

Lightweight Quantum-Convolutional Neural Network for Brain Tumor Classification

by

**Tamanna Khatun
221014063**

*Thesis (CSE4099-B) submitted in partial fulfillment of the
requirements for the degree of*

Bachelor of Science in Computer Science and Engineering

Under the supervision of

Nafees Mansoor, PhD



**DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
UNIVERSITY OF LIBERAL ARTS BANGLADESH**

SPRING 2026

© *Tamanna Khatun*
All rights reserved

DECLARATION

Project Title	Lightweight Quantum-Convolutional Neural Network for Brain Tumor Classification
Author	<i>Tamanna Khatun</i>
Student ID	221014063
Supervisor	Nafees Mansoor, PhD

I declare that this thesis entitled *Lightweight Quantum-Convolutional Neural Network for Brain Tumor Classification* is the result of my own work except as cited in the references. The thesis has not been accepted for any degree and is not concurrently submitted in candidature of any other degree.

Tamanna Khatun
221014063

Department of Computer Science and Engineering
University of Liberal Arts Bangladesh

Date: April 21, 2026



Department of Computer Science and Engineering
University of Liberal Arts Bangladesh
Mohammadpur, Dhaka - 1207

CERTIFICATE

This is to certify that the thesis entitled **Lightweight Quantum-Convolutional Neural Network for Brain Tumor Classification**, submitted by **Tamanna Khatun** (Student ID: 221014063) an undergraduate student of the **Department of Computer Science and Engineering** has been examined. Upon recommendation by the examination committee, we hereby accord our approval of it as the presented work and submitted report fulfill the requirements for its acceptance in partial fulfillment for the degree of *Bachelor of Science in Computer Science and Engineering*.

Nafees Mansoor, PhD
Professor

Department of Computer Science and Engineering
University of Liberal Arts Bangladesh

Muhammad Golam Kibria, PhD, SMIEEE
Professor & Head

Department of Computer Science and Engineering
University of Liberal Arts Bangladesh

Place: Dhaka

Date: April 21, 2026

ACKNOWLEDGEMENTS

I would like to express my deep and sincere gratitude to my research supervisor, *Nafees Mansoor, PhD*, for giving me the opportunity to do research and providing invaluable guidance throughout this work. His dynamism, vision, sincerity and motivation have deeply inspired me. He has taught me the methodology to carry out the work and to present the works as clearly as possible. It was a great privilege and honor to work and study under his guidance.

I am greatly indebted to my honorable teachers of the Department of Computer Science and Engineering at the University of Liberal Arts Bangladesh who taught me during the course of my study. Without any doubt, their teaching and guidance have completely transformed me to a person that I am today.

I am extremely thankful to my parents for their unconditional love, endless prayers, caring and immense sacrifices for educating and preparing me for my future. I would like to say thanks to my friends and relatives for their kind support and care.

Finally, I would like to thank all the people who have supported me to complete the project work directly or indirectly.

Tamanna Khatun

University of Liberal Arts Bangladesh

Date: April 21, 2026

Dedicated to my
hardworking parents,
whose support and
sacrifices brought me this
far.

– *Tamanna Khatun*

ABSTRACT

Brain tumor classification is a critical task in medical imaging because early and accurate detection directly influences treatment decisions. This thesis develops and evaluates a lightweight hybrid quantum-classical convolutional neural network for brain tumor MRI classification. The model combines a shallow grayscale CNN encoder, a six-dimensional bottleneck, batch normalization, and a six-qubit parameterized quantum circuit with ring entanglement before final classification. It is evaluated in two settings: binary tumor versus non-tumor detection and four-class classification of glioma, meningioma, pituitary, and non-tumor MRI images. To make the evaluation clearer, the hybrid model is compared with a matched lightweight classical baseline that uses the same preprocessing, encoder, bottleneck size, optimizer, data split, and training protocol, replacing only the quantum layer with a direct classical classifier. Experimental results show that the hybrid model improves binary accuracy from 85.08% to 93.29% and weighted F1-score from 0.8505 to 0.9327. In the four-class task, accuracy improves from 77.19% to 83.50% and weighted F1-score from 0.7612 to 0.8305. The results suggest that a compact parameterized quantum circuit can improve a small CNN bottleneck model under this controlled setup, while the remaining errors show that fine-grained tumor subtype classification is still more difficult than binary screening.

Keywords: Lightweight Hybrid QCNN, Parameterized Quantum Circuit, Brain Tumor Classification, Binary Classification, Multiclass Classification, Quantum Machine Learning

Contents

1	Introduction	1
1.1	Background and Motivation	1
1.2	Problem Statement	2
1.3	Scope, Significance, and Contributions	2
1.3.1	Scope of the Study	2
1.3.2	Significance of the Study	3
1.3.3	Contributions of the Thesis	3
1.4	Thesis Organization	4
2	Literature Review	5
2.1	Brain Tumor Classification in MRI Imaging	5
2.2	Deep Learning for Medical Image Classification	6
2.3	Quantum Computing Fundamentals	7
2.4	Hybrid Quantum-Classical Models and Related Work	7
2.5	Research Gap	11
2.6	Chapter Summary	12
3	Theoretical and Technical Background	13
3.1	Brain MRI and Tumor Classification Task	13
3.2	Convolutional Neural Networks	13
3.3	Parameterized Quantum Circuits	14
3.4	Hybrid CNN-QNN Architecture	15
3.5	Learning, Optimization, and Evaluation Metrics	16

4	Methodology	18
4.1	Design and Workflow Overview	18
4.2	Quantum Circuit Design	19
4.3	Hybrid Model Architecture	19
4.4	Binary and Multiclass Classification Pipelines	21
4.5	Training Configuration and Reproducibility	22
4.6	Evaluation Protocol and Constraints	23
5	Implementation and Experimental Setup	24
5.1	Dataset Description and Preprocessing	24
5.2	Data Splitting and Class Distribution	25
6	Results and Analysis	27
6.1	Binary Classification: Results and Analysis	27
6.2	Multiclass Classification: Results and Analysis	29
6.3	Comparative Discussion: Binary vs Multiclass	31
6.4	Result Reliability and Chapter Summary	31
7	Discussion	33
7.1	Interpretation of Findings	33
7.2	Relation to the Literature	33
7.3	Strengths and Limitations of the Study	34
7.4	Validity, Implications, and Directions for Improvement	34
7.5	Chapter Summary	35
8	Conclusion and Future Work	36
8.1	Summary of Findings and Achievement of Objectives	36
8.2	Research Contributions	36
8.3	Recommendations for Future Research	37
	References	38

List of Figures

4.1	Hybrid CNN-QNN architecture	20
5.1	Representative MRI samples from the multiclass dataset	25
6.1	Training curves for binary classification	27
6.2	Binary confusion matrices	28
6.3	Binary classification metric comparison	28
6.4	Training curves for multiclass classification	29
6.5	Multiclass confusion matrices	29
6.6	Multiclass classification metric comparison	30

List of Tables

2.1	Reviewed studies: model, parameters, and datasets.	9
2.2	Comparison of closely related brain tumor and medical imaging studies.	9
2.3	Concise summary of reviewed papers.	10
3.1	Classification settings used in this thesis	14
3.2	Parameterized quantum circuit used in this thesis	15
3.3	End-to-end architecture summary	16
3.4	Learning and evaluation framework	17
4.1	Architecture summary used in Chapter 4	20
4.2	Comparison of the two experiment pipelines	21
4.3	Training configuration shared by both experiments	22
4.4	Evaluation outputs and methodological constraints	23
5.1	Source datasets used in the methodology	24
5.2	Final split sizes used in the experiments	26
6.1	Binary test metrics	28
6.2	Multiclass test metrics	30
6.3	Comparison of binary and multiclass performance	31

Chapter 1

Introduction

1.1 Background and Motivation

Accurate and timely diagnosis of brain tumors is essential because treatment decisions often depend on early detection and precise identification of tumor types. Among the available imaging techniques, Magnetic Resonance Imaging (MRI) plays a crucial role due to its ability to provide detailed visualization of soft tissues without requiring invasive procedures. However, interpreting MRI scans remains challenging, particularly when tumor boundaries are unclear, visual characteristics overlap between different tumor categories, or rapid clinical assessment is required (Darwish et al., 2023; Amin et al., 2022).

