

Comparative Analysis of Hybrid Quantum Neural Networks: Efficiency, Accuracy, and Applications

Andrew Nerud



Abstract—Hybrid Quantum Neural Networks (HQNNs) integrate quantum computing with classical deep learning to improve efficiency, expressivity, and performance in high-dimensional learning tasks. While they offer theoretical advantages through quantum-enhanced feature representations and variational quantum circuits, real-world deployment remains challenging due to quantum noise, short coherence times, and hybrid processing overhead. This paper evaluates the practical viability of HQNNs by systematically reviewing comparative studies across domains such as medical imaging, quantum chemistry, and natural language processing. Using metrics like training time, parameter count, and floating-point operations (FLOPs), we find that HQNNs can achieve up to 50% fewer parameters, a 30% reduction in FLOPs, and faster convergence—while maintaining or exceeding classical model accuracy. However, these benefits are often offset by quantum-classical bottlenecks and training instabilities such as barren plateaus. Our analysis highlights the need for more robust error mitigation strategies, scalable quantum hardware, and standardized benchmarking protocols. While not yet viable replacements, HQNNs are a transitional step toward fully quantum machine learning systems, particularly in domains requiring low-resource, high-efficiency inference.

CONTENTS

I	Introduction	1
II	Background and Theoretical Foundations	2
II-A	Quantum Computing and Machine Learning	2
II-A1	Entanglement	2
II-A2	Hadamard Gate	2
II-A3	Controlled-NOT gate	2
II-B	Introduction to HQNNs	3
II-C	Quantum Feature Representations in HQNNs	3
II-D	Summary of Theoretical Foundations	3
III	Summary of Previous Research Methods	4
III-A	Datasets and Benchmarking Standards	4
III-B	Performance Evaluation Metrics	4
III-C	Quantum Hardware vs. Simulation	4
III-D	Evolution of HQNN Methodologies	4
III-E	Summary	4
IV	Findings from Previous Research	5
IV-A	HQNNs vs. Classical Neural Networks	5
IV-B	Application-Specific Findings	5
IV-B1	Medical Imaging and Disease Prediction	5
IV-B2	Quantum Chemistry Simulations	6

IV-B3	Text Processing and NLP with Quantum Feature Encoding	6
IV-B4	Handwritten Digit Recognition with HQNNs	6
IV-C	Persistent Challenges in HQNN Adoption	6
IV-D	Comparative Performance Analysis	6
V	Challenges and Future Considerations	7
V-A	Challenges in HQNN Implementation	7
V-B	Future Directions	8
V-C	The Path Forward	8
VI	Conclusion	9
	References	9

I. INTRODUCTION

In 2023, researchers at MIT demonstrated that HQNNs could classify medical images with 98.7% accuracy while using 30% fewer parameters than traditional CNNs [1]. This breakthrough highlights the potential of quantum-assisted deep learning to improve model efficiency while maintaining high accuracy.

HQNNs use **superposition and entanglement** to enhance feature extraction and reduce computational overhead. This paper reviews HQNN performance relative to classical models. We systematically review existing findings regarding:

- The computational efficiency of HQNNs versus classical neural networks.
- The impact of quantum feature representations on training performance.
- The practical limitations and challenges observed in experimental HQNN studies.

Beyond technical performance, HQNNs raise important implications in industry, ethics, and long-term viability. As quantum computing gradually shifts from academic novelty to industrial investment, HQNNs are emerging as candidates for practical AI in sectors like healthcare [2], [3], cybersecurity [4], materials science [5], [6], and environmental modeling [7]. These systems offer faster inference and lower complexity but raise equity, explainability, and access concern.

For example, privacy-aware HQNNs have been proposed for biometric tasks such as lipreading, where classical DNNs pose risks to user anonymity [8]. Meanwhile, HQNN-based pipelines are already outperforming classical models in resource-constrained tasks like corrosion inhibitor discovery [6] and battery health estimation [9], suggesting real-world feasibility even under current quantum hardware limitations.

Given this momentum, an important question arises: **Can HQNNs scale into reliable, general-purpose AI tools beyond academic benchmarks?** And if so, how should the computer science community prepare for their broader integration?

This paper examines the feasibility and limitations of HQNNs by systematically reviewing their computational efficiency, training dynamics, and real-world applications. It also expands upon the social and ethical framing of this technology, drawing from recent advancements and ongoing debates. To support this analysis, the following section first provides a foundational overview of quantum mechanics and its intersection with machine learning.

II. BACKGROUND AND THEORETICAL FOUNDATIONS

A. Quantum Computing and Machine Learning

Quantum computing uses quantum mechanics to perform computations beyond classical systems [10]. Unlike classical bits constrained to binary states (0 or 1), quantum bits (*qubits*) exist in a superposition of both states, enabling exponential computational advantages in specific domains [11].

The qubit, as the fundamental unit of quantum information, is mathematically represented as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad \text{where } \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1 \quad (1)$$

Here, α and β are complex probability amplitudes, and their squared magnitudes sum to 1. This normalization ensures that a measurement will always collapse the qubit into a definite state of $|0\rangle$ or $|1\rangle$. Multi-qubit states are expressed as tensor products of individual qubits, which enables the encoding of quantum correlations [10], [12]. The key quantum properties and definitions relevant to computing will be covered in the following sections.

