

NOISE-ROBUST HYBRID QUANTUM NEURAL NETWORKS: A FRAMEWORK FOR QUANTUM AI IN THE NISQ ERA

By

Jesus Rafael Gil Lopez

July, 2026

Director of Thesis: David Hart, PhD

Major Department: Computer Science

ABSTRACT

Quantum computing offers the potential to solve selected computational problems more efficiently than classical systems by leveraging quantum mechanical principles such as superposition and entanglement. However, current quantum devices operate in the noisy intermediate-scale quantum (NISQ) regime, where limited coherence times, gate errors, measurement instability, and hardware variability significantly constrain practical performance. As a result, many quantum machine learning architectures that appear promising under idealized simulation conditions experience substantial degradation when executed under realistic noisy environments.

This thesis investigates the problem of robustness in hybrid quantum neural network (HQNN) systems operating under realistic NISQ conditions. Rather than focusing primarily on quantum expressivity or claims of quantum advantage, the work examines whether practical hybrid quantum learning can maintain stable optimization and reliable inference in the presence of stochastic perturbations. A robustness-engineering framework is developed that integrates learned-readout decoding, multi-observable quantum feature extraction, architecture-level optimization, stability-aware training, and noise-augmented learning within a unified hybrid quantum-classical pipeline.

To evaluate the framework, a comprehensive benchmarking methodology is developed across multiple datasets, architectures, perturbation environments, software frameworks, and

hardware platforms. The experiments compare classical machine learning models, quantum kernel methods, and several HQNN architectures, including quanvolutional neural networks, quantum convolutional neural networks, and quantum transfer learning systems. Performance is evaluated under both noiseless and noisy conditions using standardized metrics, cross-noise validation, architecture-level comparisons, and real IBM Quantum hardware execution.

The results demonstrate that baseline HQNN systems are highly sensitive to optimization instability, parity-based information loss, measurement variability, and perturbation accumulation. The experiments further reveal that robustness, rather than representational expressivity alone, is the dominant factor limiting practical hybrid quantum learning under realistic NISQ conditions. Through successive robustness-engineering stages, noisy performance improves from approximately 0.37 to 0.50 in baseline architectures to approximately 0.97 under several benchmark environments while maintaining substantially improved stability and reduced degradation under perturbation.

These findings establish robustness as a central design objective for practical quantum machine learning. The thesis demonstrates that reliable hybrid quantum learning emerges through coordinated interaction between quantum feature generation, adaptive classical inference, stability-aware optimization, and hardware-aware architecture design. More broadly, the results suggest that hybrid quantum-AI systems are most effective when viewed as integrated quantum-classical computational frameworks in which quantum circuits function as probabilistic feature generators operating within larger robustness-oriented learning pipelines.

NOISE-ROBUST HYBRID QUANTUM NEURAL NETWORKS: A FRAMEWORK FOR
QUANTUM AI IN THE NISQ ERA

A Thesis

Presented to The Faculty of the Department of Computer Science

East Carolina University

In Partial Fulfillment of the Requirements for the Degree

Master of Science in Computer Science

By

Jesus Rafael Gil Lopez

July, 2026

Director of Thesis: David Hart, PhD

Thesis Committee Members:

Nasseh Tabrizi, PhD

Moritz Dannhauer, PhD

©Jesus Rafael Gil Lopez, 2026

Table of Contents

List of Tables	ix
List of Figures	x
Chapter 1: Introduction	1
1.1 Foundations of Quantum Computing in the NISQ Era	1
1.2 Qubits and Quantum Operators	3
1.3 The Noise Challenge in Quantum Systems	5
1.4 Hybrid Quantum–Classical Algorithms	8
1.5 Toward Noise-Robust Hybrid Quantum-AI Systems	10
1.6 Research Questions and Contributions	12
Chapter 2: Background and Related Work	16
2.1 Quantum Computing Fundamentals	16
2.1.1 Superposition, Measurement, and Entanglement	17
2.1.2 Quantum Circuits and Variational Models	18
2.2 Variational and Hybrid Quantum Algorithms	19
2.2.1 Variational Quantum Eigensolver (VQE)	20
2.2.2 Quantum Approximate Optimization Algorithm (QAOA)	21
2.2.3 Variational Optimization Challenges	22
2.3 Quantum Machine Learning	23
2.3.1 Quantum Kernel Methods and QSVM	24
2.3.2 Variational Quantum Neural Networks	25
2.3.3 Quantum Circuit Learning	26

2.4	Hybrid Quantum Neural Network Architectures	27
2.4.1	Quanvolutional Neural Networks	29
2.4.2	Quantum Convolutional Neural Networks	30
2.4.3	Quantum Transfer Learning	30
2.5	Quantum Noise and Error Mitigation	31
2.5.1	Noise Channels and Decoherence	32
2.5.2	Quantum Error Correction and Stabilizer Codes	34
2.5.3	Error Mitigation Techniques	35
2.5.4	Noise-Aware Learning Strategies	35
2.6	NISQ Hardware Constraints	37
2.6.1	Coherence Time and Gate Errors	37
2.6.2	Circuit Depth and Scalability Limitations	39
2.6.3	Hardware Variability and Readout Noise	40
2.7	Quantum Hardware Platforms and Scalability	41
2.7.1	Superconducting Quantum Systems	41
2.7.2	Trapped-Ion Quantum Computing	42
2.7.3	Neutral Atom Quantum Computing	43
2.7.4	Cloud Quantum Platforms and Azure Quantum	43
2.8	Comparison to Prior Work	44
2.8.1	Existing HQNN Noise Studies	45
2.8.2	Limitations of Current Robustness Evaluation	47
2.8.3	Gap Addressed by This Thesis	49
Chapter 3: Establishing Benchmark Comparisons		51
3.1	Purpose of the Benchmarking Framework	51
3.1.1	Why HQNN Robustness Requires Systematic Comparison	54
3.1.2	Limitations of Existing Noise Evaluation Methods	55
3.1.3	Role of Classical, Quantum, and Hybrid Baselines	56

3.2	Baseline HQNN Behavior and Failure Modes	57
3.2.1	Baseline Parity-Based HQNN Design	58
3.2.2	Baseline Classification Behavior	59
3.2.3	SPSA Training Instability Under Noise	60
3.2.4	Initial Threshold Limitations	62
3.3	Classical, QSVM, and HQNN Benchmark Comparisons	63
3.3.1	Quantum Kernel SVM Benchmark	64
3.3.2	Classical Baseline Performance	65
3.3.3	HQNN Performance Gap on Structured Data	66
3.3.4	Medical and Cybersecurity Classification Observations	67
3.4	Noise Robustness Benchmarking	68
3.4.1	HQNN Robustness Under Noise	69
3.4.2	Accuracy Degradation Under Increasing Noise	70
3.4.3	Noise as a Training and Evaluation Bottleneck	72
3.5	Cross-Framework and Cross-Noise Validation	73
3.5.1	Cross-Framework Noise Benchmark	74
3.5.2	Cross-Noise Robustness Heatmap	75
3.5.3	Noise-Channel Sensitivity Across Architectures	77
3.6	Architecture-Level Benchmarking	78
3.6.1	QuanNN versus QCNN Design Differences	79
3.6.2	Ideal Performance versus Noisy Performance	79
3.6.3	Robustness Tradeoffs Across HQNN Variants	80
3.7	Summary of Benchmarking Findings	82
3.7.1	What the Baseline Framework Revealed	83
3.7.2	Why Robustness Engineering Became Necessary	85
3.7.3	Transition to Improving Noisy Training	86
	Chapter 4: Improving Noisy Training	88

4.1	Motivation for Robustness Engineering	88
4.1.1	From Benchmarking to Robustness Optimization	89
4.1.2	Limitations of Baseline Parity-Based HQNNs	90
4.1.3	Hybrid Quantum-Classical Co-Design Philosophy	91
4.2	Learned Readout Hybrid Quantum Neural Networks	92
4.2.1	Replacing Parity Collapse with Learned Decoding	94
4.2.2	Multi-Shot Quantum Measurement Distributions	95
4.2.3	Classical Decoder Integration	96
4.2.4	Threshold Improvement from Approximately 0.50 to Approximately 0.82	97
4.3	Multi-Observable Quantum Feature Extraction	98
4.3.1	Z Expectations and Correlation Features	99
4.3.2	Multi-Observable Quantum Statistics	101
4.3.3	Quantum Circuits as Feature Generators	103
4.3.4	Threshold Improvement from 0.82 to 0.85	104
4.4	Architecture Robustness Optimization	105
4.4.1	Entanglement Topology Search	107
4.4.2	Linear, Ring, Full, and Local Architectures	108
4.4.3	Robustness versus Expressivity Tradeoffs	109
4.4.4	Threshold Improvement from 0.85 to 0.89	110
4.5	Noise-Aware Optimization and Stability Regularization	111
4.5.1	Dual-Loss Optimization Objectives	113
4.5.2	Stability-Regularized HQNN Training	114
4.5.3	Optimization Behavior Under Noise	115
4.5.4	Threshold Improvement from 0.89 to 0.96	117
4.6	Noise-Augmented Readout Training	118
4.6.1	Noisy Quantum Feature-Space Learning	120

4.6.2	Decoder Adaptation Under Stochastic Noise	122
4.6.3	Hybrid Robustness Co-Design	123
4.6.4	Threshold Improvement from 0.96 to 0.97	124
4.7	Framework Ablation and Mechanism Isolation	125
4.7.1	Standard Loss versus Dual-Loss Optimization	126
4.7.2	Standard Readout versus Noise-Augmented Readout	127
4.7.3	Isolating Mechanism Contributions	128
4.7.4	Full Framework versus Partial Frameworks	131
4.8	Multichannel Validation and Generalization	133
4.8.1	Depolarizing, Bit-Flip, Phase-Flip, and Amplitude Damping	134
4.8.2	Consistency Across Noise Channels	135
4.8.3	Robustness Generalization Across Architectures	136
4.9	Framework Evolution and Threshold Progression	138
4.9.1	Baseline to Robustness-Engineered HQNNs	139
4.9.2	Threshold Evolution Across Mechanisms	140
4.9.3	Final Robustness Interpretation	142
4.10	Real IBM Hardware Validation	144
4.10.1	IBM Runtime Integration	145
4.10.2	ibm_fez Hardware Execution	146
4.10.3	Simulator-to-Hardware Transferability	148
Chapter 5: Conclusion and Future Work		151
5.1	Summary of Contributions	151
5.1.1	Benchmarking Framework Contributions	153
5.1.2	Robustness-Engineering Contributions	154
5.1.3	Hybrid Quantum-Classical Co-Design Contributions	155
5.2	Final Framework Interpretation	156
5.2.1	Robustness as the Central NISQ Bottleneck	158

5.2.2	Threshold Evolution Across the Framework	159
5.2.3	Lessons for Practical Quantum AI	160
5.3	Limitations of the Present Work	161
5.3.1	NISQ Hardware Constraints	162
5.3.2	Dataset and Scalability Limitations	163
5.3.3	Simulation versus Hardware Limitations	164
5.4	Future Research Directions	166
5.4.1	Fault-Tolerant Hybrid Quantum AI	166
5.4.2	Robustness-Aware Quantum Optimization	167
5.4.3	Quantum-Classical Co-Designed AI Systems	168
5.4.4	Large-Scale Hardware Deployment	169
5.4.5	Final Conclusions	171
Bibliography		174
Appendix A: Additional Experiments and Applications		176
A.1	Nonlinear Benchmark Evaluation	176
A.1.1	Moons Dataset Benchmark	177
A.1.2	Circles Dataset Benchmark	178
A.1.3	XOR Conditional Hybrid Advantage	180
A.2	Quantum-Classical Feature Fusion	181
A.2.1	Hybrid Feature Fusion Design	182
A.2.2	Classical versus Hybrid Feature Spaces	183
A.2.3	Conditional Hybrid Complementarity	184
A.3	Data-Reuploading and Feature-Map Search	185
A.3.1	Angle Encoding versus Reuploading	186
A.3.2	Interaction-Based Feature Maps	187
A.3.3	Expressivity versus Robustness Tradeoffs	188

A.4	Noise Curriculum Learning	189
A.4.1	Fixed-Noise versus Curriculum Training	190
A.4.2	Adaptive Noise Scheduling	190
A.4.3	Optimization Stability Under Curriculum Learning	191
A.5	Variance-Reduced SPSA Optimization	192
A.5.1	SPSA Gradient Variability	193
A.5.2	Variance Reduction Techniques	194
A.5.3	Stability Improvements in Noisy Optimization	195
A.6	Applied Domain Experiments	196
A.6.1	Medical Risk Classification	197
A.6.2	Cybersecurity Anomaly Detection	198
A.6.3	Energy Optimization Scenarios	199
A.6.4	Additional Findings and Implications	200
A.6.5	Practical Lessons for NISQ-Era HQNN Design	201
A.6.6	Limits of Current Hybrid Quantum Learning Systems	202
A.7	Framework Implementation and Experimental Reproducibility	203
A.7.1	Overall Experimental Workflow	203
A.7.2	Reference Hybrid Quantum Neural Network Circuit	204
A.7.3	Baseline Parity Decoder	206
A.7.4	Learned Readout and Multi-Observable Feature Extraction	207
A.7.5	Dataset Preparation	208

List of Tables

2.1	Representative characteristics of major NISQ hardware platforms. Values represent typical ranges reported for contemporary devices and illustrate the practical constraints affecting hybrid quantum learning systems.	38
-----	--	----

List of Figures

1.1	Bloch sphere representation of a single qubit	5
1.2	NISQ hardware versus fault-tolerant quantum computing	7
1.3	Hybrid quantum-classical algorithm workflow	9
2.1	Variational quantum optimization cycle.	20
2.2	Comparison of three hybrid quantum neural network architecture families . .	28
2.3	Major quantum noise channels	33
3.1	Overall benchmarking framework used to evaluate hybrid quantum learning systems under realistic NISQ conditions.	53
3.2	Baseline parity-based hybrid quantum neural network architecture.	59
3.3	SPSA training instability under noisy HQNN execution.	61
3.4	Performance comparison between a classical machine learning model, a quantum support vector machine (QSVM), and the baseline hybrid quantum neural network (HQNN) on a representative structured classification benchmark.	67
3.5	Classification accuracy degradation under increasing noise probability for representative hybrid quantum neural network architectures.	71
3.6	Cross-noise robustness heatmap comparing representative HQNN architectures across multiple perturbation channels.	76
3.7	Robustness tradeoffs across representative HQNN architecture variants. . . .	82
4.1	Comparison between baseline parity collapse and the proposed learned-readout architecture.	94
4.2	Comparison between single-observable inference and multi-observable quantum feature extraction.	102

4.3	Comparison of representative entanglement topologies evaluated during the architecture search.	109
4.4	Optimization behavior under noisy execution before and after stability-aware optimization.	116
4.5	Distinction between noise-aware readout and stability-aware training.	120
4.6	Component-level ablation of the robustness-engineered HQNN framework.	130
4.7	Ablation comparison between the baseline HQNN, partial robustness-engineered variants, and the complete framework.	132
4.8	Threshold evolution throughout the robustness-engineering framework.	142
4.9	IBM <i>ibm_fez</i> hardware validation results.	147
4.10	Simulator-to-hardware transferability across major framework stages.	149
A.1	Moons dataset nonlinear classification benchmark.	178
A.2	Circles dataset nonlinear classification benchmark.	179
A.3	Experimental workflow used to reproduce the proposed HQNN framework.	204

Chapter 1

Introduction

1.1 Foundations of Quantum Computing in the NISQ Era

Quantum computing represents a fundamentally different computational paradigm that leverages the principles of quantum mechanics to process information in ways that are not achievable with conventional digital systems. Unlike classical computers, which operate on binary bits restricted to values of 0 or 1, quantum computers use quantum bits, or qubits, that can exist in superposition states and become entangled with one another. These quantum-mechanical properties enable quantum systems to represent and manipulate information within exponentially large state spaces, creating computational possibilities beyond those available to classical architectures [1], [3].

The foundations of quantum computing emerged from the recognition that classical computers struggle to efficiently simulate quantum systems. In the early 1980s, Richard Feynman argued that quantum phenomena could be simulated more naturally using quantum machines rather than classical digital hardware [1]. This insight motivated the development of a formal framework for quantum computation that was later expanded through advances in quantum algorithms, complexity theory, and quantum information science [2], [3]. Within this framework, quantum states evolve through unitary transformations and measurements, allowing algorithms to exploit interference and entanglement to guide computational processes.

The significance of quantum computing arises from its potential to provide computational advantages for selected classes of problems. Landmark algorithms such as Shor's factoring algorithm and Grover's search algorithm demonstrated that quantum systems can theoretically outperform classical computers in tasks involving cryptography, search, optimization,

and simulation [1], [2]. More generally, an (n) -qubit quantum register can represent a superposition of (2^n) computational basis states simultaneously, enabling quantum algorithms to explore large solution spaces in ways that are difficult to reproduce using classical methods alone. Entanglement further enhances this capability by creating nonclassical correlations between qubits that cannot be efficiently represented using conventional computational models. These correlations allow quantum systems to encode and process highly structured information across multiple variables simultaneously, forming the basis of many proposed applications in quantum simulation, optimization, machine learning, and scientific computing [3], [4]. As a result, quantum computing has become a major interdisciplinary research area spanning physics, computer science, mathematics, engineering, and artificial intelligence.

Despite this theoretical promise, practical quantum computation remains constrained by the fragility of quantum information. Real quantum systems are highly sensitive to environmental interactions, imperfect control operations, and measurement errors. Decoherence gradually destroys quantum correlations over time, while gate imperfections and readout errors introduce additional uncertainty during computation. Consequently, even relatively small perturbations can significantly affect computational accuracy and limit the complexity of executable quantum circuits.

These limitations define the current Noisy Intermediate-Scale Quantum (NISQ) era, a term introduced by Preskill to describe quantum devices containing tens to hundreds of noisy qubits without large-scale fault tolerance [6]. Although modern quantum processors are sufficiently advanced to explore computational regimes beyond straightforward classical simulation, they remain constrained by finite coherence times, imperfect gate fidelities, hardware variability, and restricted circuit depth. As a result, many idealized quantum algorithms remain impractical on present-day hardware despite their theoretical potential.

The distinction between theoretical quantum advantage and practical quantum computation motivates much of contemporary quantum computing research. Rather than assuming idealized noiseless execution, modern approaches increasingly focus on developing algorithms

and learning systems capable of operating reliably under realistic hardware constraints. This thesis is situated within that context and investigates how hybrid quantum neural network systems can be engineered to maintain stable learning behavior under noisy NISQ conditions. The central premise explored throughout the thesis is that robustness, rather than expressive capability alone, constitutes the dominant bottleneck limiting practical hybrid quantum learning on current quantum hardware.

1.2 Qubits and Quantum Operators

The fundamental unit of information in quantum computing is the quantum bit, or qubit. Unlike a classical bit, which exists exclusively in the state 0 or 1, a qubit can exist in a superposition of both computational basis states simultaneously. Mathematically, a qubit is represented within a two-dimensional complex Hilbert space and described as a linear combination of the basis states $(|0\rangle)$ and $(|1\rangle)$, where complex amplitudes determine the probabilities associated with measurement outcomes [2], [3]. This ability to represent multiple possibilities simultaneously forms one of the defining characteristics of quantum computation.

When multiple qubits are combined into a quantum register, the size of the accessible state space grows exponentially. A register containing (n) qubits can represent a superposition across (2^n) computational basis states. This scaling behavior is one of the primary reasons quantum systems are considered promising for selected classes of computational problems, including optimization, simulation, cryptography, and machine learning. Rather than evaluating a single configuration at a time, quantum systems can manipulate probability amplitudes distributed across many correlated computational states.

A second defining property of quantum systems is entanglement. Entangled qubits exhibit correlations that cannot be described through classical probability distributions alone, requiring the state of the system to be treated collectively rather than as a collection of independent components [2], [3]. These correlations allow quantum circuits to encode and process complex relationships between variables and play a central role in many quantum

algorithms and variational learning architectures.

Quantum information is manipulated through quantum gates, which are unitary transformations applied to qubit states. Single-qubit gates perform rotations and phase transformations on individual qubits, while multi-qubit gates introduce interactions and generate entanglement between different parts of the system. By arranging these operations sequentially, quantum circuits implement computational procedures that transform quantum states into forms suitable for measurement and inference.

Modern hybrid quantum learning systems frequently rely on parameterized quantum circuits (PQCs), in which selected gate operations contain trainable parameters. These parameters are optimized through classical feedback procedures, allowing the quantum circuit to function as an adaptive computational layer within a larger hybrid architecture. Parameterized circuits form the foundation of variational quantum algorithms and hybrid quantum neural networks (HQNNs), making them especially important for practical NISQ-era quantum machine learning.

Measurement represents the final stage of quantum computation. When a qubit is measured, its quantum state collapses probabilistically into one of the computational basis states, producing a classical output. Because measurement destroys the original quantum state, quantum algorithms must be designed carefully to extract useful information only after the desired quantum transformations have been performed. The relationship between state preparation, quantum evolution, and measurement therefore lies at the core of quantum algorithm design.

Figure 1.1 presents the Bloch sphere representation of a single qubit, a geometric visualization in which quantum states correspond to points on the surface of a sphere and quantum gates can be interpreted as rotations within this state space. The concepts introduced in this section—superposition, entanglement, quantum gates, parameterized circuits, and measurement—form the operational foundations of the hybrid quantum learning systems examined throughout this thesis. Their practical behavior under realistic noisy conditions becomes

increasingly important in the next section, which examines the challenges introduced by quantum noise and decoherence.




Figure 1.1: Bloch sphere representation of a single qubit and its relationship to quantum operators. A qubit can be represented as a superposition of the computational basis states $|0\rangle$ and $|1\rangle$, while quantum gates can be interpreted as transformations of the state within this geometric representation.

1.3 The Noise Challenge in Quantum Systems

Although quantum computing offers substantial theoretical advantages, its practical realization remains fundamentally constrained by noise. Current quantum processors operate within what is commonly referred to as the Noisy Intermediate-Scale Quantum (NISQ) regime, characterized by limited qubit counts, finite coherence times, imperfect gate fidelities, and stochastic measurement errors [6]. Unlike idealized theoretical models, real quantum systems are highly sensitive to environmental interactions and hardware imperfections, making reliable large-scale quantum computation difficult under present technological

conditions.

The fragility of quantum information arises primarily from decoherence and imperfect quantum control. Decoherence refers to the gradual loss of quantum phase information caused by unwanted interactions between a quantum system and its surrounding environment. As coherence is lost, superposition states and entanglement structures degrade, reducing the reliability of computation. In addition, physical implementations of quantum gates are affected by calibration inaccuracies, pulse imperfections, crosstalk, and device-specific variability, introducing errors during state transformations and measurements [6], [17].

These limitations place strict constraints on executable circuit depth. While many theoretical quantum algorithms assume long coherence times and highly accurate gate operations, real NISQ processors accumulate errors as the number of gates and entangling operations increases. Eventually, noise dominates the computation and limits the usefulness of deeper circuits. This challenge is particularly important in variational quantum learning systems, where repeated circuit execution during training amplifies the effects of hardware imperfections.

A long-term solution to this problem is fault-tolerant quantum computing. Fault-tolerant architectures use quantum error correction to encode logical qubits across many physical qubits, allowing errors to be detected and corrected during computation [18], [19]. In principle, fault tolerance could enable large-scale reliable quantum computation. In practice, however, current hardware remains far from the scale required for widespread implementation of these techniques, and most contemporary quantum algorithms must operate without full error correction.

As illustrated in Figure 1.2, a significant robustness gap exists between ideal fault-tolerant quantum computation and current NISQ hardware. While fault-tolerant systems assume stable quantum evolution over long computational durations, NISQ devices operate under severe restrictions imposed by decoherence, gate errors, readout noise, and limited circuit depth.




Figure 1.2: Comparison between current NISQ hardware and ideal fault-tolerant quantum computing. NISQ systems operate under restrictions imposed by decoherence, gate errors, readout noise, and limited circuit depth, while ideal fault-tolerant systems rely on logical qubits and quantum error correction to support longer and more reliable computation.

To study and benchmark noisy quantum systems, researchers commonly model several standard noise channels. Bit-flip noise randomly changes computational basis states, phase-flip noise alters relative phase information, amplitude damping models energy relaxation processes, and depolarizing noise partially randomizes quantum states [14], [15], [16]. Although these models simplify the behavior of real hardware, they provide a useful framework for analyzing how noise affects quantum algorithms, variational optimization procedures, and hybrid quantum machine learning systems.

The presence of noise creates one of the central tradeoffs in modern quantum computing. Increasing circuit depth and entanglement generally improves representational capability under ideal conditions, but it also increases sensitivity to hardware imperfections. Consequently, many successful NISQ-era algorithms rely on shallow, structured, and noise-aware designs that balance expressivity against robustness. This observation is particularly important for hybrid quantum machine learning systems and directly motivates the robustness-engineering approach developed throughout this thesis. Rather than treating noise as a secondary implementation detail, the work presented here approaches robustness as a primary design objective for practical hybrid quantum learning under realistic NISQ conditions.

1.4 Hybrid Quantum–Classical Algorithms

The limitations imposed by noise, finite coherence times, and restricted circuit depth in NISQ-era quantum processors have led to the emergence of hybrid quantum–classical algorithms as the dominant paradigm for practical near-term quantum computation. Rather than relying on deep fault-tolerant quantum circuits, hybrid approaches combine quantum subroutines with classical optimization, control, and post-processing systems, allowing useful computational behavior to be achieved despite the constraints of current hardware.

At the center of many hybrid approaches are variational quantum algorithms (VQAs), which use parameterized quantum circuits optimized through an iterative classical feedback loop. In this framework, a quantum circuit prepares a parameter-dependent quantum state, measurements are performed on the resulting state, and a classical optimizer updates the circuit parameters based on the observed outcomes. This repeated interaction between quantum state evolution and classical optimization forms the computational foundation of many modern NISQ-era algorithms.

Two influential examples of this paradigm are the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA). VQE was originally developed for quantum simulation and molecular energy estimation, while QAOA was designed for combinatorial optimization problems [7], [8]. Although these algorithms target different applications, both demonstrate how shallow parameterized quantum circuits can be combined with classical optimization procedures to perform useful computation without requiring large-scale fault-tolerant hardware.

As illustrated in Figure 1.3, hybrid quantum–classical computation operates through an iterative workflow in which classical and quantum systems cooperate during execution. Classical processors provide optimization, parameter control, and decision-making capabilities, while quantum circuits generate probabilistic representations and perform specialized computational transformations. This division of responsibilities allows quantum resources to be used selectively while leveraging the maturity and scalability of classical computing

infrastructure.




Figure 1.3: Hybrid quantum-classical algorithm workflow. Classical processors provide input preparation, parameter control, optimization, and post-processing, while a parameterized quantum circuit prepares and transforms quantum states before measurement. The measured outputs are returned to the classical optimizer, which updates the trainable circuit parameters through an iterative feedback loop.

The effectiveness of hybrid algorithms arises from their ability to reduce quantum resource requirements while still exploiting uniquely quantum phenomena such as superposition, interference, and entanglement. By delegating preprocessing, optimization, and post-processing tasks to classical systems, hybrid methods partially mitigate the limitations imposed by noisy hardware and enable useful computation using relatively shallow circuits compatible with present-day devices.

Hybrid architectures have become especially important in quantum machine learning. Many hybrid quantum neural network (HQNN) systems use parameterized quantum circuits as trainable feature generators embedded within larger classical learning pipelines. Classical optimizers and inference mechanisms interact continuously with quantum subroutines, creating a coordinated computational framework that combines quantum representation capabilities with the stability and scalability of classical machine learning.

However, hybridization alone does not eliminate the challenges associated with noisy quantum computation. Variational circuits remain sensitive to measurement fluctuations, hardware perturbations, optimization instability, and circuit-depth limitations. As a result,

the central challenge in practical NISQ-era quantum computing is no longer simply constructing hybrid systems, but designing hybrid architectures capable of maintaining reliable behavior under realistic noisy execution conditions. This challenge motivates the development of noise-robust hybrid quantum-AI systems, which form the focus of the remainder of this thesis.

1.5 Toward Noise-Robust Hybrid Quantum-AI Systems

Although hybrid quantum-classical algorithms provide a practical computational pathway for the NISQ era, their effectiveness remains constrained by the same hardware limitations discussed in the previous sections. Variational quantum circuits, despite being designed for shallow execution, remain sensitive to noise during state preparation, gate application, and measurement. As a result, model behavior can vary significantly across hardware platforms, perturbation environments, and optimization settings, limiting the reliability of many current hybrid quantum learning systems.

A common limitation in existing quantum machine learning research is that noise is frequently treated as a secondary implementation issue rather than as a central component of the learning process. Error mitigation techniques are often applied after model construction as corrective procedures during execution or evaluation. While these methods can improve performance under selected conditions, they do not necessarily guarantee stable optimization behavior or consistent generalization across different hardware environments [14], [15]. Consequently, many variational quantum models continue to exhibit instability when exposed to realistic NISQ conditions.

These observations motivate a different perspective adopted throughout this thesis. Rather than treating robustness as an auxiliary objective, this work considers robustness to be a primary design requirement for practical hybrid quantum-AI systems. From this viewpoint, noise should influence architecture design, feature representation, optimization procedures, decoder strategies, and evaluation methodologies from the beginning of the development

process rather than being addressed only after performance degradation has occurred.

The central contribution of this thesis is the development of a robustness-engineered hybrid quantum-AI framework designed specifically for realistic NISQ environments. The framework integrates classical preprocessing, parameterized quantum circuits, adaptive inference mechanisms, and stability-aware optimization procedures within a unified computational pipeline. Instead of maximizing circuit complexity alone, the proposed approach emphasizes reliable learning behavior under stochastic perturbation conditions and realistic hardware constraints.

The proposed framework integrates classical preprocessing, parameterized quantum circuits, adaptive inference, and robustness evaluation within a unified pipeline. Classical and quantum components interact through an iterative learning process in which robustness considerations are incorporated directly into both training and evaluation. This design allows hybrid quantum models to be analyzed systematically across multiple perturbation environments while preserving compatibility with current hardware limitations.

An important characteristic of the framework is its emphasis on comparative robustness evaluation. Rather than focusing on a single hybrid quantum neural network architecture, the thesis investigates multiple HQNN variants, including quanvolutional neural networks, quantum convolutional neural networks, and quantum transfer learning systems. These architectures are evaluated under both idealized and noisy conditions in order to identify how architectural decisions influence optimization stability, degradation behavior, and inference reliability.

The broader objective of this work is not simply to determine whether hybrid quantum models can achieve competitive performance on selected tasks. Instead, the goal is to identify the principles that enable stable and reproducible hybrid quantum learning under realistic NISQ constraints. This perspective establishes the foundation for the remainder of the thesis, which develops a systematic benchmarking framework, introduces robustness-engineering mechanisms, validates them across multiple perturbation environments and hardware plat-

forms, and ultimately argues that robustness, rather than expressive capability alone, represents the dominant bottleneck limiting practical hybrid quantum learning today.

1.6 Research Questions and Contributions

The preceding sections established that, although quantum computing offers substantial theoretical computational capabilities, practical implementation remains constrained by the limitations of NISQ-era hardware. Noise, finite coherence times, restricted circuit depth, and optimization instability collectively limit the reliability of current quantum systems. While hybrid quantum–classical algorithms provide a practical pathway for computation under these constraints, their performance remains strongly influenced by noise, architecture design, and training methodology. These challenges motivate the need for a systematic framework capable of evaluating and improving the robustness of hybrid quantum learning systems under realistic operating conditions.

This thesis addresses that challenge by investigating the role of noise in hybrid quantum neural networks and by developing a robustness-engineering framework designed specifically for NISQ-era quantum machine learning. Rather than focusing exclusively on predictive performance under idealized conditions, the work examines how hybrid quantum learning systems behave under realistic perturbation environments and how stability can be improved through coordinated architecture, optimization, and inference design.

To guide this investigation, the following research questions are considered:

- **RQ1:** How does noise affect the performance of hybrid quantum neural network models under realistic NISQ conditions?
- **RQ2:** How do different hybrid quantum architectures, including quanvolutional neural networks, quantum convolutional neural networks, and quantum transfer learning systems, compare in terms of robustness, optimization stability, and reliability?

- **RQ3:** To what extent can robustness-aware training and evaluation strategies improve the consistency and resilience of hybrid quantum learning systems?
- **RQ4:** What design principles can be derived from empirical evaluation to guide the development of practical noise-robust hybrid quantum-AI systems?

To address these research questions, this thesis develops a noise-robust Hybrid Quantum Neural Network framework that combines systematic benchmarking with robustness-engineering techniques designed for realistic NISQ hardware. Rather than proposing an entirely new HQNN architecture, this work identifies the principal factors limiting reliable hybrid quantum learning and introduces engineering strategies that improve robustness under noisy execution.

Specifically, this thesis makes three primary research contributions:

1. **Systematic Benchmarking Methodology:** Development of a unified benchmarking methodology that evaluates Hybrid Quantum Neural Networks, Quantum Support Vector Machines, and classical machine learning models under identical datasets, optimization procedures, software frameworks, evaluation metrics, and multiple representative quantum noise channels. This benchmarking framework enables direct comparison of robustness, optimization stability, degradation behavior, and reproducibility under realistic NISQ conditions.
2. **Robustness-Engineering Methodology:** Development of a coordinated robustness-engineering framework that introduces three principal technical contributions:
 - a learned-readout decoding mechanism that replaces fixed parity decoding with adaptive classical inference from quantum measurement distributions;
 - multi-observable quantum feature extraction and architecture optimization that preserve richer quantum information while balancing expressivity and robustness;

- stability-aware optimization and training objectives that improve convergence behavior under noisy quantum execution.

3. **Comparative Architecture Analysis:** Systematic evaluation of QuanNN, QCNN, and Quantum Transfer Learning architectures under both idealized and realistic noisy conditions to determine how architectural design influences robustness, optimization stability, and performance degradation across multiple quantum noise environments.

Beyond these primary technical contributions, this thesis derives practical design guidelines for developing reliable Hybrid Quantum Neural Networks under realistic NISQ hardware constraints. These guidelines are obtained through the benchmarking studies, robustness-engineering framework, ablation experiments, threshold progression analysis, and IBM Quantum hardware validation presented throughout the remainder of this dissertation.

A central outcome of this work is the identification of robustness as a primary bottleneck limiting practical hybrid quantum learning. Across the benchmarking and framework-development experiments, the results consistently demonstrate that instability under noise often limits performance before the theoretical expressive capability of variational quantum circuits can be fully utilized. This observation shifts the focus from increasing quantum complexity alone toward engineering reliable behavior under realistic execution conditions.

The scope of this thesis is restricted to near-term quantum computing systems operating without large-scale fault tolerance. Consequently, the emphasis is placed on robustness, optimization stability, and practical execution behavior rather than on demonstrating broad claims of quantum advantage. The work therefore approaches hybrid quantum machine learning primarily as a robustness-engineering problem shaped by hardware limitations, stochastic perturbations, and noisy optimization dynamics.

The remainder of the thesis develops this framework progressively. Chapter 2 reviews the theoretical foundations of quantum computing, variational algorithms, hybrid quantum neural networks, and quantum machine learning. Chapter 3 establishes benchmark comparisons and analyzes baseline HQNN behavior under noisy conditions. Chapter 4 develops

the robustness-engineering framework and evaluates its impact on noisy training stability. Chapter 5 summarizes the major findings, discusses limitations, and outlines future directions for practical hybrid quantum-AI systems. Additional experiments, nonlinear benchmarks, optimization studies, and hardware validation results are presented in Appendix A.

Chapter 2

Background and Related Work

2.1 Quantum Computing Fundamentals

Quantum computing is founded on the principles of quantum mechanics and provides a computational framework that differs fundamentally from classical digital computation. Whereas classical systems operate on binary states that are either 0 or 1, quantum systems process information using quantum states that evolve according to the laws of quantum mechanics. This allows quantum computers to exploit phenomena such as superposition, entanglement, and interference, enabling computational behaviors that cannot be reproduced directly by conventional classical architectures [2], [3].

The theoretical framework of quantum computation describes how information is encoded, transformed, and measured within quantum systems. Quantum states are represented mathematically in complex Hilbert spaces, quantum operations are implemented through unitary transformations, and measurements produce probabilistic outcomes determined by the amplitudes of the underlying quantum state [2], [3]. Together, these principles establish the foundation for quantum circuits, variational algorithms, and hybrid quantum–classical learning systems.

One of the defining characteristics of quantum computation is its ability to represent and manipulate exponentially scaling state spaces. At the same time, practical quantum systems remain constrained by noise, decoherence, finite coherence times, and imperfect hardware operations. Consequently, the effectiveness of modern quantum algorithms depends not only on theoretical expressivity but also on the ability to operate reliably under realistic NISQ conditions. The following subsections review the core concepts of superposition, measure-

ment, entanglement, and quantum circuits that form the foundation of the hybrid quantum learning architectures examined throughout this thesis.