To address these challenges, research in automated brain tumor analysis has progressively shifted from traditional image-processing pipelines toward data-driven learning approaches. Earlier studies focused on combining preprocessing techniques, segmentation methods, and quantum-inspired optimization strategies, whereas more recent work has explored transfer learning, quantum variational classifiers, and hybrid quantum-classical convolutional architectures for MRI-based tumor classification (Darwish et al., 2023; Amin et al., 2022; Khan et al., 2024; Muniasamy et al., 2025). Existing studies further indicate that the difficulty of brain tumor classification varies considerably depending on the task, with multiclass classification generally proving more challenging than binary tumor-versus-nontumor detection (Khan et al., 2024; Muniasamy et al., 2025; Nugraha et al., 2025).

The primary motivation for this thesis stems from recent advances in hybrid quantum-classical image classification. Previous research suggests that compact

parameterized quantum circuits can effectively complement classical feature extractors by capturing complex nonlinear relationships within learned representations, while maintaining relatively lightweight model architectures suitable for near-term quantum computing environments (Fan et al., 2024; Hafeez et al., 2024; Cheng et al., 2024). Building on this idea, this thesis investigates a compact CNN-QNN framework for MRI-based brain tumor classification within a controlled experimental setting. In addition, a lightweight classical baseline is incorporated to enable a direct and fair comparison of performance and efficiency.

1.2 Problem Statement

Despite the recent growth of hybrid quantum-classical learning, the practical behavior of these models in brain tumor MRI classification is still not easy to judge. Existing studies differ in label definitions, dataset composition, preprocessing choices, circuit design, and evaluation protocol, so reported results are not directly comparable (Amin et al., 2022; Khan et al., 2024; Muniasamy et al., 2025; Nugraha et al., 2025). In addition, some studies report promising hybrid performance, whereas others show that strong classical baselines can still perform better in multiclass settings (Muniasamy et al., 2025; Gencer and Gencer, 2025).

The central problem addressed in this thesis is therefore the need for a clear, controlled evaluation of a lightweight hybrid quantum-classical model across both binary and multiclass brain tumor classification tasks. To address this problem, the thesis evaluates a single hybrid CNN-QNN architecture and a matched lightweight classical baseline under the same preprocessing, training, and testing workflow. This makes it possible to discuss both task difficulty and the observed contribution of the quantum layer.

1.3 Scope, Significance, and Contributions

1.3.1 Scope of the Study

This study focuses on the MRI datasets and implementation available in the thesis workspace. Two balanced image-folder datasets are used: a binary dataset with tumor and not_tumor classes, and a multiclass dataset with glioma, meningioma,

notumor, and pituitary classes. Images are converted to grayscale, resized to 180×180 , normalized, and processed by a shallow CNN whose latent representation is projected to six features before entering a six-qubit parameterized quantum circuit. The study is implemented with PyTorch and Qiskit and is evaluated on held-out test data using accuracy, precision, recall, F1-score, weighted F1-score, and confusion matrices. A matched lightweight classical baseline is included; real quantum hardware, cross-dataset validation, and strong pretrained CNN baselines are outside the scope of this thesis.

1.3.2 Significance of the Study

The importance of this thesis lies in its systematic examination of hybrid quantum-classical learning for brain MRI classification across tasks of varying complexity. Unlike many studies that simultaneously alter datasets, architectures, or training settings, this work maintains a consistent experimental configuration by keeping the encoder design, bottleneck dimension, optimization strategy, and data-splitting protocol unchanged. This controlled setup enables a clearer comparison between the hybrid model and its classical counterpart, providing practical insight into the contribution of compact quantum layers to medical image classification. Furthermore, the study serves as a reproducible application of recent advances in hybrid quantum machine learning within the context of brain tumor diagnosis (Fan et al., 2024; Hafeez et al., 2024; Cheng et al., 2024).

1.3.3 Contributions of the Thesis

The primary contributions of this thesis are summarized as follows:

1. Development of a compact hybrid CNN-QNN architecture for brain MRI classification that combines a six-feature classical bottleneck with a six-qubit parameterized quantum circuit.
2. Design of a lightweight classical baseline with a comparable structure, enabling a direct assessment of the impact of replacing the conventional classification head with a quantum layer.
3. Comprehensive evaluation of both binary and multiclass brain tumor classification tasks within a unified experimental framework.

4. Establishment of a reproducible research pipeline through controlled data handling, checkpoint-based training, and standardized performance evaluation procedures suitable for thesis-level experimentation.

1.4 Thesis Organization

The remainder of this thesis is organized as follows. Chapter 2 reviews related work on brain tumor MRI classification and hybrid quantum-classical learning. Chapter 3 presents the theoretical background of the proposed approach. Chapter 4 describes the research methodology, including data preparation, model design, training, and evaluation. Chapter 5 outlines the implementation and experimental setup. Chapter 6 presents and analyzes the experimental results. Chapter 7 discusses the findings, limitations, and implications of the study. Finally, Chapter 8 concludes the thesis and suggests directions for future research.

Chapter 2

Literature Review

2.1 Brain Tumor Classification in MRI Imaging

Brain tumor classification using MRI has become a major area of research because MRI provides detailed, noninvasive visualization of brain anatomy and tumor characteristics. Most classification studies focus on identifying common tumor categories such as glioma, meningioma, and pituitary tumors, while some datasets also include a no-tumor class. Early approaches typically relied on multi-stage processing pipelines involving image preprocessing, segmentation, handcrafted feature extraction, and optimization techniques. For instance, [Darwish et al. \(2023\)](#) proposed a quantum-inspired level-set method to improve tumor localization prior to classification. In recent years, however, the field has increasingly adopted deep learning-based approaches that learn feature representations directly from images, reducing the need for manual feature engineering and enabling more streamlined end-to-end classification systems ([Amin et al., 2022](#); [Khan et al., 2024](#); [Muniasamy et al., 2025](#); [Gencer and Gencer, 2025](#); [Ait Haddou and Bennai, 2025](#)).

Existing research also highlights the significant influence of dataset design and task formulation on classification performance. Many studies focus on three-class classification using contrast-enhanced T1-weighted MRI scans, while others incorporate an additional no-tumor category or combine classification with segmentation tasks ([Amin et al., 2022](#); [Nugraha et al., 2025](#); [Gencer and Gencer, 2025](#); [Ajlouni et al., 2023](#)). As a result, reported outcomes are affected not only by model architecture but also by factors such as the number of classes, label definitions, and data distribution. Several studies have reported that distinguishing between glioma and meningioma remains particularly challenging due to similarities in their visual

appearance, whereas pituitary tumors are generally easier to identify (Khan et al., 2024; Muniasamy et al., 2025). Although substantial progress has been achieved in MRI-based brain tumor classification, differences in datasets, evaluation protocols, and task settings continue to make direct comparisons between studies difficult.

