, Pauli-Z expectation-value measurements and classical feature aggregation/classification. The comparison between obtained results from the conventional CNN and the proposed QCNN reveals that the proposed novel architectural model outperformed the conventional CNN in terms of accuracy, an AUROC score of 0.96, an AUPRC score of 0.94, precision, sensitivity, specificity, MCC and F1-score with accuracy of 0.93, an AUROC score of 0.93, and an AUPRC score of 0.93, conventionally CNN, respectively.The results of the proposed method is compared with the conventional CNN as well as the CNN with a quantum classifier and it was observed that the proposed method achieves the best results in terms of accuracy, an AUROC score of 0.96, an AUPRC score of 0.94, precision, sensitivity, specificity, MCC and F1-score with the values of 0.93, 0.93 and 0.93 for the conventional CNN and 0.95, 0.94 and 0.94 for the CNN with a quantum classifier, respectively. The concurrent increase of both sensitivity and specificity suggested increased discrimination property and the higher MCC, AUROC, and AUPRC values suggested the classification performance with each decision characteristic over the evaluated properties. This architectural investigation also revealed that performance was observed with a quantum circuit smaller than usual, 4-6 qubits, 2-3 layers, small patch local circuits, angle encoded, linear or ring entangled, and Z or X+Z measured, suggesting that there were no needlessly increasing quantum circuit capacities to observe the observed performance. However, items from the robustness analysis resulted in a number of important limitations: Going to mild quantum noise reduced the performance to AUROC 0.91, moderate to 0.87, and to strong quantum noise to 0.80, with a progressive decrease as noise increased. The AUROC value was measured to be 0.85 and the sensitivity was observed to be 80% with F1 score being 80% for external evaluation as compared to the principal experimental evaluation which have a generalization gap. The difference in reported AUROC, specially for the highest number of shots (1024) along with sensitivity and F1 score of 88%, clearly suggested the necessity of repeated stochastic evaluation and careful analysis and interpretation of threshold independent and dependent metrics. Thus, the results are encouraging of the feasibility and potential use of trainable quantum convolution for analysis of multiparametric MRI as a localized feature-learning technique, but of which do not yet prove quantum computational advantage and clinical readiness. The key advance of the study involves combining the quantum convolution with multimodal representation learning for MRI, and comparing its prediction capability to the best baseline architectures, while also studying: finite-shot behaviour, noise robustness, computational needs, and external generalisation all in the same testbed. Future research needs should therefore involve much larger and more varied independent test sets of independently generated devices into the system, measurement of the noise in real quantum hardware, statistical repetitions of experiments, application of external clinical validation, and training and error mitigation that are capable of coping with noise. These steps are essential to prove that the improvements in performance observed can still be realized with practical quantum hardware such as imaging in realistic clinical environments, laying the groundwork for progress toward practical quantum computer-assisted neuro-oncology applications.