2.1.1 Superposition, Measurement, and Entanglement

A quantum bit, or qubit, is the fundamental unit of information in quantum computation. Unlike a classical bit, which exists in a definite state of either 0 or 1, a qubit can exist in a superposition of both computational basis states simultaneously. Mathematically, a qubit is represented as a linear combination of the basis states ($|0\rangle$) and ($|1\rangle$), where complex-valued amplitudes determine the probabilities associated with measurement outcomes [2], [3]. This ability to represent multiple computational possibilities at once forms one of the foundational characteristics of quantum information processing and enables quantum systems to explore exponentially large state spaces as the number of qubits increases.

Measurement provides the mechanism through which quantum information is converted into classical information. When a quantum state is measured, the superposition collapses probabilistically into one of the computational basis states, producing a classical output. The probability of observing each outcome is determined by the amplitudes of the quantum state prior to measurement [3]. Because measurement destroys the original superposition, quantum algorithms must be designed to manipulate quantum information through a sequence of operations before the final observation stage. This relationship between state preparation, evolution, and measurement is central to quantum computation and quantum machine learning.

A second defining property of quantum systems is entanglement, which describes non-classical correlations between multiple qubits. When qubits become entangled, the state of the overall system can no longer be described as a collection of independent components and must instead be treated as a unified quantum state [2], [3]. Entanglement allows quantum circuits to represent highly correlated computational structures and plays an important role in optimization, simulation, and machine learning applications. At the same time, entan-

gled states are particularly sensitive to noise and decoherence, making them both a source of computational power and a source of vulnerability within practical NISQ-era quantum systems.

2.1.2 Quantum Circuits and Variational Models

Quantum computation is commonly described using the quantum circuit model, in which qubits are initialized into quantum states, transformed through sequences of quantum gates, and ultimately measured to produce classical outputs. Within this framework, computation is represented as the evolution of quantum states through unitary operations acting on one or more qubits [3]. Single-qubit gates manipulate individual quantum states through rotations and phase transformations, while multi-qubit gates generate correlations and entanglement between different parts of the system. Together, these operations form quantum circuits capable of implementing a wide variety of computational tasks.

The expressive capability of a quantum circuit depends on factors such as circuit depth, gate structure, and entanglement topology. In general, deeper circuits and more complex entanglement patterns can represent richer computational relationships, but they also become increasingly sensitive to noise and hardware imperfections. This tradeoff between expressive capability and execution reliability is particularly important in the NISQ regime, where finite coherence times and imperfect gate fidelities limit the practical complexity of executable quantum circuits.

To address these hardware constraints, many modern quantum algorithms rely on variational models based on parameterized quantum circuits. In these approaches, selected quantum gates contain trainable parameters that are optimized through classical feedback procedures. The resulting hybrid quantum–classical optimization process forms the foundation of variational quantum algorithms, quantum machine learning models, and hybrid quantum neural networks. As a result, parameterized quantum circuits serve as the computational backbone of many NISQ-era learning systems and provide the foundation for the

variational methods discussed in the following section

2.2 Variational and Hybrid Quantum Algorithms

The development of quantum algorithms has evolved from early theoretical demonstrations of quantum advantage toward more practical approaches designed specifically for noisy intermediate-scale quantum (NISQ) hardware. While foundational algorithms demonstrated that quantum systems could outperform classical computation for selected problems, many of these methods require deep and highly coherent quantum circuits that remain impractical on current devices. As a result, modern quantum computing research increasingly emphasizes variational and hybrid algorithms capable of operating under realistic hardware constraints [6], [9].

Variational quantum algorithms form the foundation of this approach. In a variational framework, a parameterized quantum circuit prepares a quantum state whose behavior depends on a set of trainable parameters. Measurements performed on the resulting state are evaluated by a classical optimizer, which updates the parameters iteratively according to an objective function. This hybrid quantum–classical feedback process allows useful computation to be performed using relatively shallow circuits while leveraging classical systems for optimization and control.

Figure 2.1 illustrates the general structure of a variational optimization loop. This computational model underlies many NISQ-era algorithms and hybrid quantum machine learning systems, including the Variational Quantum Eigensolver (VQE), the Quantum Approximate Optimization Algorithm (QAOA), and hybrid quantum neural networks. Although these methods reduce quantum resource requirements substantially, they also introduce new challenges associated with optimization stability, noisy measurements, and hardware variability. The following subsections examine these algorithms and the practical limitations that arise when variational learning systems operate under realistic NISQ conditions.




Figure 2.1: A parameter vector θ defines a parameterized quantum circuit that prepares the state $|\psi(\theta)\rangle$. Measurements estimate observables used to evaluate an objective function, and a classical optimizer updates θ iteratively until convergence. This cycle underlies VQE, QAOA, and many hybrid quantum machine learning systems.

2.2.1 Variational Quantum Eigensolver (VQE)

The Variational Quantum Eigensolver (VQE) is one of the earliest and most influential variational quantum algorithms developed for noisy intermediate-scale quantum (NISQ) hardware [7]. VQE was originally designed to approximate the ground-state energies of quantum systems while avoiding the deep circuit requirements associated with traditional fault-tolerant quantum algorithms. Its development demonstrated that useful quantum computation could be achieved through shallow parameterized quantum circuits combined with classical optimization procedures.

The central idea behind VQE is to prepare a parameterized quantum state and evaluate its expected energy with respect to a target Hamiltonian. Measurements performed on the quantum processor are used to estimate this energy, while a classical optimizer iteratively updates the circuit parameters in an attempt to minimize the objective function. This hybrid quantum–classical feedback loop allows the computational workload to be divided between quantum state preparation and classical optimization, making VQE substantially more practical for current hardware than many earlier quantum algorithms.

Beyond its original applications in quantum chemistry and simulation, VQE established

the computational paradigm that now underlies many variational quantum machine learning systems. The same optimization structure—parameterized quantum circuits combined with classical feedback procedures—appears throughout modern hybrid quantum neural networks and other NISQ-era learning architectures. As a result, VQE is often viewed not only as a quantum simulation algorithm, but also as one of the foundational models for contemporary hybrid quantum computation.

2.2.2 Quantum Approximate Optimization Algorithm (QAOA)

The Quantum Approximate Optimization Algorithm (QAOA) is a variational hybrid quantum algorithm developed for solving combinatorial optimization problems on noisy intermediate-scale quantum (NISQ) hardware [8]. Similar to VQE, QAOA combines parameterized quantum circuits with classical optimization procedures, making it compatible with the limited coherence times and circuit depths of current quantum devices. The algorithm was originally proposed for optimization tasks such as MaxCut, graph partitioning, and constraint-satisfaction problems, where the objective is to identify high-quality approximate solutions within large search spaces.

QAOA operates by alternating between two types of quantum operations. One operation encodes the optimization objective through a problem-specific Hamiltonian, while the other acts as a mixing operator that explores alternative candidate solutions. The behavior of the circuit is controlled by a set of variational parameters that are optimized iteratively through a classical feedback loop. Measurements obtained from the quantum circuit are used to evaluate solution quality, and a classical optimizer updates the parameters in an attempt to improve performance over successive iterations.

The significance of QAOA extends beyond combinatorial optimization itself. Together with VQE, it helped establish the broader paradigm of variational hybrid quantum computation in which quantum state preparation and classical optimization work cooperatively within a single computational framework. This approach has influenced a wide range of

NISQ-era algorithms, including variational quantum machine learning systems and hybrid quantum neural networks, which rely on the same fundamental interaction between parameterized quantum circuits and classical optimization procedures.

2.2.3 Variational Optimization Challenges

Although variational quantum algorithms provide a practical framework for computation on current quantum hardware, their effectiveness is strongly influenced by optimization instability and sensitivity to noise. Unlike many optimization procedures used in classical machine learning, variational quantum optimization relies on repeated evaluations of quantum circuits whose outputs are inherently probabilistic. As a result, the optimization process is often stochastic, making convergence behavior more difficult to analyze and control under realistic execution conditions.

One of the primary challenges arises from the dependence of variational algorithms on repeated quantum measurements. Even under ideal simulation conditions, measurement outcomes exhibit statistical variability due to the probabilistic nature of quantum mechanics. On real hardware, additional sources of uncertainty—including decoherence, gate errors, and readout noise—further distort objective-function evaluations. These effects can cause parameter updates to become inconsistent, producing oscillatory convergence behavior and unstable training trajectories.

A second challenge involves the scaling behavior of parameterized quantum circuits. As circuit depth and parameter count increase, optimization landscapes often become more difficult to navigate. In many variational systems, gradients diminish rapidly as circuit complexity grows, producing regions known as barren plateaus in which parameter updates become increasingly ineffective [9]. Within these regions, optimization may slow dramatically or fail to converge toward useful solutions.

The problem becomes even more severe under realistic NISQ conditions. Noise not only affects the final output of a quantum circuit, but also alters the optimization process itself

by distorting gradient estimates and objective-function evaluations. Consequently, deeper and more expressive variational circuits frequently exhibit poorer practical training behavior despite possessing greater theoretical representational capability. This creates a recurring tradeoff between expressivity and robustness that appears throughout modern variational quantum computing.

Another important consideration is the choice of optimizer. Variational quantum systems commonly rely on classical optimization methods such as gradient descent, SPSA, COBYLA, and Adam-based procedures. However, these algorithms respond differently to stochastic measurement variability and hardware noise. Optimization performance therefore depends not only on circuit design, but also on the interaction between measurement statistics, optimizer dynamics, and the underlying hardware environment. Understanding these challenges is essential for analyzing hybrid quantum learning systems and provides important context for the benchmarking and robustness-engineering studies developed in later chapters of this thesis.

2.3 Quantum Machine Learning

Quantum machine learning (QML) is an interdisciplinary field that combines concepts from quantum computing and classical machine learning to develop new approaches for data representation, prediction, and optimization. The central idea is that quantum systems may provide access to computational structures and feature spaces that are difficult to represent efficiently using classical methods alone [10], [13]. As a result, QML has emerged as one of the most active research areas within NISQ-era quantum computing.

The development of QML has generally followed two complementary directions. The first focuses on quantum-enhanced feature representations, in which classical data is encoded into quantum states and processed through quantum circuits to construct richer feature spaces for downstream learning tasks. The second focuses on trainable quantum models, where parameterized quantum circuits function as adaptive computational layers within hybrid

quantum-classical learning systems. Both approaches seek to exploit quantum phenomena such as superposition, entanglement, and interference while remaining compatible with practical hardware constraints.

Despite their theoretical promise, quantum learning systems face several important challenges. Their effectiveness depends not only on representational capability, but also on optimization stability, scalability, and robustness to noise. These considerations become especially important in the NISQ regime, where decoherence, gate errors, and limited circuit depth strongly influence practical model performance.

Most current QML systems are therefore hybrid in nature. Classical processors perform preprocessing, optimization, and inference tasks, while quantum circuits contribute feature generation, state transformation, or parameterized computation. This division of responsibilities allows quantum resources to be used selectively while leveraging the stability and scalability of classical machine learning infrastructure.

The following subsections review three major categories of quantum machine learning approaches relevant to this thesis: quantum kernel methods, variational quantum neural networks, and quantum circuit learning. Together, these approaches provide the conceptual foundation for understanding the hybrid quantum learning architectures investigated throughout the remainder of this thesis.

2.3.1 Quantum Kernel Methods and QSVM

Quantum kernel methods represent one of the earliest and most structured approaches to quantum machine learning. In these methods, classical data is encoded into quantum states through a quantum feature map, and similarities between data samples are computed using inner products within the quantum Hilbert space [10], [11]. The resulting similarity measure defines a quantum kernel that can be incorporated into classical machine learning algorithms such as support vector machines.

The Quantum Support Vector Machine (QSVM) applies this principle by replacing the

classical kernel function with a kernel generated through quantum circuit evaluation. Rather than learning a parameterized quantum model directly, the quantum processor is used to compute pairwise similarities between quantum-encoded data points. These similarities are then supplied to a classical support vector machine for optimization and classification.

One important characteristic of QSVM systems is the separation between quantum feature generation and model optimization. The quantum subsystem computes kernel values, while the classifier itself is optimized classically. Because the quantum circuit does not require repeated parameter updates during training, QSVM systems often exhibit greater optimization stability than fully variational quantum learning architectures.

The effectiveness of quantum kernel methods depends strongly on the choice of feature map and the fidelity of quantum state preparation. As hardware noise increases, the accuracy of kernel estimation can deteriorate, reducing classification reliability. Nevertheless, QSVMs remain important because they provide a conceptually clear and relatively stable benchmark for comparing quantum learning methods and understanding the tradeoffs between expressivity, trainability, and robustness in more complex hybrid architectures.

2.3.2 Variational Quantum Neural Networks

Variational quantum neural networks (VQNNs) are quantum machine learning models built from parameterized quantum circuits optimized through classical feedback procedures [12], [13]. In these systems, trainable quantum gates act as adaptive computational layers whose parameters are adjusted iteratively to minimize a learning objective, analogous to optimization in classical neural networks.

A typical VQNN begins by encoding classical input data into quantum states through an embedding procedure. These states are then transformed through parameterized quantum gates and entangling operations before measurements are performed. The resulting outputs are evaluated by a classical optimizer, which updates the circuit parameters based on the observed performance.

The flexibility of VQNNs has made them one of the dominant paradigms in NISQ-era quantum machine learning. Different circuit architectures, feature maps, entanglement structures, and optimization strategies can be combined to create models tailored for classification, regression, representation learning, and other tasks.

However, this flexibility introduces significant challenges. Because VQNNs rely on repeated execution of parameterized quantum circuits, they are highly sensitive to noise, stochastic measurement variability, and optimization instability. Barren plateaus, noisy gradients, and hardware imperfections can all degrade training reliability, particularly as circuit depth and parameter count increase.

As a result, practical VQNN design requires balancing representational flexibility against robustness and trainability. This tradeoff between expressivity and stability becomes a recurring theme throughout the remainder of this thesis and motivates the robustness-engineering framework developed in later chapters.

2.3.3 Quantum Circuit Learning

Quantum circuit learning (QCL) is a broader framework for trainable quantum computation in which parameterized quantum circuits are optimized to learn mappings between inputs and outputs [13]. Rather than focusing on a specific application, QCL treats parameterized quantum circuits as general-purpose computational models capable of approximating nonlinear functions through quantum state evolution.

In a typical QCL workflow, classical data is encoded into quantum states and processed through a sequence of trainable quantum operations. Measurements performed on the resulting quantum state generate outputs that are compared with target values using a loss function. Classical optimization algorithms then update the circuit parameters iteratively to improve performance.

The framework closely parallels the operational structure of classical neural network learning, but the transformations occur within quantum state space rather than conventional

vector spaces. Because quantum systems can represent exponentially scaling Hilbert spaces, QCL architectures may access computational representations that are difficult to reproduce explicitly using classical methods alone.

An important characteristic of QCL is its flexibility. Different encoding strategies, measurement operators, variational ansätze, and entanglement structures can produce significantly different learning behavior. Consequently, the effectiveness of QCL systems depends heavily on architecture design, optimization dynamics, and compatibility with hardware constraints.

Like other variational approaches, QCL inherits many of the challenges associated with noisy quantum computation, including unstable gradients, measurement variability, and circuit-depth limitations. These issues have motivated growing interest in shallow circuit design, adaptive optimization strategies, and noise-aware learning procedures, all of which play an important role in practical hybrid quantum learning systems.

2.4 Hybrid Quantum Neural Network Architectures

Hybrid quantum neural networks (HQNNs) represent an important class of quantum machine learning systems that combine parameterized quantum circuits with classical learning components. These architectures seek to integrate the representational capabilities of quantum systems with the scalability, optimization stability, and computational maturity of classical machine learning methods. Within the NISQ era, hybrid approaches have become especially attractive because they allow quantum resources to be incorporated into learning pipelines while remaining compatible with current hardware limitations.

Unlike purely quantum learning systems, HQNNs distribute computation across both quantum and classical subsystems. Classical processors typically perform preprocessing, optimization, parameter updates, and inference management, while quantum circuits provide feature generation, nonlinear transformations, or trainable computational layers. This division of responsibilities allows quantum computation to be used selectively while leveraging

the reliability of classical infrastructure.

A wide variety of HQNN architectures have been proposed, differing primarily in how quantum circuits are integrated into the learning process. Some approaches emphasize localized quantum feature extraction through shallow modular circuits, while others rely on hierarchical entangling structures or pretrained classical feature extractors combined with quantum inference layers. These differences create important tradeoffs between representational capability, trainability, scalability, and robustness under noise.

Figure 2.2 compares the three major HQNN architectures considered throughout this thesis: quanvolutional neural networks (QuanNNs), quantum convolutional neural networks (QCNNs), and quantum transfer learning (QTL) systems. These architectures represent distinct strategies for integrating quantum computation into hybrid learning pipelines and provide useful reference points for evaluating robustness under realistic NISQ conditions.




Figure 2.2: c

considered in this thesis Quanvolutional neural networks use local quantum filters for feature extraction, quantum convolutional neural networks apply hierarchical quantum convolution and pooling operations, and quantum transfer learning systems combine classical feature extraction with trainable quantum inference layers.

The following subsections examine each architecture individually and discuss how their structural differences influence optimization behavior, representational flexibility, and robustness under noisy execution environments.

2.4.1 Quanvolutional Neural Networks

Quannevolutional Neural Networks (QuanNNs) extend the concept of classical convolutional neural networks by replacing conventional convolutional filters with shallow quantum circuits [10], [12], [13]. Input data is divided into localized regions, and each region is processed independently through a quantum feature-generation circuit. The resulting outputs are then aggregated and supplied to downstream classical learning layers.

The defining characteristic of QuanNNs is their emphasis on localized quantum computation. Rather than constructing a large globally entangled system, the architecture distributes computation across many relatively independent quantum filters. This organization naturally limits circuit depth and reduces the amount of entanglement required during execution.

One advantage of this design is compatibility with current NISQ hardware. Because each quantum circuit remains relatively small and shallow, noise accumulation is partially constrained within localized computational regions. This often improves execution reliability compared to architectures that rely heavily on deep entangling structures.

The expressive capability of QuanNNs depends strongly on the design of the quantum filters and the strategy used to combine the resulting feature maps. While the architecture can capture useful local patterns, additional classical processing is often required to model more complex global relationships within the data.

Because of their modular structure and relatively strong noise tolerance, QuanNNs are frequently regarded as one of the most practical HQNN architectures for near-term quantum hardware and serve as an important robustness benchmark throughout this thesis.

2.4.2 Quantum Convolutional Neural Networks

Quantum Convolutional Neural Networks (QCNNs) are hybrid architectures inspired by the hierarchical organization of classical convolutional neural networks but implemented using quantum operations and entangling circuit layers [10], [13]. Unlike QuanNNs, which emphasize localized independent quantum filters, QCNNs construct increasingly compressed feature representations through correlated quantum transformations.

QCNN architectures typically consist of alternating layers of parameterized quantum operations and pooling-like procedures that progressively reduce the dimensionality of the quantum state. Through this hierarchical structure, the architecture can capture broader relationships across the input representation while maintaining relatively efficient parameter usage.

A major advantage of QCNNs is their ability to model global correlations through entanglement. Because information is distributed across multiple qubits simultaneously, the architecture can construct richer quantum representations than many purely localized approaches.

However, this increased representational capability comes with a cost. Entanglement allows noise and errors to propagate throughout the computational system, making QCNNs more sensitive to decoherence, gate imperfections, and optimization instability. As circuit depth increases, maintaining reliable execution becomes increasingly challenging under realistic NISQ conditions.

Consequently, QCNNs illustrate one of the recurring themes of this thesis: architectures with greater theoretical expressive capability often exhibit reduced robustness under noisy execution environments, creating a practical tradeoff between expressivity and stability.

2.4.3 Quantum Transfer Learning

Quantum Transfer Learning (QTL) represents a hybrid approach in which pretrained classical learning systems are combined with quantum computational layers [10], [13]. In most

implementations, a classical neural network first extracts features from the input data, after which a quantum circuit is used as a trainable classification or inference component.

The primary advantage of QTL is its compatibility with existing machine learning ecosystems. Rather than replacing classical feature extraction entirely, the quantum subsystem operates on representations that have already been processed by a powerful classical model. This reduces the burden on quantum hardware and allows relatively small circuits to be used effectively within larger learning pipelines.

Because the quantum component typically appears near the end of the computational workflow, QTL systems often exhibit greater optimization stability than more deeply integrated HQNN architectures. At the same time, the contribution of the quantum subsystem may be more limited because it operates on compressed feature representations generated primarily by classical models.

Although QTL generally provides less direct quantum influence than QuanNN or QCNN architectures, it remains an important reference model because it illustrates one of the most practical strategies for integrating quantum computation into contemporary machine learning systems under realistic NISQ constraints.

2.5 Quantum Noise and Error Mitigation

The practical performance of quantum algorithms is fundamentally constrained by the presence of noise. Unlike classical computing systems, where errors can often be addressed through deterministic redundancy and mature hardware engineering, quantum systems are highly sensitive to environmental interactions, imperfect control operations, and measurement uncertainty. As a result, noise plays a central role in determining the reliability, scalability, and usefulness of quantum computation under realistic hardware conditions.

Quantum noise arises from several physical and operational sources, including decoherence, gate imperfections, calibration instability, thermal interactions, and readout errors. These effects cause quantum states to deviate from ideal evolution, reducing state fidelity

and degrading computational accuracy over time [6], [14], [15], [16]. In the absence of large-scale fault-tolerant quantum error correction, most current quantum algorithms must operate within these imperfect execution environments.

Understanding noise requires both a physical and algorithmic perspective. Physically, noise describes how environmental interactions and imperfect controls alter quantum states. Algorithmically, noise affects optimization behavior, measurement reliability, convergence stability, and final computational performance. This distinction is especially important for quantum machine learning systems, where noisy measurements and repeated circuit evaluations can directly influence training dynamics.

This section reviews the major forms of quantum noise, the principles of quantum error correction, modern error mitigation techniques, and noise-aware learning strategies relevant to hybrid quantum machine learning. These concepts provide the foundation for the robustness-oriented framework developed later in the thesis.

2.5.1 Noise Channels and Decoherence

Quantum noise is commonly modeled through noise channels, which describe how quantum states are altered by imperfect operations and environmental interactions [6], [14], [15]. These channels provide a mathematical and conceptual framework for studying the probabilistic transformations that occur during realistic quantum computation. They are especially useful for benchmarking because they allow different types of hardware disturbance to be represented in controlled experimental settings.

Among the most widely studied noise channels are bit-flip noise, phase-flip noise, amplitude damping, phase damping, and depolarizing noise. Bit-flip noise randomly changes a qubit from one computational basis state to the other. Phase-flip noise alters relative phase information without directly changing the measured classical value. Amplitude damping models energy relaxation, in which excited states decay toward lower-energy configurations. Phase damping describes loss of phase coherence without direct energy loss, while depolar-

izing noise partially randomizes the quantum state and is often one of the most disruptive stochastic error models.

These noise processes correspond to important behaviors observed in real quantum hardware. Environmental interactions, imperfect control signals, thermal fluctuations, calibration drift, and readout instability all contribute to error formation during computation. Because quantum algorithms rely on superposition, phase coherence, and entanglement, even relatively small perturbations can significantly affect computational outcomes.

Decoherence is one of the most important limitations in practical quantum systems. As quantum states interact with the surrounding environment, phase information gradually degrades, causing the system to lose the quantum properties required for reliable computation [3], [6]. This process limits the duration over which quantum information can be preserved and therefore restricts the depth and complexity of executable quantum circuits.

Figure 2.3 summarizes the major quantum noise channels considered throughout this thesis and their qualitative effects on quantum state evolution.




Figure 2.3: considered throughout this thesis and their qualitative effects on quantum state evolution. Bit-flip noise changes computational basis populations, phase-flip noise alters relative phase information, amplitude damping models energy relaxation, and depolarizing noise represents stochastic state randomization. These mechanisms reduce the reliability of quantum computation by disturbing the state information used during measurement and learning.

Noise accumulation becomes especially severe in deep or highly entangled circuits. As the number of quantum operations increases, errors can propagate through the system and progressively reduce computational reliability. This sensitivity to accumulated noise is one of

the defining practical challenges of NISQ-era quantum computation and directly motivates the robustness analysis developed in later chapters.

2.5.2 Quantum Error Correction and Stabilizer Codes

Quantum error correction provides the theoretical foundation for reliable large-scale quantum computation in the presence of noise [18], [19]. Unlike classical error correction, quantum error correction cannot simply copy quantum information because of the no-cloning theorem. Instead, logical quantum states are encoded across multiple physical qubits so that errors can be detected indirectly without measuring and destroying the encoded quantum information.

The core idea is to distribute quantum information across an entangled multi-qubit structure. Error syndromes are then measured to identify whether particular error patterns have occurred, while preserving the logical quantum state. Stabilizer codes provide one of the most influential mathematical frameworks for this process. They define protected subspaces using sets of commuting operators and form the basis for many modern fault-tolerant quantum computing architectures, including surface-code approaches [18], [19].

A central distinction in quantum error correction is the difference between physical qubits and logical qubits. Physical qubits are the actual hardware-level quantum systems, while logical qubits are protected units of quantum information encoded across many physical qubits. Because current hardware error rates remain relatively high, practical fault-tolerant systems require substantial physical-qubit overhead to maintain reliable logical qubits.

Although quantum error correction provides a long-term pathway toward scalable quantum computation, present-day NISQ hardware remains far from the scale required for broad fault-tolerant execution. Current devices generally lack the qubit counts, gate fidelities, and coherence times necessary to support full error-corrected computation. Consequently, most near-term quantum algorithms must operate under partially noisy conditions without complete logical protection.

2.5.3 Error Mitigation Techniques

Because full quantum error correction is not yet practical on current hardware, significant research has focused on error mitigation techniques that reduce the observed effects of noise without requiring logical encoding [14], [15], [16]. Error mitigation differs from error correction because it does not attempt to remove errors completely during computation. Instead, it compensates statistically or algorithmically for their influence on measured outputs.

Common mitigation strategies include measurement calibration, zero-noise extrapolation, probabilistic error cancellation, randomized compiling, and virtual distillation. Measurement calibration attempts to reduce systematic readout errors by characterizing the measurement process and adjusting observed outcomes accordingly. Zero-noise extrapolation estimates an ideal output by evaluating circuits at different effective noise levels and extrapolating toward a zero-noise limit.

Other methods attempt to reduce the impact of hardware imperfections through sampling or circuit transformation procedures. Probabilistic error cancellation uses noise characterization to statistically reconstruct ideal expectation values, while randomized compiling converts coherent errors into more stochastic forms that are easier to analyze. Virtual distillation uses multiple noisy copies of quantum states to improve estimates of target quantities.

Although error mitigation can improve computational accuracy, it usually introduces additional execution overhead and often depends on assumptions about the underlying noise model. For this reason, mitigation should be understood as a practical near-term tool rather than a complete replacement for fault-tolerant quantum computation. In hybrid quantum learning systems, mitigation can be useful, but it does not by itself resolve deeper challenges involving optimization instability and robustness under repeated noisy training.

2.5.4 Noise-Aware Learning Strategies

Beyond post-processing mitigation techniques, recent research has increasingly explored the integration of noise awareness directly into the design and training of quantum machine

learning systems. In these approaches, noise is not treated only as an external disturbance to be corrected after execution. Instead, it is treated as part of the computational environment that shapes model design, training behavior, and evaluation.

In variational quantum learning systems, noise influences both optimization dynamics and final inference behavior. Noisy measurements introduce variability into objective-function evaluations and gradient estimates, potentially destabilizing convergence during training. At the same time, certain forms of stochastic perturbation may act similarly to regularization mechanisms in classical machine learning by encouraging parameter configurations that remain stable under repeated evaluation [12], [13], [16].

Noise-aware learning strategies attempt to align training conditions with the conditions under which the model will ultimately be evaluated. Common approaches include training directly under simulated noise, selecting hardware-efficient circuit structures, reducing unnecessary entanglement, using shallow variational architectures, and incorporating robustness-oriented terms into the optimization objective. These methods shift the focus from correcting noise after computation toward designing systems that can tolerate noise during learning.

This perspective is particularly important for hybrid quantum neural networks. HQNN systems rely on repeated quantum circuit evaluations throughout training, meaning that noise affects not only the final prediction but also the path taken by the optimizer. A model may appear expressive under ideal simulation while becoming unstable under realistic noisy execution. For this reason, robustness must be evaluated as a property of the full training process rather than only as a final accuracy measurement.

Noise-aware learning therefore provides a conceptual bridge between traditional error mitigation and the robustness-engineering framework developed in this thesis. While mitigation attempts to compensate for errors after they occur, robustness-oriented learning seeks to construct hybrid quantum systems that remain stable despite stochastic perturbations. This distinction becomes central in later chapters, where HQNN architectures are benchmarked, modified, and evaluated according to their ability to maintain reliable behavior

under realistic NISQ conditions.

2.6 NISQ Hardware Constraints

The practical capabilities of quantum computation must ultimately be evaluated in the context of the physical systems that implement it. In the noisy intermediate-scale quantum (NISQ) era, quantum devices are characterized by limited qubit counts, finite coherence times, imperfect gate operations, and non-negligible measurement errors [6], [17]. These limitations constrain the complexity and reliability of executable quantum circuits and therefore play a central role in determining the practical usefulness of quantum algorithms.

Unlike classical computing platforms, which benefit from decades of engineering refinement and highly reliable digital logic, current quantum hardware remains inherently sensitive to environmental interactions and control imperfections. Consequently, the performance of quantum systems depends not only on algorithmic design, but also on the physical characteristics of the underlying hardware.

These constraints are particularly important for variational quantum algorithms and hybrid quantum learning systems because they rely on repeated circuit execution during training and inference. As a result, hardware limitations directly influence optimization stability, measurement reliability, and scalability. The following subsections examine three of the most important constraints affecting practical NISQ-era computation: coherence time, circuit scalability, and hardware variability.

2.6.1 Coherence Time and Gate Errors

Two of the most significant limitations in current quantum hardware are coherence time and gate error rates. Coherence time refers to the duration over which a quantum system can preserve its quantum state before interactions with the environment cause information loss. Once coherence degrades, superposition and entanglement become increasingly unreliable, reducing the accuracy of quantum computation [6].

In addition to decoherence, quantum gate operations are subject to implementation errors arising from imperfect physical control mechanisms, calibration drift, and hardware variability. Even small inaccuracies in gate execution accumulate throughout a circuit and can significantly affect final measurement outcomes [17]. These errors become particularly important in algorithms that rely on repeated circuit evaluations, such as variational quantum learning systems.

The combined effects of decoherence and gate imperfections place practical limits on the complexity of executable quantum computations. As circuit depth increases, maintaining an accurate quantum state becomes progressively more difficult. Consequently, many NISQ-era algorithms are intentionally designed to remain shallow in order to preserve execution reliability.

Table 2.1 summarizes representative hardware characteristics commonly observed in current quantum devices, including coherence times, gate error rates, and measurement fidelities. These values illustrate the practical constraints under which near-term quantum algorithms must operate.

Table 2.1: Representative characteristics of major NISQ hardware platforms. Values represent typical ranges reported for contemporary devices and illustrate the practical constraints affecting hybrid quantum learning systems.

Platform	Coherence	Gate Fidelity	Readout
Superconducting	50–300 μ s	1Q: 99.8–99.99% 2Q: 98–99.8%	95–99%
Trapped-Ion	1–100 s	1Q: 99.9–99.99% 2Q: 99–99.9%	99–99.9%
Neutral Atom	100 ms–10 s	1Q: 99–99.9% 2Q: 97–99.5%	98–99%
Photonic	No intrinsic decoherence	Platform dependent	95–99%

An important observation is that multi-qubit operations generally exhibit substantially higher error rates than single-qubit gates. Because entangling operations are central to many

quantum algorithms and machine learning architectures, their reliability often becomes a major determinant of overall system performance. As a result, practical circuit design frequently involves balancing representational capability against accumulated hardware error.

2.6.2 Circuit Depth and Scalability Limitations

The limitations imposed by coherence time and gate fidelity directly constrain the scalability of quantum computation. In principle, increasing the number of qubits and circuit depth allows quantum systems to represent more complex functions and computational structures. In practice, however, deeper circuits accumulate noise more rapidly, reducing the reliability of both computation and measurement [6], [9].

This creates a fundamental tradeoff between computational expressivity and execution stability. Additional circuit layers may increase representational flexibility, but they also increase exposure to decoherence, gate errors, and stochastic perturbations. Consequently, circuits that appear highly expressive under idealized simulation conditions may become impractical when executed on real hardware.

Scalability limitations extend beyond circuit depth alone. Increasing the number of qubits introduces additional calibration complexity, control challenges, and measurement instability. Larger systems require more operations, more qubit interactions, and more repeated measurements, all of which increase the probability of accumulated error during execution.

These constraints are especially important in variational quantum learning systems, where optimization requires repeated evaluations of parameterized quantum circuits. The combination of deep circuits, noisy measurements, and iterative optimization can significantly reduce convergence reliability and practical performance.

For this reason, many successful NISQ-era algorithms emphasize shallow, hardware-efficient circuit structures rather than maximizing circuit complexity. This design philosophy appears repeatedly throughout modern hybrid quantum learning research and forms an important foundation for the robustness-oriented framework developed later in this thesis.

2.6.3 Hardware Variability and Readout Noise

In addition to coherence limitations and gate errors, current quantum hardware exhibits substantial variability across devices, executions, and calibration states. Quantum processors are not perfectly stationary systems; their behavior changes over time due to fluctuations in temperature, calibration drift, environmental interference, and device-specific operating conditions.

This variability introduces additional uncertainty into quantum computation. Identical circuits executed on different hardware backends—or even on the same backend at different times—may produce measurably different outputs. Such fluctuations complicate reproducibility and make performance evaluation more challenging than in conventional deterministic computing environments.

Measurement reliability represents another major challenge. Readout noise occurs during the final stage of computation, when quantum states are converted into classical outcomes. Imperfect measurement processes can produce incorrect bit values even when the underlying quantum state has been prepared correctly [14]–[16]. Because many quantum algorithms depend heavily on repeated measurements, readout noise can significantly influence computational accuracy.

The impact of readout noise is especially important in variational algorithms and hybrid quantum learning systems. Measurement outcomes are used repeatedly during optimization to estimate objective functions and guide parameter updates. Variability in these measurements can therefore contribute directly to optimization instability and inconsistent convergence behavior.

Taken together, coherence limitations, gate errors, hardware variability, and readout noise define many of the practical boundaries of current quantum computation. Understanding these constraints is essential for evaluating both the capabilities and limitations of quantum algorithms operating under realistic NISQ-era conditions and provides important context for the robustness analyses presented later in this thesis.

2.7 Quantum Hardware Platforms and Scalability

The practical capabilities of quantum algorithms are shaped not only by their mathematical formulation but also by the physical hardware on which they are implemented. Different quantum computing platforms vary substantially in terms of qubit implementation, coherence behavior, gate fidelity, connectivity, scalability, and noise characteristics. These differences directly influence the reliability and complexity of executable quantum circuits and therefore play an important role in determining the practical performance of quantum algorithms.

Several hardware architectures have emerged as leading candidates for scalable quantum computation in the NISQ era. Among the most prominent are superconducting quantum systems, trapped-ion quantum computers, and neutral atom platforms. Each architecture offers distinct advantages and limitations regarding coherence time, gate fidelity, connectivity, and scalability, creating different tradeoffs between computational capability and hardware reliability.

At the same time, cloud-based quantum computing ecosystems have expanded access to quantum hardware by allowing researchers to execute circuits remotely through integrated software platforms. These environments have accelerated the development of hybrid quantum learning systems by enabling experimentation across multiple hardware backends without requiring direct access to specialized laboratory infrastructure.

The following subsections briefly review several major quantum hardware architectures and discuss their relevance to practical NISQ-era quantum computation.

2.7.1 Superconducting Quantum Systems

Superconducting quantum systems represent one of the most widely developed hardware architectures for quantum computation and currently form the basis of several major industrial quantum platforms, including those developed by IBM, Google, and Rigetti [20], [21]. In these systems, qubits are implemented using superconducting electrical circuits operating at

temperatures near absolute zero, where quantum behavior can be maintained and controlled through microwave-based operations.

One of the primary advantages of superconducting architectures is their fast gate execution speed and compatibility with large-scale fabrication techniques derived from semiconductor manufacturing. These characteristics have enabled rapid growth in qubit counts and have made superconducting systems the dominant platform for many cloud-accessible quantum computing services.

However, superconducting qubits also exhibit relatively short coherence times and remain vulnerable to gate errors, calibration drift, crosstalk, and readout instability. As a result, practical algorithms executed on these systems frequently rely on shallow circuit structures, variational optimization methods, and noise-aware strategies designed to operate within current NISQ limitations.

2.7.2 Trapped-Ion Quantum Computing

Trapped-ion quantum computing represents an alternative hardware approach in which qubits are encoded within the internal electronic states of charged ions confined by electromagnetic fields [17]. Quantum operations are performed using precisely controlled laser pulses that manipulate the quantum states and create interactions between ions.

A major advantage of trapped-ion systems is their exceptionally long coherence times and high gate fidelities. In addition, many trapped-ion architectures provide flexible qubit connectivity, allowing interactions between qubits that would require additional routing operations in other hardware platforms. These properties make trapped-ion systems attractive for experiments requiring high precision and reliable execution.