2.2 Deep Learning for Medical Image Classification

Deep learning, particularly Convolutional Neural Networks (CNNs), has become the leading approach for medical image classification due to its ability to automatically learn hierarchical feature representations directly from image data. In the context of brain MRI analysis, CNN-based models have been widely applied to tasks such as tumor detection, tumor subtype classification, and feature extraction for clinical decision support (Amin et al., 2022; Muniasamy et al., 2025; Gencer and Gencer, 2025; Ajlouni et al., 2023). Transfer learning and lightweight CNN architectures are especially popular because medical imaging datasets are typically smaller than large-scale natural image datasets and often suffer from class imbalance.

The literature reviewed in this study indicates that modern deep-learning models continue to serve as strong benchmarks for evaluating newer approaches. For example, Gencer and Gencer (2025) combined EfficientNetB0 with a quantum genetic feature selection method and achieved competitive performance on multiclass brain tumor classification tasks. Likewise, Muniasamy et al. (2025) reported that DenseNet121 outperformed a hybrid quantum-assisted model under their experimental setting. Although Ajlouni et al. (2023) observed that hybrid architectures may offer benefits such as improved training behavior and convergence, their findings also suggest that the classical feature extraction component remains responsible for much of the model’s predictive capability. These observations highlight the importance of comparing quantum-enhanced approaches against carefully designed deep-learning baselines rather than evaluating them independently.

Another recurring theme in the literature is the influence of data-related factors on model performance. Brain MRI datasets are often limited in size, class distributions are frequently imbalanced, and preprocessing protocols vary considerably across studies. For instance, Nugraha et al. (2025) demonstrated that oversampling techniques significantly improved QCNN performance, particularly for binary

classification tasks. Such findings emphasize that successful medical image classification depends not only on model architecture but also on the quality of data preparation, class balancing strategies, and evaluation procedures.

2.3 Quantum Computing Fundamentals

The quantum computing concepts most relevant to this thesis include qubits, superposition, entanglement, parameterized quantum circuits (PQCs), and measurement-based output generation. Unlike classical bits, which exist only in one of two states, qubits can represent information in a superposition of states. Furthermore, entanglement enables interactions between qubits, allowing complex relationships to be modeled within the quantum system. In quantum machine learning applications, classical features are typically compressed into a lower-dimensional representation and encoded into qubits using feature maps or rotation gates. The encoded quantum state is then processed through trainable variational gates and entangling operations, after which measurements are performed to produce outputs that can be used by classical layers (Fan et al., 2024; Hafeez et al., 2024; Ranga et al., 2024; Prajapat et al., 2025; Cheng et al., 2024).

This general workflow is common across many hybrid quantum-classical learning models. For example, Hafeez et al. (2024) and Prajapat et al. (2025) employ rotation-based data encoding together with shallow PQCs, while Ranga et al. (2024) and Cheng et al. (2024) utilize variational quantum circuits whose measured expectation values are passed back to classical components for further processing. From a practical perspective, these design choices are largely influenced by current hardware limitations. Since existing quantum devices and simulators support only a limited number of qubits and relatively shallow circuits, most effective hybrid architectures rely on compact latent representations, lightweight quantum modules, and classical post-processing stages.

2.4 Hybrid Quantum-Classical Models and Related Work

Hybrid quantum-classical learning has emerged as one of the most practical approaches for image classification because it combines the strengths of both com-

puting paradigms. In these architectures, high-dimensional feature extraction is performed by a classical neural network, while a low-dimensional quantum module is used to further transform the extracted features. This strategy reduces qubit requirements and makes the models more suitable for current quantum hardware and simulation environments (Fan et al., 2024; Hafeez et al., 2024; Shi et al., 2024; Ranga et al., 2024; Prajapat et al., 2025; Cheng et al., 2024). Studies on general image-classification tasks have reported promising results, particularly for binary classification and reduced-input problems. However, they also indicate that performance tends to become less consistent as the number of classes and the complexity of the feature space increase (Hafeez et al., 2024; Shi et al., 2024; Ranga et al., 2024; Prajapat et al., 2025).

A similar trend can be observed in medical imaging applications. For instance, Amin et al. (2022) integrated InceptionV3, SegNet, and a quantum variational classifier for brain tumor analysis. Likewise, Khan et al. (2024) and Nugraha et al. (2025) applied QCNN-based architectures to MRI brain tumor classification and found that these models generally performed better on simpler classification tasks than on more challenging multiclass problems. Other studies have continued to extend the hybrid-learning paradigm. Ajlouni et al. (2023) introduced an adaptive hybrid quantum CNN, while Ait Haddou and Bennai (2025) explored explainable hybrid models for brain tumor classification. In contrast, Gencer and Gencer (2025) incorporated quantum-inspired optimization techniques into an otherwise classical framework, illustrating that the term “quantum” in the literature may refer either to an actual quantum processing layer or to optimization methods inspired by quantum principles. Collectively, these studies suggest that the focus of current research has shifted from demonstrating the feasibility of hybrid models to understanding the conditions under which they can remain compact, trainable, and practically useful for real-world medical imaging tasks.

Table 2.1 summarizes the main reviewed studies, with emphasis on model type, reported parameter information, and dataset choice. Tables 2.2 and 2.3 then narrow the focus to the closest MRI-related studies and the broader paper set discussed in this chapter.

Table 2.1: Reviewed studies: model, parameters, and datasets.

Citation	Model	Parameters	Dataset
Amin et al., 2022	InceptionV3 + QVC + SegNet	4 qubits; 6 layers	Kaggle brain MRI; BraTS 2020
Khan et al., 2024	HQCNN	4 qubits	3064 CE-T1 MRI slices; 3-class
Muniasamy et al., 2025	HQC-CNN; DenseNet121 baseline	NR	3064 MRI images from 233 patients; 3-class
Nugraha et al., 2025	QCNN	4 qubits; depth 4	3264 MRI images; binary and 4-class
Gencer and Gencer, 2025	EfficientNetB0 + QGA-FS	NR	Figshare MRI (3-class); Kaggle MRI (4-class)
Ajilouni et al., 2023	Adaptive HQCNN	8256 params; 4 qubits	Kaggle brain tumor MRI; REMBRANDT
Ait Haddou and Bennai, 2025	HQCM-EBTC	5 qubits; depth 2	7576 MRI images; 4-class
Hafeez et al., 2024	H-QNN	2 qubits	Car/Bike; Football/Rugby; Waste
Shi et al., 2024	PCA-QNN; CNN-QNN	NR	MNIST; FashionMNIST
Ranga et al., 2024	Hybrid QNN	NR	MNIST 0 vs. 1
Prajapat et al., 2025	QCNN	45 params	MNIST 0 vs. 1; Wine
Fan et al., 2024	QC-CNN	10 qubits	Overhead-MNIST; So2Sat LCZ42; PatternNet; RSI-CB256; NaSC-TG2

Note. In Table 2.1, NR = not reported in the paper.

Table 2.2: Comparison of closely related brain tumor and medical imaging studies.