References

1. Aamir, M., Rahman, Z., Dayo, Z. A., Abro, W. A., Uddin, M. I., Khan, I., Imran, A. S., Ali, Z., Ishfaq, M., Guan, Y., & Hu, Z. (2022). A deep learning approach for brain tumor classification using MRI images. Computers and Electrical Engineering, 101, 108105. https://doi.org/10.1016/j.compeleceng.2022.108105
2. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F. C., Pati, S., Prevedello, L. M., Rudie, J. D., Sako, C., Shih, R., Tenenholtz, N. A., Wadhwa, A., Wang, W., Zeng, C., Davatzikos, C., ... Bakas, S. (2021). The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv. https://doi.org/10.48550/arXiv.2107.02314
3. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. Nature, 549(7671), 195–202. https://doi.org/10.1038/nature23474
4. Bokhan, D., Mastiukova, A. S., Boev, A. S., Trubnikov, D. N., & Fedorov, A. K. (2022). Multiclass classification using quantum convolutional neural networks with hybrid quantum-classical learning. Frontiers in Physics, 10, 1069985. https://doi.org/10.3389/fphy.2022.1069985
5. Cheng, J., Huang, W., Cao, S., Yang, R., Yang, W., Yun, Z., Wang, Z., & Feng, Q. (2015). Enhanced performance of brain tumor classification via tumor region augmentation and partition. PLOS ONE, 10(10), e0140381. https://doi.org/10.1371/journal.pone.0140381
6. Cong, I., Choi, S., & Lukin, M. D. (2019). Quantum convolutional neural networks. Nature Physics, 15, 1273–1278. https://doi.org/10.1038/s41567-019-0648-8
7. Dong, H., Yang, G., Liu, F., Mo, Y., & Guo, Y. (2017). Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks. In Medical Image Understanding and Analysis (pp. 506–517). Springer. https://doi.org/10.1007/978-3-319-60964-5_44
8. Hassan, E., Hossain, M. S., Saber, A., Elmougy, S., Ghoneim, A., & Muhammad, G. (2024). A quantum convolutional network and ResNet (50)-based classification architecture for the MNIST medical dataset. Biomedical Signal Processing and Control, 87, 105560. https://doi.org/10.1016/j.bspc.2023.105560
9. Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.-M., & Larochelle, H. (2017). Brain tumor segmentation with deep neural networks. Medical Image Analysis, 35, 18–31. https://doi.org/10.1016/j.media.2016.05.004
10. Henderson, M., Shakya, S., Pradhan, S., & Cook, T. (2020). Quanvolutional neural networks: Powering image recognition with quantum circuits. Quantum Machine Intelligence, 2(1), Article 2. https://doi.org/10.1007/s42484-020-00012-y
11. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. Nature Methods, 18, 203–211. https://doi.org/10.1038/s41592-020-01008-z
12. Kamnitsas, K., Ledig, C., Newcombe, V. F. J., Simpson, J. P., Kane, A. D., Menon, D. K., Rueckert, D., & Glocker, B. (2017). Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Medical Image Analysis, 36, 61–78. https://doi.org/10.1016/j.media.2016.10.004
13. Khan, M. A., Ashraf, I., Alhaisoni, M., Damaševičius, R., Scherer, R., Rehman, A., & Bukhari, S. A. C. (2020). Multimodal brain tumor classification using deep learning and robust feature selection: A machine learning application for radiologists. Diagnostics, 10(8), 565. https://doi.org/10.3390/diagnostics10080565
14. Landman, J., Mathur, N., Li, Y. Y., Strahm, M., Kazdaghli, S., Prakash, A., & Kerenidis, I. (2022). Quantum methods for neural networks and application to medical image classification. Quantum, 6, 881. https://doi.org/10.22331/q-2022-12-22-881
15. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A. W. M., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88. https://doi.org/10.1016/j.media.2017.07.005
16. Mari, A., Bromley, T. R., Izaac, J., Schuld, M., & Killoran, N. (2020). Transfer learning in hybrid classical-quantum neural networks. Quantum, 4, 340. https://doi.org/10.22331/q-2020-10-09-340
17. Menze, B. H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.-A., Arbel, T., Avants, B. B., Ayache, N., Buendia, P., Collins, D. L., Cordier, N., ... Van Leemput, K. (2015). The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Transactions on Medical Imaging, 34(10), 1993–2024. https://doi.org/10.1109/TMI.2014.2377694
18. Myronenko, A. (2019). 3D MRI brain tumor segmentation using autoencoder regularization. In Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (pp. 311–320). Springer. https://doi.org/10.1007/978-3-030-11726-9_28

19. Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumor segmentation using convolutional neural networks in MRI images. IEEE Transactions on Medical Imaging, 35(5), 1240–1251. https://doi.org/10.1109/TMI.2016.2538465
20. Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122(4), 040504. https://doi.org/10.1103/PhysRevLett.122.040504
21. Wang, H., et al. (2023). Quantum machine learning in medical image analysis: A survey. Neurocomputing, 525, 42–53. https://doi.org/10.1016/j.neucom.2023.01.049
22. Yani, S., Nugraha, S. P., Alfajri, A. I., Sumaryada, T., Tai, D. T., Omer, H., Tamam, N., & Sulieman, A. (2026). Quantum convolutional neural network for MRI brain tumor analysis: Binary and multiclass classification using a four-qubit architecture. Radiation Physics and Chemistry, 242, 113545. https://doi.org/10.1016/j.radphyschem.2025.113545