Despite these strengths, trapped-ion hardware faces scalability challenges. Gate operations are generally slower than those used in superconducting architectures, and maintaining stable control over increasingly large ion chains introduces additional engineering complexity. Consequently, trapped-ion systems often provide an important contrast between execution

fidelity and scalability within NISQ-era quantum computing.

2.7.3 Neutral Atom Quantum Computing

Neutral atom quantum computing has emerged as another promising approach to scalable quantum computation. In these systems, individual atoms are trapped and manipulated using optical tweezers or laser-generated trapping arrays, while quantum operations are implemented through controlled atomic interactions.

A major advantage of neutral atom architectures is their scalability potential. Large arrays of atoms can be arranged in highly configurable geometries, providing flexibility for constructing programmable quantum systems. Many implementations rely on Rydberg interactions, which enable controllable entanglement between nearby atoms and support the construction of complex quantum circuits.

Although neutral atom systems continue to improve rapidly, they remain experimentally challenging and face limitations related to gate fidelity, control precision, and large-scale calibration stability. As with other NISQ platforms, practical performance ultimately depends not only on scalability but also on reliability and robustness under realistic operating conditions.

2.7.4 Cloud Quantum Platforms and Azure Quantum

The emergence of cloud-based quantum computing platforms has significantly expanded access to quantum hardware and accelerated research in hybrid quantum machine learning. Rather than requiring direct access to specialized laboratory systems, researchers can now execute quantum circuits remotely through cloud services provided by major technology companies and research organizations.

Platforms such as IBM Quantum, Amazon Braket, Google Quantum AI, and Microsoft Azure Quantum provide integrated ecosystems combining hardware access, simulation tools, software development frameworks, and runtime management services [20], [23]. These en-

vironments allow researchers to evaluate algorithms across multiple hardware architectures while maintaining consistent development workflows.

Cloud platforms are particularly important for hybrid quantum–classical systems because they naturally support distributed computational workflows. Classical preprocessing, optimization, and post-processing can be performed on conventional computing infrastructure while quantum circuits are executed remotely on specialized hardware. In this thesis, cloud-accessible platforms play an important role in benchmarking, reproducibility, and hardware-validation experiments involving hybrid quantum learning systems.

2.8 Comparison to Prior Work

Recent research in Hybrid Quantum Neural Networks has increasingly investigated the impact of realistic Noisy Intermediate-Scale Quantum (NISQ) hardware on quantum machine learning. Previous studies have demonstrated that decoherence, gate errors, measurement uncertainty, and limited circuit depth can significantly reduce predictive performance and optimization stability. These works have established the importance of robustness as a practical challenge for near-term quantum computing and provide the foundation upon which this thesis builds.

Although this body of literature has produced important advances, existing studies generally focus on individual architectures, isolated robustness mechanisms, or selected quantum noise models. Consequently, direct comparison across different Hybrid Quantum Neural Network approaches remains difficult because experiments frequently use different datasets, software frameworks, evaluation metrics, and training methodologies. In many cases, robustness is evaluated under only one or two perturbation sources, making it difficult to determine whether the observed behavior generalizes across representative NISQ environments.

This thesis extends previous work by introducing a unified robustness characterization methodology that evaluates multiple Hybrid Quantum Neural Network architectures, Quantum Support Vector Machines, and classical machine learning models under identical bench-

mark conditions. The proposed framework systematically analyzes optimization dynamics, cross-noise robustness, architecture sensitivity, degradation behavior, and simulator-to-hardware transferability using standardized datasets, evaluation metrics, and software frameworks. This methodology enables direct comparison of robustness across different learning paradigms while providing a consistent experimental basis for evaluating practical Hybrid Quantum Machine Learning systems.

The following subsections compare representative HQNN robustness studies with the methodology developed in this thesis, identify the limitations of current evaluation strategies, and explain how the proposed framework extends existing work through systematic robustness characterization and validation.

2.8.1 Existing HQNN Noise Studies

A substantial portion of recent Hybrid Quantum Neural Network research has focused on understanding how parameterized quantum models behave under realistic Noisy Intermediate-Scale Quantum conditions. Several foundational studies established the computational structures upon which modern HQNN systems are built, while more recent work has examined how those structures respond to noise, optimization instability, and hardware limitations.

Havlíček et al. [11] demonstrated how quantum circuits can construct quantum-enhanced feature spaces for supervised learning. Their work established an important foundation for quantum kernel methods, but it primarily focused on the representational capability of quantum feature maps rather than on systematic evaluation of training robustness across multiple HQNN architectures and noise channels. In contrast, the present thesis uses quantum kernel methods as one comparative baseline within a broader framework that also evaluates classical models, variational HQNNs, optimization behavior, and progressive degradation under noise.

Mitarai et al. [12] introduced quantum circuit learning as a general variational framework in which parameterized quantum circuits are optimized through classical feedback. Their

work established the trainable circuit model that underlies many contemporary HQNN architectures. However, the primary objective was to demonstrate the learning capability of parameterized quantum circuits rather than to characterize how measurement interpretation, architecture topology, and noisy optimization interact across standardized benchmark environments. This thesis extends that foundation by analyzing these factors jointly and by treating robustness as a systems-level design objective.

Benedetti et al. [13] provided a broader analysis of parameterized quantum circuits as machine learning models and emphasized the influence of encoding, circuit structure, optimization, and measurement on learning behavior. Their work offers an important conceptual framework for understanding variational quantum models. The present thesis builds upon that perspective but differs in experimental scope and purpose. Rather than surveying parameterized circuit design generally, this work evaluates multiple Hybrid Quantum Neural Network architectures, Quantum Support Vector Machines, and classical baselines under identical datasets, metrics, software environments, and representative quantum noise channels.

Kashif et al. [24] investigated the effect of quantum noise on the training performance of Hybrid Quantum Neural Networks and directly demonstrated that phase-flip, bit-flip, phase-damping, amplitude-damping, and depolarizing errors can influence convergence and classification behavior. This study is the closest prior work to the noise-analysis component of the present thesis. However, its primary emphasis is the effect of different noise channels on HQNN training within a selected experimental architecture. The present work extends this direction by combining baseline failure analysis, classical and quantum benchmarking, cross-framework evaluation, progressive noise sweeps, cross-noise heatmaps, architecture comparison, robustness-mechanism ablation, and IBM hardware validation within one unified methodology.

Prior research has also shown that robustness is strongly architecture dependent. Quantum convolutional, convolutional, and other modular quantum architectures may respond differently

to perturbation because of differences in circuit depth, entanglement topology, feature encoding, and measurement strategy. These findings motivate architecture-aware evaluation, but direct comparison across studies remains difficult when each investigation uses different datasets, preprocessing procedures, optimizers, metrics, and software frameworks.

The contribution of this thesis is therefore not the isolated observation that quantum noise affects learning. That conclusion is already well established in the literature. The contribution is the development of a standardized robustness-characterization process that evaluates multiple computational paradigms under controlled conditions and then uses the identified failure modes to guide the development of learned-readout decoding, multi-observable feature extraction, architecture optimization, and stability-aware training.

Collectively, the cited studies establish the foundations of quantum feature spaces, parameterized quantum circuit learning, and noise-sensitive HQNN training. This thesis builds upon those foundations by integrating them into a unified systems-level methodology for identifying, comparing, and engineering robustness under realistic NISQ conditions.

2.8.2 Limitations of Current Robustness Evaluation

Although Havlíček et al. [11], Mitarai et al. [12], Benedetti et al. [13], and Kashif et al. [24] provide important foundations for quantum feature learning, parameterized circuit models, and noisy HQNN training, their experimental objectives differ from the unified robustness-characterization objective of this thesis. Havlíček et al. primarily examine quantum-enhanced feature spaces, Mitarai et al. establish a trainable quantum circuit learning model, Benedetti et al. analyze parameterized quantum circuits as machine learning models, and Kashif et al. investigate the influence of representative noise channels on HQNN training. Together, these studies establish the principal components of the research area, but they do not provide one standardized methodology combining classical, kernel-based quantum, and variational hybrid models across identical datasets, software environments, optimization conditions, and noise channels.

One major limitation within the broader literature is therefore the absence of standardized benchmarking across computational paradigms. Because individual studies often use different datasets, feature maps, circuit structures, optimization procedures, shot counts, metrics, and perturbation assumptions, numerical results cannot always be compared directly. A model reported as robust in one study may have been evaluated under substantially different conditions from a model evaluated elsewhere.

The benchmarking methodology proposed in this thesis addresses this limitation by evaluating classical machine learning models, Quantum Support Vector Machines, and multiple Hybrid Quantum Neural Network architectures under shared preprocessing, training, testing, and perturbation conditions. This controlled design makes it possible to attribute observed differences more directly to architecture, optimization behavior, measurement interpretation, and noise sensitivity rather than to incompatible experimental protocols. A second limitation concerns the scope of robustness analysis itself. Existing work often evaluates predictive accuracy under noisy conditions while devoting comparatively less attention to optimization stability, convergence behavior, architecture structure, and the interaction between noise and training dynamics. Consequently, robustness is frequently treated as a performance metric rather than as a broader systems-level property emerging from multiple interacting factors. Unlike many previous studies that primarily report final predictive accuracy, this thesis additionally analyzes optimization dynamics, convergence stability, degradation trends, and cross-noise robustness, providing a more comprehensive characterization of HQNN behavior under realistic noisy execution.

Another challenge involves the heavy reliance on simplified simulation environments. While simulator-based studies provide useful experimental control, they do not always capture calibration drift, backend variability, crosstalk, measurement uncertainty, and execution-time fluctuations that influence practical quantum hardware behavior. As a result, conclusions drawn from idealized simulations do not always transfer directly to real NISQ devices. To complement simulator-based evaluation, the proposed framework validates the

final robustness-engineered HQNN directly on IBM Quantum hardware, allowing simulator performance to be compared with practical quantum execution.

Cross-platform validation also remains relatively limited. Quantum machine learning systems implemented in different software ecosystems can exhibit differences arising from transpilation procedures, backend execution characteristics, and noise-model assumptions. Without systematic cross-framework analysis, it is difficult to determine whether observed behavior reflects genuine architectural properties or implementation-specific artifacts. The framework developed in this thesis addresses this issue through validation across Qiskit, Cirq, and PennyLane, reducing the likelihood that observed differences originate solely from software implementation.

Taken together, these limitations suggest that the field requires more integrated methodologies capable of combining architecture comparison, optimization analysis, multichannel perturbation evaluation, reproducible benchmarking, and hardware validation within one unified experimental framework.

2.8.3 Gap Addressed by This Thesis

The limitations identified throughout the existing literature motivate the need for a more systematic methodology for evaluating robustness in Hybrid Quantum Machine Learning. Rather than investigating isolated architectures, individual noise models, or single experimental configurations, there remains a need for a unified framework capable of comparing multiple learning paradigms under standardized benchmark conditions while simultaneously analyzing optimization stability, robustness degradation, and practical hardware behavior.

This thesis addresses that gap through the development of a systematic robustness characterization methodology and a robustness-engineering framework for Hybrid Quantum Neural Networks. The proposed methodology combines unified benchmarking, cross-noise evaluation, optimization analysis, architecture comparison, component-level ablation, threshold progression analysis, and IBM Quantum hardware validation within one reproducible exper-

imental framework.

Unlike previous studies that typically introduce one new HQNN architecture or one isolated robustness mechanism, this thesis adopts a systems-engineering perspective. Robustness is treated as a property emerging from the coordinated interaction between measurement interpretation, quantum feature representation, circuit architecture, optimization strategy, and noise-aware learning rather than from any single architectural modification. This perspective motivates the staged robustness-engineering framework developed throughout Chapter 4.

A central premise of this thesis is that robustness, rather than quantum expressive capability alone, represents the dominant bottleneck limiting practical Hybrid Quantum Neural Networks under realistic NISQ conditions. Consequently, the framework developed in later chapters treats robustness as a first-class design objective instead of a secondary evaluation metric.

By combining systematic robustness characterization, unified cross-noise benchmarking, staged robustness engineering, component-level validation, and simulator-to-hardware transferability, this thesis contributes a comprehensive methodology for evaluating, designing, and improving Hybrid Quantum Neural Networks under realistic quantum hardware constraints. These contributions establish the experimental foundation for the benchmarking studies presented in Chapter 3 and the robustness-engineering framework developed in Chapter 4.

Chapter 3

Establishing Benchmark Comparisons

3.1 Purpose of the Benchmarking Framework

The practical evaluation of Hybrid Quantum Neural Networks requires more than isolated demonstrations of model performance under idealized conditions. In the NISQ era, quantum machine learning systems operate within environments shaped by noise, optimization instability, hardware variability, and architectural constraints. Consequently, understanding the practical reliability of HQNN systems requires a systematic methodology capable of evaluating how different models behave across a range of computational and perturbation environments.

Previous quantum machine learning studies have frequently evaluated individual architectures, selected datasets, or isolated quantum noise models using different software frameworks and experimental protocols. While these investigations have provided valuable insight into specific aspects of HQNN behavior, direct comparison between studies remains difficult because of differences in datasets, evaluation metrics, optimization procedures, and perturbation assumptions. As a result, robustness has often been evaluated within isolated experimental settings rather than through a standardized comparative methodology.

To address this limitation, this chapter develops a structured benchmarking framework for evaluating hybrid quantum learning systems under realistic NISQ conditions. Rather than treating benchmarking as a secondary validation step, comparative evaluation is positioned as a central component of the research methodology. The objective is not only to determine whether a model performs well under a selected configuration, but also to understand how performance changes as noise environments, optimization settings, architectural

structures, and computational constraints vary.

Unlike many previous HQNN studies, the proposed framework evaluates classical machine learning models, Quantum Support Vector Machines, and multiple Hybrid Quantum Neural Network architectures under identical datasets, software frameworks, evaluation metrics, optimization procedures, and representative quantum noise channels. This standardized methodology enables direct comparison of predictive performance, optimization stability, robustness degradation, and reproducibility across different learning paradigms while providing a consistent experimental basis for evaluating practical Hybrid Quantum Machine Learning systems.

The relationship to prior work is important. Havlíček et al. [11] provide a quantum-kernel model that separates quantum feature generation from classical optimization, while Mitarai et al. [12] and Benedetti et al. [13] establish parameterized quantum circuits as trainable learning models. Kashif et al. [24] further demonstrate that representative quantum errors can substantially alter HQNN training behavior. The benchmarking framework developed here brings these previously separate perspectives into one controlled evaluation structure by comparing classical baselines, quantum kernels, and variational HQNNs under shared datasets, metrics, optimization settings, software frameworks, and noise environments.

Figure 3.1 summarizes the overall benchmarking methodology developed throughout this chapter. The framework serves as the experimental foundation for the remainder of the thesis and provides the structure through which robustness behavior is analyzed systematically across multiple architectures and perturbation environments.

Benchmark Data	Model Families	Training Setup	Noise Tests
Iris, WDBC, synthetic datasets	SVM, MLP, QSVM, baseline, HQNN	SPSA, circuit depth, shot count	Depolarizing, bit-flip, phase-flip, damping

Evaluation Metrics	Cross-Framework Validation	Robustness Analysis
Accuracy, AUC, loss variance	Qiskit, Cirq, PennyLane	Degradation trends, stability, reproducibility

Systems-Level Interpretation: predictive accuracy, convergence behavior, noise sensitivity, and reproducibility under realistic NISQ conditions.

Figure 3.1: The proposed benchmarking framework compares standardized benchmark datasets, classical and quantum learning models, training configurations, representative quantum noise channels, evaluation metrics, and cross-framework validation results under identical experimental conditions. This unified methodology enables systematic comparison of predictive performance, optimization stability, robustness degradation, and reproducibility across multiple Hybrid Quantum Machine Learning paradigms operating under realistic NISQ environments.

An important characteristic of the proposed framework is its interpretation of robustness as a systems-level property rather than solely a final predictive accuracy metric. In this thesis, robustness encompasses optimization stability, convergence consistency, degradation behavior, reproducibility, and resilience to realistic quantum noise. This broader perspective distinguishes the proposed methodology from many previous HQNN evaluations, which often emphasize predictive accuracy under isolated experimental conditions, and provides the experimental foundation for the robustness-engineering framework developed in Chapter 4.

3.1.1 Why HQNN Robustness Requires Systematic Comparison

Hybrid quantum neural networks exhibit highly variable behavior across different architectures, datasets, optimization strategies, and noise environments. Small changes in circuit depth, entanglement structure, feature encoding, or optimizer configuration can produce substantial differences in convergence stability and noisy inference performance. Consequently, isolated evaluations of individual HQNN systems often provide only limited insight into the broader reliability of hybrid quantum learning under realistic execution conditions.

This variability becomes especially important in the NISQ regime, where hardware imperfections strongly influence computational behavior. Variational quantum systems are sensitive not only to decoherence and gate errors, but also to optimization instability arising from noisy measurements and stochastic parameter updates [9], [13]. As a result, architectures that appear similar under noiseless simulation may behave very differently once realistic perturbation environments are introduced.

Recent HQNN studies have demonstrated this sensitivity explicitly. Comparative analyses involving QuanNNs, QCNNs, and QTL architectures reported substantial differences in degradation behavior under depolarizing noise, amplitude damping, bit-flip noise, and phase-flip noise. Some architectures maintained relatively stable behavior under moderate perturbation conditions, whereas others experienced rapid performance collapse as noise intensity increased.

These observations suggest that robustness cannot be evaluated reliably through isolated experiments alone. Meaningful analysis requires systematic comparison across multiple architectures, datasets, optimization conditions, and perturbation environments. Only through such comparative evaluation can broader robustness patterns and architectural tradeoffs be identified.

Systematic benchmarking therefore serves as the foundation for understanding how robustness emerges within practical hybrid quantum learning systems. By evaluating architectures consistently across controlled environments, the framework developed in this chapter

seeks to identify the design principles that support reliable behavior under realistic NISQ constraints.

3.1.2 Limitations of Existing Noise Evaluation Methods

Although noise analysis has become increasingly common in hybrid quantum machine learning research, many existing evaluation methodologies remain limited in scope and consistency. A common limitation is the reliance on isolated noise studies that evaluate model performance under only one or two perturbation environments without examining broader robustness behavior systematically.

In many cases, robustness is assessed primarily through final predictive accuracy measurements. While such measurements provide useful information, they often fail to capture other important aspects of practical reliability, including optimization stability, convergence consistency, sensitivity to initialization, and degradation dynamics during training. Consequently, two models with similar noisy accuracy may exhibit substantially different behavior during optimization.

Another important limitation concerns reproducibility. Different studies frequently employ incompatible datasets, noise models, optimization procedures, and software frameworks, making direct comparison difficult. Furthermore, many evaluations rely heavily on idealized simulation environments that do not fully reflect the variability of real quantum hardware systems.

Cross-framework consistency is also rarely examined systematically. Implementations in Qiskit, Cirq, PennyLane, and hardware-specific runtimes may exhibit differences arising from compilation procedures, noise-model assumptions, and backend execution characteristics. Without cross-platform validation, it becomes difficult to determine whether observed robustness behavior reflects genuine architectural properties or implementation-specific effects.

These limitations motivate the need for a more integrated evaluation methodology ca-

pable of combining standardized datasets, multichannel perturbation analysis, architecture comparison, optimization assessment, and reproducible benchmarking within a unified experimental framework.

3.1.3 Role of Classical, Quantum, and Hybrid Baselines

A central component of the benchmarking framework is the use of comparative baselines spanning classical machine learning systems, quantum kernel methods, and hybrid quantum neural network architectures. These baselines provide essential reference points for interpreting the practical behavior of hybrid quantum learning systems under realistic noisy conditions.

Classical machine learning models serve as the first major baseline category. Because modern classical algorithms remain highly effective across many structured learning tasks, comparison against strong classical methods is necessary for evaluating whether observed quantum behavior reflects genuinely useful computational advantages or simply differences in evaluation methodology.

Quantum kernel methods, particularly Quantum Support Vector Machines (QSVMs), provide a second important baseline. Unlike fully variational quantum neural networks, QSVM systems separate quantum feature generation from classical optimization and therefore often exhibit greater optimization stability under noisy conditions. These methods help distinguish challenges associated with variational optimization from those associated with quantum feature representation itself.

Hybrid quantum neural network architectures form the third major baseline category. Throughout this work, multiple HQNN variants—including QuanNNs, QCNNs, and QTL systems—are evaluated systematically. These architectures differ substantially in circuit structure, entanglement behavior, modularity, and optimization dynamics, allowing the benchmarking framework to analyze how architectural decisions influence robustness and stability under noise.

Together, these baseline categories provide the comparative foundation necessary for evaluating practical NISQ-era quantum machine learning. By examining classical, quantum, and hybrid approaches under identical experimental conditions, the framework establishes the reference structure used throughout the remainder of the benchmarking studies.

3.2 Baseline HQNN Behavior and Failure Modes

Before developing robustness-oriented hybrid quantum learning systems, it is necessary to establish how baseline HQNN architectures behave under realistic NISQ conditions. Although hybrid quantum neural networks are often presented as promising candidates for near-term quantum machine learning, their practical behavior is frequently constrained by noise, measurement variability, and optimization instability. Understanding these limitations is essential because they motivate the robustness-engineering framework developed throughout the remainder of this thesis.

The experiments presented in this section evaluate baseline HQNN systems prior to the introduction of learned-readout decoding, multi-observable feature extraction, architecture optimization, and stability-aware training procedures. These baseline models rely primarily on shallow variational quantum circuits, parity-based measurement strategies, and conventional optimization methods commonly used in early NISQ-era quantum machine learning studies [12], [13].

While such architectures can perform meaningful learning tasks under idealized simulation conditions, their behavior often changes significantly when realistic perturbation environments are introduced. Noise affects not only final predictive performance, but also the optimization process itself. Training trajectories become increasingly unstable, convergence behavior becomes inconsistent, and degradation frequently appears even under moderate perturbation levels.

The experiments reveal several recurring failure modes. Variational optimization often produces oscillatory behavior under noisy conditions, measurement variability amplifies in-

ference instability, and deeper circuits accumulate perturbations more rapidly than they gain expressive capability. These observations suggest that practical HQNN performance is constrained by robustness and optimization reliability rather than by expressivity alone.

The following subsections examine the baseline parity-based architecture, evaluate its classification behavior, analyze SPSA optimization instability, and establish the initial robustness thresholds that motivate the framework evolution developed in Chapter 4.

3.2.1 Baseline Parity-Based HQNN Design

The baseline HQNN architecture used throughout the initial benchmarking experiments follows a conventional variational quantum learning design commonly employed in early NISQ-era quantum machine learning research [12], [13]. The system combines classical feature preprocessing with a shallow parameterized quantum circuit whose outputs are interpreted through a parity-based measurement procedure.

Classical input features are first encoded into quantum states using parameterized rotation gates. These encoded states are then processed through variational layers composed of trainable single-qubit rotations and entangling operations. Following circuit execution, measurements are performed on the quantum register and the resulting bitstrings are converted into class predictions through parity-based decoding.

The parity-readout mechanism represents one of the simplest methods for transforming quantum measurement outcomes into classical predictions. In this approach, the parity of a measured bitstring determines the predicted class label. Although computationally straightforward, the procedure introduces a substantial information bottleneck because many distinct quantum measurement distributions collapse into the same binary output.

Figure 3.2 shows the baseline parity-based HQNN architecture.




Figure 3.2: Classical features are encoded into a parameterized quantum circuit, measured as bitstring samples, and converted into binary predictions through parity decoding. Although computationally efficient and compatible with shallow NISQ circuits, parity decoding compresses rich quantum measurement distributions into a limited binary representation, creating an information bottleneck.

One advantage of this architecture is its compatibility with NISQ hardware constraints. The circuit remains shallow and requires relatively few quantum resources. However, the simplicity of parity-based decoding also limits the amount of information extracted from the quantum system and increases sensitivity to measurement fluctuations. These characteristics make the architecture an effective baseline for studying the practical limitations of hybrid quantum learning under noisy conditions.

3.2.2 Baseline Classification Behavior

The first benchmark evaluates the behavior of the baseline HQNN architecture on a controlled classification task designed to examine hybrid quantum learning under simplified conditions. The objective is not to demonstrate state-of-the-art predictive performance, but rather to establish how a standard variational HQNN behaves before the introduction of robustness-oriented modifications.

Under noiseless simulation conditions, the baseline model is capable of learning simple decision boundaries and achieving moderate classification performance. However, even in

these idealized settings, training behavior exhibits sensitivity to parameter initialization and optimizer configuration. Different runs frequently produce noticeably different convergence trajectories.

When realistic perturbation channels are introduced, the behavior of the model becomes substantially less stable. Measurement fluctuations and noisy parameter updates distort optimization trajectories, causing classification accuracy to oscillate throughout training. In several runs, the model exhibits inconsistent convergence behavior and unstable decision-boundary formation.

An important observation is that degradation does not occur uniformly throughout training. In many cases, the model temporarily approaches reasonable performance before later destabilizing as optimization proceeds. This suggests that optimization dynamics themselves become a major source of performance degradation in variational quantum learning systems.

The experiment also highlights a limitation of parity-based decoding. Because predictions depend on compressed parity outputs rather than richer measurement distributions, relatively small perturbations in measured bitstrings can produce disproportionately large changes in final predictions. These observations provide the first indication that robustness challenges originate not only from hardware noise, but also from the way quantum information is interpreted within the learning pipeline.

3.2.3 SPSA Training Instability Under Noise

To further investigate optimization behavior, the next benchmark evaluates baseline HQNN training using Simultaneous Perturbation Stochastic Approximation (SPSA), one of the most widely used optimization methods for variational quantum algorithms [9]. SPSA is attractive in NISQ environments because it estimates gradients efficiently using only a small number of circuit evaluations.

Despite this efficiency, SPSA optimization becomes highly sensitive to stochastic perturbations under realistic noisy conditions. Because gradient estimates are constructed from

repeated noisy measurements, parameter updates may fluctuate substantially between optimization steps. These fluctuations become increasingly severe as hardware noise and measurement variability distort the underlying objective-function evaluations.

The experiments performed in this work demonstrate that baseline HQNN systems trained with SPSA frequently exhibit oscillatory convergence behavior. Rather than converging smoothly toward stable minima, optimization trajectories often fluctuate around unstable regions of the loss landscape. In several runs, classification performance approaches near-random behavior despite extended optimization.

Figure 3.3 shows the SPSA training instability under noise.




Figure 3.3: The noisy SPSA trajectory oscillates around the binary cross-entropy random-classification region rather than converging smoothly, illustrating how stochastic measurement perturbations can destabilize variational optimization under realistic NISQ conditions.

Another important observation is that instability increases as circuit complexity grows. Additional variational layers and entangling operations increase measurement variability and make the optimization landscape more difficult to navigate. Consequently, deeper circuits often produce poorer practical behavior despite possessing greater theoretical expressive

capability.

These results suggest that optimization instability represents one of the primary practical bottlenecks in NISQ-era hybrid quantum learning. The challenge is not simply that noise reduces final predictive accuracy, but that noise fundamentally alters the optimization dynamics governing variational learning itself.

3.2.4 Initial Threshold Limitations

The baseline experiments establish an important reference point for the remainder of the thesis. Across multiple noisy evaluation environments, parity-based HQNN architectures frequently exhibit noisy classification accuracy in the approximate range of 0.37 to 0.50, depending on optimizer configuration, circuit structure, and perturbation conditions.

These results indicate that many baseline variational HQNN systems operate only marginally above random guessing under realistic noisy environments. While isolated runs occasionally achieve higher temporary performance, optimization instability and stochastic degradation frequently prevent the models from maintaining stable behavior over time.

A particularly important observation is that increasing circuit complexity does not reliably improve practical performance. In many cases, additional entanglement layers and deeper variational structures increase instability faster than they improve representational capability. This finding challenges the assumption that greater quantum complexity necessarily produces better learning behavior under realistic NISQ conditions.

The results therefore suggest that the dominant limitation is not a lack of expressive capability, but rather the inability to maintain reliable optimization and inference under perturbation. This distinction becomes one of the central themes of the thesis and directly motivates the robustness-engineering framework developed in Chapter 4.

The baseline experiments provide a clear starting point for that progression. They demonstrate that parity-based decoding introduces severe information bottlenecks, variational optimization becomes unstable under stochastic perturbations, and practical hybrid quantum

learning requires more than simply increasing circuit depth. These findings establish the experimental foundation upon which the subsequent framework evolution is constructed.

3.3 Classical, QSVM, and HQNN Benchmark Comparisons

A central objective of the benchmarking framework is to evaluate hybrid quantum neural networks relative to both classical machine learning baselines and alternative quantum learning approaches under controlled and reproducible conditions. Evaluating HQNN architectures in isolation provides only limited insight into their practical usefulness because performance must ultimately be interpreted relative to competing computational paradigms operating under the same constraints.

For this reason, the framework integrates three major model categories: classical machine learning systems, quantum kernel methods, and hybrid quantum neural network architectures. Classical models provide strong reference baselines representing mature and highly optimized learning techniques. Quantum kernel methods such as QSVMs provide a more structured quantum learning paradigm, while HQNN architectures offer adaptive variational models capable of learning quantum feature representations directly.

This comparative structure allows the framework to examine not only whether hybrid quantum systems achieve meaningful performance, but also under what conditions they exhibit advantages, limitations, or instability relative to alternative approaches. Particular attention is given to the influence of noise sensitivity, optimization behavior, architectural complexity, and feature representation on practical model performance under realistic NISQ conditions. While Havlíček et al. primarily evaluated quantum kernel methods, this thesis compares quantum kernel models directly with multiple HQNN architectures and classical baselines under identical benchmark conditions.

The benchmark results reveal several recurring patterns. Classical machine learning systems remain highly competitive across many structured learning tasks. Quantum kernel methods generally provide stable behavior but limited adaptability, while variational HQNN

architectures offer greater flexibility at the cost of increased sensitivity to optimization instability and noise.

The following subsections examine these comparisons in greater detail and establish the performance reference points used throughout the remainder of the thesis.

3.3.1 Quantum Kernel SVM Benchmark

The Quantum Support Vector Machine (QSVM) serves as one of the primary comparative baselines within the benchmarking framework. Unlike variational HQNN systems that require repeated parameter optimization during training, QSVM architectures separate quantum feature generation from classifier optimization. This distinction allows the framework to evaluate how much instability originates from variational optimization rather than from quantum feature representation itself [10], [11].

Within the QSVM pipeline, classical data is encoded into quantum states through a feature map, and pairwise similarities between samples are computed using quantum kernel estimation. These kernel values are then supplied to a classical support vector machine that performs the optimization and classification stages. Because optimization remains entirely classical, QSVM systems often exhibit greater stability than fully variational learning architectures. Across the benchmark experiments, QSVM systems demonstrated relatively strong and consistent classification performance. Under controlled simulation conditions, several experiments achieved classification accuracy near 0.84 while maintaining stable optimization behavior under moderate perturbation levels.

An important observation is that QSVM systems often degrade more gracefully under noise than variational HQNN architectures. Because the quantum circuit is not updated repeatedly during training, optimization instability is reduced substantially relative to trainable quantum models.

At the same time, QSVM performance depends strongly on the choice of feature map and the quality of quantum kernel estimation. As hardware noise increases, kernel fidelity

deteriorates, limiting classification reliability and reducing practical performance.

3.3.2 Classical Baseline Performance

Classical machine learning models provide essential reference points for evaluating practical hybrid quantum learning systems. Modern classical algorithms remain highly optimized, scalable, and stable across a wide range of structured learning tasks. Consequently, meaningful assessment of HQNN performance requires direct comparison against strong classical baselines operating under equivalent experimental conditions.

The classical benchmarks used throughout this work were evaluated using the same preprocessing procedures, training-validation protocols, and performance metrics applied to the quantum models. This consistency ensures that observed differences arise from the computational approach itself rather than from methodological variations.

Across many structured datasets, the classical baselines consistently demonstrated strong predictive performance together with highly stable optimization behavior. In several benchmark tasks, classical models achieved classification accuracy near or above 0.95 while maintaining smooth convergence trajectories under conditions where variational quantum systems exhibited substantial instability.

These results indicate that practical NISQ-era hybrid quantum learning should not be interpreted as a universal replacement for classical machine learning. Instead, classical systems remain highly competitive across many conventional learning environments, particularly when optimization stability and robustness are considered.

For this reason, the goal of the benchmarking framework is not to demonstrate universal quantum superiority, but rather to identify the conditions under which hybrid quantum representations may provide useful complementary computational behavior.

3.3.3 HQNN Performance Gap on Structured Data

The comparative benchmarking experiments reveal that baseline HQNN architectures frequently exhibit a measurable performance gap relative to strong classical machine learning systems on structured datasets. While HQNN systems are capable of meaningful classification behavior under selected conditions, their performance often degrades substantially once realistic noise and optimization instability are introduced.

Several factors contribute to this gap. Variational quantum systems rely heavily on repeated noisy circuit execution during optimization, causing stochastic perturbations to accumulate throughout training. In addition, shallow circuit-depth constraints imposed by NISQ hardware limit the representational flexibility of practical HQNN architectures. Parity-based decoding strategies further compress quantum measurement distributions into simplified outputs, reducing inference sensitivity.

The experiments also demonstrate that increasing circuit complexity does not necessarily improve practical HQNN performance. In many cases, deeper circuits with stronger entanglement produced greater optimization instability and noisier convergence behavior despite possessing greater theoretical expressive capability.

Figure 3.4 shows Classical vs QSVM vs HQNN Benchmark Comparison.




Figure 3.4: While the QSVM achieves competitive performance, the baseline HQNN exhibits substantially lower accuracy under realistic NISQ conditions due to optimization instability, limited circuit depth, and parity-based information loss.

Importantly, the existence of this performance gap does not imply that hybrid quantum learning lacks value. Rather, it highlights a central challenge of NISQ-era quantum machine learning: practical performance is frequently constrained more by instability and noise sensitivity than by lack of expressivity. This observation becomes one of the primary motivations for the robustness-engineering framework developed in Chapter 4.

3.3.4 Medical and Cybersecurity Classification Observations

To extend the benchmarking analysis beyond synthetic datasets, the framework also evaluated hybrid quantum learning systems on domain-inspired classification tasks involving medical risk prediction and cybersecurity anomaly detection. These experiments were designed to determine whether the patterns observed in earlier benchmarks persisted within more realistic application scenarios.

The results demonstrated that HQNN systems could learn meaningful decision boundaries under controlled conditions in both domains. However, the same limitations observed in

earlier experiments remained visible. Optimization instability, noisy measurements, and sensitivity to perturbation environments continued to influence model behavior during training and inference.

An important observation is that the instability patterns identified in synthetic benchmarks also appeared in these more realistic environments. This suggests that the challenges observed throughout the benchmarking framework are not artifacts of simplified datasets, but instead reflect broader characteristics of variational quantum learning under realistic NISQ conditions.

These experiments therefore reinforce a central conclusion of Chapter 3: practical hybrid quantum learning depends as much on stability and robustness as it does on representational capability. This conclusion directly motivates the robustness-oriented methods introduced in the following chapter.

3.4 Noise Robustness Benchmarking

A central objective of this thesis is to evaluate how hybrid quantum neural networks behave under realistic noisy execution conditions rather than under idealized simulation assumptions alone. While many quantum machine learning studies emphasize performance in noiseless environments, practical NISQ systems operate under persistent decoherence, gate errors, measurement instability, and stochastic hardware perturbations. Consequently, understanding the robustness of HQNN architectures requires systematic evaluation across controlled perturbation environments.

The benchmarking framework developed in this work evaluates HQNN systems under several widely studied quantum noise channels, including depolarizing noise, bit-flip noise, phase-flip noise, and amplitude damping. These perturbation models reflect experimentally observed behavior in real quantum hardware and provide a controlled environment for studying how hybrid quantum learning systems respond to increasing levels of disturbance [14], [15], [16].

Rather than focusing exclusively on final classification accuracy, the framework evaluates robustness as a broader systems-level property involving convergence stability, degradation behavior, optimization reliability, and sensitivity to perturbation intensity. This perspective is important because noisy quantum systems may exhibit unstable optimization trajectories even when isolated accuracy measurements appear acceptable.

The experiments reveal several important trends. Different HQNN architectures respond differently to distinct noise channels, degradation behavior depends strongly on optimization stability, and deeper variational circuits frequently become unstable under moderate perturbation levels. Most importantly, the results demonstrate that noise influences not only inference quality, but also the underlying training dynamics governing variational learning itself.

The following subsections examine HQNN behavior under controlled perturbation environments, analyze degradation patterns as noise intensity increases, and establish noise as one of the primary bottlenecks limiting practical hybrid quantum learning under realistic NISQ conditions.

3.4.1 HQNN Robustness Under Noise

The first robustness benchmark evaluates how baseline HQNN architectures behave under controlled perturbation environments. Several common quantum noise channels are introduced into variational circuit execution in order to examine how different forms of hardware instability affect hybrid quantum learning behavior.

The evaluated perturbation channels include depolarizing noise, bit-flip noise, phase-flip noise, and amplitude damping. These channels were selected because they correspond closely to experimentally observed error processes in real NISQ hardware systems [14], [15], [16]. By applying these perturbations systematically, the framework can evaluate whether observed behavior reflects general robustness characteristics or sensitivity to specific noise models.