Citation	Task	Data	Approach	Main finding
Amin et al., 2022	Multiclass MRI	4-class; multi-dataset	TL + QVC + segmentation	Strong, but multi-stage.
Khan et al., 2024	3-class MRI	CE-T1; 3-class	HQCNN	Glioma and meningioma remain harder to separate.
Muniasamy et al., 2025	3-class MRI	3064 MRI; 233 patients	HQC-CNN vs. DenseNet121	The classical baseline still performs better.
Nugraha et al., 2025	Binary and 4-class MRI	3264 MRI; oversampled	QCNN	Binary classification is easier than multiclass.

Continued on next page

Citation	Task	Data	Approach	Main finding
Gencer and Gencer, 2025	3-class and 4-class MRI	Figshare + Kaggle	EfficientNetB0 + QGA-FS	Quantum support is mainly in feature selection.
Ajlouni et al., 2023	Medical image classification	MRI and medical sets	Adaptive HQCNN	Faster convergence; competitive accuracy.

Table 2.3: Concise summary of reviewed papers.

Citation	Focus	Model	Task	Relevance
Darwish et al., 2023	MRI tumor analysis	Quantum-inspired framework	Detection	Early non-deep MRI baseline.
Amin et al., 2022	Brain tumor detection	Hybrid DL + QVC	Multiclass MRI	Hybrid medical pipeline.
Khan et al., 2024	Brain tumor diagnosis	HQCNN	Multiclass MRI	Close MRI comparator.
Fan et al., 2024	Image classification	QC-CNN	General image	Hybrid design rationale.
Hafeez et al., 2024	Binary classification	H-QNN	Binary image	Low-qubit binary evidence.
Muniasamy et al., 2025	MRI tumor classification	HQC-CNN	Multiclass MRI	Hybrid vs. classical baseline.
Nugraha et al., 2025	MRI detection and classification	QCNN	Binary + multiclass MRI	Same domain, both settings.
Gencer and Gencer, 2025	Brain tumor classification	EffNetB0 + QGA	Multiclass MRI	Quantum-inspired selection.
Shi et al., 2024	Image multi-classification	CNN-QNN	Multiclass image	Shows multiclass difficulty.
Ranga et al., 2024	MNIST binary classification	Hybrid QNN	Binary image	Compact hybrid baseline.
Prajapat et al., 2025	Efficient image classification	QCNN	Binary image	Parameter-efficient PQC.
Ajlouni et al., 2023	Medical image diagnosis	Adaptive HQCNN	Medical image	Medical hybrid convergence.
Cheng et al., 2024	Learnability analysis	Hybrid QC-CNN	Image classification	Training-behavior support.
Ait Haddou and Bennai, 2025	Explainable tumor classification	HQCM-EBTC	Multiclass MRI	Adds explainability to hybrid MRI classification.

Taken together, Tables 2.1–2.3 show that only a minority of the reviewed studies explicitly report parameter counts, even when efficiency is presented as an advan-

tage. They also show that comparison is often complicated by changing datasets, task settings, and pipeline complexity.

2.5 Research Gap

The literature reviewed in this chapter reveals several limitations that motivate the present study. First, many existing works focus exclusively on either binary or multiclass classification, making it difficult to assess how a single hybrid architecture performs as classification tasks become more complex (Hafeez et al., 2024; Nugraha et al., 2025; Shi et al., 2024; Ranga et al., 2024; Prajapat et al., 2025). Second, comparisons across studies are often challenging because of differences in datasets, class definitions, image resolutions, preprocessing procedures, and class-balancing techniques (Amin et al., 2022; Nugraha et al., 2025; Gencer and Gencer, 2025; Ajlouni et al., 2023; Ait Haddou and Bennai, 2025). These variations make it difficult to determine whether performance differences arise from the model itself or from the experimental setup.

Another limitation is that several studies employ relatively complex frameworks involving transfer learning, segmentation modules, explainability components, or feature-selection techniques. While such additions can improve performance, they also make it harder to isolate the specific contribution of the hybrid quantum-classical classifier (Amin et al., 2022; Gencer and Gencer, 2025; Ait Haddou and Bennai, 2025). Furthermore, despite the growing interest in lightweight hybrid architectures, relatively few studies investigate whether a compact model with a fixed quantum bottleneck can maintain its effectiveness across both binary and multiclass brain tumor classification tasks. Existing work also reports model complexity and computational requirements inconsistently, making efficiency-focused comparisons difficult.

These gaps highlight the need for a controlled and reproducible study that evaluates the same compact hybrid architecture across different classification settings while maintaining consistent preprocessing, balanced data handling, and standardized evaluation metrics.

2.6 Chapter Summary

This chapter reviewed the key literature related to MRI-based brain tumor classification, deep learning for medical image analysis, fundamental quantum computing concepts, and recent developments in hybrid quantum-classical learning. The reviewed studies demonstrate that CNN-based models continue to provide strong benchmarks for medical image classification, while hybrid approaches show promise when quantum components are kept compact and integrated within carefully designed workflows. The literature also indicates that multiclass brain tumor classification remains more challenging than binary classification and that variations in datasets and preprocessing procedures often complicate direct comparisons between studies. These findings motivate the development of a compact, reproducible hybrid CNN-QNN framework evaluated under a consistent experimental setting for both binary and multiclass MRI classification tasks.

Chapter 3

Theoretical and Technical Background

3.1 Brain MRI and Tumor Classification Task

MRI is the main imaging modality used in this thesis because it provides strong soft-tissue contrast and supports noninvasive brain examination. In tumor classification, MRI is useful for recognizing structural abnormalities, but interpretation remains difficult because tumor boundaries, texture, and contrast patterns vary across cases ([Amin et al., 2022](#); [Khan et al., 2024](#); [Muniasamy et al., 2025](#)).

- **Binary task:** decide whether an MRI image belongs to tumor or not_tumor.
- **Multiclass task:** classify an MRI image as glioma, meningioma, notumor, or pituitary.
- **Why both tasks matter:** binary classification measures tumor detection ability, while multiclass classification tests finer subtype discrimination.
- **Expected difficulty:** multiclass classification is harder because some tumor types, especially glioma and meningioma, share similar visual patterns ([Khan et al., 2024](#); [Muniasamy et al., 2025](#); [Nugraha et al., 2025](#)).

3.2 Convolutional Neural Networks

CNNs are designed for image data and learn local spatial patterns through trainable filters. In medical imaging, they are effective because they capture edges, texture

Table 3.1: Classification settings used in this thesis

Setting	Classes	Main objective	Difficulty
Binary	tumor, not_tumor	Detect tumor presence	Lower
Multiclass	glioma, meningioma, notumor, pituitary	Separate tumor subtypes and normal cases	Higher

changes, and higher-level visual structure without handcrafted features ([Amin et al., 2022](#); [Gencer and Gencer, 2025](#); [Ajlouni et al., 2023](#)).

- **Convolution:** learns local pattern detectors by sliding kernels across the image.
- **ReLU activation:** introduces nonlinearity so the network can model complex image structure.
- **Max pooling:** reduces spatial resolution while preserving strong feature responses.
- **Dropout:** reduces overfitting by randomly deactivating part of the representation during training.
- **Role in this thesis:** compress the 180×180 grayscale MRI input into a compact feature vector before quantum processing.

The classical feature extractor in this thesis is intentionally shallow. It uses two small convolution blocks and a single projection layer so that most of the model capacity remains compact rather than shifting toward a large classical backbone.

3.3 Parameterized Quantum Circuits

Parameterized quantum circuits (PQCs) are variational circuits with encoded inputs and trainable parameters. In hybrid learning, classical features are mapped to qubits, transformed through trainable gates, and then measured to produce outputs for prediction ([Fan et al., 2024](#); [Hafeez et al., 2024](#); [Ranga et al., 2024](#); [Prajapat et al., 2025](#); [Cheng et al., 2024](#)).

- **Qubits:** hold the encoded latent features in quantum state form.
- **Encoding layer:** maps classical values into the circuit through rotation gates.
- **Entanglement:** allows interactions among encoded features rather than independent processing.
- **Trainable block:** adapts the quantum transformation during optimization.
- **Measurement:** returns expectation values that can be consumed by a classical layer.

Table 3.2: Parameterized quantum circuit used in this thesis

Stage	Operation	Purpose
Initialization	Hadamard on 6 qubits	Create superposition before data encoding
Data encoding	R_Y rotations	Map the 6 classical features to qubit states
Entanglement	Ring CZ gates	Model interactions among neighboring features
Trainable layer	R_Y rotations	Learn a task-specific quantum transformation
Readout	Pauli-Z expectation values	Return 6 outputs to the classical classifier

This PQC is shallow but practical. It matches the six-dimensional latent vector produced by the classical projection layer and keeps the quantum block small enough for simulation-based hybrid training.

3.4 Hybrid CNN-QNN Architecture

The thesis model is hybrid because it combines a classical CNN front end with a quantum layer and a final classical classifier. This follows the common hybrid design principle that classical layers handle high-dimensional image extraction, while the quantum block operates on a compact latent representation (Fan et al., 2024; Hafeez et al., 2024; Shi et al., 2024; Cheng et al., 2024).

- **Input:** grayscale MRI image resized to 180×180 .

- **Classical front end:** two convolution layers, ReLU, max pooling, and dropout.
- **Compression stage:** flattening followed by a linear layer from 7056 features to 6 latent values and batch normalization.
- **Quantum connector:** tanh-bounded latent vector passed into a 6-qubit QNN layer wrapped for PyTorch training.
- **Output head:** final linear layer producing 2 logits for binary classification or 4 logits for multiclass classification.
- **Baseline comparison:** a matched classical version keeps the same CNN encoder and six-dimensional bottleneck but replaces the QNN layer with a direct linear classifier.

Table 3.3: End-to-end architecture summary

Component	Configuration	Role
Image input	$1 \times 180 \times 180$	Standardized grayscale MRI input
Conv block 1	$1 \rightarrow 2$, kernel 5×5 , ReLU, max pool	Extract low-level spatial features
Conv block 2	$2 \rightarrow 4$, kernel 5×5 , ReLU, max pool	Build higher-level feature maps
Regularization	Dropout = 0.25	Reduce overfitting
Projection	Flatten 7056 $\rightarrow$ 6, batch normalization, tanh	Create a compact latent vector
Quantum layer	6-qubit EstimatorQNN	Transform latent features quantum mechanically
Classifier head	Linear $6 \rightarrow 2$ or $6 \rightarrow 4$	Produce task-specific logits

3.5 Learning, Optimization, and Evaluation Metrics

The model is trained as a supervised classifier. The quantum layer is differentiable through the Qiskit-PyTorch interface, so classical and quantum parameters can be optimized in one end-to-end training loop. The matched baseline uses the same learning setup, which makes the comparison focused on the QNN stage rather than on preprocessing or optimizer changes.

- **Loss:** cross-entropy over class logits.
- **Gradient flow:** backpropagation through CNN layers, projection layer, quantum layer, and classifier head.
- **Optimizer:** AdamW with learning rate 0.001 and weight decay 10^{-4} .
- **Scheduler:** ReduceLROnPlateau with factor 0.5 and patience 3.
- **Training control:** checkpointing and early stopping with patience 5.
- **Main metrics:** accuracy, precision, recall, F1-score, weighted F1-score, and confusion matrix analysis.
- **Why multiple metrics matter:** accuracy alone can hide class-wise errors, especially in multiclass MRI classification ([Amin et al., 2022](#); [Khan et al., 2024](#); [Nugraha et al., 2025](#); [Ajlouni et al., 2023](#)).

Table 3.4: Learning and evaluation framework

Aspect	Choice	Purpose
Loss function	Cross-entropy	Optimize categorical classification
Optimization	AdamW	Stable adaptive updates with regularization
Learning-rate control	ReduceLROnPlateau	Slow the update rate when validation loss stalls
Training stopping	Early stopping + checkpoints	Preserve the best model state and reduce overfitting
Primary metrics	Accuracy, precision, recall, F1-score, weighted F1-score	Evaluate overall and class-wise performance
Error analysis	Confusion matrix	Reveal class confusion patterns
Baseline control	Matched classical bottleneck model	Isolate the effect of the QNN layer

In the multiclass setting, macro and weighted averages are both informative because they summarize class-level performance from complementary perspectives. The confusion matrix is particularly important because it shows whether the model fails uniformly or struggles with specific tumor pairs such as glioma and meningioma.

Chapter 4

Methodology

4.1 Design and Workflow Overview

This thesis follows an implementation-oriented experimental design. A hybrid CNN-QNN framework is trained and evaluated under two related settings: binary tumor detection and multiclass tumor-type classification. A matched lightweight classical CNN baseline is also trained for each setting. The design keeps preprocessing, CNN encoder, bottleneck size, optimizer, data split, and training controls consistent so that observed differences can be interpreted in relation to task difficulty and to the added QNN layer (Fan et al., 2024; Ajlouni et al., 2023; Cheng et al., 2024).

- **Research mode:** Controlled implementation, training, and test-set evaluation.
- **Experiment tracks:** Binary (tumor vs. not_tumor) and multiclass (glioma, meningioma, notumor, pituitary).
- **Core design choice:** Reuse the same compact backbone in both the hybrid model and the classical baseline.
- **Outcome focus:** Compare predictive behavior, class-wise errors, model compactness, and reproducibility under a lightweight setting.

4.2 Quantum Circuit Design

Both experiment tracks use the same six-qubit parameterized quantum circuit. The circuit receives the six-dimensional latent vector produced by the classical encoder, applies feature encoding and entanglement, and returns six expectation values for the final classifier. This follows the compact hybrid pattern commonly used in image-oriented quantum-classical models (Fan et al., 2024; Hafeez et al., 2024; Cheng et al., 2024).

- **Qubits:** 6
- **Input parameters:** $x = (x_1, \dots, x_6)$
- **Trainable parameters:** $\theta = (\theta_1, \dots, \theta_6)$
- **Encoding gate:** $R_Y(x_i)$
- **Entanglement:** ring-connected CZ gates
- **Trainable block:** $R_Y(\theta_i)$
- **Readout:** one Pauli-Z expectation value per qubit
- **Execution backend:** `StatevectorEstimator`

The ring entanglement pattern is used because it keeps the circuit shallow while allowing every qubit to participate in a symmetric interaction structure. Since the circuit is simulated with a statevector estimator, the methodology evaluates idealized hybrid learning behavior rather than noisy hardware execution.