The experiments reveal that HQNN robustness depends strongly on both architecture

design and perturbation type. Some architectures exhibit relatively stable behavior under selected low-noise conditions, whereas others degrade rapidly even under moderate perturbation levels. Architectures relying heavily on entangling operations frequently become more unstable as stochastic errors propagate across correlated quantum states.

An important observation is that different forms of noise affect hybrid quantum learning systems in qualitatively different ways. Bit-flip and phase-flip noise often produce gradual degradation patterns, whereas depolarizing noise tends to generate more severe instability and rapid performance collapse. Amplitude damping introduces additional degradation by driving quantum states toward lower-energy configurations and reducing the diversity of accessible quantum representations.

The results also demonstrate that noise sensitivity is strongly architecture dependent. More modular and shallow structures frequently maintain greater stability under perturbation, whereas deeper highly entangled architectures exhibit more severe degradation under identical conditions. These findings suggest that meaningful robustness analysis requires systematic evaluation across multiple perturbation environments rather than reliance on a single noise model.

3.4.2 Accuracy Degradation Under Increasing Noise

To quantify robustness more systematically, the benchmarking framework evaluates how HQNN performance changes as perturbation intensity increases progressively. Rather than analyzing isolated noise conditions, these experiments examine degradation trajectories across multiple perturbation levels in order to identify stability thresholds and failure regions within variational hybrid learning systems.

The results demonstrate that baseline HQNN architectures experience substantial degradation as perturbation probabilities increase. Under low-noise conditions, several models maintain moderate classification capability and relatively stable optimization behavior. However, as noise intensity grows, optimization trajectories become increasingly unstable, mea-

sured outputs become more variable, and predictive accuracy begins to deteriorate rapidly.

Figure 3.5 shows HQNN accuracy degradation under increasing noise.




Figure 3.5: Although all models experience performance degradation as perturbation intensity increases, shallower modular architectures maintain greater stability, while the baseline HQNN approaches near-random classification performance under moderate to high noise conditions. The dashed horizontal line indicates random binary classification accuracy.

An important observation is that degradation is not always smooth or monotonic. In many experiments, systems exhibit oscillatory behavior in which temporary performance improvements are followed by abrupt instability and collapse. This suggests that perturbations influence not only final prediction quality, but also the optimization dynamics governing parameter evolution during training.

The degradation curves also reveal important differences between architectures. QuanNN systems frequently maintain more stable behavior at moderate perturbation levels due to their shallow and modular structure, whereas more globally entangled architectures often experience sharper degradation as circuit complexity increases.

Another significant finding is that degradation thresholds occur substantially earlier than

would be expected from idealized simulation assumptions. In several experiments, noisy classification accuracy drifted toward approximately 0.37 to 0.50 under moderate perturbation conditions, approaching near-random classification behavior. These observations establish important practical limitations of baseline HQNN systems and provide quantitative evidence that robustness becomes a defining factor in hybrid quantum learning under NISQ conditions.

3.4.3 Noise as a Training and Evaluation Bottleneck

The experiments presented throughout this section demonstrate that noise is not merely a secondary implementation issue in hybrid quantum machine learning, but rather one of the defining bottlenecks limiting practical NISQ-era computation. Noise affects not only final inference accuracy, but also optimization stability, convergence behavior, measurement consistency, and architectural scalability.

A particularly important observation is that many variational HQNN systems become unstable during training before their theoretical representational capacity can be fully utilized. As circuit complexity increases, stochastic perturbations distort optimization landscapes and amplify measurement variability, often causing optimization procedures to oscillate or collapse before meaningful convergence can occur.

These findings suggest that practical hybrid quantum learning is constrained less by the absence of expressive quantum representations and more by the difficulty of maintaining stable optimization under noisy execution conditions. In other words, many baseline HQNN systems fail not because quantum circuits cannot represent useful computational structures, but because realistic NISQ environments prevent reliable training and inference from emerging consistently.

This observation becomes one of the central conclusions of the benchmarking phase of the thesis. The results indicate that practical hybrid quantum learning should be approached not solely as a representational problem, but also as a robustness problem. Improving

performance therefore requires more than increasing circuit complexity; it requires methods capable of stabilizing optimization and reducing sensitivity to stochastic perturbations.

These findings establish the primary motivation for the robustness-engineering methods developed in Chapter 4. The next chapter introduces a sequence of techniques designed specifically to address the instability, degradation, and optimization bottlenecks identified throughout the baseline benchmarking experiments.

3.5 Cross-Framework and Cross-Noise Validation

A major challenge in evaluating Hybrid Quantum Machine Learning systems is determining whether observed performance reflects genuine architectural behavior or implementation-specific artifacts. Differences in software frameworks, transpilation procedures, simulator assumptions, and backend execution can influence experimental results, making it difficult to compare robustness consistently across studies.

Previous HQNN research has frequently evaluated robustness within a single software ecosystem or under one or two selected quantum noise models. Although these studies have provided valuable insight into architecture-specific behavior, they generally offer limited evidence regarding whether the observed robustness trends generalize across different implementation environments and representative NISQ perturbation processes.

This validation extends the scope of Kashif et al. [24], who examine the effects of several representative noise channels on HQNN training, by asking an additional reproducibility question: whether the same qualitative degradation patterns persist across Qiskit, Cirq, and PennyLane implementations. It also extends the parameterized-circuit perspective described by Benedetti et al. [13] by evaluating whether architecture-level robustness conclusions remain stable when implementation environments and perturbation channels change.

To address this limitation, the benchmarking framework developed in this thesis incorporates both cross-framework and cross-noise validation. Hybrid Quantum Neural Networks were evaluated using Qiskit, Cirq, and PennyLane while maintaining identical datasets, op-

timization procedures, evaluation metrics, and representative quantum noise channels. This methodology allows robustness trends to be compared consistently across multiple software ecosystems and perturbation environments.

The objective of this validation process is not merely to compare implementation outputs, but to determine whether the robustness characteristics observed throughout the benchmarking study represent genuine properties of the learning architectures rather than artifacts introduced by individual software frameworks. In addition, evaluating multiple quantum noise channels provides a broader characterization of HQNN robustness than evaluations performed under isolated perturbation conditions.

The results demonstrate that although absolute performance values vary slightly across software frameworks and quantum noise models, the overall robustness trends remain highly consistent. These observations provide additional confidence that the degradation patterns identified throughout Chapter 3 reflect meaningful computational behavior rather than implementation-specific effects and establish a reliable experimental foundation for the robustness-engineering methodology introduced in Chapter 4.

3.5.1 Cross-Framework Noise Benchmark

The first validation experiment evaluates HQNN behavior across multiple quantum software ecosystems in order to examine the reproducibility of robustness observations under different execution environments. The experiments were implemented using Qiskit, Cirq, and PennyLane, each of which provides distinct approaches to circuit construction, simulation, transpilation, and backend execution.

Although these frameworks implement many common quantum operations, differences in simulator configuration, transpilation procedures, noise-model implementation, and measurement handling can produce measurable variation in final performance outcomes. Evaluating HQNN architectures across multiple frameworks therefore provides an important test of whether observed robustness behavior is implementation-specific or architecture-driven.

The experiments demonstrate that while small quantitative differences exist between frameworks, the broader qualitative robustness patterns remain largely consistent. Architectures that exhibited instability under noisy conditions within one framework generally displayed similar degradation behavior across alternative implementations. Likewise, architectures that maintained relatively stable behavior under selected perturbation environments tended to preserve these trends across multiple software ecosystems.

An important observation is that framework-dependent variation becomes more pronounced as circuit complexity increases. Differences in compilation procedures, qubit mapping, and backend optimization can alter the effective execution characteristics of deeper variational circuits. However, shallow and modular architectures often exhibit stronger consistency because they are less sensitive to compilation complexity and accumulated gate noise.

These findings reinforce the importance of reproducibility within hybrid quantum machine learning research. Cross-framework validation provides evidence that the robustness patterns identified throughout the benchmarking analysis reflect meaningful architectural behavior rather than isolated software-specific effects.

3.5.2 Cross-Noise Robustness Heatmap

The second validation experiment evaluates HQNN architectures across multiple perturbation channels simultaneously using cross-noise robustness heatmaps. Rather than examining performance under only a single perturbation environment, this analysis evaluates behavior under depolarizing noise, bit-flip noise, phase-flip noise, amplitude damping, and phase damping using comparable benchmark conditions.

The resulting heatmaps provide a broader view of robustness behavior across perturbation environments and allow stability regions and degradation patterns to be identified more systematically. This approach reveals that robustness is highly nonuniform across noise channels and that different architectures respond differently depending on the nature of the

perturbation.

Figure 3.6 shows a cross-noise robustness heatmap.




Figure 3.6: The heatmap shows that robustness is not uniform across noise models: some architectures remain stable under bit-flip or phase-flip perturbations while degrading more sharply under depolarizing or amplitude damping noise.

Several important trends emerge from this analysis. Some architectures maintain relatively stable behavior under selected low-noise conditions but degrade rapidly under depolarizing or amplitude damping environments. Others display greater resilience to localized perturbations such as bit-flip or phase-flip noise while remaining sensitive to more globally disruptive stochastic channels.

The heatmap analysis demonstrates that evaluating HQNN systems under a single perturbation model is insufficient for understanding practical robustness. Meaningful characterization requires multichannel evaluation capable of capturing the diverse ways in which stochastic perturbations influence variational quantum learning systems.

3.5.3 Noise-Channel Sensitivity Across Architectures

The final stage of the validation framework examines how different HQNN architectures respond to specific categories of noise. Whereas the previous subsection visualized broad robustness patterns across perturbation environments, this analysis focuses directly on architecture-specific sensitivity and degradation mechanisms.

The experiments reveal that different architectures possess distinct robustness profiles depending on their structural organization and optimization behavior. QuanNN architectures often exhibit comparatively greater resilience under moderate perturbation conditions due to their shallow and modular design. Because localized quantum filters operate relatively independently, perturbation accumulation and error propagation remain partially constrained within smaller computational regions.

QCNN architectures frequently exhibit stronger sensitivity to stochastic perturbations as entanglement complexity increases. Errors affecting one portion of the circuit can propagate throughout correlated quantum states, producing more severe degradation during optimization and inference. Quantum Transfer Learning systems occupy an intermediate position, often maintaining relatively stable behavior because the quantum subsystem operates primarily as a final inference layer built on top of pretrained classical representations.

These observations reinforce one of the central conclusions emerging from the benchmarking framework: robustness is strongly architecture dependent in hybrid quantum learning systems. Circuit structure, entanglement topology, modularity, and optimization dynamics all interact to determine how perturbations propagate through the computational pipeline.

Collectively, the cross-framework and cross-noise validation experiments demonstrate that practical HQNN robustness cannot be characterized adequately through isolated performance measurements alone. Meaningful evaluation requires systematic comparison across architectures, perturbation environments, software ecosystems, and optimization conditions in order to identify the structural properties that support reliable hybrid quantum learning under realistic NISQ constraints.

3.6 Architecture-Level Benchmarking

The robustness and practical performance of hybrid quantum neural networks depend strongly on architectural design choices. Although many HQNN systems share common variational and hybrid optimization principles, differences in circuit organization, entanglement structure, modularity, and quantum–classical integration can produce substantially different behavior under realistic NISQ conditions. Consequently, architecture-level benchmarking represents an important component of the evaluation framework developed throughout this thesis.

The experiments presented in this section compare several representative HQNN architectures, including Quantvolutional Neural Networks (QuanNNs), Quantum Convolutional Neural Networks (QCNNs), and Quantum Transfer Learning (QTL) systems. These architectures represent distinct strategies for integrating quantum computation into learning pipelines and therefore provide an opportunity to examine how architectural decisions influence performance under both idealized and noisy execution conditions.

Rather than evaluating architectures independently, the benchmarking framework applies common datasets, preprocessing procedures, optimization settings, and evaluation metrics across all models. This controlled structure allows differences in behavior to be attributed more directly to architectural design choices rather than to variations in experimental configuration.

The results reveal a recurring pattern throughout the benchmark analysis. Architectures that exhibit strong expressive capability under idealized simulation conditions often become more sensitive to noise as circuit complexity increases, whereas simpler and more modular architectures frequently maintain more stable behavior under realistic perturbation environments.

Understanding this relationship between architecture and robustness is essential for interpreting the practical capabilities of hybrid quantum learning systems and provides important context for the robustness-engineering methods developed in Chapter 4.

3.6.1 QuanNN versus QCNN Design Differences

Quanvolutional Neural Networks and Quantum Convolutional Neural Networks represent two distinct strategies for incorporating quantum computation into hybrid learning systems [10], [12], [13]. Although both architectures draw inspiration from classical convolutional neural networks, they differ substantially in how quantum operations are organized and how information is processed throughout the computational pipeline.

QuanNN architectures emphasize localized quantum filtering. Small regions of the input are processed independently through shallow quantum circuits that act as quantum feature extractors. The resulting local features are then aggregated and passed to downstream learning components. This design naturally promotes modularity and limits the extent to which perturbations can propagate across the overall system.

QCNN architectures instead emphasize hierarchical feature extraction through entangling quantum operations. Information is processed through correlated quantum states that evolve across multiple layers, allowing the architecture to capture broader global relationships within the data. This structure often provides greater theoretical expressive capability but also introduces increased sensitivity to accumulated noise and optimization instability.

From a design perspective, the contrast between the two architectures reflects a broader tradeoff. QuanNN systems prioritize locality and modularity, whereas QCNN systems prioritize global quantum correlations. As a result, the two architectures frequently exhibit different robustness profiles when evaluated under realistic NISQ conditions.

These structural differences make QuanNN and QCNN architectures useful benchmarks for studying how circuit organization influences both learning performance and robustness under noise.

3.6.2 Ideal Performance versus Noisy Performance

The architecture-level experiments reveal a substantial distinction between HQNN behavior under idealized simulation conditions and under realistic noisy execution environments.

Several architectures that perform strongly in noiseless simulations experience significant degradation once stochastic perturbations are introduced into the computational pipeline.

Under idealized conditions, QCNN architectures often achieve stronger representational capability due to their hierarchical entanglement structure and global feature extraction behavior. In several experiments, deeper QCNN configurations demonstrated improved classification performance relative to shallower modular architectures when hardware noise was absent or minimal.

However, once realistic perturbation channels were introduced, the performance characteristics changed substantially. Increased entanglement complexity and deeper variational structures caused QCNN systems to accumulate errors more rapidly, producing unstable optimization trajectories and noticeable degradation in predictive performance. By contrast, QuanNN architectures frequently maintained more stable behavior because their modular structure reduced error propagation and limited the impact of stochastic perturbations on the overall computational pipeline.

These observations demonstrate that idealized evaluation alone provides an incomplete picture of practical HQNN capability. Architectures that appear highly expressive under simulation assumptions may become significantly less effective once realistic NISQ perturbations are introduced.

Consequently, robustness evaluation under noisy conditions becomes essential for understanding the practical usefulness of hybrid quantum learning systems.

3.6.3 Robustness Tradeoffs Across HQNN Variants

The comparative benchmarking experiments demonstrate that different HQNN architectures occupy distinct positions within the tradeoff space between representational expressivity, optimization complexity, and robustness under noise. No single architecture consistently dominates across all evaluation environments. Instead, practical performance depends heavily on how architectural structure interacts with perturbation accumulation, optimization stability,

circuit depth, and hardware variability.

QuanNN architectures generally exhibit stronger robustness under moderate perturbation environments due to their shallow and modular organization. Because quantum operations are distributed across relatively independent local circuits, optimization instability and perturbation propagation remain partially localized. This frequently allows QuanNN systems to maintain more stable convergence behavior under noisy conditions.

QCNN systems, by contrast, often provide greater theoretical expressive capability through hierarchical entanglement and global feature extraction. However, the same structural characteristics that increase expressive power also increase sensitivity to decoherence and stochastic perturbations. Under realistic NISQ conditions, deeper QCNN circuits frequently experience sharper degradation as noise accumulates throughout the computational state.

QTL architectures occupy a different position within this tradeoff space. Because the quantum subsystem typically operates as a final inference layer built on top of pretrained classical feature extractors, QTL systems often maintain relatively stable optimization behavior. However, the influence of the quantum component is generally more limited than in fully integrated variational architectures.

Figure 3.7 shows robustness tradeoffs across HQNN architectures.

Robustness Tradeoffs Across HQNN Variants

Architecture	Strength	Limitation	Tradeoff Position
QuanNN	shallow modular circuits	moderate expressivity	higher robustness
QCNN	hierarchical entanglement	higher noise sensitivity	higher expressivity
QTL	stable classical features	limited quantum role	stable optimization
Baseline HQNN	simple shallow design	parity information bottleneck	low robustness baseline

Practical HQNN design requires balancing expressivity, circuit depth, optimization stability, and robustness under noisy execution.

Figure 3.7: QuanNN systems generally provide stronger robustness under moderate noise due to shallow modular structure, QCNN systems provide greater expressive capacity but often exhibit stronger sensitivity to accumulated perturbations, and QTL systems maintain stable optimization by limiting the role of the quantum subsystem.

Taken together, these results indicate that practical HQNN performance cannot be understood solely through theoretical notions of quantum expressivity. Successful hybrid quantum learning systems must balance representational flexibility, circuit complexity, optimization stability, and robustness under noisy execution conditions. This observation becomes one of the key motivations for the robustness-engineering framework introduced in the following chapter.

3.7 Summary of Benchmarking Findings

The benchmarking experiments presented throughout this chapter establish several important observations regarding the practical behavior of hybrid quantum neural networks under realistic NISQ conditions. Across multiple datasets, optimization settings, architectures, software frameworks, and perturbation environments, the results consistently demonstrate

that robustness and optimization stability are central determinants of practical HQNN performance.

The experiments revealed that baseline variational HQNN systems frequently exhibit substantial instability once realistic noise channels are introduced. While several architectures demonstrated meaningful learning capability under idealized simulation conditions, noisy execution environments often produced oscillatory optimization behavior, unstable convergence dynamics, and significant degradation in predictive performance.

The benchmarking framework also demonstrated that robustness behavior is strongly architecture dependent. Modular architectures such as QuanNN systems frequently maintained greater stability under noisy execution conditions, whereas more globally entangled systems often experienced sharper degradation as circuit complexity increased. Quantum kernel methods provided comparatively stable optimization behavior, but typically lacked the adaptive representational flexibility available to trainable variational HQNN architectures.

Another important finding was that noise affects not only final predictive accuracy, but also the optimization process itself. Variational quantum systems repeatedly exhibited sensitivity to noisy measurements, stochastic gradient estimation, and unstable parameter updates. These effects frequently prevented deeper architectures from realizing their theoretical expressive advantages under practical NISQ conditions.

Collectively, the experiments demonstrate that practical hybrid quantum learning cannot be evaluated meaningfully through idealized simulation accuracy alone. Instead, robustness must be interpreted as a broader systems-level property involving optimization stability, degradation behavior, architectural resilience, and reproducibility across multiple perturbation environments.

3.7.1 What the Baseline Framework Revealed

The benchmarking framework developed throughout this chapter produced a comprehensive view of hybrid quantum neural network behavior under realistic NISQ conditions. By evalu-

ating classical baselines, quantum kernel methods, and multiple HQNN architectures across controlled perturbation environments, the framework revealed several recurring patterns that remained consistent throughout the experimental analysis.

The first major observation was that baseline HQNN systems frequently exhibited substantial instability under noisy execution conditions. Although many architectures demonstrated meaningful learning capability under idealized simulations, performance often deteriorated once realistic perturbation channels were introduced. Optimization trajectories became increasingly variable, convergence behavior became less predictable, and classification performance frequently degraded as perturbation intensity increased.

A second important finding was that robustness behavior proved highly architecture dependent. Modular architectures generally maintained stronger stability under noisy conditions than more globally entangled systems. While deeper and more connected circuits often provided greater theoretical representational flexibility, they also accumulated perturbations more rapidly and exhibited greater optimization instability. This observation revealed a recurring tradeoff between expressivity and robustness that appeared throughout multiple benchmark environments.

The comparative evaluation also demonstrated that classical machine learning systems remained highly competitive across many structured learning tasks. Quantum kernel methods frequently exhibited greater optimization stability than fully variational HQNN architectures, while classical baselines often maintained superior performance under identical benchmark conditions. These results suggest that practical hybrid quantum learning should not be interpreted through assumptions of universal superiority over classical approaches. Instead, the usefulness of quantum representations appears conditional and strongly influenced by architecture design, optimization stability, and perturbation resilience.

Another significant observation involved the role of optimization itself. The experiments repeatedly showed that noisy variational optimization represents one of the primary challenges facing practical HQNN systems. SPSA training frequently produced oscillatory con-

vergence behavior, noisy parameter updates distorted objective-function evaluations, and repeated circuit execution amplified stochastic variability throughout training. In many cases, instability emerged before the representational capacity of the quantum circuit could be fully utilized.

The benchmarking analysis further revealed that several commonly used HQNN design choices contributed directly to instability. In particular, parity-based measurement decoding introduced severe information bottlenecks by collapsing rich quantum measurement distributions into simplified binary outputs. This process amplified sensitivity to measurement fluctuations and reduced the amount of useful information available for downstream inference. The results suggested that improving how quantum information is interpreted may be just as important as improving the quantum circuit itself.

Collectively, these findings lead to a central conclusion. The dominant limitation of practical hybrid quantum learning under realistic NISQ conditions is not a lack of theoretical expressivity. Instead, the primary challenge is maintaining reliable optimization and inference behavior in the presence of stochastic perturbations, measurement variability, and hardware noise. Baseline HQNN systems often fail before their theoretical representational capabilities can be fully exploited.

This conclusion provides the foundation for the remainder of the thesis. Having identified the principal failure modes limiting baseline HQNN architectures, the next chapter investigates whether these limitations can be reduced systematically through robustness-oriented design. The transition from Chapter 3 to Chapter 4 therefore marks a shift from diagnosis to intervention, moving from the analysis of instability toward the development of a structured robustness-engineering framework for practical hybrid quantum learning.

3.7.2 Why Robustness Engineering Became Necessary

The instability patterns observed throughout the benchmarking experiments demonstrated that robustness cannot be treated as a secondary implementation issue in practical hybrid

quantum machine learning systems. Instead, robustness emerged as one of the defining operational constraints governing whether variational HQNN architectures could function reliably under realistic NISQ conditions.

Several experiments showed that increasing circuit complexity or entanglement depth often worsened practical performance rather than improving it. Although deeper circuits possess greater theoretical expressive capability, they also accumulate perturbations more rapidly and destabilize optimization behavior. This revealed a fundamental disconnect between theoretical expressivity and practical trainability in variational quantum systems.

The benchmarking results also demonstrated that standard optimization procedures alone were insufficient for maintaining reliable learning behavior under noise. SPSA optimization frequently produced unstable convergence trajectories, noisy parameter updates distorted objective-function evaluations, and repeated circuit execution amplified stochastic perturbations throughout training.

These observations motivated a broader shift in how hybrid quantum learning systems were interpreted. Rather than viewing HQNN performance primarily as a representational problem, the benchmark results suggested that practical NISQ-era quantum machine learning should be approached as a robustness problem shaped by optimization stability, architecture structure, measurement reliability, and resistance to stochastic perturbations.

Consequently, the need for robustness engineering emerged naturally from the experimental evidence. The benchmark results demonstrated that practical HQNN systems require improved decoding strategies, more stable optimization procedures, architecture-aware design choices, and noise-aware evaluation methodologies in order to maintain reliable behavior under realistic NISQ constraints.

3.7.3 Transition to Improving Noisy Training

The findings presented throughout this benchmarking chapter establish the primary motivation for the robustness-engineering framework introduced in Chapter 4. The experiments

revealed that baseline variational HQNN systems suffer from several interconnected limitations, including noisy optimization instability, degradation under stochastic perturbations, information bottlenecks in parity-based decoding, and sensitivity to circuit-depth accumulation effects.

Importantly, the results suggest that these limitations are not isolated failures of particular implementations, but rather structural characteristics of variational hybrid quantum learning under NISQ conditions. This observation shifts the focus of the thesis from identifying failure modes to developing mechanisms capable of overcoming them.

Chapter 4 therefore begins a staged progression aimed at improving noisy training behavior through a sequence of coordinated robustness-oriented mechanisms. The first stage introduces learned-readout decoding strategies designed to replace parity-based measurement collapse with richer interpretation of quantum measurement distributions. Subsequent sections develop multi-observable feature extraction, architecture-level optimization, stability-aware training procedures, and noise-adaptive learning strategies.

This transition represents one of the central conceptual developments of the thesis. Rather than treating noise solely as an unavoidable limitation, the following chapter investigates whether robustness can be engineered systematically through coordinated interaction between architecture design, optimization procedures, decoder adaptation, and noise-aware learning mechanisms.

The remainder of the thesis demonstrates that substantial improvements in noisy performance can be achieved when hybrid quantum learning systems are designed explicitly around realistic NISQ constraints rather than idealized noiseless assumptions.

Chapter 4

Improving Noisy Training

4.1 Motivation for Robustness Engineering

The benchmarking experiments presented in Chapter 3 demonstrated that baseline hybrid quantum neural network systems exhibit substantial instability under realistic NISQ-era execution conditions. Although several HQNN architectures showed meaningful learning capability under idealized simulation environments, noisy perturbations frequently produced unstable optimization behavior, rapid degradation in predictive accuracy, and inconsistent convergence dynamics. These observations revealed that practical hybrid quantum learning is constrained not only by representational capability, but also by the ability of variational systems to maintain stable behavior under stochastic hardware conditions.

A particularly important conclusion emerging from the benchmarking framework was that many baseline HQNN systems failed before their theoretical expressive potential could be fully utilized. Increased circuit depth and stronger entanglement often amplified optimization instability faster than they improved representational flexibility. As a result, architectures that appeared promising under noiseless assumptions frequently became unreliable once realistic perturbation environments were introduced.

The experiments further showed that instability arose from multiple interacting sources simultaneously. Noise distorted measurement outcomes, variational optimization produced oscillatory parameter updates, parity-based decoding compressed rich measurement distributions into unstable binary outputs, and deeper entangling circuits accumulated errors rapidly during execution. Consequently, no single isolated modification appeared sufficient to resolve the broader robustness limitations observed throughout the benchmarking process.

These findings motivated a transition from isolated model evaluation toward a robustness-oriented perspective. Rather than treating noise solely as an unavoidable implementation issue, this thesis approaches robustness as a systems-level design problem involving the interaction of architecture structure, optimization methodology, decoder design, feature representation, and evaluation strategy.

The remainder of this chapter develops a sequence of mechanisms designed to improve noisy training behavior under realistic NISQ conditions. These mechanisms include learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware training procedures, and noise-adaptive learning strategies. Together, they form a staged progression aimed at improving both optimization stability and predictive reliability under noisy execution environments.

More importantly, the chapter reframes the central challenge of practical hybrid quantum learning. Rather than asking how expressive a quantum model can become under ideal conditions, the focus shifts toward understanding how reliable that model remains when operating under realistic hardware constraints. This change in perspective forms the foundation of the robustness-engineering framework developed throughout the remainder of the thesis.

4.1.1 From Benchmarking to Robustness Optimization

The benchmarking framework developed in Chapter 3 revealed that evaluating HQNN systems under realistic NISQ conditions requires more than measuring predictive accuracy under isolated perturbation environments. The results demonstrated that robustness itself must become a central design objective within practical hybrid quantum learning systems.

Several observations motivated this transition. Noisy degradation patterns proved highly architecture dependent. Some modular architectures maintained relatively stable behavior under moderate perturbation conditions, whereas more globally entangled systems frequently experienced rapid instability as noise accumulated. Optimization instability also emerged

repeatedly across multiple variational systems, particularly when stochastic perturbations distorted gradient estimation and objective-function evaluation.

These observations suggested that practical HQNN performance depends not only on the expressive power of a quantum circuit, but also on how effectively the overall hybrid system maintains stability during training and inference. Consequently, robustness could no longer be interpreted merely as an evaluation metric. Instead, it became necessary to optimize directly for reliable behavior under noisy execution conditions.

This shift represents an important conceptual transition within the thesis. The benchmarking experiments focused primarily on identifying failure modes in baseline HQNN systems. The objective of Chapter 4 is different: to investigate whether those failure modes can be reduced systematically through coordinated design choices and robustness-aware learning strategies.

The sections that follow therefore move from diagnosis to intervention. Rather than documenting instability, they examine how architecture design, decoder adaptation, feature representation, and optimization procedures can be engineered explicitly to improve robustness under realistic NISQ constraints.

4.1.2 Limitations of Baseline Parity-Based HQNNs

One of the most significant limitations identified during the benchmarking experiments was the instability associated with parity-based decoding strategies. In many baseline architectures, quantum measurement outcomes were reduced to binary class predictions through parity-collapse procedures in which measured bitstrings were compressed into simplified parity values.

Although computationally straightforward, this decoding strategy introduced a severe information bottleneck within the learning pipeline. Rich probabilistic measurement distributions generated by the quantum circuit were reduced to highly compressed binary outputs, causing substantial loss of representational information during inference. As a result,

small stochastic perturbations in measured bitstrings could produce disproportionately large changes in final predictions.

The experiments revealed that parity-based systems frequently exhibited unstable behavior even under moderate perturbation environments. Noisy measurement fluctuations directly affected parity outputs, increasing classification instability and reducing optimization reliability during training. In many cases, instability originated not only from hardware noise itself, but also from the way measurement outcomes were interpreted.

Another important limitation involved optimization sensitivity. Because parity-based decoding produces highly compressed outputs, the optimization landscape often becomes noisy and discontinuous under stochastic perturbations. This behavior amplified SPSA instability and contributed to the oscillatory convergence patterns observed throughout the benchmarking analysis.

These observations suggested that improving robustness would require more than reducing hardware noise alone. It would also require redesigning the interface between quantum computation and classical inference so that richer measurement information could be preserved and interpreted more effectively during training and prediction.

4.1.3 Hybrid Quantum-Classical Co-Design Philosophy

The limitations observed throughout the benchmarking analysis motivated a broader shift in how hybrid quantum learning systems are interpreted conceptually. Rather than viewing the quantum subsystem as an isolated computational engine operating independently from classical computation, this thesis adopts a hybrid quantum-classical co-design philosophy in which robustness emerges through coordinated interaction between quantum and classical components.

Within this perspective, the quantum subsystem is responsible primarily for generating probabilistic feature representations and nonlinear computational transformations through variational quantum circuits. Classical systems provide optimization stability, adaptive de-

coding, parameter control, preprocessing, and feedback mechanisms. Practical hybrid quantum learning therefore depends not on maximizing quantum complexity alone, but on balancing quantum expressivity against classical stabilization and control.

This philosophy differs from many earlier approaches that emphasized quantum circuit expressivity as the primary determinant of model capability. The benchmarking experiments demonstrated that increasing circuit depth and entanglement often worsened practical performance under realistic NISQ conditions due to optimization instability and noise accumulation. Consequently, the framework developed in this chapter prioritizes coordinated interaction between variational circuit structure, classical optimization, adaptive decoding, and noise-aware learning procedures.

An important implication of this perspective is that robustness becomes an emergent property of the overall computational system rather than a feature of the quantum circuit alone. Stability depends on how effectively quantum feature generation, classical optimization, measurement interpretation, and adaptive learning mechanisms operate together within the hybrid pipeline.

This co-design philosophy provides the conceptual foundation for the robustness-oriented methods developed throughout the remainder of the chapter. Rather than attempting to eliminate noise entirely, the objective is to construct hybrid systems capable of maintaining reliable learning behavior despite realistic stochastic hardware perturbations.

4.2 Learned Readout Hybrid Quantum Neural Networks

One of the most significant limitations identified during the baseline benchmarking experiments was the instability introduced by parity-based decoding. In the baseline HQNN systems, quantum measurement outputs were compressed into binary parity values before classification decisions were produced. Although this approach is computationally simple, it discards much of the probabilistic information generated by the quantum circuit and creates a severe information bottleneck within the learning pipeline.

The experiments presented in Chapter 3 demonstrated that parity-based readout systems frequently became unstable under stochastic perturbation environments. Small fluctuations in measured bitstrings often produced disproportionately large changes in final predictions, reducing optimization stability and increasing sensitivity to measurement noise. These observations suggested that a significant portion of the observed instability originated not only from the quantum circuit itself, but also from the way quantum information was interpreted during inference.

The learned-readout framework addresses this limitation by replacing fixed parity collapse with adaptive classical decoding. Rather than mapping all quantum outputs directly into simplified binary parity states, the learned-readout architecture treats the quantum circuit as a probabilistic feature generator. Its measurement distributions are then processed through trainable classical decoding layers that learn how to interpret noisy quantum outputs more effectively.

This modification represents an important conceptual shift in the interpretation of hybrid quantum learning systems. In the baseline framework, the quantum subsystem effectively performed both feature generation and direct classification through parity collapse. In the learned-readout architecture, these responsibilities are separated. The quantum circuit generates probabilistic measurement features, while the classical decoder learns how to transform those features into stable class predictions.

Figure 4.1 shows the differences between parity collapse and learned readout.




Figure 4.1: The baseline HQNN converts quantum measurement outcomes directly into binary parity values, creating information loss and sensitivity to noisy measurements. The proposed framework instead interprets complete quantum measurement distributions using a trainable classical decoder, preserving probabilistic information and learning more stable decision boundaries under noisy NISQ execution.

The experiments developed throughout this section demonstrate that learned-readout architectures substantially improve noisy inference behavior relative to baseline parity systems. More importantly, the results show that decoder adaptation itself becomes a major contributor to robustness under realistic NISQ conditions. This stage therefore marks the first major step in the robustness-engineering progression developed throughout Chapter 4.

4.2.1 Replacing Parity Collapse with Learned Decoding

The baseline HQNN architectures evaluated in Chapter 3 relied on parity-based measurement decoding, in which measured quantum bitstrings were compressed into simplified binary outputs according to their parity value. Although this approach allowed straightforward classification, it discarded most of the information contained within the underlying quantum measurement distribution.

Under realistic noisy conditions, this information loss became highly problematic. Small

stochastic perturbations affecting measured bitstrings frequently altered parity outcomes directly, causing unstable prediction behavior and increasing sensitivity to measurement variability. Instead of filtering noise, parity decoding often amplified noise-induced instability.

To address this limitation, the learned-readout framework retained the measurement distribution produced by the quantum circuit and processed it through trainable classical decoding layers. Instead of reducing the quantum output to a single parity value, the decoder received probabilistic information representing the distribution of observed measurement outcomes across repeated circuit executions.

This modification allowed the hybrid learning system to interpret noisy quantum outputs adaptively rather than relying on fixed deterministic parity rules. The classical decoder could learn which measurement patterns remained informative under stochastic perturbation conditions and adjust inference behavior accordingly.

An important consequence of this change was the reduction of the information bottleneck introduced by parity-based inference. Richer measurement distributions preserved more information about the underlying quantum state evolution, improving both the stability and flexibility of the hybrid learning pipeline.

4.2.2 Multi-Shot Quantum Measurement Distributions

A key component of the learned-readout framework is the use of multi-shot quantum measurement distributions rather than single-shot parity outcomes. Repeated circuit execution produces probabilistic distributions over measurement outcomes, reflecting both the intended quantum state and the stochastic effects of noisy execution environments.

In the baseline architecture, most of this information was discarded during inference. The learned-readout framework instead preserves the multi-shot distribution generated by repeated circuit evaluations and uses it as an input representation for downstream decoding. Rather than relying only on binary parity outputs, the decoder receives structured

probabilistic patterns describing how frequently different measurement states occur.

This approach provides two important advantages. First, it increases the amount of information available to the hybrid learning system by preserving a richer description of quantum state evolution. Second, it improves resilience to isolated noisy measurements. Single-shot outcomes can fluctuate substantially due to stochastic perturbations, whereas aggregated multi-shot distributions provide a more stable statistical representation of the underlying quantum behavior.

The use of multi-shot distributions also aligns with the broader hybrid quantum-classical co-design philosophy developed throughout this thesis. The quantum subsystem acts as a probabilistic feature generator, while the classical subsystem learns how to interpret these probabilistic representations during training.

For this reason, multi-shot measurement distributions provide the bridge between quantum feature generation and adaptive classical inference. They allow the model to retain useful quantum information without forcing the circuit to produce a direct binary decision through parity collapse.

4.2.3 Classical Decoder Integration

The learned-readout framework integrates a trainable classical decoder directly into the HQNN inference pipeline. This decoder operates as a lightweight classical component positioned after the quantum measurement stage and before the final classification output.

The decoder receives the multi-shot quantum measurement distribution as its input representation and learns adaptive mappings between probabilistic quantum outputs and target class labels. Because the decoder parameters are optimized together with the variational quantum circuit, the hybrid system evolves cooperatively during training.

This design introduces additional representational flexibility without increasing quantum circuit depth. Instead of relying entirely on deeper entangling quantum structures for expressivity, part of the learning burden is transferred to the classical subsystem. This reduces

pressure on noisy quantum hardware while preserving useful nonlinear inference capability.