4.3 Hybrid Model Architecture

The implemented hybrid model combines a shallow CNN front end with a quantum bottleneck and a task-specific output head. The classical encoder reduces the 180×180 MRI input to a six-dimensional latent vector, normalizes it with `BatchNorm1d`, bounds it with `tanh`, and passes it into the QNN through `EstimatorQNN` and `TorchConnector`. The baseline keeps the same encoder and bottleneck but removes the QNN stage.

Table 4.1: Architecture summary used in Chapter 4

Stage	Configuration	Role
Input	$1 \times 180 \times 180$	Standardized image input grayscale MRI
Conv block 1	Conv2d(1, 2, 5) ReLU + pool	Low-level feature extraction
Conv block 2	Conv2d(2, 4, 5) ReLU + pool	Higher-level feature extraction
Regularization	Dropout = 0.25	Reduce overfitting
Projection	Flatten 7056 $\rightarrow$ 6 batch normalization + tanh	Create compact latent features
Quantum block	6-qubit EstimatorQNN	Nonlinear feature transformation
Classifier head	Linear 6 $\rightarrow$ C	Output class logits




Figure 4.1: Hybrid CNN-QNN architecture

This architecture is intentionally compact. The same backbone is reused in both experiment tracks, and only the final output dimension changes with the number of classes. In the matched baseline, the six-dimensional bottleneck is connected directly to the output head. The result is a lightweight methodology that supports both a binary-versus-multiclass comparison and a hybrid-versus-classical comparison.

4.4 Binary and Multiclass Classification Pipelines

The binary and multiclass pipelines share the same preprocessing, optimizer, scheduler, and evaluation logic. Their main differences lie in the dataset root, class labels, and output-head dimension. Within each task, the hybrid and baseline models use the same CNN encoder and bottleneck; only the QNN transformation is replaced in the baseline.

Table 4.2: Comparison of the two experiment pipelines

Track	Classes	Data root	Head	Main objective
Binary	tumor, not_tumor	binary	$6 \rightarrow 2$	Detect tumor presence
Multiclass	glioma, meningioma, notumor, pituitary	multiclass	$6 \rightarrow 4$	Separate tumor subtypes and normal cases

The common pipeline can be summarized as follows:

- load the selected dataset root and define class names
- apply training and evaluation transforms
- build train, validation, and test loaders
- instantiate the shared hybrid model with the required output size
- instantiate the matched classical baseline with the same bottleneck and output size
- train with validation monitoring and checkpoint saving
- restore the best validation-loss weights
- evaluate on the hold-out test set and export artifacts

4.5 Training Configuration and Reproducibility

Training settings are kept fixed across both tracks so that performance differences reflect the classification task rather than a changed optimization setup.

Table 4.3: Training configuration shared by both experiments

Aspect	Setting	Purpose
Batch size	$B = 10$	Small-batch supervised training
Optimizer	AdamW	Stable adaptive updates with weight decay
Learning rate	$\eta = 0.001$	Base optimization step size
Weight decay	10^{-4}	Mild regularization
Loss function	Cross-entropy	Multi-class logit optimization
Scheduler	ReduceLROnPlateau	Lower learning rate after validation stagnation
Scheduler patience	3 epochs	Control scheduler sensitivity
Maximum epochs	20	Upper bound for training duration
Early stopping patience	5 epochs	Stop after sustained non-improvement

Reproducibility is strengthened through:

- a fixed seed of 42 for Python, NumPy, and PyTorch
- seeded data splitting and worker initialization
- automatic checkpoint loading when a resumable run exists
- checkpoint saving after every epoch
- restoration of the best validation-loss weights before final testing

This strategy makes interrupted runs recoverable and keeps the experimental path consistent across repeated executions.

4.6 Evaluation Protocol and Constraints

Final evaluation is performed only after training finishes or early stopping is triggered. The test set remains unseen during optimization.

- switch the model to evaluation mode
- disable gradient computation with `torch.no_grad()`
- generate logits for each test batch
- convert logits to predicted labels using `argmax`
- compute accuracy, class-wise precision, recall, and F1-score
- generate confusion matrices and save output artifacts

Table 4.4: Evaluation outputs and methodological constraints

Component	Used in this thesis	Methodological note
Primary metrics	Accuracy, precision, recall, F1-score	Supports overall and class-wise analysis
Error analysis	Confusion matrix	Exposes class-specific confusion patterns
Saved artifacts	Checkpoints, weights, curves, confusion matrices, metric charts	Preserves the best run state and outputs
Quantum execution	StatevectorEstimator simulation	Idealized quantum behavior, not hardware noise
Data paths	Absolute Windows paths in scripts	Convenient locally, weaker portability
Baseline comparison	Matched lightweight classical baseline	Isolates the QNN layer under the same pipeline
Implementation record	Notebook-derived scripts and saved artifacts	Some naming inconsistency remains in exported scripts

The methodology is therefore reproducible and controlled, but it should be interpreted within the limits of simulation-based quantum execution, a compact shallow encoder, a single split, and a codebase that prioritizes experiment completion over packaging. These constraints do not invalidate the study; they define the scope within which the results should be read.

Chapter 5

Implementation and Experimental Setup

5.1 Dataset Description and Preprocessing

The experiments use two folder-structured MRI datasets stored as JPG images. The binary and multiclass datasets are separate, but both are balanced at the source-folder level. Across the two experiment tracks, the workspace contains 18,000 MRI images.

Table 5.1: Source datasets used in the methodology

Track	Classes	Training source	Test source	Total
Binary	tumor, not_tumor	4200 / class	1200 / class	10,800
Multiclass	glioma, meningioma, notumor, pituitary	1400 / class	400 / class	7,200

The data-preparation pipeline is identical in structure for both experiment tracks:

- convert each image to grayscale with one channel
- resize all images to 180×180
- convert images to tensors
- normalize tensors with mean $[0.5]$ and standard deviation $[0.5]$
- apply augmentation only to the training subset

The training transform uses random horizontal flipping with probability 0.5 and random rotation within $\pm 10^\circ$. Validation and test samples use deterministic preprocessing only. This separation improves generalization during learning while keeping evaluation stable and reproducible.




Figure 5.1: Representative MRI samples from the multiclass dataset

5.2 Data Splitting and Class Distribution

The original folder structure already provides a separate hold-out test set for each track. Validation data are created from the training folders using a seeded random split with:

$$\text{VAL_RATIO} = 0.15 \quad \text{SEED} = 42$$

- 85% of the source training data are used for model fitting.

- 15% are reserved for validation.
- Test folders remain untouched until final evaluation.
- An explicit assertion checks that train and validation indices do not overlap.

Class counts are audited after splitting to confirm that the balanced source folders remain reasonably balanced in the derived subsets. This is important because the split is random rather than stratified.

Table 5.2: Final split sizes used in the experiments


Dataset Type	Train	Validation	Test	Class balance
Binary	7140	1260	2400	Balanced at source and audited after split
Multiclass	4760	840	1600	Balanced at source and audited after split

Chapter 6

Results and Analysis

6.1 Binary Classification: Results and Analysis

The binary task compares the matched classical baseline with the hybrid CNN-QNN model on the same 2400-image hold-out test set. The baseline reached 85.08% test accuracy, while the hybrid model reached 93.29%, correctly classifying 2239 of 2400 images. Weighted F1-score increased from 0.8505 to 0.9327.



(a) Baseline

(b) Hybrid CNN-QNN

Figure 6.1: Training curves for binary classification

The baseline and hybrid confusion matrices are:

$$\text{Baseline} = \begin{bmatrix} 1081 & 119 \\ 239 & 961 \end{bmatrix}, \quad \text{Hybrid} = \begin{bmatrix} 1188 & 12 \\ 149 & 1051 \end{bmatrix}$$

with class order not_tumor, tumor. The hybrid model sharply reduces false positives from 119 to 12 and false negatives from 239 to 149. This improves tumor recall from 80.08% to 87.58% and preserves very high tumor precision at 98.87%.

- The hybrid model gives an 8.21-point accuracy gain over the matched baseline.
- The largest practical improvement is the reduction of tumor false negatives from 239 to 149.




Figure 6.2: Binary confusion matrices




Figure 6.3: Binary classification metric comparison

Table 6.1: Binary test metrics

Class / Average	Precision	Recall	F1-score	Support
Classical baseline				
not_tumor	0.8189	0.9008	0.8579	1200
tumor	0.8898	0.8008	0.8430	1200
Accuracy		0.8508		2400
Weighted average	0.8544	0.8508	0.8505	2400
Hybrid CNN-QNN				
not_tumor	0.8886	0.9900	0.9365	1200
tumor	0.9887	0.8758	0.9289	1200
Accuracy		0.9329		2400
Weighted average	0.9386	0.9329	0.9327	2400

- The binary result supports the usefulness of the compact QNN layer for coarse tumor detection under this controlled setup.

6.2 Multiclass Classification: Results and Analysis

The multiclass task is more difficult because the model must separate glioma, meningioma, notumor, and pituitary. On the 1600-image hold-out test set, the baseline reached 77.19% accuracy, while the hybrid model reached 83.50%, correctly classifying 1336 images. Weighted F1-score increased from 0.7612 to 0.8305.




Figure 6.4: Training curves for multiclass classification




Figure 6.5: Multiclass confusion matrices

The baseline and hybrid confusion matrices are:

$$\text{Baseline} = \begin{bmatrix} 238 & 79 & 74 & 9 \\ 49 & 220 & 98 & 33 \\ 3 & 6 & 388 & 3 \\ 6 & 5 & 0 & 389 \end{bmatrix}, \quad \text{Hybrid} = \begin{bmatrix} 265 & 69 & 61 & 5 \\ 23 & 285 & 74 & 18 \\ 2 & 5 & 393 & 0 \\ 2 & 5 & 0 & 393 \end{bmatrix}$$

with class order glioma, meningioma, notumor, pituitary. The hybrid model improves glioma recall from 59.50% to 66.25% and meningioma recall from 55.00% to 71.25%. Notumor and pituitary remain the easiest classes, each reaching 98.25% recall in the hybrid model.




Figure 6.6: Multiclass classification metric comparison

Table 6.2: Multiclass test metrics

Class / Average	Precision	Recall	F1-score	Support
Classical baseline				
glioma	0.8041	0.5950	0.6839	400
meningioma	0.7097	0.5500	0.6197	400
notumor	0.6929	0.9700	0.8083	400
pituitary	0.8963	0.9725	0.9329	400
Accuracy		0.7719		1600
Weighted average	0.7757	0.7719	0.7612	1600
Hybrid CNN-QNN				
glioma	0.9075	0.6625	0.7659	400
meningioma	0.7830	0.7125	0.7461	400
notumor	0.7443	0.9825	0.8470	400
pituitary	0.9447	0.9825	0.9632	400
Accuracy		0.8350		1600
Weighted average	0.8449	0.8350	0.8305	1600

- The hybrid model gives a 6.31-point accuracy gain over the matched baseline.
- The largest class-wise gain is in meningioma, where recall improves by 16.25 points.

- The hardest remaining class is glioma, mainly because it is still confused with meningioma and notumor.

6.3 Comparative Discussion: Binary vs Multiclass

The binary task remains easier than the multiclass task for both models. The hybrid model performs better than the matched baseline in both settings, but the four-class task still has lower accuracy and more class-specific errors.

Table 6.3: Comparison of binary and multiclass performance

Measure	Binary	Multiclass
Output classes	2	4
Test samples	2400	1600
Baseline accuracy	85.08%	77.19%
Hybrid accuracy	93.29%	83.50%
Accuracy gain	8.21 points	6.31 points
Hybrid weighted F1-score	0.9327	0.8305
Strongest hybrid class behavior	not_tumor recall 99.00%	pituitary F1- score 0.9632
Main remaining weakness	Tumor false nega- tives	Glioma recall and class overlap

The 9.79-point difference between the hybrid binary and hybrid multiclass accuracies shows the cost of moving from coarse detection to subtype discrimination. The baseline and hybrid results both support the same pattern: the compact architecture handles broad tumor detection more easily than fine-grained separation among visually similar tumor classes.

6.4 Result Reliability and Chapter Summary

Several features support result reliability:

- both tracks use balanced datasets and a separate hold-out test set
- the train-validation split is seeded and checked for overlap
- the baseline and hybrid models use the same preprocessing, optimizer, split, and training controls

- checkpointing and best-weight restoration reduce accidental reporting from a weak epoch

The findings should still be read within the study scope:

- the quantum layer is evaluated with ideal statevector simulation
- the baseline is lightweight and matched, not a strong pretrained CNN benchmark
- results are based on a single split rather than repeated multi-seed evaluation

Overall, the chapter shows that the hybrid CNN-QNN model improves over the matched lightweight baseline in both binary and multiclass MRI classification. The strongest result is binary tumor detection, while the main remaining challenge is fine-grained confusion among tumor subtypes.

Chapter 7

Discussion

7.1 Interpretation of Findings

The findings of this thesis indicate that the proposed hybrid CNN-QNN model consistently outperforms the matched lightweight classical baseline in both binary and multiclass brain tumor MRI classification. In the binary task, accuracy improved from (85.08%) to (93.29%), while in the multiclass task it increased from (77.19%) to (83.50%). As expected, binary tumor detection proved less challenging than four-class tumor classification.

A closer examination of class-level results provides additional insight. In the binary setting, the hybrid model reduced the number of missed tumor cases from (239) to (149). For multiclass classification, the highest performance was observed for the pituitary and notumor classes, whereas glioma remained the most difficult category to distinguish. These results suggest that the six-qubit quantum bottleneck enhances overall class separation but still struggles to capture subtle differences between visually similar tumor types.

Overall, the objectives of this thesis were successfully achieved through the development, evaluation, and analysis of both the hybrid model and its matched classical counterpart across two classification settings.

7.2 Relation to the Literature

The results align with several observations reported in previous studies. Consistent with earlier work, multiclass brain tumor classification was more challenging than binary classification (Khan et al., 2024; Nugraha et al., 2025). Similarly, pituitary

tumors were classified more reliably than glioma and meningioma, reflecting trends reported in the literature (Khan et al., 2024; Muniasamy et al., 2025).

Some studies have reported higher overall performance (Amin et al., 2022; Gencer and Gencer, 2025; Ajlouni et al., 2023); however, direct comparison is difficult because of differences in datasets, preprocessing methods, and model complexity. Within its controlled experimental setting, this thesis provides additional evidence that compact hybrid quantum-classical models can effectively support medical image classification and can improve upon a comparable lightweight classical baseline (Fan et al., 2024; Hafeez et al., 2024; Cheng et al., 2024).