The decoder also improves robustness by learning to compensate partially for recurring noisy measurement patterns. During training, it adapts to the statistical structure of the quantum measurement distributions and learns more stable decision boundaries than fixed parity rules can provide.

As a result, the decoder becomes more than a simple post-processing layer. It functions as an adaptive robustness-aware inference mechanism that helps stabilize the interface between quantum feature generation and classical prediction.

4.2.4 Threshold Improvement from Approximately 0.50 to Approximately 0.82

The introduction of learned-readout decoding produced one of the first major robustness improvements observed throughout the framework evolution. Across multiple benchmark environments, replacing parity collapse with adaptive learned decoding substantially improved noisy inference behavior relative to the baseline HQNN architectures.

In the baseline parity-based systems, noisy classification accuracy frequently drifted within approximately the 0.37 to 0.50 range under realistic perturbation environments. In several unstable optimization runs, performance approached near-random classification behavior. After integrating learned-readout decoding and multi-shot measurement interpretation, noisy accuracy improved substantially, reaching approximately 0.82 under several benchmark conditions.

This improvement represents one of the most important findings in the early stages of the robustness-engineering framework. The results demonstrate that a large portion of the instability observed in baseline HQNN systems originated not solely from the quantum circuit itself, but from how quantum information was decoded and interpreted during inference.

The experiments therefore show that decoder design plays a critical role in practical hybrid quantum learning robustness. Instead of treating the quantum subsystem as a standalone classifier, the learned-readout framework interprets the quantum circuit as a proba-

bilistic feature generator whose outputs require adaptive classical interpretation under noisy conditions.

Another important implication is that robustness improvements can emerge without dramatically increasing quantum circuit complexity. The threshold improvement was achieved primarily through improved hybrid inference design rather than deeper quantum architectures. This observation reinforces one of the central conclusions of the thesis: practical NISQ-era hybrid learning depends heavily on coordinated quantum-classical co-design rather than on quantum complexity alone.

4.3 Multi-Observable Quantum Feature Extraction

The learned-readout framework developed in the previous section demonstrated that preserving richer measurement information can significantly improve robustness under realistic NISQ conditions. By replacing parity collapse with adaptive decoding, the framework reduced one of the major information bottlenecks identified during the baseline benchmarking experiments. However, despite these improvements, the underlying quantum representation remained relatively constrained. The quantum subsystem continued to rely on a limited collection of measurement outputs, restricting the amount of information available to the downstream learning process.

This observation motivated the next stage of the robustness-engineering framework. If decoder design could improve robustness by preserving more measurement information, it became reasonable to ask whether the feature representation itself could also be expanded. Rather than relying on a single observable or a compressed parity output, the framework was modified to extract multiple observables simultaneously from the variational quantum circuit.

The motivation for this approach arises from the nature of quantum state evolution. Different observables capture different aspects of the computational behavior of a quantum system. Some measurements reflect the behavior of individual qubits, while others reveal

correlations generated through entanglement and interaction patterns within the circuit. Restricting the learning process to a single observable inevitably discards useful information that may contribute to classification stability and robustness.

The multi-observable framework therefore expands the role of quantum measurement within the hybrid learning pipeline. Instead of producing a single compressed output, the quantum subsystem generates a collection of statistical features describing multiple aspects of the quantum state. These features are then provided to the classical learning components for interpretation and optimization.

This modification also represents an important conceptual transition within the thesis. The quantum circuit is no longer viewed primarily as a classifier that directly produces predictions. Instead, it increasingly functions as a probabilistic feature-generation mechanism whose outputs are interpreted through adaptive classical processing. This interpretation aligns naturally with the hybrid quantum-classical co-design philosophy introduced earlier in the chapter.

The results presented throughout this section demonstrate that richer quantum feature representations improve both noisy inference behavior and optimization stability. More importantly, they establish that robustness can be improved not only through better decoding, but also through more informative quantum feature extraction.

4.3.1 Z Expectations and Correlation Features

The first stage of the multi-observable framework focused on the extraction of expectation values and correlation statistics from variational quantum circuits. Rather than relying exclusively on binary measurement outcomes, the framework incorporated measurements that preserve additional information about the statistical behavior of the quantum state during execution.

Particular attention was given to Pauli-Z expectation values measured across multiple qubits. These measurements provide information about the average behavior of the quantum

state across repeated circuit executions. Unlike parity-based decoding, which compresses measurement outcomes into a single binary quantity, expectation values preserve a more detailed description of the probabilistic structure generated by the circuit.

An important advantage of expectation values is their ability to reduce sensitivity to isolated measurement fluctuations. Individual bitstrings can vary substantially due to stochastic perturbations, especially under noisy execution conditions. Expectation values instead aggregate information across repeated measurements, producing a more stable representation of the underlying quantum behavior.

In addition to single-qubit measurements, the framework incorporated correlation features describing relationships between pairs of qubits. These correlations capture information regarding how entanglement and variational interactions influence different regions of the quantum system. Because quantum circuits generate complex dependency structures through entangling operations, correlation statistics often contain information that is not visible through independent measurements alone.

The inclusion of correlation features expanded the representational capacity of the hybrid learning system. Rather than describing qubits independently, the framework could capture interaction patterns generated during circuit evolution. These patterns frequently remained informative even when individual measurement outcomes were affected by noise.

The experiments suggested that expectation values and correlation features improved the consistency of downstream classification behavior across several benchmark environments. By preserving richer information regarding quantum state evolution, the framework reduced sensitivity to the measurement variability that had previously destabilized parity-based systems.

This stage therefore established an important principle that remains central throughout the remainder of the thesis: preserving more information from quantum state evolution generally produces more stable and robust hybrid learning behavior under realistic NISQ conditions.

4.3.2 Multi-Observable Quantum Statistics

Building upon expectation-value extraction, the framework next incorporated multiple observable statistics simultaneously into the feature-generation process. Rather than relying on a single measurement operator, the architecture aggregated information from several observables in order to construct a richer representation of the quantum state.

This approach significantly increased the informational capacity of the quantum feature space. Different observables capture different aspects of quantum behavior, including local state evolution, entanglement structure, and interaction dynamics between qubits. By combining these measurements, the framework obtained a more complete description of the computational patterns generated by the variational circuit.

The resulting feature vectors consisted of multiple expectation values, correlation measurements, and statistical summaries derived from repeated circuit execution. Together, these measurements formed a higher-dimensional representation that preserved substantially more information than parity-based decoding or single-observable inference.

A key advantage of this design is the distribution of information across multiple measurement channels. When a representation depends heavily on a single observable, perturbations affecting that observable can have a disproportionately large impact on model behavior. Multi-observable representations reduce this vulnerability by distributing information more broadly throughout the feature space.

As shown in Figure 4.2, the transition from single-observable inference to multi-observable feature extraction substantially expands the amount of information available to the hybrid learning pipeline.

Single-Observable Inference	Multi-Observable Feature Extraction
Measurement Strategy One observable, such as a single Z expectation value	Measurement Strategy Multiple observables, including Z_0 , Z_1 , ZZ , XX , and correlation terms
Feature Representation Low-dimensional feature vector with limited quantum information	Feature Representation Higher-dimensional feature vector using expectations, correlations, and statistical summaries
Noise Behavior A single noisy measurement channel can strongly affect prediction	Noise Behavior Information is distributed across several measurement channels
Robustness Effect More sensitive to isolated stochastic perturbations	Robustness Effect Improves noisy inference stability by preserving more quantum-state information

Key Insight: Multi-observable extraction preserves more information from the quantum circuit than single-observable inference and produces a richer feature representation for noisy hybrid learning.

Figure 4.2: Instead of relying on a single measurement channel, the proposed representation aggregates local expectation values, correlation observables, and statistical summaries from repeated circuit execution, producing a richer feature vector that distributes information across multiple measurement channels and improves robustness under noisy conditions.

The results indicate that richer measurement representations improve both robustness and optimization stability. Because the learning process no longer depends on a narrow subset of measurement outcomes, the framework becomes less sensitive to isolated fluctuations and more resilient under stochastic perturbation conditions.

This observation further supports the broader argument developed throughout Chapter 4: robustness improvements often emerge from preserving and utilizing quantum information more effectively rather than from increasing quantum complexity alone.

4.3.3 Quantum Circuits as Feature Generators

The introduction of multi-observable extraction produced an important reinterpretation of the role of quantum circuits within hybrid learning systems. Earlier HQNN architectures frequently treated the quantum subsystem as a direct classification mechanism responsible for producing final predictions. The results obtained throughout this chapter suggest a different perspective.

Within the multi-observable framework, the quantum subsystem increasingly behaves as a probabilistic feature generator rather than as a standalone classifier. Through superposition, entanglement, and variational state evolution, the quantum circuit generates structured statistical representations that are subsequently interpreted by classical learning components.

This interpretation provides a useful explanation for several observations obtained during the benchmarking experiments. Many baseline HQNN architectures struggled under noisy conditions despite possessing theoretically expressive quantum circuits. Increasing circuit complexity alone often failed to improve practical performance because useful information generated by the quantum subsystem was either discarded or interpreted inefficiently.

By treating the quantum circuit as a feature-generation mechanism, the learning process becomes distributed across the hybrid system. The quantum subsystem focuses on generating informative representations, while the classical subsystem performs optimization, decoding, stabilization, and inference. This division of responsibility aligns closely with the practical strengths of each computational paradigm.

An additional benefit of this perspective is improved compatibility with NISQ hardware constraints. Useful feature representations can often be generated using relatively shallow circuits, reducing sensitivity to decoherence and accumulated gate errors. Rather than re-

quiring increasingly complex quantum architectures, the framework extracts greater value from existing quantum resources through more effective representation design.

This interpretation also supports the broader hybrid-first philosophy adopted throughout the thesis. Practical hybrid quantum learning emerges not from maximizing quantum computation in isolation, but from coordinating quantum and classical components in a manner that exploits the strengths of both systems.

The multi-observable framework therefore establishes one of the central conceptual conclusions of the thesis: under realistic NISQ conditions, quantum circuits function most effectively as probabilistic feature generators embedded within larger hybrid computational pipelines.

4.3.4 Threshold Improvement from 0.82 to 0.85

The introduction of multi-observable quantum feature extraction produced another measurable improvement in noisy HQNN robustness relative to the learned-readout baseline. Although the numerical gain was smaller than the earlier transition from parity collapse to learned decoding, the improvement remained consistent across multiple benchmark environments.

Following the introduction of learned-readout decoding, the framework achieved noisy classification performance near 0.82 under several evaluation conditions. While this represented a substantial improvement relative to baseline parity-based systems, analysis suggested that additional information generated by the quantum subsystem was still being discarded during feature extraction.

The incorporation of multiple observables expanded the available feature space and provided a richer description of quantum state behavior. As a result, the hybrid learning system gained access to additional information that could be used to distinguish signal from noise during both training and inference.

Across multiple benchmark environments, the expanded feature representation increased

noisy classification performance to approximately 0.85. The improvement was accompanied by more stable optimization behavior and reduced sensitivity to isolated measurement fluctuations.

Although the magnitude of the improvement was smaller than the earlier learned-readout transition, its significance extends beyond the numerical threshold itself. The results demonstrate that robustness can be improved through more informative feature extraction even when the underlying quantum circuit remains largely unchanged.

A second implication concerns the interpretation of quantum computation within hybrid learning systems. The improvement emerged because the framework preserved more information generated by the quantum subsystem, not because additional quantum complexity was introduced. This finding further supports the view that practical NISQ-era performance depends heavily on information utilization and hybrid design rather than on circuit expressivity alone.

The multi-observable framework therefore represents the second major robustness improvement within Chapter 4. Having established the importance of richer feature representations, the next section investigates how architecture design itself influences robustness, leading to the architecture optimization mechanisms that increase noisy performance from 0.85 to 0.89.

4.4 Architecture Robustness Optimization

The previous sections demonstrated that substantial robustness improvements can be achieved through improved interpretation of quantum information and richer feature representations. Learned-readout decoding reduced the information bottleneck associated with parity collapse, while multi-observable feature extraction preserved a broader description of quantum state evolution. Together, these mechanisms increased noisy performance from baseline levels near 0.50 to approximately 0.85 under several benchmark environments.

Despite these improvements, the benchmarking analysis revealed that circuit architecture

remained one of the strongest determinants of practical HQNN behavior. Models with similar parameter counts and comparable feature representations frequently exhibited very different optimization trajectories once realistic noise sources were introduced. Some architectures maintained relatively stable performance under perturbation, while others degraded rapidly despite possessing greater theoretical expressive capability.

These observations motivated the next stage of the robustness-engineering framework: architecture robustness optimization. Rather than treating circuit topology as a fixed design decision, this stage investigates whether architecture selection itself can be optimized for stability under realistic NISQ conditions. The objective is not simply to identify the most expressive circuit, but to determine which architectural structures maintain reliable behavior when noise, decoherence, and optimization instability are present simultaneously.

The experiments presented throughout this section focus on the role of entanglement topology, connectivity structure, and information propagation within variational HQNN systems. Multiple architectures were evaluated under identical datasets, training procedures, and noise environments in order to isolate the influence of circuit organization on robustness behavior.

A consistent pattern emerged throughout the analysis. Architectures designed primarily for expressivity frequently became more sensitive to perturbation accumulation as circuit complexity increased. By contrast, architectures that constrained information flow and limited entanglement propagation often achieved more stable optimization behavior despite possessing lower theoretical representational capacity.

This finding represents an important development in the broader thesis narrative. Earlier sections established that decoder design and feature representation influence robustness. The results presented here demonstrate that robustness is also shaped directly by the structure of the quantum circuit itself. Consequently, architecture selection becomes a central component of practical HQNN design under realistic NISQ constraints.

4.4.1 Entanglement Topology Search

To investigate how circuit structure influences robustness, the framework performed a systematic search across multiple entanglement topologies. The objective was to determine whether specific connectivity patterns provide greater resilience to noise accumulation and optimization instability than others.

Entanglement topology plays a critical role in variational quantum learning because it determines how information propagates throughout the circuit. Different connectivity structures create different interaction patterns between qubits, producing distinct optimization landscapes and noise-propagation behaviors even when the total number of qubits remains unchanged.

The topology search evaluated a range of architectures spanning highly localized interaction patterns to more globally connected structures. Each topology was trained using identical datasets, optimization procedures, and evaluation metrics. This controlled design ensured that observed performance differences could be attributed primarily to architectural organization rather than to unrelated experimental factors.

The results revealed that robustness depends strongly on how entanglement is distributed throughout the circuit. Architectures with extensive global connectivity frequently accumulated perturbations more rapidly because errors introduced in one region of the system could influence larger portions of the computational state. In contrast, more localized interaction patterns often restricted error propagation and improved optimization stability.

Another important observation was that architectural quality cannot be evaluated under noiseless conditions alone. Several topologies that appeared effective during idealized simulation exhibited substantial instability once realistic noise models were introduced. Consequently, architecture selection must be evaluated directly within noisy execution environments if practical robustness is the primary objective.

The topology search therefore established that circuit organization is not merely an implementation detail. It is a major factor influencing optimization behavior, noise sensitivity,

and overall learning stability within hybrid quantum systems.

4.4.2 Linear, Ring, Full, and Local Architectures

Four representative entanglement topologies emerged as particularly important during the architecture search: linear architectures, ring architectures, fully connected architectures, and localized modular structures.

Linear architectures employ sequential nearest-neighbor interactions that constrain information flow throughout the circuit. Because interactions occur only between adjacent qubits, these architectures remain relatively simple and often exhibit stable behavior under noisy execution conditions. However, their limited connectivity can restrict representational flexibility.

Ring architectures extend this concept by connecting the endpoints of the linear structure, creating a cyclic interaction pattern. This modification increases the ability of information to propagate throughout the circuit while maintaining moderate architectural complexity. In many benchmark environments, ring structures provided a useful balance between expressivity and stability.

Fully connected architectures maximize entanglement interactions by allowing extensive communication between qubits. Under idealized conditions, these systems often achieved strong representational capability because they could model highly correlated quantum states. However, they also proved substantially more sensitive to noise accumulation, optimization instability, and perturbation propagation.

Localized modular architectures organized interactions into partially independent computational regions. Rather than maximizing global connectivity, these designs emphasized controlled information flow and limited error propagation. Across several benchmark environments, modular structures frequently produced the most stable optimization behavior under realistic NISQ conditions.

As shown in Figure 4.3, the relative performance of these architectures depended strongly

on the interaction between expressivity and robustness.




Figure 4.3: Linear and ring structures provide controlled connectivity, fully connected circuits increase expressive capacity but amplify noise sensitivity, and localized modular designs restrict perturbation propagation by organizing interactions into smaller computational regions.

The comparative analysis demonstrated that no single architecture dominates universally across all environments. Instead, practical effectiveness depends on the balance between representational flexibility, optimization stability, and resistance to perturbation accumulation.

4.4.3 Robustness versus Expressivity Tradeoffs

One of the most important findings emerging from the architecture search was the existence of a recurring tradeoff between expressivity and robustness. This relationship appeared consistently across multiple datasets, architectures, and noise environments and became one of the central conceptual conclusions of the robustness-engineering framework.

Under idealized conditions, deeper and more densely connected architectures frequently

exhibited stronger representational capability. Their ability to generate complex entanglement structures allowed them to model highly correlated relationships and produce rich quantum feature representations.

However, these same characteristics often became liabilities under realistic NISQ conditions. Increased connectivity amplified noise propagation, deeper circuits accumulated more gate errors, and optimization landscapes became increasingly difficult to navigate as perturbations distorted objective-function evaluations.

By contrast, architectures that intentionally limited entanglement and constrained information flow often achieved greater stability during training. Although these systems possessed lower theoretical expressive capability, they frequently maintained more reliable optimization behavior and produced more consistent classification performance under noisy execution conditions.

This observation suggests that maximizing expressivity alone is not an effective design strategy for practical hybrid quantum learning. The most useful architectures are not necessarily those capable of representing the largest class of functions, but rather those capable of maintaining stable learning behavior under realistic hardware constraints.

The expressivity-robustness tradeoff therefore becomes a central design consideration throughout the remainder of the thesis. Robustness-aware architecture selection emerges as a practical requirement for NISQ-era HQNN systems.

4.4.4 Threshold Improvement from 0.85 to 0.89

The architecture optimization stage produced another measurable improvement in noisy HQNN performance. Following the introduction of learned-readout decoding and multi-observable feature extraction, the framework had previously achieved noisy performance near 0.85 under several benchmark environments.

After selecting architectures optimized specifically for robustness, noisy classification performance increased to approximately 0.89 while simultaneously exhibiting more stable

convergence behavior during training. The improvement was observed across multiple benchmark conditions and remained consistent throughout repeated optimization runs.

The significance of this improvement extends beyond the numerical threshold itself. Unlike the learned-readout stage, which focused on measurement interpretation, and the multi-observable stage, which focused on feature representation, architecture optimization altered the way quantum information was generated within the circuit. Robustness improvements therefore emerged directly from structural changes to the quantum subsystem.

The experiments also demonstrated that the strongest architectures were often not the most complex architectures. In many benchmark environments, moderately connected and modular designs outperformed more densely entangled systems once realistic noise conditions were introduced. This result reinforces the broader conclusion that practical performance depends heavily on stability rather than on raw quantum complexity.

Another important implication concerns the role of NISQ constraints. Hardware limitations fundamentally influence which architectures are useful in practice. Designs that perform well under idealized simulation may not remain effective when decoherence, gate errors, and stochastic perturbations are present. Robust architecture selection must therefore be guided by realistic execution conditions rather than theoretical capability alone.

The architecture optimization stage represents the third major robustness improvement within the framework evolution. Having demonstrated that decoder design, feature representation, and circuit architecture all influence robustness, the next section investigates how optimization itself can be modified to further improve learning stability and increase noisy performance from 0.89 to 0.96.

4.5 Noise-Aware Optimization and Stability Regularization

The previous stages of the robustness-engineering framework demonstrated that substantial improvements in noisy HQNN performance can be achieved through better decoding strategies, richer feature representations, and robustness-oriented architecture design. These

modifications increased noisy performance from baseline levels near 0.50 to approximately 0.89 while significantly improving learning stability under realistic NISQ conditions.

Despite these gains, the benchmarking experiments continued to reveal an important limitation. Even when architectures were carefully designed and richer feature representations were available, optimization behavior frequently remained sensitive to stochastic perturbations. Training trajectories often exhibited oscillatory convergence patterns, noisy parameter updates, and substantial variability across repeated runs. These observations suggested that optimization itself remained a major source of instability within hybrid quantum learning systems.

This finding motivated the next stage of the framework evolution: noise-aware optimization and stability regularization. Rather than treating optimization as a conventional loss-minimization problem, this stage introduces robustness directly into the training objective. The central idea is that HQNN systems should optimize not only for predictive performance, but also for stable behavior under realistic noisy execution conditions.

From this perspective, robustness becomes an explicit optimization target rather than an indirect consequence of architecture design. The learning process is encouraged to identify parameter configurations that maintain both predictive accuracy and consistency when exposed to stochastic perturbations. This modification shifts the role of optimization from simple performance maximization toward robustness-aware learning.

The experiments presented throughout this section demonstrate that optimization stability plays a central role in practical hybrid quantum learning. More importantly, the results show that substantial robustness gains can be achieved by modifying the optimization process itself rather than relying exclusively on changes to architecture or feature representation.

The findings also reinforce a broader conclusion emerging throughout the thesis. Reliable HQNN behavior depends on coordinated interaction between multiple components of the learning pipeline. Decoder design, feature extraction, architecture selection, and optimization stability all contribute to robustness, and improvements in any one component can

influence the effectiveness of the entire system.

4.5.1 Dual-Loss Optimization Objectives

The baseline HQNN systems evaluated in earlier chapters optimized primarily for predictive performance. Standard supervised learning objectives encouraged the model to improve classification accuracy, but they did not explicitly account for robustness or stability under noisy execution conditions. As a result, optimization frequently converged toward parameter regions that performed well under isolated evaluations while remaining highly sensitive to stochastic perturbations.

To address this limitation, the framework introduced a dual-loss optimization strategy designed to balance predictive performance with robustness. Instead of relying solely on a conventional classification objective, the training process incorporated an additional stability-oriented component that evaluated the consistency of model behavior under noisy conditions.

The purpose of the additional objective was not to replace the classification loss, but rather to complement it. The framework continued to reward accurate predictions while simultaneously encouraging parameter configurations that maintained reliable behavior when exposed to perturbations. This modification allowed robustness to influence the optimization process directly.

An important consequence of the dual-loss formulation was the reduction of unstable parameter exploration. Highly sensitive parameter regions that produced inconsistent predictions became less attractive during training because instability contributed directly to the optimization objective. The optimizer therefore favored solutions exhibiting both predictive quality and robustness.

The modified objective also reduced the tendency of the system to exploit noisy measurement fluctuations. In several baseline configurations, optimization occasionally converged toward parameter settings that appeared favorable under limited evaluation but deteriorated

rapidly when tested repeatedly under noise. The dual-loss framework reduced this behavior substantially by rewarding consistency across noisy execution environments.

Another important benefit was improved generalization across perturbation conditions. Parameter configurations selected through stability-aware objectives frequently maintained stronger performance when evaluated under unseen noise levels than configurations optimized solely for predictive accuracy.

The introduction of dual-loss optimization therefore represents an important conceptual transition within the thesis. Optimization is no longer interpreted solely as a mechanism for maximizing performance, but also as a mechanism for directly engineering robustness into the learning process.

4.5.2 Stability-Regularized HQNN Training

Building upon the dual-loss objective, the framework next introduced stability-regularized training procedures designed specifically for noisy quantum learning environments. Rather than exposing the model only to idealized execution conditions during optimization, the training process incorporated controlled perturbation environments throughout learning.

This modification allowed the optimizer to observe how parameter configurations behaved under realistic execution conditions while training was still occurring. Instead of discovering instability only during post-training evaluation, the framework incorporated robustness information directly into the learning process itself.

One of the most important effects of stability regularization was improved convergence consistency. Optimization trajectories became substantially smoother across repeated runs, and the severe oscillatory behavior observed in earlier benchmark experiments was reduced considerably. Parameter updates became less sensitive to isolated measurement fluctuations and more representative of broader performance trends.

The regularization process also altered the structure of the optimization landscape experienced during training. Highly unstable regions that previously attracted optimization

trajectories became less favorable, while smoother and more reproducible regions became increasingly attractive. This encouraged convergence toward parameter configurations exhibiting greater resilience to perturbation.

The experiments further suggested that stability regularization improved reproducibility. Baseline HQNN systems often produced noticeably different outcomes depending on initialization conditions and perturbation ordering. Regularized systems exhibited far more consistent training behavior across repeated evaluations performed under identical benchmark settings.

An additional benefit was reduced sensitivity to stochastic gradient variability. Variational quantum optimization frequently suffers from noisy objective evaluations, particularly under realistic hardware conditions. Stability regularization helped moderate the impact of these fluctuations and produced more reliable optimization behavior throughout training.

These findings reinforce an increasingly important theme within the thesis: optimization stability is not merely a secondary implementation detail. It is one of the primary factors determining whether practical hybrid quantum learning can function reliably under realistic NISQ constraints.

4.5.3 Optimization Behavior Under Noise

The integration of dual-loss objectives and stability regularization substantially altered the behavior of variational optimization under noisy execution conditions. Compared with baseline systems, the modified framework produced smoother convergence trajectories, reduced optimization variance, and more consistent learning behavior across repeated runs.

The benchmarking experiments revealed that noise influences optimization through several interconnected mechanisms. Measurement fluctuations distort objective-function evaluations, stochastic perturbations alter effective gradient information, and repeated execution variability introduces uncertainty into parameter updates. Together, these effects create optimization landscapes that are substantially more difficult to navigate than those encountered

under idealized simulation conditions.

Within baseline HQNN systems, these effects frequently produced unstable convergence. Training trajectories often oscillated between competing parameter regions, and optimization occasionally failed to maintain consistent progress toward stable solutions. These behaviors became increasingly severe as noise intensity increased.

After introducing stability-aware optimization, however, convergence behavior improved substantially. Training trajectories became smoother, optimization variance decreased, and repeated runs exhibited far greater consistency. The framework maintained more stable parameter evolution throughout training despite continued exposure to stochastic perturbations.

As shown in Figure 4.4, stability-aware optimization significantly reduced oscillatory training behavior and improved convergence reliability under noisy execution conditions.




Figure 4.4: The baseline trajectory exhibits oscillatory loss behavior caused by noisy measurements, unstable objective-function evaluations, and stochastic parameter updates. The stability-aware framework produces smoother convergence, lower variance, and more consistent progress under comparable noisy conditions.

Another important observation was that optimization improvements amplified the effectiveness of earlier framework mechanisms. Learned-readout decoding, multi-observable

feature extraction, and architecture optimization all benefited from more stable training dynamics. Consequently, robustness gains accumulated across multiple framework stages rather than emerging from any single isolated modification.

The results therefore indicate that optimization behavior represents one of the most influential determinants of practical HQNN performance. Even highly expressive architectures cannot perform reliably if optimization remains unstable under realistic execution conditions.

4.5.4 Threshold Improvement from 0.89 to 0.96

The introduction of noise-aware optimization and stability regularization produced one of the largest robustness improvements observed throughout the framework evolution. Following architecture robustness optimization, the framework had achieved noisy classification performance near 0.89 across several benchmark environments. After incorporating stability-aware optimization procedures, noisy performance increased to approximately 0.96.

The significance of this improvement extends beyond the numerical threshold itself. Earlier framework stages improved how information was decoded, represented, and generated. This stage improved how that information was learned. The resulting gains demonstrate that optimization stability is capable of influencing performance as strongly as architecture design and feature representation.

A particularly important observation was the dramatic reduction in performance degradation between clean and noisy evaluation environments. Earlier benchmark systems frequently experienced substantial drops when perturbations were introduced. Following stability-aware optimization, several benchmark configurations exhibited only minimal differences between clean and noisy performance.

The experiments also demonstrated that the observed gains did not arise from a single mechanism operating in isolation. Instead, the improvement emerged from coordinated interaction between all previously introduced robustness mechanisms. Learned-readout decoding preserved richer measurement information, multi-observable extraction generated more in-

formative features, architecture optimization improved robustness at the circuit level, and stability-aware optimization enabled those improvements to be learned more effectively.

This layered interaction represents one of the most important scientific conclusions of the thesis. Practical HQNN robustness is not the result of a single architectural innovation. It emerges progressively through coordinated robustness engineering across the entire hybrid learning pipeline.

The improvement from 0.89 to 0.96 therefore marks a major turning point in the framework evolution. For the first time, the system begins approaching consistently reliable learning behavior under realistic NISQ conditions. This result establishes the foundation for the final robustness stage introduced in the following section, where noise-augmented readout training further increases noisy performance to approximately 0.97.

4.6 Noise-Augmented Readout Training

The stability-aware optimization framework introduced in the previous section substantially improved noisy HQNN performance and produced one of the largest robustness gains observed throughout the framework evolution. By incorporating stability directly into the optimization objective, the framework achieved noisy performance near 0.96 while significantly reducing degradation under stochastic perturbation environments.

Despite these improvements, an important limitation remained. The decoder system continued to learn primarily from feature representations generated under conditions that did not fully reflect the variability encountered during realistic execution. Although optimization became more stable, the learning process itself was still influenced by a mismatch between training conditions and deployment conditions. As a result, the framework remained vulnerable to distributional shifts between clean and noisy quantum feature representations.

To address this limitation, the robustness-engineering framework introduced noise-augmented readout training. Rather than exposing the decoder only to idealized quantum outputs, this approach incorporated noisy quantum feature distributions directly into the training process.

The objective was to allow the decoder to learn how realistic quantum representations behave under perturbation rather than requiring it to adapt only after performance degradation occurred during evaluation.

This modification represents an important transition within the framework evolution. Earlier stages focused on preserving information, improving architecture design, and stabilizing optimization. Noise-augmented training instead focuses on adaptation. The learning system is trained directly within the conditions it is expected to encounter during practical deployment.

The underlying principle is straightforward. If noise is an unavoidable characteristic of NISQ hardware, then training procedures should incorporate that reality rather than attempting to learn exclusively from idealized representations. By exposing the decoder to realistic perturbation environments during training, the framework encourages the development of inference strategies that remain effective when noise is present.

The experiments presented throughout this section demonstrate that direct exposure to noisy quantum feature distributions improves inference reliability and robustness under realistic execution conditions. More importantly, the results reveal that the strongest performance emerges when feature generation, decoder adaptation, architecture design, and optimization stability operate cooperatively within a unified hybrid framework.

Noise-Aware Readout	Stability-Aware Training
Goal: Decode noisy quantum measurement outputs into stable predictions.	Goal: Stabilize optimization when loss values and gradient estimates fluctuate under noise.
Acts on: Measurement distributions produced by quantum circuit execution.	Acts on: Parameter updates produced during variational training.
Input: Bitstrings, shot counts, expectation values, and noisy quantum feature distributions.	Input: Loss values, gradient estimates, update variance, and noisy objective evaluations.
Method: Learned classical decoder interprets noisy quantum features during inference.	Method: Dual-loss objective and stability regularization reduce unstable parameter movement.
Primary benefit: Reliable inference from noisy measurements.	Primary benefit: Smoother convergence and stable training dynamics.
Key distinction: readout improves inference, while training improves optimization.	

Figure 4.5: Noise-aware readout improves how noisy quantum measurement outputs are decoded into predictions, while stability-aware training improves how model parameters are updated when loss evaluations and gradient estimates fluctuate under noisy execution.

4.6.1 Noisy Quantum Feature-Space Learning

Earlier stages of the framework established that variational quantum circuits function effectively as probabilistic feature generators. However, the statistical properties of these feature representations often differ substantially between clean and noisy execution environments. Measurement distributions collected under realistic perturbation conditions frequently exhibit altered probability structures, shifted expectation statistics, and increased stochastic variability relative to idealized simulation outputs.

These differences create an important challenge for downstream inference systems. A decoder trained primarily on cleaner feature representations may learn decision boundaries that perform well under ideal conditions but become less effective when exposed to realistic measurement variability. Consequently, a gap can emerge between training behavior and deployment behavior.

To reduce this gap, the framework introduced noisy quantum feature-space learning. During training, the decoder was exposed directly to feature representations generated under controlled perturbation environments. Rather than learning only from idealized quantum outputs, the system learned from the same types of noisy distributions encountered during practical execution.

An important consequence of this approach was improved tolerance to feature-space variability. Because the decoder repeatedly observed perturbed measurement distributions during optimization, it became less sensitive to isolated fluctuations and more capable of identifying persistent statistical structure within noisy representations.

The experiments also suggested improved generalization across perturbation conditions. Decoders trained using noise-augmented feature distributions frequently maintained more consistent performance when evaluated under previously unseen noise environments than decoders trained primarily on cleaner representations.

Another benefit involved improved alignment between training and evaluation. Earlier architectures occasionally experienced performance degradation because the statistical properties of the evaluation environment differed from those encountered during optimization. Noise-augmented learning reduced this mismatch substantially.

These findings support a broader conclusion emerging throughout the thesis. Robustness depends not only on preserving information and stabilizing optimization, but also on ensuring that the learning process reflects realistic execution conditions from the beginning.

4.6.2 Decoder Adaptation Under Stochastic Noise

A central component of the noise-augmented framework was adaptive decoder training under stochastic perturbation environments. Although learned-readout decoding introduced earlier in the chapter significantly improved robustness relative to parity-based inference, the decoder still lacked explicit adaptation to dynamically changing feature distributions during training.

The noise-augmented framework addressed this limitation by integrating stochastic perturbation exposure directly into decoder optimization. During training, feature representations were generated under varying noise conditions, allowing the decoder to learn adaptive mappings between noisy quantum outputs and target labels.

This procedure transformed the decoder from a static inference mechanism into an explicitly robustness-aware component. Rather than interpreting noisy outputs using decision boundaries learned primarily under cleaner conditions, the decoder evolved continuously in response to realistic feature-space variability.

The experiments demonstrated that adaptive decoder training reduced inference instability substantially. The decoder learned to identify informative patterns that persisted across perturbation environments while becoming less sensitive to transient fluctuations generated by noise.

An important result emerged from later ablation experiments. When evaluated independently of the full framework, the adaptive decoder produced measurable improvements in noisy performance relative to conventional decoding approaches. This observation indicates that decoder adaptation functions as a meaningful robustness mechanism in its own right rather than simply serving as a secondary consequence of other framework modifications.

The adaptation process also improved consistency across repeated evaluations. Earlier baseline systems frequently produced noticeably different outputs under identical conditions because decision boundaries remained highly sensitive to measurement variability. Adaptive decoders maintained more reproducible behavior across repeated noisy executions.

These findings reinforce the broader argument that robustness depends not only on the quality of quantum representations, but also on how effectively those representations are interpreted under realistic perturbation conditions.

4.6.3 Hybrid Robustness Co-Design

The experiments conducted throughout Chapter 4 consistently revealed that no single modification fully explained the observed improvements in HQNN robustness. Learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware training, and noise-augmented decoder adaptation all contributed to improved performance, but none alone accounted for the complete framework behavior.

This observation motivated the broader concept of hybrid robustness co-design. Within this perspective, robustness is interpreted not as a property of a single subsystem, but as an emergent characteristic arising from coordinated interaction between quantum and classical components throughout the learning pipeline.

The quantum subsystem serves primarily as a probabilistic feature-generation mechanism operating under realistic noisy conditions. The classical subsystem provides optimization, decoder adaptation, stabilization, and inference capabilities that allow useful information to be extracted from those quantum representations. Robustness emerges through cooperation between these components rather than through improvements to either subsystem in isolation.

This perspective differs substantially from earlier interpretations of quantum machine learning that emphasized quantum expressivity as the dominant factor influencing performance. The experiments presented throughout this thesis repeatedly demonstrated that increasing quantum complexity alone often failed to improve practical behavior and frequently increased sensitivity to noise.

Instead, the strongest performance emerged when architecture design, feature extraction, optimization strategy, and decoder adaptation were engineered jointly around realistic

NISQ constraints. This coordinated design process allowed the framework to preserve useful information while simultaneously maintaining stability throughout training and inference.

The concept of hybrid robustness co-design therefore becomes one of the central interpretations of the thesis. Practical HQNN systems should be viewed as integrated computational ecosystems in which robustness is engineered across the entire learning pipeline rather than being assigned to any single component.

4.6.4 Threshold Improvement from 0.96 to 0.97

The introduction of noise-augmented readout training produced the highest noisy performance achieved throughout the robustness-engineering framework. Following the stability-aware optimization stage, the framework had achieved noisy classification performance near 0.96 under several benchmark environments. After incorporating noisy feature-space learning and adaptive decoder training, performance improved further to approximately 0.97.

Although the numerical improvement was smaller than several earlier threshold transitions, its significance should not be evaluated solely through the magnitude of the increase. At this stage of framework evolution, the system was already operating near the upper limits observed throughout the benchmark environments. Further gains therefore became progressively more difficult to achieve.

The improvement emerged primarily through better alignment between training conditions and evaluation conditions. By exposing the decoder directly to noisy feature distributions during optimization, the framework reduced sensitivity to distributional shifts and improved reliability under realistic execution environments.

Another important observation was the emergence of highly consistent inference behavior across repeated noisy evaluations. Earlier baseline systems frequently experienced severe degradation under perturbation conditions, while the final framework maintained remarkably stable performance despite continued exposure to realistic noise environments.

The results also reinforce one of the central conclusions of the thesis. Robustness im-

provements do not emerge from deeper quantum circuits alone. Instead, they arise through coordinated interaction between feature generation, architecture design, optimization procedures, decoder adaptation, and training methodology.

The transition from 0.96 to 0.97 therefore represents more than a numerical improvement. It marks the culmination of the robustness-engineering progression developed throughout Chapter 4 and demonstrates that highly reliable hybrid quantum learning behavior can emerge when robustness is incorporated systematically throughout the entire learning pipeline.

Having established the full robustness-engineering framework, the next section investigates which mechanisms contribute most strongly to the observed performance gains through controlled ablation analysis and mechanism isolation experiments.

4.7 Framework Ablation and Mechanism Isolation

The robustness-engineering framework developed throughout the previous sections introduced a sequence of coordinated modifications designed to improve HQNN behavior under realistic NISQ conditions. These modifications included learned-readout decoding, multi-observable feature extraction, architecture robustness optimization, stability-aware optimization, and noise-augmented readout training. Collectively, these mechanisms increased noisy performance from baseline levels near 0.50 to approximately 0.97 while substantially improving optimization stability and inference reliability.

Despite these improvements, an important scientific question remained. To what extent did each mechanism contribute independently to the observed gains, and how much of the final performance emerged through interaction between framework components? Without controlled ablation analysis, the observed improvements could potentially be interpreted as a collection of empirical observations rather than as a systematically validated robustness methodology.

For this reason, the framework incorporated a series of mechanism-isolation experiments

designed to evaluate individual contributions under controlled benchmark conditions. The objective was not merely to verify that the final framework performed well, but to determine why it performed well. By isolating specific components while holding others fixed, the experiments provided insight into the causal role played by each robustness mechanism.

The resulting analysis revealed that multiple components contributed independently to robustness improvements. At the same time, the strongest performance consistently emerged when these mechanisms operated together. No single modification explained the full progression observed throughout Chapter 4.

This observation is important because it transforms the interpretation of the framework. Rather than representing a collection of unrelated improvements, the results indicate that robustness emerges progressively through coordinated interaction between feature generation, decoder adaptation, architecture design, optimization stabilization, and noise-aware training.

The experiments presented throughout this section therefore provide the strongest evidence supporting the broader robustness-engineering philosophy developed throughout the thesis. They demonstrate that practical HQNN robustness is fundamentally a systems-level property rather than the consequence of any individual modification.

4.7.1 Standard Loss versus Dual-Loss Optimization

The first ablation study focused on evaluating the contribution of stability-aware optimization independently from the remainder of the framework. Earlier sections demonstrated that dual-loss optimization substantially improved noisy convergence behavior, but it remained necessary to determine whether those gains originated directly from the stability objective itself.

To isolate this effect, the experiments compared standard supervised optimization against the dual-loss formulation while maintaining identical architectures, datasets, decoder structures, and noise conditions. In the baseline configuration, optimization focused exclusively

on predictive performance. In the dual-loss configuration, optimization incorporated an additional stability-oriented objective designed to encourage consistent behavior under perturbation.

The results demonstrated that the dual-loss framework independently improved robustness. Compared with standard optimization, the modified objective produced smoother convergence trajectories, reduced oscillatory behavior, and improved consistency across repeated noisy evaluations. These improvements occurred even when architecture and decoder design remained unchanged.

An important implication of this finding is that optimization instability itself represents a major source of performance degradation within variational quantum learning systems. Earlier sections established that noise affects measurement outcomes and feature representations. The ablation results further demonstrate that optimization dynamics amplify these effects when stability is not incorporated directly into the learning objective.

The dual-loss objective therefore functions as more than a regularization technique. It fundamentally changes how the optimization process navigates noisy parameter landscapes by favoring solutions that remain reliable under perturbation rather than solutions that merely perform well under isolated evaluations.

These observations establish optimization stability as a primary robustness mechanism within the overall framework and confirm that meaningful improvements can be achieved by modifying the learning objective itself.

4.7.2 Standard Readout versus Noise-Augmented Readout

The second ablation study examined the contribution of noise-augmented decoder training independently from stability-aware optimization. Although learned-readout decoding and noise-aware adaptation produced substantial improvements throughout the framework evolution, it remained important to determine whether decoder adaptation itself represented an independent robustness mechanism.

To isolate this effect, the experiments compared standard learned-readout decoding against noise-augmented decoder training while maintaining identical architectures, optimization procedures, and benchmark conditions. In the standard configuration, the decoder learned primarily from cleaner feature representations. In the noise-augmented configuration, the decoder was exposed directly to noisy feature-space distributions during training.

The results revealed that adaptive decoder training independently improved noisy inference behavior. Even without the complete framework, exposure to noisy feature representations produced more stable classification performance and reduced sensitivity to stochastic perturbations during evaluation.

One particularly important observation was the improvement in reproducibility. Earlier baseline systems frequently produced inconsistent predictions across repeated noisy evaluations because decision boundaries remained highly sensitive to measurement variability. Noise-augmented decoders maintained substantially more consistent behavior under identical conditions.

The experiments also demonstrated that decoder adaptation improves the ability of the learning system to interpret noisy quantum information effectively. Rather than relying on decision boundaries learned primarily under idealized conditions, the decoder learns directly from realistic execution behavior and therefore becomes better aligned with deployment environments.

These findings confirm that decoder adaptation is not simply a secondary consequence of optimization stabilization. It is a meaningful robustness mechanism in its own right and contributes independently to the overall framework progression.

4.7.3 Isolating Mechanism Contributions

Beyond pairwise comparisons, the framework next performed broader mechanism-isolation experiments designed to evaluate the contribution of each robustness component individually. These experiments systematically enabled and disabled selected framework mechanisms while

maintaining all remaining conditions constant.

The resulting analysis revealed that different mechanisms contributed to robustness in different ways. Learned-readout decoding primarily reduced information loss during measurement interpretation. Multi-observable feature extraction improved the richness and stability of quantum representations. Architecture optimization reduced noise propagation and improved convergence reliability. Stability-aware optimization improved learning dynamics, while noise-augmented training improved adaptation to realistic execution conditions.

Although each mechanism produced measurable gains independently, their effects were not equivalent. Some mechanisms primarily improved inference reliability, while others influenced optimization stability or feature quality. Consequently, robustness improvements emerged through multiple pathways rather than through a single dominant intervention.

An important conclusion emerged from this analysis. No individual mechanism was capable of reproducing the full robustness progression observed throughout Chapter 4. Each component addressed a different limitation within the hybrid learning pipeline, and removing any one of them reduced overall framework effectiveness.

This observation significantly strengthens the scientific interpretation of the framework. The observed gains cannot be explained as the consequence of a single optimization trick, decoder modification, or architectural adjustment. Instead, robustness emerges progressively through layered interaction between multiple coordinated components.

The mechanism-isolation experiments therefore provide direct evidence that practical HQNN robustness is inherently multidimensional and cannot be reduced to a single design decision.

Framework Variant	Read-out	Multi-Obs.	Architecture	Stability	Noise Aug.	Accuracy
Baseline HQNN	–	–	–	–	–	0.50
Learned Readout Only	ON	–	–	–	–	0.82
Multi-Observable Only	–	ON	–	–	–	0.79
Architecture Only	–	–	ON	–	–	0.77
Stability-Aware Only	–	–	–	ON	–	0.84
Complete Framework	ON	ON	ON	ON	ON	0.97

Figure 4.6: Component-level ablation of the proposed HQNN framework. Each partial configuration activates one principal robustness mechanism while leaving the remaining mechanisms inactive. All isolated configurations improve noisy classification accuracy relative to the baseline, but none matches the complete framework. The final result therefore reflects coordinated interaction across the hybrid learning pipeline. Noise-augmented training is shown only within the complete framework because a separate noise-augmentation-only result was not reported in this comparison.

The ablation comparison demonstrates that the mechanisms address distinct weaknesses within the hybrid learning pipeline. Learned-readout decoding produces the strongest isolated improvement, while multi-observable extraction, architecture optimization, and stability-

aware training provide complementary benefits. The complete framework achieves the highest accuracy because these mechanisms reinforce one another rather than functioning as interchangeable modifications.

4.7.4 Full Framework versus Partial Frameworks

The final ablation experiments compared the complete robustness-engineered framework against partially integrated variants in order to evaluate how different combinations of mechanisms influenced final performance. Several intermediate configurations were constructed by combining subsets of the robustness mechanisms introduced throughout the chapter.

The results demonstrated that partial frameworks consistently outperformed the original baseline architecture. Systems incorporating only learned-readout decoding, only architecture optimization, or only stability-aware optimization all produced measurable improvements relative to baseline HQNN behavior. However, none achieved the same level of performance observed in the complete framework.

The strongest robustness emerged when all mechanisms operated together. Decoder adaptation improved interpretation of noisy outputs, multi-observable extraction improved feature quality, architecture optimization reduced perturbation propagation, stability-aware optimization improved convergence behavior, and noise-augmented training improved adaptation to realistic execution conditions. The combined effect exceeded the contribution of any individual mechanism.

As shown in Figure 4.7, the complete framework consistently outperformed all partial variants across the evaluated benchmark environments.




Figure 4.7: Each partial mechanism improves noisy performance relative to the baseline, but the complete framework achieves the strongest result by combining learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware optimization, and noise-aware training into a coordinated robustness-engineered pipeline.

One particularly important observation was that framework behavior appeared cooperative rather than merely additive. The interaction between robustness mechanisms produced improvements greater than would be expected if each component contributed independently. This suggests that robustness emerges through reinforcement between components operating throughout the hybrid learning pipeline.

The full-framework comparison therefore provides the strongest validation of the robustness-engineering philosophy developed throughout Chapter 4. Practical HQNN robustness is not achieved through isolated modifications or increased quantum complexity alone. It emerges through coordinated design across feature generation, architecture selection, optimization procedures, decoder adaptation, and noise-aware training.

This conclusion completes the causal validation stage of the framework. Having demonstrated both the individual contributions and the cooperative interaction of the robustness mechanisms, the next section investigates whether these improvements generalize across

multiple noise channels and architecture families.

4.8 Multichannel Validation and Generalization

The previous section established that the robustness improvements observed throughout Chapter 4 emerged from identifiable and measurable mechanisms rather than from isolated experimental artifacts. Through controlled ablation analysis, the framework demonstrated that decoder adaptation, feature extraction, architecture optimization, stability-aware training, and noise-augmented learning each contributed meaningfully to overall performance. An important question nevertheless remained unresolved. Do these improvements generalize across different perturbation environments, or do they depend primarily on the specific noise conditions used during framework development?

To address this question, the robustness-engineering framework incorporated a series of multichannel validation experiments. Rather than evaluating performance under a single perturbation model, the framework was tested across several fundamentally different noise processes commonly used in NISQ benchmarking studies. This approach allowed the analysis to determine whether the observed gains represented genuine robustness improvements or merely adaptation to isolated experimental conditions.

The distinction is important because practical quantum hardware rarely experiences only one source of noise. Real quantum systems operate under combinations of decoherence, gate imperfections, measurement variability, and environmental disturbances. Consequently, a framework that performs well under one perturbation model but degrades substantially under another provides limited practical value.

The experiments presented throughout this section evaluate framework behavior under depolarizing noise, bit-flip noise, phase-flip noise, and amplitude damping. These perturbation channels influence quantum states through different physical mechanisms and therefore provide a useful basis for assessing generalization.

The results demonstrate that the robustness-engineered HQNN framework maintains

substantially greater stability across multiple perturbation environments than baseline architectures. More importantly, the improvements generalize consistently across several fundamentally different noise channels rather than depending on a single perturbation model.

This observation provides strong evidence that the framework is capturing broader robustness principles rather than overfitting to a specific experimental setting. The improvements therefore extend beyond isolated benchmarks and support the development of more reliable hybrid quantum learning systems under realistic NISQ conditions.

4.8.1 Depolarizing, Bit-Flip, Phase-Flip, and Amplitude Damping

The multichannel validation experiments evaluated framework performance under four widely studied quantum noise channels: depolarizing noise, bit-flip noise, phase-flip noise, and amplitude damping. Each channel affects quantum computation differently and therefore provides unique insight into the behavior of hybrid learning systems under perturbation.

Depolarizing noise represents one of the most disruptive stochastic error processes because it partially randomizes the quantum state during execution. Earlier benchmark experiments demonstrated that baseline HQNN architectures frequently experienced severe degradation under depolarizing conditions, often approaching near-random classification performance as perturbation intensity increased.

Bit-flip and phase-flip channels introduce different forms of state corruption. Bit-flip noise alters computational basis states directly, while phase-flip noise modifies phase relationships without necessarily changing measurement outcomes immediately. Although these channels affect quantum states through different mechanisms, both frequently produced unstable behavior within earlier baseline architectures.

Amplitude damping introduces yet another form of perturbation by modeling energy relaxation processes within the quantum system. Unlike depolarizing noise, which introduces generalized randomness, amplitude damping drives quantum states toward lower-energy configurations and can gradually alter computational behavior over time.

Despite these substantial differences between perturbation models, the robustness-engineered framework maintained significantly improved behavior across all evaluated environments. Although the degree of degradation varied between channels, the framework consistently exhibited stronger performance, improved convergence stability, and more reliable inference behavior than baseline HQNN systems.

One particularly important observation was that the framework improvements persisted even when perturbations altered quantum states through fundamentally different physical mechanisms. This suggests that the robustness gains emerged from broader stabilization principles rather than from adaptation to one specific form of noise.

The results therefore support the conclusion that the framework improves general system reliability rather than merely compensating for isolated perturbation characteristics.

4.8.2 Consistency Across Noise Channels

Beyond evaluating each perturbation environment independently, the framework also examined consistency across multiple noise channels. This analysis focused on whether the improvements observed during framework evolution remained stable as the perturbation environment changed.

Consistency is an important robustness indicator because practical quantum systems operate under varying hardware conditions. Models that perform well under one perturbation channel but degrade substantially under another may achieve strong benchmark results while still exhibiting poor real-world reliability. Consequently, stable behavior across heterogeneous environments represents an important objective for practical hybrid quantum learning.

The experiments revealed that the robustness-engineered framework maintained substantially greater consistency across perturbation environments than baseline architectures. Earlier parity-based systems frequently exhibited unpredictable degradation patterns when the perturbation channel changed, whereas the robustness-engineered framework produced

smoother and more reproducible behavior across all evaluated conditions.

An important finding was that degradation became more predictable as the framework evolved. Early architectures often responded dramatically to changes in perturbation type, making performance difficult to anticipate. In contrast, the robustness-engineered systems exhibited more gradual and consistent responses across different noise environments.

This improvement suggests that the framework stabilized the overall learning process rather than merely improving performance within individual benchmark settings. Decoder adaptation, richer feature extraction, architecture optimization, and stability-aware training collectively reduced sensitivity to perturbation-specific fluctuations throughout the computational pipeline.

Another notable observation was that consistency improved progressively throughout the Chapter 4 framework evolution. Early mechanisms produced localized gains under selected conditions, while later robustness-oriented modifications increasingly generalized across heterogeneous perturbation environments.

The consistency analysis therefore provides additional evidence that the framework is improving fundamental stability properties within the hybrid learning system rather than optimizing narrowly for selected benchmark conditions.

4.8.3 Robustness Generalization Across Architectures

The final stage of multichannel validation examined whether the robustness-engineering framework generalized across different HQNN architectures. Earlier benchmarking experiments demonstrated that architectural structure strongly influences degradation behavior under noise. Consequently, it was necessary to determine whether the framework improvements remained effective across multiple classes of hybrid quantum models.

The experiments evaluated the framework across several representative HQNN architectures, including Quantvolutional Neural Networks, Quantum Convolutional Neural Networks, and Quantum Transfer Learning systems. The robustness mechanisms introduced through-

out Chapter 4 were integrated into each architecture while maintaining comparable datasets, optimization procedures, and perturbation environments.

The results demonstrated that framework improvements generalized consistently across all evaluated architecture classes. Although the magnitude of improvement varied depending on circuit structure and connectivity patterns, every architecture benefited from the robustness-oriented modifications introduced throughout the framework evolution.

Modular architectures such as QuanNNs continued to exhibit particularly strong robustness behavior due to their constrained perturbation propagation and shallow circuit organization. More globally entangled architectures such as QCNNs remained comparatively more sensitive to perturbation accumulation, but nevertheless achieved substantial improvements relative to their baseline counterparts.

An important implication of this finding is that the framework does not depend on one specific architectural configuration. The robustness mechanisms target broader properties of hybrid learning systems and therefore remain effective even when circuit organization changes significantly.

The generalization observed across architecture families further strengthens the scientific interpretation of the framework. Rather than overfitting to a particular circuit topology or benchmark environment, the robustness-engineering methodology captures principles that remain applicable across multiple classes of hybrid quantum systems.

Collectively, the multichannel validation experiments provide strong evidence that the framework represents a generalized approach to improving hybrid quantum learning under realistic NISQ conditions. The observed gains extend across different perturbation environments, different physical noise processes, and different HQNN architectures. This broad generalization establishes the foundation for the final section of Chapter 4, which synthesizes the complete framework evolution and interprets the threshold progression observed throughout the robustness-engineering process.

4.9 Framework Evolution and Threshold Progression

The robustness-engineering framework developed throughout this chapter evolved through a sequence of coordinated modifications designed to improve hybrid quantum learning under realistic NISQ conditions. Rather than emerging from a single architectural innovation or optimization procedure, the observed robustness improvements accumulated progressively across multiple stages of framework development. Each stage addressed a different source of instability while simultaneously contributing to broader improvements throughout the hybrid learning pipeline.

The benchmarking experiments presented earlier in the thesis established that baseline HQNN systems exhibit substantial instability under realistic noisy execution environments. Optimization trajectories frequently oscillated, parity-based decoding discarded useful measurement information, and stochastic perturbations often caused predictive performance to deteriorate significantly. These limitations established the practical motivation for the robustness-oriented methods introduced throughout Chapter 4.

The subsequent framework evolution addressed these limitations systematically. Learned-readout decoding improved interpretation of noisy quantum measurements. Multi-observable feature extraction preserved richer statistical information from quantum state evolution. Architecture optimization reduced perturbation propagation through more robust entanglement structures. Stability-aware optimization improved convergence behavior, while noise-augmented training aligned the learning process more closely with realistic execution environments.

An important consequence of this progression was the emergence of measurable threshold improvements throughout the framework evolution. Each mechanism contributed to robustness in a different way, yet all participated in a broader process of reducing instability and improving reliability under perturbation.

The resulting progression transformed the interpretation of the thesis itself. What began as an investigation into HQNN limitations evolved into a robustness-engineering framework

capable of maintaining stable hybrid quantum learning behavior under realistic NISQ constraints.

The following subsections summarize this evolution, examine the threshold progression observed throughout the framework, and provide a final interpretation of robustness within practical hybrid quantum learning systems.

4.9.1 Baseline to Robustness-Engineered HQNNs

The earliest HQNN architectures evaluated throughout the benchmarking framework relied on shallow variational circuits, parity-based decoding procedures, and conventional optimization objectives. Although these systems demonstrated meaningful learning capability under idealized simulation environments, their behavior under realistic noisy conditions proved considerably less stable.

Several recurring limitations emerged. Optimization frequently exhibited oscillatory convergence behavior, parity-based measurement interpretation introduced severe information bottlenecks, and deeper entanglement structures amplified perturbation propagation throughout the quantum circuit. In multiple benchmark environments, noisy performance drifted between approximately 0.37 and 0.50, approaching near-random classification behavior.

The first major improvement occurred through the introduction of learned-readout decoding. By preserving richer measurement information and replacing parity collapse with adaptive decoder interpretation, the framework substantially improved noisy performance and increased stability during inference. This stage increased noisy performance to approximately 0.82 and established the first major robustness breakthrough within the framework evolution.

The next stage introduced multi-observable feature extraction. Rather than relying on a narrow collection of measurements, the framework incorporated expectation values and correlation statistics that preserved a broader description of quantum state evolution. This

richer representation increased noisy performance to approximately 0.85 while improving feature-space stability.

Architecture robustness optimization further improved performance by identifying circuit structures better suited to realistic NISQ environments. The experiments demonstrated that shallow and modular architectures frequently maintained greater stability than more densely entangled systems. This stage increased noisy performance to approximately 0.89 and reinforced the importance of architecture selection in practical quantum learning.

The largest improvement emerged through stability-aware optimization and dual-loss training. By incorporating robustness directly into the learning objective, the framework substantially improved convergence behavior and reduced degradation under perturbation. This stage increased noisy performance to approximately 0.96 while producing highly stable optimization trajectories.

The final robustness mechanism introduced noise-augmented readout training. By exposing the learning system directly to noisy feature distributions during optimization, the framework improved adaptation to realistic execution environments and increased noisy performance further to approximately 0.97.

Taken together, these stages demonstrate that practical HQNN robustness emerged through progressive engineering across multiple interacting components rather than through a single isolated innovation.

4.9.2 Threshold Evolution Across Mechanisms

One of the most important findings emerging from Chapter 4 is that robustness improvements can be traced quantitatively throughout the framework evolution. The threshold progression provides direct evidence that the observed gains accumulated systematically as additional robustness mechanisms were introduced.

The baseline parity-based systems established the initial performance threshold under realistic noisy conditions. In these architectures, noisy behavior frequently deteriorated to-

ward the 0.37 to 0.50 range due to optimization instability, information loss during decoding, and sensitivity to perturbation.

The introduction of learned-readout decoding produced the first major transition, increasing noisy performance from approximately 0.50 to approximately 0.82. This result demonstrated that decoder design itself plays a major role in practical HQNN robustness.

Multi-observable feature extraction increased performance further from 0.82 to 0.85 by preserving additional information from quantum state evolution. Architecture robustness optimization subsequently increased performance from 0.85 to 0.89 through improved entanglement organization and reduced perturbation propagation.

The introduction of stability-aware optimization produced one of the largest single improvements within the framework evolution, increasing performance from 0.89 to 0.96. This result highlighted the importance of optimization stability as a primary determinant of practical HQNN behavior.

Finally, noise-augmented readout training increased performance from 0.96 to 0.97 by improving adaptation to realistic execution environments and reducing sensitivity to distributional shifts between training and deployment conditions.

As shown in Figure 4.8, the complete framework evolution exhibits a clear and systematic progression across all major robustness mechanisms.




Figure 4.8: Beginning from the baseline parity-based HQNN, successive robustness mechanisms—including learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware optimization, and noise-augmented training—produce systematic improvements in noisy classification performance. The progression illustrates that the final robustness gains emerge through the cumulative interaction of multiple coordinated mechanisms rather than from any single modification.

An important observation emerging from this progression is that no single mechanism explains the final framework behavior. The improvements accumulated incrementally as multiple components interacted throughout the hybrid learning pipeline. The threshold evolution therefore provides quantitative evidence supporting the broader robustness-engineering philosophy developed throughout the chapter.

4.9.3 Final Robustness Interpretation

The complete framework evolution developed throughout Chapter 4 leads to a broader interpretation of practical hybrid quantum learning under realistic NISQ conditions. The experiments consistently demonstrated that instability in HQNN systems does not origi-

nate from one isolated source. Instead, degradation emerges through interaction between measurement variability, information loss, architecture design, optimization instability, and stochastic hardware perturbations.

Consequently, the strongest improvements were achieved not through isolated corrections, but through coordinated robustness engineering across the entire hybrid computational pipeline. Decoder adaptation, feature extraction, architecture selection, optimization stabilization, and noise-aware training each addressed different aspects of the overall problem.

This observation significantly reframes how practical hybrid quantum learning systems should be interpreted. Earlier perspectives often emphasized quantum expressivity and circuit complexity as the primary determinants of performance. The experiments presented throughout this thesis suggest a different conclusion. Under realistic NISQ conditions, robustness and optimization stability are equally important and frequently more influential than additional representational complexity.

Another important finding concerns the role of classical computation within hybrid systems. Several of the strongest robustness improvements emerged through adaptive decoder design, stability-aware optimization, and noise-aware learning procedures. These mechanisms demonstrate that classical components play a central role in enabling reliable hybrid quantum learning.

The framework therefore supports the hybrid quantum-classical co-design philosophy developed throughout the chapter. Practical HQNN systems achieve their strongest performance not by maximizing quantum complexity alone, but by balancing quantum feature generation with classical stabilization, adaptation, and inference.

The final interpretation emerging from Chapter 4 is therefore clear. Robustness, rather than expressivity alone, represents the dominant practical bottleneck limiting hybrid quantum learning under realistic NISQ conditions. The framework developed throughout this chapter demonstrates that substantial improvements can be achieved through systematic

robustness engineering across the entire learning pipeline.

4.10 Real IBM Hardware Validation

The previous sections demonstrated that the proposed robustness-engineering framework substantially improved Hybrid Quantum Neural Network performance under controlled simulator environments. Through learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware training, and noise-aware learning, the framework consistently achieved greater robustness than the baseline HQNN. An important question nevertheless remained: do these improvements transfer beyond simulation and remain effective on physical quantum hardware?

Previous HQNN studies have frequently concluded with simulator-based evaluation or limited hardware experiments, making it difficult to determine whether proposed robustness improvements generalize to practical NISQ execution. In contrast, this thesis validates the complete robustness-engineering framework directly on IBM Quantum hardware, allowing simulator performance to be compared with physical execution under realistic hardware constraints.

The objective of this evaluation is not simply to reproduce simulator accuracy, but to determine whether the robustness principles developed throughout this thesis remain effective when exposed to gate imperfections, readout variability, calibration drift, decoherence, and other sources of stochastic hardware behavior.

The experimental results demonstrate that hardware execution introduces measurable degradation relative to idealized simulation, as expected for current NISQ devices. Nevertheless, the robustness-engineered HQNN maintains substantially greater stability and predictive performance than the baseline implementation while preserving the same overall improvement trends observed throughout simulator-based evaluation.

More importantly, these experiments demonstrate that the proposed robustness mechanisms transfer meaningfully from simulation to physical IBM Quantum hardware. Although

practical quantum execution remains challenging, the framework consistently improves reliability under realistic hardware conditions, indicating that the proposed methodology addresses practical NISQ limitations rather than simulator-specific behavior.

These findings provide an important bridge between theoretical robustness analysis and practical quantum machine learning deployment. By demonstrating simulator-to-hardware transferability, this thesis establishes that robustness engineering represents a practical systems-level methodology for developing reliable Hybrid Quantum Neural Networks on current quantum hardware.

4.10.1 IBM Runtime Integration

To support practical hardware execution, the HQNN framework was integrated with the IBM Quantum runtime ecosystem using Qiskit-based execution pipelines. The IBM runtime environment provides access to cloud-based superconducting quantum hardware together with infrastructure for transpilation, scheduling, circuit execution, and measurement management.

The integration process required adapting the framework to realistic hardware constraints. Variational circuits developed during simulator-based experimentation had to be mapped onto hardware-specific qubit layouts while accounting for connectivity restrictions, finite shot budgets, and backend-specific calibration conditions.

An important challenge involved balancing representational capability against execution reliability. Earlier sections demonstrated that deeper and more heavily entangled architectures frequently become less stable under noisy conditions. This challenge became even more significant during hardware execution because transpilation often introduces additional gate operations beyond those present in the original circuit design.

The robustness-engineered framework benefited directly from the design principles developed throughout Chapter 4. Shallow architectures, adaptive decoding procedures, and stability-aware optimization all contributed to improved hardware compatibility by reducing

sensitivity to execution variability and limiting unnecessary circuit complexity.

Another important observation was that hardware conditions themselves vary over time. Calibration updates, changing backend availability, and fluctuating noise profiles can influence execution behavior even when the same circuit is executed repeatedly. Consequently, practical robustness requires resilience not only to noise but also to temporal variability in hardware conditions.

These observations reinforce the importance of evaluating hybrid quantum learning systems beyond simulator environments. Practical deployment depends on the ability of a framework to maintain stable behavior despite the additional complexities introduced by real hardware execution.

4.10.2 `ibm_fez` Hardware Execution

Following runtime integration, the robustness-engineered HQNN framework was executed on the IBM superconducting backend `ibm_fez`. The experiments evaluated classification behavior, inference stability, and robustness under real hardware conditions using benchmark tasks previously examined in simulation.

The hardware experiments demonstrated that the robustness-engineered framework maintained substantially greater stability than earlier baseline HQNN architectures. Although hardware noise reduced overall performance relative to idealized simulations, the framework preserved meaningful learning behavior and avoided the severe degradation patterns frequently observed in parity-based systems.

As shown in Figure 4.9, the hardware results remained broadly consistent with the trends observed throughout the simulator-based evaluation.




Figure 4.9: Although real hardware execution reduces performance relative to simulator-based evaluation, the robustness-engineered HQNN framework preserves meaningful classification behavior and remains substantially more stable than the baseline parity-based hardware implementation.

One particularly important observation was that learned-readout decoding and multi-observable feature extraction transferred effectively to hardware execution. The richer measurement representations generated by these mechanisms remained more stable than parity-based approaches even when realistic readout variability and calibration effects were present.

The experiments also demonstrated that stability-aware optimization and noise-aware training improved consistency across repeated hardware evaluations. Earlier baseline systems frequently produced highly variable outputs under repeated execution, whereas the robustness-engineered framework maintained more predictable behavior despite backend variability.

Although hardware execution remains significantly more challenging than simulator evaluation, the results provide strong evidence that the framework improvements are not limited to simulation alone.

4.10.3 Simulator-to-Hardware Transferability

One of the most important questions addressed by the hardware-validation experiments concerns simulator-to-hardware transferability. Many quantum machine learning systems perform adequately in simulation but experience substantial degradation when executed on physical devices. Understanding whether robustness improvements survive this transition is therefore essential for evaluating practical usefulness.

The experiments demonstrated that several of the robustness mechanisms developed throughout the thesis transferred successfully from simulator environments to real hardware conditions. Learned-readout decoding, architecture optimization, stability-aware training, and noise-aware learning all contributed to improved hardware behavior relative to earlier baseline systems.

As shown in Figure 4.10, the overall relationship between simulator performance and hardware performance remained qualitatively consistent despite differences in absolute performance values.




Figure 4.10: Although hardware execution reduces absolute performance relative to simulator evaluation, the overall robustness trend remains consistent: learned-readout decoding, architecture optimization, stability-aware training, and the full robustness-engineered framework each improve hardware behavior relative to the baseline parity-based HQNN.

An important observation was that transferability depended strongly on architecture design. Shallow modular architectures generalized more effectively to hardware environments than highly entangled circuits. This finding aligns closely with the broader conclusions developed throughout the thesis regarding the relationship between robustness and architectural simplicity under realistic NISQ constraints.

The experiments also revealed that simulator environments cannot capture every aspect of hardware behavior perfectly. Physical devices introduce additional perturbation sources, calibration effects, and backend variability that are difficult to model completely through simplified stochastic simulations. Consequently, hardware performance frequently differed quantitatively from simulator predictions even when overall trends remained consistent.

Despite these limitations, the results demonstrate that simulator-based robustness engineering can produce meaningful improvements in practical hardware behavior. This find-

ing represents one of the most important outcomes of the thesis because it connects the robustness-engineering framework directly to real-world quantum computing environments.

More broadly, the hardware-validation experiments reinforce the central conclusion developed throughout this work. Reliable hybrid quantum learning under realistic NISQ conditions emerges not through maximizing circuit complexity, but through coordinated interaction between architecture design, decoder adaptation, optimization stability, and robustness-aware learning procedures operating across both simulator and hardware environments.

Chapter 5

Conclusion and Future Work

5.1 Summary of Contributions

This thesis investigated the problem of robustness in Hybrid Quantum Neural Networks operating under realistic Noisy Intermediate-Scale Quantum (NISQ) conditions. While previous research has primarily focused on increasing quantum circuit expressivity and evaluating predictive performance under idealized execution, this work demonstrates that practical HQNN performance is limited primarily by robustness under realistic quantum noise.

The contributions of this thesis combine original technical mechanisms, systematic evaluation methodology, and the integration of established Hybrid Quantum Machine Learning techniques into a unified robustness-engineering framework. Rather than claiming that every individual component was developed for the first time, this work introduces several original contributions while incorporating meaningful concepts that build upon previous research.

The principal original contributions of this thesis are summarized as follows:

- **Systematic Robustness Characterization Framework:** A unified methodology for identifying robustness limitations through baseline HQNN analysis, optimization dynamics, classical versus quantum benchmarking, progressive noise degradation, and cross-noise robustness evaluation under standardized experimental conditions.
- **Unified Cross-Noise Benchmarking Methodology:** A benchmarking framework that evaluates Hybrid Quantum Neural Networks, Quantum Support Vector Machines, and classical machine learning models under identical datasets, software frameworks, evaluation metrics, and multiple realistic quantum noise channels, enabling direct comparison of robustness and degradation behavior.

- **Learned-Readout Decoding Mechanism:** An adaptive decoding strategy that replaces conventional parity-based measurement with a trainable decoder operating on the complete quantum measurement distribution, substantially reducing information loss during inference.
- **Robustness-Engineering Framework:** A staged framework integrating learned-readout decoding, multi-observable feature extraction, architecture optimization, stability-aware optimization, and noise-aware learning into a coordinated Hybrid Quantum Neural Network pipeline, where each mechanism addresses a specific robustness bottleneck identified during benchmarking.
- **Stability-Aware Loss and Training:** A stability-aware optimization methodology that explicitly incorporates optimization robustness into the learning objective, producing significantly more stable convergence under realistic quantum noise.
- **Comprehensive Framework Validation:** Systematic validation through component-level ablation studies, threshold progression analysis, cross-noise evaluation, IBM Quantum hardware experiments, and simulator-to-hardware transferability analysis, demonstrating that the proposed robustness improvements generalize consistently across simulation and physical quantum hardware.
- **Scientific Finding:** Experimental evidence demonstrating that robustness, rather than quantum expressivity alone, represents the dominant bottleneck limiting practical Hybrid Quantum Neural Networks under realistic NISQ conditions.

These original contributions are integrated with established concepts from the Hybrid Quantum Machine Learning literature, including variational quantum circuits, hybrid quantum-classical optimization, multi-observable quantum representations, architecture optimization strategies, and existing HQNN models. The resulting framework therefore represents both a set of original technical contributions and a broader systems-engineering methodology for

developing reliable Hybrid Quantum Neural Networks under realistic quantum hardware constraints.

The following subsections summarize these contributions in greater detail, beginning with the benchmarking methodology, followed by the robustness-engineering framework, experimental validation, and the practical design principles established throughout this research.

5.1.1 Benchmarking Framework Contributions

One of the primary contributions of this thesis is the development of a structured benchmarking framework for evaluating hybrid quantum neural network systems under realistic NISQ conditions. Earlier HQNN studies frequently relied on isolated experiments, architecture-specific evaluations, or limited simulator environments, making systematic comparison difficult across datasets, perturbation conditions, and implementation frameworks.

The benchmarking methodology developed throughout this work addressed these limitations by integrating standardized evaluation procedures, multichannel perturbation analysis, architecture-level comparison, and cross-framework validation within a unified experimental pipeline. This structure enabled consistent comparison between HQNN architectures, quantum kernel methods, and classical baselines under shared conditions.

An important contribution of the framework involves the evaluation of robustness as a systems-level property rather than focusing exclusively on predictive performance. The experiments examined optimization stability, degradation behavior, convergence consistency, reproducibility, and perturbation resilience in addition to conventional accuracy metrics.

The framework also incorporated multichannel perturbation analysis. Rather than evaluating performance under a single noise model, the experiments examined behavior across depolarizing noise, bit-flip perturbations, phase-flip perturbations, amplitude damping, and hardware-style execution variability. These results demonstrated that robustness behavior depends strongly on perturbation structure and cannot be characterized reliably through isolated evaluations alone.

Another important contribution involved cross-platform validation. The framework incorporated experiments spanning Qiskit, Cirq, PennyLane, simulator environments, and IBM hardware execution systems. These studies demonstrated that several robustness trends generalized consistently across software ecosystems and hardware platforms.

More broadly, the benchmarking methodology contributes a shift in how practical HQNN systems should be evaluated. Rather than emphasizing idealized performance alone, the framework promotes robustness, reproducibility, and stability as central evaluation objectives for NISQ-era quantum learning systems.

5.1.2 Robustness-Engineering Contributions

The central technical contribution of this thesis is the development of a staged robustness-engineering framework for hybrid quantum neural network systems operating under realistic NISQ constraints.

The framework evolved through a sequence of coordinated mechanisms designed to improve stability under noisy execution conditions. The first major contribution involved replacing parity-based measurement collapse with learned-readout decoding. This modification preserved richer measurement information and significantly reduced the information bottlenecks present in earlier architectures.

The framework next introduced multi-observable feature extraction procedures that incorporated expectation values and correlation statistics into the learning process. These richer representations improved feature-space stability and reduced sensitivity to isolated perturbation events during inference.