7.3 Strengths and Limitations of the Study

This study offers several strengths. A single hybrid architecture was evaluated across both binary and multiclass tasks, and a matched lightweight classical baseline was used under the same experimental conditions. In addition, balanced datasets, separate test sets, seeded data handling, checkpointed training, and saved evaluation artifacts contributed to a reproducible workflow.

Despite these strengths, several limitations should be acknowledged. The classical encoder remains relatively shallow, and the baseline, while appropriate for controlled comparison, does not represent a strong pretrained CNN benchmark. Furthermore, the quantum component was evaluated using an ideal simulator rather than real quantum hardware, and results were obtained from a single train-test split rather than repeated multi-seed experiments. These factors limit the extent to which the findings can be generalized.

7.4 Validity, Implications, and Directions for Improvement

The internal validity of the study is supported by controlled data splitting and leakage checks, although the absence of repeated runs or cross-validation limits statistical confidence. External validity is also restricted because the model was not evaluated on independent datasets or hardware-based quantum backends. Moreover, standard classification metrics alone cannot fully capture clinical utility or robustness.

Nevertheless, the study demonstrates that a compact hybrid CNN-QNN architecture can be trained end-to-end on MRI data and can outperform a matched lightweight classical model, particularly in binary tumor detection. As such, the work serves as a practical demonstration of lightweight hybrid quantum learning rather than a clinically deployable solution.

Future research could strengthen the findings by incorporating stronger pre-trained CNN baselines, improving the classical feature extractor, exploring more expressive quantum circuits, and conducting evaluations on external datasets, noisy quantum simulations, and repeated experimental runs.

7.5 Chapter Summary

This chapter discussed the experimental findings and their relationship to existing literature. The results show that the proposed hybrid model provides consistent improvements over a matched lightweight baseline, achieving particularly strong performance in binary classification while maintaining competitive multiclass results. Although several limitations remain, the study demonstrates the potential of compact hybrid quantum-classical architectures and provides a reproducible foundation for future research in MRI-based brain tumor classification.

Chapter 8

Conclusion and Future Work

8.1 Summary of Findings and Achievement of Objectives

This thesis presented a hybrid CNN-QNN framework for brain tumor classification using MRI images. The proposed model was evaluated on both binary and multiclass classification tasks and compared against a matched lightweight classical baseline. Experimental results showed consistent improvements with the hybrid approach, increasing test accuracy from (85.08%) to (93.29%) in the binary setting and from (77.19%) to (83.50%) in the multiclass setting. The findings indicate that the quantum layer contributes positively to classification performance, while also confirming that multiclass tumor classification remains more challenging than binary tumor detection.

The objectives of the study were successfully achieved. A complete hybrid classification pipeline was designed and implemented, evaluated alongside a comparable classical baseline, and analyzed using standard performance metrics and confusion matrices. Overall, the results demonstrate the potential of hybrid quantum-classical learning for MRI classification while highlighting areas that require further improvement.

8.2 Research Contributions

The main contributions of this thesis are summarized as follows:

- Development of an end-to-end hybrid CNN-QNN framework for brain MRI classification.
- Evaluation of the same hybrid architecture on both binary and multiclass classification tasks within a unified experimental setting.
- Inclusion of a matched lightweight classical baseline to facilitate a clearer assessment of the quantum layer’s contribution.
- Establishment of a reproducible workflow accompanied by a balanced analysis of the model’s strengths and limitations.

The contribution of this work is primarily methodological and empirical. Rather than claiming superiority over advanced deep-learning architectures, it provides insight into the practical behavior of a compact hybrid model and its effectiveness relative to a comparable lightweight classical alternative.

8.3 Recommendations for Future Research

Future studies should compare the proposed approach against stronger classical and pretrained CNN models to better assess its competitiveness. Additional improvements may be achieved through enhanced feature extraction, more effective latent representations, and systematic exploration of quantum circuit configurations, including qubit count, circuit depth, and entanglement strategies.

Further validation should also include repeated experiments, cross-validation, external datasets, and evaluations on noisy simulators or real quantum hardware. From a reproducibility perspective, improved code organization, portable implementations, and standardized experimental environments would strengthen the reliability and accessibility of future work. Collectively, these directions can provide a more comprehensive understanding of the capabilities and limitations of hybrid quantum-classical models for medical image classification.

References

- Ait Haddou, M. and Bennai, M. (2025). Hqcm-ebtc: A hybrid quantum-classical model for explainable brain tumor classification. *arXiv preprint arXiv:2506.21937*.
- Ajlouni, N., Ozyavas, A., Takaoglu, M., Takaoglu, F., and Ajlouni, F. (2023). Medical image diagnosis based on adaptive hybrid quantum cnn. *BMC Medical Imaging*, 23(1):126.
- Amin, J., Anjum, M. A., Sharif, M., Jabeen, S., Kadry, S., and Moreno Ger, P. (2022). A new model for brain tumor detection using ensemble transfer learning and quantum variational classifier. *Computational Intelligence and Neuroscience*, 2022:3236305.
- Cheng, T., Zhao, R.-S., Wang, S., Wang, R., and Ma, H.-Y. (2024). Analysis of learnability of a novel hybrid quantum-classical convolutional neural network in image classification. *Chinese Physics B*, 33(4):040303.
- Darwish, S. M., Abu Shaheen, L. J., and Elzoghbi, A. A. (2023). A new medical analytical framework for automated detection of mri brain tumor using evolutionary quantum inspired level set technique. *Bioengineering*, 10(7):819.
- Fan, F., Shi, Y., Guggemos, T., and Zhu, X. X. (2024). Hybrid quantum-classical convolutional neural network model for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 35(12):18145–18159.
- Gencer, K. and Gencer, G. (2025). Hybrid deep learning approach for brain tumor classification using efficientnetb0 and novel quantum genetic algorithm. *PeerJ Computer Science*, 11:e2556.
- Hafeez, M. A., Munir, A., and Ullah, H. (2024). H-qnn: A hybrid quantum-classical neural network for improved binary image classification. *AI*, 5(3):1462–1481.

- Khan, M. A. Z., Galib, A. A. O., Innan, N., and Bennai, M. (2024). Brain tumor diagnosis using quantum convolutional neural networks. *arXiv preprint arXiv:2401.15804*.
- Muniasamy, A., Alquhtani, S. A. S., Alshehri, A. H., Begum, A., and Sabahath, A. (2025). Investigating hybrid quantum-assisted classical and deep learning model for mri brain tumor classification. *Journal of Image and Graphics*, 13(1):123–129.
- Nugraha, S. P., Alfajri, A. I., Sumaryada, T., Tai, D. T., Tamam, N., Sulieman, A., and Yani, S. (2025). Application of quantum convolutional neural networks for mri-based brain tumor detection and classification. *arXiv preprint arXiv:2509.02582*.
- Prajapat, S., Tomar, M., Kumar, P., Kumar, R., and Vasilakos, A. V. (2025). Quantum computing meets deep learning: A qcn model for accurate and efficient image classification. *Mathematics*, 13(19):3148.
- Ranga, D., Prajapat, S., Akhtar, Z., Kumar, P., and Vasilakos, A. V. (2024). Hybrid quantum-classical neural networks for efficient mnist binary image classification. *Mathematics*, 12(23):3684.
- Shi, M., Situ, H., and Zhang, C. (2024). Hybrid quantum neural network structures for image multi-classification. *Physica Scripta*, 99(5):056012.