Additional contributions emerged through architecture robustness optimization. The experiments demonstrated that shallow and modular architectures frequently maintain greater stability under realistic perturbation conditions than more heavily entangled variational systems. These findings established an important relationship between architecture design and practical robustness.

The framework also introduced stability-aware optimization procedures, dual-loss learning objectives, noise-augmented decoder adaptation, curriculum-based perturbation scheduling, and variance-reduced SPSA optimization. Together, these mechanisms substantially improved convergence behavior and reduced instability during noisy training.

A particularly important contribution involved the ablation and mechanism-isolation studies. These experiments demonstrated that robustness gains emerged progressively through interaction between multiple coordinated mechanisms rather than through any single modification operating independently.

The framework ultimately improved noisy performance from approximately 0.37 to 0.50 in the earliest baseline systems to approximately 0.97 under several benchmark environments while simultaneously reducing degradation between clean and noisy execution conditions. These results establish robustness engineering as a practical strategy for improving hybrid quantum learning under realistic NISQ constraints.

5.1.3 Hybrid Quantum-Classical Co-Design Contributions

A broader conceptual contribution of this thesis is the development of a hybrid quantum-classical co-design philosophy for practical quantum machine learning systems.

Earlier perspectives frequently emphasized quantum circuit expressivity and entanglement complexity as primary determinants of performance. The experiments conducted throughout this work suggest a different interpretation. Practical HQNN behavior depends heavily on coordinated interaction between quantum and classical computational components rather than on quantum complexity alone.

Within the proposed co-design framework, quantum circuits function primarily as probabilistic feature generators. Through variational state evolution, they produce nonlinear and high-dimensional representations that can be incorporated into broader learning pipelines. Classical systems provide optimization stability, adaptive inference, preprocessing, decoder interpretation, and robustness-aware control.

The experiments repeatedly demonstrated that the strongest improvements emerged when quantum feature generation was combined with adaptive classical learning procedures. Decoder adaptation, stability-aware optimization, feature fusion, and noise-aware training all contributed substantially to practical performance.

This interpretation also contributes a more balanced perspective regarding quantum usefulness in the NISQ era. The experiments do not support broad claims of universal superiority over classical machine learning systems. Instead, the results suggest that quantum representations may provide conditional benefits under selected nonlinear and probabilistic learning environments when integrated within carefully engineered hybrid architectures.

Another important contribution involves the reinterpretation of variational quantum circuits themselves. Throughout the framework evolution, quantum circuits consistently performed most effectively when treated as probabilistic feature generators rather than standalone classifiers. This perspective improved stability while preserving useful representational flexibility under realistic perturbation conditions.

The hybrid quantum-classical co-design philosophy developed throughout this thesis therefore reframes practical quantum machine learning as a systems-level engineering problem operating at the intersection of quantum feature generation, optimization stability, adaptive inference, hardware-aware design, and robustness engineering. This perspective serves as the conceptual foundation for the conclusions presented in the remainder of this chapter.

5.2 Final Framework Interpretation

The experiments developed throughout this thesis support a broader reinterpretation of practical hybrid quantum learning under realistic noisy intermediate-scale quantum conditions. Earlier work in quantum machine learning frequently emphasized representational power, circuit complexity, and the possibility of achieving quantum advantages through increasingly expressive variational architectures. The results presented throughout this work suggest a different conclusion. Practical HQNN performance depends just as strongly on optimization

stability, perturbation resilience, decoder adaptation, and architecture robustness as it does on representational capability.

The robustness-engineering framework evolved progressively through a sequence of coordinated modifications beginning with unstable parity-based architectures and culminating in substantially more reliable hybrid learning systems capable of maintaining strong noisy performance under realistic execution conditions. Importantly, the strongest improvements did not emerge from dramatic increases in quantum complexity. Instead, they emerged through systematic improvements in how quantum information was represented, interpreted, optimized, and stabilized throughout the learning process.

The experiments also demonstrated that idealized simulator performance provides an incomplete picture of practical HQNN behavior. Several architectures that appeared highly expressive under noiseless conditions became unstable when realistic perturbations were introduced. Conversely, shallower and more modular architectures frequently maintained stronger optimization stability and more reliable inference despite possessing lower theoretical representational complexity.

These observations motivate a broader conceptual shift. Practical hybrid quantum learning should not be viewed primarily as a problem of maximizing quantum depth or entanglement. Instead, it should be understood as a robustness-engineering challenge operating at the intersection of noisy optimization, probabilistic feature generation, adaptive inference, and hardware-aware design.

A second important conclusion concerns the role of hybrid computation itself. Throughout the framework evolution, the strongest improvements emerged through cooperation between quantum and classical components rather than through isolated quantum processing. Quantum circuits generated useful feature representations, while classical systems provided optimization stability, adaptive decoding, and robustness-aware learning behavior.

The framework therefore contributes not only a collection of experimental results, but also a broader interpretation of how practical hybrid quantum learning should be designed

and evaluated under realistic NISQ constraints.

The following subsections examine robustness as the dominant practical bottleneck, summarize the threshold progression observed throughout the framework evolution, and discuss broader lessons for future quantum-AI systems.

5.2.1 Robustness as the Central NISQ Bottleneck

One of the clearest conclusions emerging from this thesis is that robustness represents the dominant practical bottleneck limiting hybrid quantum learning under current NISQ conditions. Throughout the benchmarking experiments, HQNN systems repeatedly exhibited sensitivity to noisy optimization, measurement variability, perturbation accumulation, and unstable convergence behavior. These limitations frequently prevented variational architectures from fully utilizing their theoretical representational capabilities.

An important observation is that the primary limitation was not simply insufficient expressivity. In many benchmark environments, increasing circuit depth and entanglement complexity actually reduced practical performance because instability increased more rapidly than representational flexibility improved. This finding significantly changes how practical hybrid quantum learning should be interpreted.

The experiments demonstrated that robustness failures emerge through interaction between multiple sources of instability. Parity-based decoding discarded useful information, noisy optimization distorted parameter updates, perturbations propagated through entangled circuits, and stochastic measurement fluctuations altered feature representations. No single mechanism explained the observed degradation independently.

These observations suggest that practical progress in hybrid quantum learning cannot depend solely on larger circuits, additional qubits, or increased expressivity. Improvements in robustness, optimization reliability, adaptive inference, and architecture stability are equally important for practical deployment under realistic hardware conditions.

The findings also align closely with broader NISQ literature emphasizing the limitations

imposed by decoherence, gate errors, and hardware variability. However, the results extend these observations by demonstrating how such limitations interact directly with optimization and learning dynamics in HQNN systems.

More broadly, the experiments suggest that future advances in practical hybrid quantum learning will depend heavily on robustness-aware design methodologies capable of managing instability throughout the entire computational pipeline rather than focusing exclusively on representational capability.

5.2.2 Threshold Evolution Across the Framework

A second major conclusion of this thesis involves the threshold progression observed throughout the robustness-engineering framework. Rather than presenting isolated improvements independently, the experiments revealed a systematic sequence of transitions in which each mechanism contributed incrementally to improved noisy performance.

The baseline parity-based architectures established the initial threshold under realistic perturbation conditions. In these systems, noisy performance frequently operated between approximately 0.37 and 0.50 due to unstable optimization, information loss during decoding, and sensitivity to stochastic perturbations.

The first major transition occurred after replacing parity collapse with learned-readout decoding, increasing performance from approximately 0.50 to approximately 0.82. Multi-observable feature extraction then increased performance from 0.82 to 0.85 through richer probabilistic feature representations and improved measurement utilization.

Architecture robustness optimization further increased performance from 0.85 to 0.89 by identifying circuit structures better suited for realistic NISQ environments. Stability-aware optimization produced one of the largest improvements, increasing performance from 0.89 to 0.96 while substantially reducing degradation under perturbation.

Finally, noise-augmented training increased performance from 0.96 to 0.97 by aligning the learning process more closely with realistic execution conditions. Each transition addressed

a different source of instability while contributing to broader improvements throughout the hybrid learning pipeline.

An important observation emerging from this progression is that no single mechanism explains the final framework behavior. The improvements accumulated systematically as multiple robustness mechanisms interacted throughout the learning system. This provides direct support for the robustness-engineering philosophy developed throughout the thesis.

The threshold evolution therefore represents more than a sequence of numerical improvements. It provides evidence that practical HQNN robustness can be engineered progressively through coordinated modifications that improve stability at multiple levels of the learning process.

5.2.3 Lessons for Practical Quantum AI

The framework evolution developed throughout this thesis suggests several broader lessons regarding the future of practical quantum-AI systems under realistic NISQ constraints.

The first lesson is that hybrid quantum learning should not be interpreted as a direct replacement for classical machine learning. Across many benchmark environments and application studies, classical systems remained highly competitive and frequently outperformed unstable baseline HQNN architectures. Quantum usefulness instead appeared conditional and dependent on problem structure, optimization stability, and robustness under perturbation.

A second lesson concerns the importance of hybrid quantum-classical cooperation. The strongest improvements consistently emerged when quantum and classical systems operated cooperatively rather than competitively. Quantum circuits contributed probabilistic feature representations, while classical systems provided optimization, adaptive inference, and stability management. This cooperative interaction became one of the defining characteristics of the framework.

The experiments also reinforced the value of shallow and modular architectures. Through-

out the benchmarking and optimization studies, increasing complexity frequently reduced practical reliability under noisy conditions. More localized architectures repeatedly maintained stronger stability and more reproducible behavior than heavily entangled alternatives.

Another important lesson involves the necessity of hardware-aware evaluation. Several robustness mechanisms generalized successfully from simulation to IBM hardware execution, but the experiments also revealed additional variability introduced through transpilation, calibration drift, and backend-specific perturbation effects. These findings highlight the importance of evaluating quantum learning systems under realistic hardware conditions whenever possible.

Collectively, the results suggest that the future of practical quantum-AI systems will depend less on demonstrations of quantum complexity and more on advances in robustness-aware optimization, adaptive inference, hardware-aware design, and hybrid quantum-classical co-design. In this sense, the thesis contributes not only a robustness-engineering framework, but also a broader perspective regarding how practical quantum learning should evolve under realistic NISQ constraints.

The conclusions developed throughout this section establish the foundation for the remaining discussion of limitations, future research directions, and the final conclusions of the thesis.

5.3 Limitations of the Present Work

Although the robustness-engineering framework developed throughout this thesis demonstrated substantial improvements in HQNN stability and noisy performance, several important limitations remain. These limitations arise from both the physical constraints of current NISQ hardware and the methodological boundaries of the experimental framework itself. Consequently, the results should be interpreted as evidence regarding practical hybrid quantum learning under present technological conditions rather than as definitive conclusions regarding future large-scale quantum learning systems.

A major limitation concerns the capabilities of existing quantum hardware. Current superconducting, trapped-ion, and neutral-atom systems continue to operate under significant constraints involving decoherence, gate errors, readout variability, and calibration instability. These limitations restrict the depth and scale of variational circuits that can be executed reliably and therefore limit the practical complexity of HQNN architectures.

Another limitation involves scalability. The benchmark environments examined throughout this thesis were intentionally selected to remain compatible with realistic NISQ-era execution constraints. Although the framework demonstrated meaningful robustness improvements, the experiments do not establish that current HQNN systems can scale directly to the size and complexity of modern industrial machine learning systems.

The work is also constrained by the relationship between simulator-based evaluation and physical hardware execution. While several robustness mechanisms generalized successfully to real IBM hardware environments, simulations cannot fully reproduce all sources of variability present in practical quantum systems.

The following subsections discuss these limitations in greater detail and identify several challenges that remain open for future research.

5.3.1 NISQ Hardware Constraints

One of the most significant limitations of the present work arises from the constraints imposed by current NISQ hardware. Existing quantum processors continue to operate within environments characterized by finite coherence times, noisy gate operations, imperfect measurements, and backend-specific variability. These factors directly influence the stability and scalability of variational quantum learning systems.

Throughout the experiments, perturbation accumulation frequently increased as circuit complexity grew. More expressive architectures often became increasingly difficult to optimize because additional operations introduced greater opportunities for error propagation. Although the robustness-engineering framework substantially improved stability relative to

earlier baseline systems, practical circuit depth remained strongly constrained by hardware limitations.

Another important challenge concerns calibration variability. During hardware validation experiments, repeated executions occasionally produced different performance characteristics because calibration conditions changed over time. Variations in gate fidelity, measurement quality, and qubit coherence introduced additional uncertainty beyond the perturbation environments modeled during simulation.

The experiments were also limited by the number of qubits that could be used reliably in practical hardware settings. Although modern devices contain growing qubit counts, effective utilization remains constrained by connectivity restrictions, calibration quality, and backend-specific noise characteristics.

A further limitation is the absence of fault-tolerant quantum error correction. The framework was developed specifically for realistic pre-fault-tolerant NISQ environments. Consequently, the results describe robustness under current hardware conditions rather than behavior in future fault-tolerant quantum systems.

These observations reinforce one of the central conclusions of the thesis: practical hybrid quantum learning remains fundamentally influenced by noisy hardware behavior, making robustness-aware design essential under present technological conditions.

5.3.2 Dataset and Scalability Limitations

Another important limitation concerns dataset scale and experimental scalability. The benchmark environments used throughout the thesis were intentionally selected to remain compatible with realistic circuit-depth limitations and practical execution constraints. As a result, the evaluated datasets remain considerably smaller than those commonly used in large-scale industrial machine learning applications.

The experiments demonstrated meaningful robustness improvements and nonlinear learning behavior under controlled benchmark conditions. However, these results should not be

interpreted as evidence that current HQNN systems can immediately scale to machine learning workloads involving millions of parameters or extremely large datasets.

Feature-space dimensionality also remains a practical constraint. Although quantum systems operate within exponentially large state spaces in theory, practical variational circuits remain limited by shot noise, hardware fidelity, optimization stability, and circuit depth restrictions. Consequently, the framework focused primarily on low-dimensional and medium-dimensional feature representations compatible with stable execution.

The application-oriented experiments were likewise simplified relative to production-level systems. Medical risk classification, cybersecurity anomaly detection, and energy optimization tasks were designed to evaluate robustness behavior under structured environments rather than to compete directly with highly optimized industrial machine learning pipelines.

Another limitation is the scope of learning paradigms examined. The present work focused primarily on supervised hybrid learning systems. Reinforcement learning, generative modeling, unsupervised representation learning, and transformer-based quantum architectures were beyond the scope of the framework.

These limitations indicate that substantial additional research remains necessary before large-scale quantum-AI systems become practical under realistic hardware conditions.

5.3.3 Simulation versus Hardware Limitations

A further limitation concerns the relationship between simulator-based experimentation and physical hardware execution. Much of the robustness-engineering framework was developed initially through controlled simulation environments that allowed systematic benchmarking across architectures, perturbation models, and optimization procedures. While these environments provide valuable insight, they cannot fully reproduce every characteristic of real quantum hardware.

The IBM hardware validation experiments demonstrated that several robustness mechanisms generalized successfully from simulation to physical execution. Learned-readout de-

coding, stability-aware optimization, and modular architecture design all contributed to improved hardware behavior relative to earlier baseline systems. However, hardware execution also introduced additional variability that was not captured completely through simplified perturbation models.

One source of discrepancy involved transpilation overhead. Circuits optimized within abstract simulation environments frequently required additional gate decompositions and routing operations when executed on hardware. These modifications increased perturbation exposure and occasionally altered practical behavior relative to the original simulator configuration.

Temporal variability also represents a significant challenge. Real hardware systems experience changing calibration conditions, fluctuating coherence characteristics, and evolving noise profiles. As a result, repeated executions under nominally identical settings may encounter different operating conditions over time. Simulator environments generally assume more stable perturbation behavior and therefore cannot reproduce these effects perfectly.

The perturbation models used throughout the framework also remain simplified relative to the complexity of real hardware noise. While common models such as depolarizing noise, phase damping, amplitude damping, and bit-flip perturbations capture important error mechanisms, practical devices may exhibit correlated errors, crosstalk effects, leakage behavior, and backend-specific artifacts that are difficult to model accurately.

Despite these limitations, the experiments demonstrated that simulator-based robustness engineering can still provide meaningful benefits for hardware execution. The qualitative trends observed throughout the framework remained broadly consistent across both environments even when absolute performance values differed. These findings support the continued use of simulation as a valuable development tool while emphasizing the importance of direct hardware validation whenever practical.

5.4 Future Research Directions

The robustness-engineering framework developed throughout this thesis demonstrates that practical hybrid quantum learning can maintain substantially improved stability under realistic NISQ conditions when architecture design, optimization procedures, decoder adaptation, and perturbation-aware training are engineered cooperatively. At the same time, the experiments revealed several challenges that remain unresolved and therefore motivate important directions for future research.

Some of these directions depend on continued advances in quantum hardware. Others involve algorithmic and architectural questions that remain relevant regardless of future hardware improvements. Collectively, they suggest that progress in practical hybrid quantum learning will likely emerge through continued interaction between robustness engineering, hardware-aware design, adaptive optimization, and hybrid quantum-classical computation.

The results developed throughout this thesis indicate that future quantum machine learning research should move beyond isolated demonstrations of expressive capability and focus increasingly on reliability, reproducibility, scalability, and robustness under realistic operating conditions.

An important implication of the framework is that practical HQNN systems should be viewed as evolving computational ecosystems rather than static learning architectures. Future systems will likely require tighter integration between quantum and classical components, more adaptive optimization procedures, and greater awareness of hardware-specific behavior.

The following subsections discuss several of the most promising research directions emerging from the results of this work.

5.4.1 Fault-Tolerant Hybrid Quantum AI

One of the most important future directions involves the integration of robustness-engineering principles with fault-tolerant quantum computing architectures. The present framework was

developed specifically for realistic NISQ conditions in which full-scale quantum error correction remains impractical. Future quantum processors may eventually support logical qubits protected through fault-tolerant encoding schemes capable of maintaining coherent quantum computation over substantially longer time scales.

The results of this thesis suggest that fault tolerance alone will not eliminate the need for robustness-aware learning procedures. Even if physical error rates are reduced dramatically, practical hybrid quantum learning systems will still depend on stable optimization, adaptive inference, effective feature representations, and scalable learning pipelines.

An especially important research question concerns how the relationship between expressivity and robustness changes when variational circuits can operate reliably at depths far beyond current NISQ limitations. Future fault-tolerant systems may enable more expressive architectures while simultaneously introducing new optimization and scalability challenges.

Future investigations may also examine the interaction between robustness-aware learning and logical-level error-corrected quantum computation. Such studies could determine whether the robustness principles identified throughout this thesis remain relevant when hardware-level errors are reduced substantially.

Another promising direction involves combining fault-tolerant quantum processing with adaptive classical learning systems. The hybrid co-design principles developed throughout this work suggest that even highly reliable quantum hardware may continue to benefit from coordinated classical stabilization and inference mechanisms.

These directions indicate that fault-tolerant quantum AI should be viewed not as a replacement for robustness engineering, but as a future environment in which robustness-aware methodologies can continue to evolve.

5.4.2 Robustness-Aware Quantum Optimization

A second major research direction involves the continued development of robustness-aware optimization methods for variational quantum learning systems. Throughout the thesis,

optimization instability repeatedly emerged as one of the most significant factors limiting practical HQNN performance.

Although the framework introduced several stabilization mechanisms, including dual-loss optimization, curriculum-based perturbation scheduling, and variance-reduced SPSA procedures, many optimization challenges remain unresolved. Variational quantum systems continue to exhibit noisy gradient estimation, unstable parameter evolution, and sensitivity to increasingly expressive circuit structures.

Future research may therefore focus on adaptive optimization procedures designed specifically for noisy quantum environments. Examples include perturbation-aware optimizers, uncertainty-driven parameter updates, dynamic regularization methods, and optimization procedures capable of responding directly to hardware conditions during training.

Another promising direction involves integrating machine learning techniques into the optimization process itself. Meta-learning, adaptive scheduling, reinforcement-based optimization control, and AI-assisted optimizer selection may provide new approaches for improving convergence reliability under realistic execution conditions.

The experiments also suggest that future optimization systems may increasingly incorporate robustness as a primary objective rather than treating stability as a secondary concern. Such approaches would align closely with the central conclusions developed throughout this thesis.

These directions reinforce a recurring theme throughout the framework evolution: optimization stability is not merely an implementation detail, but a defining component of practical hybrid quantum learning.

5.4.3 Quantum-Classical Co-Designed AI Systems

The hybrid quantum-classical co-design philosophy developed throughout this thesis suggests several important future directions involving increasingly integrated learning systems. Earlier perspectives often treated quantum computation as an independent computational

engine with classical processing serving only a supporting role. The results of this work suggest a different interpretation.

Future hybrid AI systems may evolve toward deeply integrated architectures in which quantum and classical computation operate as coordinated components within a unified learning framework. Quantum processors may specialize in probabilistic feature generation, combinatorial exploration, and nonlinear representation learning, while classical systems provide optimization, inference, uncertainty estimation, and robustness management.

Another important direction involves dynamic allocation of computational responsibilities. Rather than assigning fixed roles to quantum and classical subsystems, future architectures may adapt continuously according to hardware conditions, optimization behavior, and task-specific requirements.

The experiments also suggest that future systems may increasingly emphasize quantum feature generation rather than direct quantum classification. Throughout the framework evolution, variational circuits consistently performed most effectively when acting as probabilistic feature generators embedded within larger classical learning pipelines.

This perspective may guide future hybrid AI research toward architectures that emphasize cooperation rather than competition between computational paradigms. Such systems could combine the strengths of both domains while reducing the limitations associated with each individually.

More broadly, the co-design philosophy developed throughout this thesis suggests that practical quantum AI will likely emerge through increasingly sophisticated hybrid computational ecosystems rather than through isolated quantum systems operating independently.

5.4.4 Large-Scale Hardware Deployment

A final major research direction concerns the deployment of robustness-engineered HQNN systems on increasingly capable quantum hardware platforms. Although current NISQ devices remain constrained by noise, limited coherence times, and execution variability, rapid

progress continues across superconducting, trapped-ion, neutral-atom, and emerging quantum technologies.

Future hardware improvements may support larger qubit counts, higher gate fidelities, enhanced connectivity, and reduced calibration variability. These advances could enable more expressive variational architectures while reducing several of the constraints encountered throughout the present work.

An important future challenge involves scaling robustness-aware HQNN systems beyond the relatively small benchmark environments explored in this thesis. Future studies may evaluate hybrid quantum learning on larger scientific datasets, more complex optimization tasks, high-dimensional classification problems, and increasingly realistic application domains.

Another promising direction involves cloud-based quantum infrastructures. As quantum cloud ecosystems continue to mature, future HQNN systems may operate across heterogeneous environments that combine multiple quantum processors with large-scale classical AI infrastructure. Such systems will require new approaches to robustness, scalability, and resource coordination.

The experiments developed throughout this thesis also suggest that hardware-aware adaptation will become increasingly important as systems scale. Future architectures may dynamically adjust feature maps, optimization strategies, circuit organization, and inference procedures according to backend-specific characteristics and changing hardware conditions.

Ultimately, large-scale deployment will require continued integration between robustness engineering and hardware-aware design. Even as hardware improves, practical quantum learning systems will still need to manage optimization reliability, execution variability, scalability constraints, and adaptive learning behavior. The framework developed throughout this thesis therefore provides not only a methodology for current NISQ systems, but also a foundation for future generations of scalable hybrid quantum-AI architectures.

5.4.5 Final Conclusions

This thesis investigated the problem of robustness in hybrid quantum neural network systems operating under realistic noisy intermediate-scale quantum conditions. While previous research demonstrated that variational quantum learning architectures possess meaningful theoretical representational capability, the experiments conducted throughout this work revealed that practical hybrid quantum learning is constrained fundamentally by optimization instability, perturbation sensitivity, noisy inference behavior, and hardware variability under realistic execution environments.

To address these challenges, the thesis developed a robustness-engineering framework designed specifically for practical NISQ-era hybrid quantum learning. Rather than treating noise as a secondary implementation concern, the framework incorporated robustness directly into architecture design, feature representation, decoder adaptation, optimization behavior, perturbation-aware training, and evaluation methodology. The resulting system evolved progressively from unstable parity-based HQNN architectures toward substantially more reliable hybrid learning systems capable of maintaining strong performance under noisy conditions.

One of the most important findings of this work is that practical HQNN robustness emerges through coordinated interaction between multiple mechanisms rather than through isolated architectural modifications. Learned-readout decoding reduced information bottlenecks associated with parity-based inference. Multi-observable feature extraction improved representation quality and measurement utilization. Architecture optimization identified circuit structures better suited to realistic perturbation environments. Stability-aware optimization and noise-augmented training improved convergence reliability and inference consistency throughout the learning process.

The experiments also demonstrated that practical hybrid quantum learning should not be interpreted primarily as a problem of maximizing quantum complexity. In many benchmark environments, deeper and more heavily entangled architectures degraded more rapidly un-

der perturbation than shallower and more modular alternatives. These observations revealed a recurring tradeoff between expressivity and robustness under realistic NISQ constraints. Consequently, the strongest practical behavior emerged not from increasing quantum complexity independently, but from balancing quantum feature generation, adaptive classical inference, modular architecture design, and robustness-aware optimization within unified hybrid learning pipelines.

A broader contribution of this thesis is the development of a hybrid quantum-classical co-design perspective for practical quantum machine learning. The experiments repeatedly demonstrated that quantum and classical systems operate most effectively when functioning cooperatively rather than competitively. Within this interpretation, variational quantum circuits act primarily as probabilistic feature generators, while classical systems provide optimization stability, adaptive decoder interpretation, and scalable inference support. The resulting framework reframes practical hybrid quantum learning as a systems-level engineering problem rather than solely as a question of circuit complexity.

The threshold progression observed throughout the framework evolution further reinforces this interpretation. Baseline architectures frequently operated between 0.37 and 0.50 under realistic perturbation conditions. Through successive robustness-engineering stages, the framework progressed from 0.50 to 0.82 through learned-readout decoding, from 0.82 to 0.85 through multi-observable feature extraction, from 0.85 to 0.89 through architecture optimization, from 0.89 to 0.96 through stability-aware optimization, and from 0.96 to 0.97 through noise-augmented training. These improvements emerged not through dramatic increases in hardware capability or circuit depth, but through systematic improvements in how information was represented, interpreted, optimized, and stabilized throughout the hybrid learning process.

At the same time, the experiments revealed important limitations associated with current NISQ-era systems. Present-day hardware remains constrained by finite coherence times, gate variability, calibration instability, and limited fault-tolerant capability. Clas-

sical machine learning systems also remain highly competitive across many benchmark environments, suggesting that the usefulness of quantum representations remains conditional rather than universal. These limitations reinforce the importance of continued research into robustness-aware optimization, hardware-aware architecture design, and scalable hybrid learning methodologies.

Ultimately, the central conclusion of this thesis is clear. Practical hybrid quantum learning in the NISQ era is fundamentally a robustness problem. The future of quantum-AI systems will likely depend less on isolated demonstrations of quantum complexity and more on the ability to engineer stable, adaptive, and reliable learning systems capable of operating effectively under realistic noisy conditions. In this sense, the present work contributes not only a robustness-engineering framework for HQNN systems, but also a broader perspective on how practical hybrid quantum learning may evolve as quantum hardware and quantum-AI technologies continue to mature.

Bibliography

mkbothBIBLIOGRAPHYBIBLIOGRAPHYopenbib@code enumivempty

clubpenalty 4000m R. P. Feynman, “Simulating physics with computers,” *International Journal of Theoretical Physics*, vol. 21, no. 6–7, pp. 467–488, 1982.

biblabelarabic@ enumiv

biblabelarabic@ enumiv A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*. Providence, RI: American Mathematical Society, 2002.

biblabelarabic@ enumiv M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, UK: Cambridge University Press, 2010.

biblabelarabic@ enumiv P. W. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proc. 35th Annual Symposium on Foundations of Computer Science (FOCS)*, 1994, pp. 124–134.

biblabelarabic@ enumiv L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proc. 28th Annual ACM Symposium on Theory of Computing (STOC)*, 1996, pp. 212–219.

biblabelarabic@ enumiv J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.

biblabelarabic@ enumiv A. Peruzzo et al., “A variational eigenvalue solver on a quantum processor,” *Nature Communications*, vol. 5, 2014.

biblabelarabic@ enumiv E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” arXiv:1411.4028, 2014.

biblabelarabic@ enumiv J. R. McClean et al., “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics*, vol. 18, no. 2, 2016.

biblabelarabic@ enumiv M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*. Cham, Switzerland: Springer, 2018.

biblabelarabic@ enumiv V. Havlíček et al., “Supervised learning with quantum-enhanced feature spaces,” *Nature*, vol. 567, pp. 209–212, 2019.

biblabelarabic@ enumiv K. Mitarai et al., “Quantum circuit learning,” *Physical Review A*, vol. 98, no. 3, 2018.

biblabelarabic@ enumiv M. Benedetti et al., “Parameterized quantum circuits as machine learning models,” *Quantum Science and Technology*, vol. 4, 2019.

- biblabelarabic@enumiv K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Physical Review Letters*, vol. 119, 2017.
- biblabelarabic@enumiv S. Endo, S. C. Benjamin, and Y. Li, “Practical quantum error mitigation for near-future applications,” *Physical Review X*, vol. 8, 2018.
- biblabelarabic@enumiv Y. Li and S. C. Benjamin, “Efficient variational quantum simulator incorporating active error minimization,” *Physical Review X*, vol. 7, 2017.
- biblabelarabic@enumiv F. Arute et al., “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, pp. 505–510, 2019.
- biblabelarabic@enumiv H. J. Kimble, “The quantum internet,” *Nature*, vol. 453, pp. 1023–1030, 2008.
- biblabelarabic@enumiv Nature Editorial, “Quantum computing reaches for real-world impact,” *Nature*, vol. 574, pp. 505–506, 2019.
- biblabelarabic@enumiv IBM Quantum, “Qiskit: An open-source framework for quantum computing,” 2023. [Online]. Available: <https://qiskit.org>
- biblabelarabic@enumiv Cirq Developers, “Cirq: A Python framework for quantum circuits,” 2023. [Online]. Available: <https://quantumai.google/cirq>
- biblabelarabic@enumiv Xanadu Inc., “PennyLane: A cross-platform library for quantum machine learning,” 2023. [Online]. Available: <https://pennylane.ai>
- biblabelarabic@enumiv Microsoft Azure, “Azure Quantum platform,” 2023. [Online]. Available: <https://azure.microsoft.com/en-us/products/quantum/>
- biblabelarabic@enumiv A. Kashif et al., “Impact of quantum noise on hybrid quantum neural network training,” arXiv preprint, arXiv:2304.05420, 2023.

Appendix A

Additional Experiments and Applications

A.1 Nonlinear Benchmark Evaluation

The robustness-engineering framework developed throughout Chapter 4 demonstrated substantial improvements in HQNN stability under realistic NISQ conditions. Through coordinated improvements in decoder design, feature extraction, architecture selection, optimization stability, and noise-aware training, the framework transformed highly unstable baseline systems into substantially more reliable hybrid learning models. An important question nevertheless remained. Do these improvements remain useful when the learning task itself becomes more challenging?

To investigate this issue, the framework was evaluated on several nonlinear benchmark datasets commonly used in machine learning research, including the moons dataset, circles dataset, and XOR classification problem. These datasets were selected because they contain decision structures that cannot be separated effectively through simple linear boundaries and therefore provide useful environments for evaluating nonlinear representation learning.

The objective of these experiments was not to demonstrate universal quantum advantage over classical machine learning systems. Instead, the goal was to determine whether robustness-engineered HQNN systems could maintain stable behavior while learning more complex decision structures than those examined in earlier benchmark studies.

The experiments revealed several important observations. First, the framework successfully learned nonlinear decision boundaries across all evaluated benchmarks. Second, the robustness improvements introduced throughout Chapter 4 transferred effectively to these tasks and produced substantially more stable behavior than earlier baseline HQNN architec-

tures. Third, the results suggested that hybrid quantum representations may provide useful complementary structure under selected nonlinear learning conditions.

Taken together, these experiments indicate that the benefits of robustness engineering extend beyond the original benchmark environments. Although classical models remain highly competitive, the framework demonstrates that stable hybrid quantum learning can remain effective across a broader class of nonlinear problems.

The following subsections examine the moons, circles, and XOR benchmarks in greater detail.

A.1.1 Moons Dataset Benchmark

The moons dataset provides a useful benchmark for evaluating nonlinear classification because it consists of two interleaving crescent-shaped clusters that cannot be separated effectively using a simple linear decision boundary. Successful classification therefore requires the learning system to construct a curved representation of the underlying data geometry.

In the experiments conducted throughout this thesis, the robustness-engineered HQNN framework was trained on the moons dataset under both idealized and noisy execution conditions. The objective was to evaluate whether the hybrid quantum feature representations developed in Chapter 4 could support stable nonlinear decision-boundary formation.

The results demonstrated that the framework successfully learned the curved structure of the dataset while maintaining stable optimization behavior. Earlier parity-based architectures frequently produced fragmented or unstable decision regions under noise, whereas the robustness-engineered framework generated smoother and more consistent classification boundaries.

As shown in Figure A.1, the resulting decision structures closely followed the underlying geometry of the moons dataset and remained considerably more stable under perturbation than those produced by earlier benchmark architectures.




Figure A.1: The two interleaving crescent-shaped classes require a curved decision boundary, making the task useful for evaluating whether the robustness-engineered HQNN framework can support nonlinear feature representation and stable decision-boundary formation under noisy execution conditions.

An important observation was that the improvements introduced throughout Chapter 4 contributed directly to nonlinear learning behavior. Richer feature representations, adaptive decoding, and stability-aware optimization all reduced the instability that previously distorted decision-boundary formation. Although classical baselines remained highly competitive, the results indicate that hybrid quantum representations can support effective nonlinear learning when sufficient robustness is present.

A.1.2 Circles Dataset Benchmark

The circles dataset extends the nonlinear evaluation by introducing concentric geometric structures that require the learning system to separate nested regions rather than simply constructing curved boundaries. This benchmark therefore provides a different test of nonlinear representation capability and allows the framework to be evaluated on a distinct class of geometric relationships.

The experiments demonstrated that the robustness-engineered HQNN framework successfully learned radial decision structures under both idealized and noisy conditions. Compared with earlier baseline architectures, the framework produced substantially more stable boundaries and maintained greater consistency throughout optimization.

One important observation was that perturbations affected radial decision structures differently than the crescent-shaped patterns observed in the moons benchmark. Baseline systems frequently distorted circular boundaries under noise, producing irregular classification regions and unstable convergence behavior. The robustness-oriented mechanisms introduced throughout Chapter 4 reduced these effects considerably.

As shown in Figure A.2, the framework generated smoother nonlinear boundaries and maintained stronger geometric consistency throughout the classification process.

Circles Dataset Nonlinear Classification Benchmark




Figure A.2: The concentric class structure requires a radial nonlinear decision boundary, making the task useful for evaluating whether the robustness-engineered HQNN framework can preserve geometric consistency under noisy execution conditions.

The results also reinforced earlier observations regarding architecture design. Modular

HQNN structures frequently maintained more stable nonlinear behavior than highly entangled architectures, further supporting the relationship between robustness and circuit organization established in Chapter 4.

A.1.3 XOR Conditional Hybrid Advantage

The XOR benchmark represents one of the most widely studied nonlinear classification problems because its structure cannot be solved through a simple linear decision boundary. The problem therefore provides a useful environment for examining whether hybrid quantum representations can capture higher-order interactions within the data.

The robustness-engineered HQNN framework successfully learned the XOR decision structure under both idealized and noisy conditions. Earlier baseline architectures frequently exhibited unstable nonlinear behavior when noise was introduced, whereas the robustness-oriented framework maintained substantially more reliable inference performance throughout training and evaluation.

An important observation emerging from the XOR experiments was that the usefulness of hybrid quantum learning appears conditional rather than universal. In certain nonlinear environments, quantum feature representations provided useful structure that supported classification performance. However, these benefits depended strongly on architecture design, optimization stability, decoder adaptation, and robustness-aware training.

The results therefore do not support claims of universal superiority over classical machine learning. Instead, they suggest that hybrid quantum representations may provide complementary advantages within selected nonlinear learning environments when the framework is sufficiently robust to operate reliably under realistic NISQ conditions.

This interpretation aligns closely with the broader conclusions of the thesis. Practical hybrid quantum learning depends less on raw quantum complexity and more on the ability of the overall system to maintain stable behavior under noisy execution conditions. The nonlinear benchmark experiments therefore reinforce the central argument developed

throughout the thesis: robustness remains the dominant practical requirement for effective hybrid quantum learning in the NISQ era.

A.2 Quantum-Classical Feature Fusion

One of the central themes emerging throughout this thesis is that practical hybrid quantum learning systems derive much of their effectiveness from coordinated interaction between quantum and classical computation rather than from quantum processing alone. The robustness-engineering framework developed in Chapter 4 repeatedly demonstrated that quantum circuits function most effectively as probabilistic feature generators embedded within larger hybrid learning pipelines. Classical systems provide optimization stability, adaptive inference, and scalable learning behavior, while quantum circuits contribute non-linear feature representations that may capture useful structure under selected conditions.

These observations naturally motivate the exploration of quantum-classical feature fusion. Rather than evaluating quantum representations independently, the feature-fusion framework combines quantum-generated features with conventional classical representations within a unified learning system. The objective is to determine whether the two forms of representation provide complementary information that improves overall learning behavior.

This perspective differs from approaches that frame quantum and classical computation as competing alternatives. Within the hybrid framework developed throughout this thesis, the objective is not to replace classical machine learning with quantum models. Instead, the goal is to identify situations in which quantum and classical representations can cooperate effectively within the same computational pipeline.

The experiments presented throughout this section investigate how feature fusion influences representation quality, optimization behavior, and robustness under realistic NISQ conditions. The results suggest that quantum representations often provide their greatest practical value when integrated with strong classical feature spaces rather than evaluated in isolation.

More broadly, the feature-fusion experiments reinforce the hybrid quantum-classical co-design philosophy that has guided the thesis from the beginning. Practical HQNN systems derive their strongest performance from cooperation between computational paradigms rather than from reliance on one paradigm alone.

A.2.1 Hybrid Feature Fusion Design

The hybrid feature-fusion framework integrates quantum-generated representations directly into conventional machine learning pipelines. Classical input data is first processed through standard preprocessing procedures before being encoded into variational quantum circuits. The resulting quantum representations are then combined with classical feature vectors and supplied jointly to downstream learning systems.

This design allows quantum and classical information to coexist throughout the inference process. Rather than requiring the quantum subsystem to perform complete end-to-end classification independently, the framework distributes representational responsibility across both computational domains. This reduces dependence on noisy quantum outputs while preserving access to potentially useful quantum feature structures.

An important advantage of this approach is flexibility. Classical systems continue to provide stable optimization and mature feature-processing capabilities, while quantum circuits contribute nonlinear probabilistic representations that may capture relationships not represented efficiently within the original feature space.

The experiments also suggested that feature fusion improves robustness under several noisy execution conditions. Because representational information is distributed across both quantum and classical components, perturbations affecting quantum measurements have less influence on overall system behavior than in architectures relying heavily on quantum representations alone.

Another important observation was that richer representations could be obtained without substantially increasing circuit complexity. Earlier experiments demonstrated that deeper

variational circuits often amplified instability under realistic NISQ conditions. Feature fusion instead allows additional representational diversity to emerge through cooperation between subsystems while maintaining relatively shallow quantum architectures.

These findings support the broader interpretation that practical hybrid quantum learning systems benefit most when quantum computation augments classical learning rather than attempting to replace it.

A.2.2 Classical versus Hybrid Feature Spaces

To evaluate feature fusion more systematically, the framework compared purely classical feature spaces against combined quantum-classical representations across multiple benchmark environments. The objective was to determine whether quantum-generated features contribute meaningful information beyond that already available through conventional pre-processing techniques.

The experiments demonstrated that classical feature spaces remain highly effective across many structured learning tasks. Conventional machine learning models frequently achieved strong performance while maintaining lower optimization complexity and greater stability than purely variational quantum systems operating independently.

At the same time, the inclusion of quantum-generated features occasionally produced measurable improvements in nonlinear representation behavior. In several benchmark environments, quantum representations introduced additional structure that increased decision-boundary flexibility and improved the ability of the learning system to model complex relationships within the data.

An important observation was that the usefulness of quantum features depended strongly on dataset characteristics. For relatively simple tasks, classical representations frequently dominated the learning process. In more nonlinear environments, however, quantum representations sometimes contributed information that complemented the existing classical feature space.

The experiments also demonstrated that hybrid feature spaces remained more robust than purely quantum representations under noisy execution conditions. Because the classical subsystem continued to provide stable representational support, perturbations affecting quantum features did not destabilize the overall learning process as severely as in earlier baseline HQNN architectures.

These findings support a conditional interpretation of quantum usefulness. Quantum-generated features do not universally outperform classical representations, but they may provide complementary structure under selected learning conditions.

A.2.3 Conditional Hybrid Complementarity

One of the most important conclusions emerging from the feature-fusion experiments is that quantum and classical representations frequently behave complementarily rather than competitively. The strongest performance often occurred when each subsystem contributed different forms of information within a coordinated learning framework.

This complementarity was particularly visible in nonlinear benchmark environments. Classical systems provided stable optimization behavior, scalable feature processing, and reliable convergence characteristics. Quantum feature representations occasionally contributed additional nonlinear structure and probabilistic diversity that enhanced the flexibility of the resulting decision boundaries.

An important observation was that the usefulness of quantum representations remained highly conditional. The experiments did not reveal universal advantages across all datasets or learning environments. Instead, the contribution of quantum features depended strongly on problem structure, optimization stability, feature-space geometry, and robustness under noise.

This interpretation aligns closely with the broader conclusions developed throughout the thesis. Practical hybrid quantum learning should not be viewed primarily as a competition between classical and quantum computation. Rather, it should be understood as a coop-

erative framework in which each subsystem contributes strengths that compensate for the limitations of the other.

Another important implication concerns robustness itself. Because feature fusion distributes representational responsibility across multiple computational components, the resulting systems frequently maintain greater resilience under perturbation than architectures relying heavily on quantum representations alone. This cooperative structure reduces sensitivity to isolated sources of instability while preserving access to richer feature-generation capabilities.

The conditional complementarity observed throughout these experiments therefore reinforces one of the central conclusions of this work: the strongest practical HQNN behavior emerges through coordinated interaction between quantum feature generation, classical optimization, adaptive inference, and robustness-aware learning rather than through increasing quantum complexity in isolation.

A.3 Data-Reuploading and Feature-Map Search

One of the recurring limitations observed throughout the earlier HQNN experiments was the restricted representational capacity of shallow variational circuits operating under realistic NISQ constraints. Although the robustness-engineering framework substantially improved stability and noisy performance, practical quantum systems remain limited by circuit depth, coherence time, and optimization reliability. Consequently, increasing representational flexibility without significantly increasing circuit complexity remains an important research objective.

To address this challenge, the framework explored data-reuploading strategies and feature-map search procedures designed to improve nonlinear representation capability while maintaining compatibility with realistic hardware constraints. Rather than increasing circuit depth directly, these approaches modify how classical information is embedded and transformed within the quantum system. Data reuploading repeatedly injects classical informa-

tion into the circuit throughout variational evolution, while feature-map design investigates alternative encoding structures capable of generating richer quantum representations. Together, these techniques attempt to increase expressive capability without relying exclusively on deeper or more heavily entangled circuits.

The experiments presented throughout this section demonstrate that both approaches can improve nonlinear representation behavior under selected conditions. At the same time, they reveal that increases in expressive capability frequently introduce additional sensitivity to noise and optimization instability. Consequently, these methods reinforce the broader tradeoff between expressivity and robustness that has emerged repeatedly throughout the thesis.

A.3.1 Angle Encoding versus Reuploading

The first set of experiments compared conventional angle encoding against repeated data-reuploading strategies. In standard angle encoding, classical input features are embedded into the quantum circuit once during state preparation. This approach is computationally efficient and compatible with shallow architectures, but it may limit representational flexibility when learning complex nonlinear relationships.

Data reuploading extends this process by repeatedly introducing classical information throughout circuit evolution. Instead of relying solely on the initial encoding stage, the circuit receives additional input-dependent transformations across multiple variational layers. This repeated embedding allows shallow circuits to generate richer nonlinear representations without requiring substantially deeper architectures.

The experiments demonstrated that reuploading frequently improved nonlinear representation capability relative to single-stage encoding. Several benchmark tasks exhibited smoother decision boundaries and greater feature-space flexibility when repeated embedding strategies were employed. These results are consistent with previous variational quantum learning studies suggesting that repeated feature injection increases the approximation ca-

pability of shallow circuits.

However, the experiments also revealed an important limitation. Reuploading increases the number of encoding operations performed during execution and therefore introduces additional opportunities for perturbation accumulation. In several noisy benchmark environments, aggressive reuploading strategies produced less stable optimization behavior than simpler encoding procedures.

These findings reinforce a recurring theme throughout the thesis. Techniques that improve expressivity can also increase sensitivity to noise, requiring careful balance between representational flexibility and robustness.

A.3.2 Interaction-Based Feature Maps

Building upon the reuploading experiments, the framework next investigated interaction-based feature maps designed to encode nonlinear relationships directly into quantum state evolution. Unlike simple angle-encoding procedures that primarily operate on individual qubits, interaction-based feature maps incorporate transformations that capture relationships between multiple input variables simultaneously.

The motivation for this approach is straightforward. Many nonlinear learning problems depend heavily on interactions between variables rather than on independent feature values alone. Interaction-based feature maps attempt to represent these relationships directly through quantum transformations, potentially generating richer feature spaces than independent encoding strategies.

The experiments demonstrated that interaction-based feature maps frequently produced more flexible representations than simpler encoding procedures. Several nonlinear benchmarks exhibited improved decision-boundary formation and increased representational diversity when interaction terms were incorporated into the encoding process.

At the same time, these richer representations often introduced additional robustness challenges. Because interaction-based maps relied more heavily on coordinated quantum

transformations, perturbations frequently propagated more aggressively throughout the system. As a result, optimization stability became increasingly important for maintaining useful performance.

The results therefore suggest that feature-map design should not focus exclusively on maximizing representational complexity. Effective feature maps must balance expressive capability against stability under realistic execution conditions.

A.3.3 Expressivity versus Robustness Tradeoffs

The data-reuploading and feature-map experiments provide another example of the recurring tradeoff between expressivity and robustness observed throughout this thesis. Increasing representational flexibility often improved nonlinear learning capability under controlled conditions, but these same modifications frequently increased sensitivity to perturbation and optimization instability.

More expressive feature maps generated richer quantum representations and occasionally improved nonlinear classification behavior. However, they also introduced additional computational complexity and increased opportunities for noise accumulation. Simpler encoding procedures often produced more stable behavior despite possessing lower theoretical representational capacity.

An important conclusion emerging from these experiments is that expressive capability alone does not determine practical HQNN usefulness. Under realistic NISQ conditions, representational quality must be evaluated alongside robustness, optimization stability, and hardware compatibility. Maximizing one objective without considering the others frequently leads to diminished overall performance.

The experiments therefore reinforce one of the central conclusions of the thesis. Practical hybrid quantum learning requires balancing expressivity against robustness rather than maximizing quantum complexity indiscriminately. The strongest systems are those that combine useful representations with stable optimization and reliable behavior under realistic

execution conditions.

A.4 Noise Curriculum Learning

The robustness-engineering framework developed throughout Chapter 4 demonstrated that decoder adaptation, architecture optimization, stability-aware training, and noise-aware learning can substantially improve HQNN behavior under realistic NISQ conditions. Despite these improvements, the optimization process itself remained sensitive to how perturbation environments were introduced during training. Several experiments suggested that immediate exposure to strong perturbation conditions could destabilize learning before useful feature representations had fully developed.

This observation motivated the exploration of noise curriculum learning. Rather than exposing the model to a fixed perturbation level throughout optimization, the curriculum framework gradually increased noise intensity as training progressed. The objective was to allow the system to establish stable representations under relatively mild conditions before adapting to more challenging perturbation environments.

The underlying hypothesis is that robustness develops progressively. Just as feature representations evolve during training, the ability of a model to tolerate perturbation may also improve when challenges are introduced gradually rather than all at once. Curriculum learning therefore treats perturbation exposure as a controllable component of the optimization process rather than as a static property of the training environment.

The experiments presented throughout this section demonstrate that progressive perturbation scheduling improves convergence behavior, increases optimization stability, and frequently produces more robust parameter configurations than fixed-noise training procedures.

These findings suggest that training strategy itself can function as an important robustness mechanism within practical hybrid quantum learning systems.

A.4.1 Fixed-Noise versus Curriculum Training

The first set of experiments compared conventional fixed-noise training against curriculum-based perturbation scheduling. In the fixed-noise configuration, the model was exposed to the same perturbation intensity throughout optimization. In the curriculum configuration, perturbation levels increased gradually across successive training stages.

Under fixed-noise conditions, several HQNN architectures exhibited unstable convergence behavior during the early stages of training. Noisy objective evaluations and stochastic measurement variability frequently disrupted parameter updates before stable feature representations could emerge. As a result, optimization occasionally converged toward unstable regions or exhibited persistent oscillatory behavior.

The curriculum-learning framework produced noticeably different behavior. Early optimization occurred under relatively mild perturbation conditions, allowing the model to establish useful representations before stronger perturbations were introduced. Noise intensity increased progressively as training advanced, exposing the system to increasingly difficult environments while preserving previously learned stability.

The experiments demonstrated that curriculum-trained systems frequently achieved smoother convergence behavior and improved final robustness relative to fixed-noise baselines. Models trained through progressive perturbation exposure also tended to generalize more effectively across unseen noise environments during evaluation.

These results suggest that robustness is influenced not only by architecture and optimization design, but also by how perturbation conditions are presented during training.

A.4.2 Adaptive Noise Scheduling

Building upon the curriculum-learning framework, the experiments next investigated adaptive noise scheduling strategies. Rather than following a predetermined perturbation schedule, the adaptive framework adjusted noise intensity dynamically according to observed optimization behavior.

The motivation for adaptive scheduling arises from the variability commonly observed in variational quantum optimization. Different architectures, initialization conditions, and datasets often exhibit different convergence characteristics. Consequently, a fixed progression may not always provide the most effective training environment.

Within the adaptive framework, perturbation intensity increased when optimization remained stable and decreased temporarily when excessive instability emerged. This allowed the learning process to respond directly to the behavior of the model rather than following a rigid schedule.

The experiments demonstrated that adaptive scheduling frequently improved convergence consistency and reduced sensitivity to initialization variability. In several benchmark environments, the framework maintained smoother optimization trajectories and more reproducible training behavior than either fixed-noise or fixed-curriculum procedures.

These findings suggest that effective perturbation management can improve robustness without requiring additional circuit complexity or more elaborate quantum architectures.

A.4.3 Optimization Stability Under Curriculum Learning

The final curriculum-learning experiments focused on optimization stability itself. Earlier chapters demonstrated that noisy gradients, measurement variability, and stochastic perturbations frequently destabilize variational quantum optimization. Curriculum learning was therefore evaluated primarily as a mechanism for improving convergence reliability.

The experiments demonstrated that progressive perturbation exposure consistently improved optimization stability across multiple benchmark environments. Compared with fixed-noise procedures, curriculum-trained systems exhibited smoother convergence trajectories, reduced variability across repeated runs, and more stable inference behavior during evaluation.

An important observation was that curriculum learning reduced the tendency of optimization to converge prematurely toward unstable parameter regions. By allowing useful

representations to emerge under milder conditions before introducing stronger perturbations, the framework produced more reliable optimization behavior throughout training.

The experiments also suggested that curriculum learning interacted positively with several robustness mechanisms introduced earlier in the thesis. Stability-aware optimization, adaptive decoder training, and modular architectures all benefited from more controlled perturbation exposure during optimization.

These findings reinforce one of the central conclusions of the thesis: practical HQNN robustness is fundamentally tied to optimization stability under realistic NISQ conditions. Effective hybrid quantum learning therefore depends not only on representational capability, but also on carefully managing how the learning process encounters perturbation throughout training.

A.5 Variance-Reduced SPSA Optimization

The previous sections demonstrated that robustness-aware training procedures can improve HQNN behavior under realistic NISQ conditions. Curriculum learning, adaptive perturbation scheduling, and stability-aware optimization all contributed to smoother convergence and more reliable learning dynamics. Despite these improvements, one source of instability remained persistent throughout many experiments: stochastic variability in SPSA-based optimization.

Simultaneous Perturbation Stochastic Approximation (SPSA) is widely used in variational quantum learning because it estimates gradients efficiently while requiring relatively few circuit evaluations. This characteristic makes SPSA particularly attractive for NISQ systems, where circuit execution remains expensive and noisy. However, the same stochastic properties that make SPSA computationally efficient also make it sensitive to perturbation-induced variability.

Earlier benchmark experiments repeatedly demonstrated that SPSA optimization can produce oscillatory convergence behavior, unstable parameter updates, and inconsistent

learning trajectories when operating under noisy conditions. These effects frequently appeared even when architecture design, feature extraction, and decoder adaptation had already been improved.

To address this limitation, the framework explored variance-reduction strategies designed specifically to improve the stability of SPSA optimization. Rather than modifying the quantum architecture directly, these methods focused on improving the reliability of gradient estimation and reducing the influence of stochastic perturbations on parameter updates.

The experiments demonstrated that variance reduction substantially improved optimization consistency and convergence reliability across multiple benchmark environments. More importantly, the results reinforced a recurring conclusion throughout the thesis: optimization stability remains one of the most important determinants of practical hybrid quantum learning performance under realistic NISQ conditions.

A.5.1 SPSA Gradient Variability

SPSA optimization estimates parameter updates through stochastic perturbations applied simultaneously across multiple dimensions of the parameter space. Rather than computing exact gradients, the algorithm constructs approximate update directions from differences between noisy objective-function evaluations. This approach dramatically reduces computational overhead and makes optimization feasible for larger variational quantum systems.

The stochastic nature of SPSA, however, introduces substantial variability into the optimization process. Quantum measurements are already subject to probabilistic fluctuations, decoherence effects, gate imperfections, and readout noise. When these sources of uncertainty interact with SPSA perturbations, gradient estimates can vary significantly from one optimization step to the next.

The experiments revealed that this variability frequently destabilized optimization trajectories in baseline HQNN systems. Parameter updates occasionally overshoot stable regions of the optimization landscape, while repeated perturbations produced oscillatory convergence

patterns that slowed learning and reduced reproducibility.

An important observation was that SPSA variability increased as circuit complexity and perturbation intensity increased. More expressive variational systems often generated noisier optimization landscapes, making reliable gradient estimation increasingly difficult. Consequently, improvements in representational capability were sometimes offset by reduced optimization stability.

The results therefore indicate that SPSA variability is not simply a minor implementation issue. It represents a meaningful source of instability that can limit practical performance even when the underlying learning architecture is otherwise well designed.

A.5.2 Variance Reduction Techniques

To improve optimization reliability, the framework introduced several variance-reduction procedures designed to stabilize SPSA updates while preserving the computational efficiency of the underlying algorithm. These methods focused on reducing the sensitivity of gradient estimates to stochastic perturbation fluctuations.

One of the most effective strategies involved gradient averaging across repeated perturbation evaluations. Rather than updating parameters using a single noisy estimate, the framework aggregated information from multiple evaluations to construct smoother update directions. This process reduced the influence of isolated outliers and improved consistency across optimization runs.

Additional stabilization procedures included adaptive step-size management, perturbation smoothing, and constrained update scaling. Together, these methods limited abrupt parameter movements and reduced the likelihood of large optimization oscillations generated by highly noisy gradient estimates.

The experiments demonstrated that variance reduction improved convergence reliability across several benchmark environments. Optimization trajectories became smoother, repeated runs exhibited lower variability, and parameter evolution remained more consistent

throughout training.

An important finding was that variance reduction improved more than optimization smoothness alone. Stable gradient estimates also improved decoder adaptation, feature-space learning, and architecture utilization by providing a more reliable foundation for the learning process as a whole.

These observations suggest that optimization stabilization can influence multiple components of the hybrid learning pipeline simultaneously rather than acting as an isolated improvement.

A.5.3 Stability Improvements in Noisy Optimization

The integration of variance-reduction techniques produced substantial improvements in noisy optimization behavior across multiple HQNN architectures. Compared with conventional SPSA training procedures, the stabilized framework exhibited smoother convergence trajectories, reduced oscillatory behavior, and greater reproducibility across repeated evaluations.

One particularly important observation was that variance reduction allowed the framework to exploit richer feature representations without immediately destabilizing optimization. Earlier experiments involving data reuploading and interaction-based feature maps occasionally produced unstable learning dynamics under noisy conditions. The stabilized optimization framework substantially improved the ability of the system to benefit from these richer representations while maintaining reliable convergence.

The experiments also demonstrated that optimization improvements generalized across different architectures and perturbation environments. Modular QuanNN systems continued to exhibit particularly stable behavior, while more globally connected architectures also benefited from reduced optimization variability.

Another important finding was that optimization stability often improved before measurable performance gains became visible. In several benchmark environments, convergence trajectories became noticeably smoother long before classification accuracy changed substan-

tially. This observation suggests that optimization behavior can serve as an early indicator of broader robustness improvements within hybrid learning systems.

More broadly, the results reinforce one of the central conclusions of the thesis. Practical hybrid quantum learning depends not only on feature representation, architecture design, and decoder adaptation, but also on maintaining reliable optimization dynamics under realistic perturbation conditions.

The variance-reduced SPSA framework therefore demonstrates that optimization reliability can be engineered systematically through targeted stabilization strategies. This finding strengthens the broader robustness-engineering philosophy developed throughout the thesis and provides another example of how practical HQNN performance emerges through coordinated improvements across the entire hybrid learning pipeline.

A.6 Applied Domain Experiments

The previous sections demonstrated that the robustness-engineered HQNN framework maintained stable behavior across nonlinear benchmarks, feature-fusion experiments, optimization studies, and real hardware validation environments. While these evaluations provide important insight into robustness and learning behavior, practical hybrid quantum systems must ultimately be examined within application-oriented settings where classification, anomaly detection, and optimization tasks resemble real-world computational problems more closely.

To investigate this question, the framework was evaluated across several representative application domains, including medical risk classification, cybersecurity anomaly detection, and energy optimization scenarios. These domains were selected because they contain nonlinear relationships, probabilistic uncertainty, and optimization complexity that align naturally with many of the challenges explored throughout the thesis.

The objective of these experiments was not to demonstrate production-level superiority over mature classical systems. Instead, the goal was to determine whether the robustness-

engineered HQNN framework could maintain stable and interpretable behavior while operating within more realistic application environments.

The experiments demonstrated that the framework successfully performed meaningful learning and optimization tasks across multiple domains while maintaining substantially greater stability than earlier baseline architectures. More importantly, the results reinforced a recurring conclusion throughout the thesis: practical usefulness depends primarily on robustness, optimization reliability, and adaptive hybrid learning rather than on quantum complexity alone.

These application studies therefore provide additional evidence that robustness engineering remains central to the development of practical hybrid quantum learning systems under realistic NISQ constraints.

A.6.1 Medical Risk Classification

Medical risk classification provides a useful application environment for evaluating hybrid quantum learning because healthcare data frequently contains nonlinear feature interactions, probabilistic uncertainty, and complex decision boundaries. In addition, medical applications place a strong emphasis on reliability and consistency, making robustness particularly important.

The experiments evaluated the HQNN framework on structured medical-style classification tasks involving numerical risk indicators and simplified diagnostic features. These benchmark environments were designed to emulate probabilistic decision-making scenarios while remaining compatible with practical NISQ-era architectures.

The results demonstrated that the robustness-engineered framework successfully learned meaningful classification behavior while maintaining substantially greater stability than earlier baseline HQNN systems. Adaptive decoding, multi-observable feature extraction, and stability-aware optimization all contributed to more reliable inference under noisy conditions.

An important observation was that hybrid quantum feature representations occasionally

captured nonlinear relationships that were difficult to separate through simple linear decision boundaries. At the same time, classical machine learning systems remained highly competitive across many benchmark configurations.

These findings support a conditional interpretation of quantum usefulness. The experiments do not suggest universal advantages for medical inference tasks. Instead, they indicate that hybrid quantum representations may provide useful complementary structure when integrated within robustness-aware learning pipelines.

A.6.2 Cybersecurity Anomaly Detection

Cybersecurity anomaly detection represents another useful application environment because anomaly-detection tasks often involve sparse events, nonlinear statistical relationships, and evolving distributions. These characteristics provide a challenging setting for evaluating hybrid quantum representations.

The experiments examined structured anomaly-detection tasks designed to simulate intrusion detection and irregular activity classification. The objective was to determine whether the HQNN framework could identify subtle deviations within probabilistic feature distributions while maintaining stable behavior under noisy execution conditions.

The results demonstrated that the framework successfully learned meaningful anomaly-separation structures and maintained more reliable classification behavior than earlier baseline architectures. In several benchmark environments, the hybrid quantum representations produced smoother decision boundaries and improved stability during optimization.

An important observation was that the probabilistic nature of quantum feature representations occasionally aligned well with the irregular statistical structure characteristic of anomaly-detection problems. However, these benefits depended strongly on stability-aware optimization and adaptive decoding procedures.

The cybersecurity experiments therefore provide additional evidence that hybrid quantum learning may offer conditional usefulness within selected nonlinear environments when

supported by a robust and stable learning framework.

A.6.3 Energy Optimization Scenarios

The final application experiments examined simplified energy optimization scenarios involving constrained resource-allocation and optimization-style tasks. These benchmarks were motivated by the broader relationship between variational quantum methods and combinatorial optimization problems.

The experiments evaluated whether the robustness-engineered HQNN framework could maintain stable optimization behavior while operating within nonlinear objective landscapes under noisy conditions. The benchmark tasks were designed to emulate resource-balancing and energy-distribution problems while remaining compatible with practical circuit-depth limitations.

The results demonstrated that the framework maintained substantially more stable optimization behavior than earlier baseline architectures. Stability-aware optimization, variance-reduced SPSA procedures, and adaptive decoder mechanisms all contributed to smoother convergence and more reliable optimization trajectories.

Another important observation was that richer quantum feature representations occasionally improved the exploration of nonlinear objective landscapes when combined with strong classical optimization procedures. However, these benefits remained dependent on robustness-aware training and careful management of perturbation effects throughout optimization.

The energy optimization experiments therefore reinforce one of the central conclusions of the thesis. Practical hybrid quantum learning derives its greatest usefulness not from increasing quantum complexity alone, but from coordinated interaction between quantum feature generation, classical optimization, adaptive inference, and robustness-oriented design principles operating under realistic NISQ constraints.

A.6.4 Additional Findings and Implications

The experiments developed throughout Chapters 3, 4, and in this chapter collectively reveal several broader implications regarding the future of hybrid quantum learning under realistic NISQ conditions. Although the robustness-engineered HQNN framework demonstrated substantial improvements relative to earlier baseline architectures, the results also exposed important constraints that continue to limit practical deployment. At the same time, the experiments identified several design principles that consistently improved stability, reliability, and learning performance across multiple benchmark environments.

Perhaps the most important observation is that practical hybrid quantum learning should not be interpreted primarily as a problem of maximizing quantum complexity. Throughout the framework evolution, improvements in optimization stability, decoder adaptation, feature representation, and architecture design frequently produced greater benefits than increases in circuit depth or entanglement complexity. This observation appeared consistently across benchmark studies, hardware validation experiments, and applied-domain evaluations.

The results also reinforce the importance of hybrid quantum-classical cooperation. The strongest performance rarely emerged from quantum computation operating independently. Instead, robust behavior developed through coordinated interaction between quantum feature generation, classical optimization, adaptive inference procedures, and stability-aware learning mechanisms. This finding supports the hybrid-first interpretation adopted throughout the thesis.

At the same time, the experiments reveal that current NISQ-era quantum learning systems remain constrained by hardware limitations, optimization challenges, and scalability concerns. Although robustness engineering can substantially improve practical behavior, several fundamental limitations remain unresolved and continue to influence performance under realistic conditions.

The following subsections summarize the most important practical lessons emerging from the framework, discuss the limitations of current hybrid quantum learning systems, and

identify promising directions for future research.

A.6.5 Practical Lessons for NISQ-Era HQNN Design

One of the clearest lessons emerging from the experimental framework is that shallow and modular architectures frequently provide the most reliable behavior under realistic NISQ conditions. Earlier benchmarking experiments repeatedly demonstrated that increasing circuit depth and entanglement complexity often amplified instability more rapidly than it improved representational capability. More localized architectures consistently maintained stronger optimization stability and more reliable inference behavior under perturbation.

Another important lesson concerns decoder design. The experiments showed that parity-based measurement interpretation introduced severe information bottlenecks that amplified instability throughout the learning pipeline. Learned-readout decoding and multi-observable feature extraction substantially improved robustness by preserving richer information and allowing adaptive interpretation of noisy quantum outputs.

The results also emphasize the importance of optimization stability. Stability-aware training, variance-reduced SPSA procedures, and curriculum-based learning all contributed meaningfully to improved convergence behavior. These findings suggest that optimization reliability should be treated as a primary design objective rather than as a secondary implementation concern.

A further lesson involves the role of quantum representations themselves. Across multiple experiments, quantum-generated features occasionally provided useful nonlinear structure and representational diversity. However, these benefits were highly dependent on robustness-aware optimization, architecture design, and adaptive decoding. Consequently, quantum usefulness appears conditional rather than universal.

The collective evidence therefore supports a practical design philosophy centered on modular architectures, adaptive decoding, stability-aware optimization, multichannel validation, and coordinated hybrid quantum-classical integration rather than maximizing variational

complexity alone.

A.6.6 Limits of Current Hybrid Quantum Learning Systems

Although the robustness-engineered framework demonstrated substantial improvements relative to earlier baseline systems, the experiments also revealed several important limitations affecting current hybrid quantum learning architectures.

One major limitation concerns scalability. Even with improved optimization procedures and modular architecture design, larger variational systems frequently remained vulnerable to perturbation accumulation and optimization instability. Increasing the number of qubits, parameters, or entangling operations often introduced additional sensitivity to noise and reduced training reliability.

Another limitation involves hardware dependence. The experiments demonstrated that performance frequently varied across simulator environments, transpilation procedures, backend calibration conditions, and hardware execution settings. Although the framework generalized more effectively than earlier architectures, practical behavior remained influenced by backend-specific characteristics that are difficult to control completely.

Optimization reliability also remains a significant challenge. Variational quantum learning continues to depend heavily on stochastic optimization procedures that are sensitive to noisy objective-function evaluations and perturbation-induced variability. While the framework reduced these effects substantially, optimization remains one of the principal bottlenecks limiting practical deployment.

A further limitation concerns interpretability. Although several experiments demonstrated meaningful nonlinear feature-generation behavior, understanding precisely how quantum representations contribute to inference remains difficult. Compared with many classical machine learning systems, quantum feature spaces remain challenging to analyze rigorously.

Finally, the experiments consistently demonstrated that current HQNN systems do not universally outperform strong classical baselines. Instead, their usefulness depends strongly

on dataset structure, optimization stability, noise conditions, and the effectiveness of the surrounding hybrid learning pipeline.

A.7 Framework Implementation and Experimental Reproducibility

This section provides the technical details required to reconstruct the principal Hybrid Quantum Neural Network architecture and reproduce the experiments presented in Chapters 3 and 4. The purpose is to make the experimental methodology understandable independently of the accompanying source code. The circuit structure, preprocessing procedure, measurement strategy, classical readout, training configuration, noise models, robustness metrics, software organization, and execution workflow are therefore documented explicitly.

The complete implementation is available in the accompanying public repository. However, the information in this section provides the minimum specification necessary for another researcher to reimplement the framework using an equivalent quantum computing environment.

A.7.1 Overall Experimental Workflow

The experimental framework follows a common sequence across the benchmarking and robustness-engineering pipelines:

1. load or generate the selected benchmark dataset;
2. divide the data into training and test subsets using a fixed random seed;
3. standardize the classical input features using statistics calculated from the training subset;
4. reduce the feature dimension when the number of original features exceeds the available number of logical qubits;
5. encode the resulting feature vector into a parameterized quantum circuit;

6. apply trainable single-qubit rotations and the selected entanglement topology;
7. execute the circuit under noiseless simulation, simulated quantum noise, or physical hardware conditions;
8. convert measurement counts into quantum-derived features;
9. train a fixed or learned classical readout;
10. calculate predictive and robustness metrics;
11. repeat the experiment when statistical validation is required.

The quantum circuit therefore functions primarily as a trainable quantum feature generator. In the baseline configuration, classification is performed using fixed parity decoding. In the proposed framework, the measurement distribution is preserved and interpreted through a learned classical readout.




Figure A.3: Experimental workflow used to reproduce the proposed HQNN framework. Classical data are preprocessed and encoded into a parameterized quantum circuit. The resulting measurement distribution is converted into quantum-derived features and interpreted by a classical readout before predictive and robustness metrics are calculated.

A.7.2 Reference Hybrid Quantum Neural Network Circuit

The principal simulated HQNN configuration uses four logical qubits for datasets reduced to four classical input features. Let the standardized input vector be

$$\mathbf{x} = \begin{bmatrix} x_0 & x_1 & x_2 & x_3 \end{bmatrix}^T.$$

Each retained classical feature is encoded into one qubit using a rotation around the Y -axis:

$$U_{\text{enc}}(\mathbf{x}) = \bigotimes_{i=0}^3 R_Y(s_x x_i),$$

where s_x is a feature-scaling coefficient used to map standardized classical features into suitable rotation angles.

The encoded state is

$$\psi_{\text{enc}}(\mathbf{x}) = U_{\text{enc}}(\mathbf{x})0000.$$

After feature encoding, each qubit receives trainable rotations around the X - and Z -axes:

$$U_{\text{var}}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \bigotimes_{i=0}^3 R_Z(\phi_i) R_X(\theta_i).$$

The complete one-layer circuit can therefore be written as

$$\psi(\mathbf{x}, \boldsymbol{\theta}, \boldsymbol{\phi}) = U_{\text{ent}} U_{\text{var}}(\boldsymbol{\theta}, \boldsymbol{\phi}) U_{\text{enc}}(\mathbf{x})0000,$$

where U_{ent} represents the selected entanglement topology.

The architecture-search stage compares four circuit structures:

No entanglement: only feature-encoding and trainable single-qubit rotations are applied.

Linear entanglement: controlled- Z gates connect adjacent qubits q_i and q_{i+1} .

Ring entanglement: the linear topology is extended by connecting the final qubit to the first qubit.

Fully connected entanglement: every logical qubit is entangled with every other logical qubit.

The linear architecture was selected for the principal robustness experiments because it provided a favorable balance between quantum feature interaction, circuit depth, trainability, and accumulated noise.

A.7.3 Baseline Parity Decoder

For an input sample $\mathbf{x}$, circuit execution produces a measurement-count dictionary

$$C(\mathbf{x}) = \{b : n_b\},$$

where $b \in \{0, 1\}^n$ is a measured bitstring and n_b is its observed count.

The empirical probability of bitstring b is

$$p(b \mid \mathbf{x}) = \frac{n_b}{N_{\text{shots}}},$$

where N_{shots} is the total number of circuit executions.

The baseline HQNN compresses the entire measurement distribution into a single parity expectation:

$$f_{\text{parity}}(\mathbf{x}) = \sum_b (-1)^{w(b)} p(b \mid \mathbf{x}),$$

where $w(b)$ is the Hamming weight of bitstring b .

A binary prediction is then produced using a fixed sign or threshold rule. Although this decoder is computationally simple, it maps many distinct measurement distributions to the same scalar output. It therefore discards information that may be valuable for classification.

A.7.4 Learned Readout and Multi-Observable Feature Extraction

The learned-readout configuration preserves the quantum measurement distribution rather than immediately compressing it into one parity value.

For a four-qubit circuit, the full computational-basis probability vector contains

$$2^4 = 16$$

entries.

The principal multi-observable representation combines:

1. full computational-basis bitstring probabilities;
2. single-qubit Z -expectation values;
3. pairwise ZZ -correlation values;
4. the global parity expectation;
5. probability-distribution summary statistics.

For logical qubit i , the local Z -expectation is reconstructed from the measurement distribution as

$$\langle Z_i \rangle = \sum_b (-1)^{b_i} p(b).$$

For logical qubits i and j , the pairwise correlation is

$$\langle Z_i Z_j \rangle = \sum_b (-1)^{b_i + b_j} p(b).$$

The global parity expectation is

$$\langle Z_0 Z_1 \cdots Z_{n-1} \rangle = \sum_b (-1)^{w(b)} p(b).$$

The distribution summaries include concentration and uncertainty statistics derived from the empirical probability vector. In the principal four-qubit implementation, the complete extraction procedure produces a quantum-derived feature vector containing the full probability distribution, local expectations, pairwise correlations, parity information, and distribution statistics.

The learned prediction is written as

$$\hat{y} = g_{\omega}(\Phi_Q(\mathbf{x})),$$

where $\Phi_Q(\mathbf{x})$ is the extracted quantum feature vector and g_{ω} is a trainable classical readout.

The framework evaluates multiple readout families, including:

- logistic regression;
- radial-basis-function support vector machines;
- random forests;
- extra-trees classifiers;
- gradient-boosting classifiers.

The strongest reported simulated configuration used a random-forest learned readout.

A.7.5 Dataset Preparation

The framework supports the following benchmark data:

- synthetic binary classification datasets;
- the Iris dataset;
- the Wisconsin Diagnostic Breast Cancer dataset;

- nonlinear synthetic datasets such as moons, circles, and XOR-style benchmarks.

The common preprocessing sequence is:

1. load or generate the complete dataset;
2. select the target classes required by the experiment;
3. divide the data into training and test subsets;
4. use a fixed random seed;
5. stratify the split when class labels permit stratification;
6. fit a standard scaler using the training data only;
7. apply the fitted scaler to the training and test data;
8. reduce dimensionality when the number of features exceeds the number of available logical qubits;
9. use the same processed data for all models included in the comparison.

The general preprocessing procedure can be reproduced as follows:

[language=Python, caption=Pseudocode for reproducible preprocessing.] X