Key Quantum Terms: A Practical Glossary

To improve clarity, we define essential terms relevant to HQNNs:

- **Qubit:** A quantum bit that can exist in a superposition of 0 and 1, enabling parallel information processing [10].
- **Superposition:** A quantum property allowing simultaneous existence in multiple states, giving rise to computational parallelism [13].
- **Entanglement:** A phenomenon where two or more qubits are interdependent; a change in one affects the state of the other. This is critical for learning feature correlations [10].
- **Quantum Gate:** A transformation applied to a qubit, analogous to classical logic gates. Gates like Hadamard and CNOT manipulate quantum states during computation [12].
- **Variational Quantum Circuit (VQC):** A parameterized quantum model trained via optimization, serving as the core quantum layer in HQNNs [14].
- **Hilbert Space:** A high-dimensional vector space where quantum states live. Feature encodings into Hilbert space enable HQNNs to capture complex data patterns [15].

- **Quantum Kernel:** A similarity measure computed in quantum feature space. Used in tasks like classification with quantum-enhanced SVMs [12].
- **Measurement:** The act of collapsing a quantum state into classical information. This step is probabilistic and introduces latency in hybrid systems [16].

1) *Entanglement:* Entanglement is a quantum phenomenon that links qubits such that measuring one immediately determines the other's state [10], [13]. This correlation is independent of distance, enabling non-local information encoding and efficient representation of feature dependencies in HQNNs.

A two-qubit entangled state, known as a Bell state, is represented as:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \quad (2)$$

This state forms the theoretical basis for quantum teleportation, quantum key distribution, and improved expressivity in variational quantum circuits.

2) *Hadamard Gate:* The Hadamard gate (H) is a foundational single-qubit gate that places a qubit into an equal superposition of states. It is often used at the beginning of quantum algorithms to initiate parallel exploration of solution spaces.

Mathematically, it is defined as:

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (3)$$

Applied to the initial basis state $|0\rangle$, it produces:

$$H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \quad (4)$$

which encodes both computational paths simultaneously for later quantum operations.

3) *Controlled-NOT gate:* The Controlled-NOT (CNOT) gate is a two-qubit quantum operation that flips the target qubit if the control qubit is in the $|1\rangle$ state. It is a key component in generating entangled states within quantum circuits.

Its matrix form is:

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (5)$$

When applied to a superposed input state:

$$CNOT \left(\frac{1}{\sqrt{2}}(|00\rangle + |10\rangle) \right) = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \quad (6)$$

it produces the Bell state from Equation (2), thus entangling the two qubits.

Quantum Machine Learning (QML) explores how quantum computing can accelerate machine learning tasks, such as classification, clustering, and generative modeling [13].

B. Introduction to HQNNs

HQNNs integrate quantum computing layers within classical deep learning models, aiming to enhance computational efficiency while leveraging quantum properties [12], [17]. These models use **variational quantum circuits (VQCs)** that leverage entanglement and superposition for feature extraction.

As shown in **Figure 1**, HQNNs leverage an **N-qubit variational circuit** where each qubit undergoes Hadamard transformations (H), controlled interactions, and parameterized rotation gates ($R_y(\theta)$) [18]. This setup enables feature representations beyond classical models, improving pattern recognition. VQCs are typically trained using hybrid optimization methods, where classical optimizers adjust the quantum gate parameters (θ) based on loss minimization techniques [14].




Fig. 1. One of the various forms of a N-qubit VQCs used in HQNNs. Each qubit begins in $|0\rangle$ and undergoes Hadamard gates, controlled operations, and $R_y(\theta)$ rotations. Adapted from [14], [15], [16], [18], [19].

HQNN Architecture Overview:

- A classical neural network is used for feature extraction and data preprocessing.
- A quantum layer (often implemented using VQCs) replaces one or more classical layers [14].
- The quantum circuit is parameterized and trained using gradient-based optimization, similar to classical deep learning models [20].




Fig. 2. General architecture of a HQNN. The model integrates classical convolutional layers with a quantum variational circuit before producing the final output [1], [21], [22].

HQNNs follow a hybrid structure where:

- A **classical neural network** is used for initial feature extraction and preprocessing.
- A **quantum variational circuit** replaces certain layers within the model, leveraging quantum gates to process data [15].
- Optimization is performed through **hybrid training methods**, combining classical gradient-based techniques with quantum variational parameter tuning [12].

Figure 2 provides an overview of a typical HQNN architecture, where a convolutional neural network (CNN) extracts features before processing them through a quantum variational circuit. This hybrid approach enables improved feature transformations while maintaining classical efficiency.

C. Quantum Feature Representations in HQNNs

Recent research shows quantum encodings can outperform classical embeddings in **high-dimensional tasks** with redundancy or sparsity, which often pose challenges to conventional deep learning approaches [23]. Quantum feature representations **efficiently encode high-dimensional data**, reducing computation and overfitting.

Other recent experimental studies have demonstrated the practical applications of HQNNs in real-world scenarios. For instance, research on **handwritten digit recognition** has shown that HQNNs can achieve comparable accuracy to classical convolutional neural networks (CNNs) while requiring substantially fewer trainable parameters, thereby reducing computational complexity [24]. Similarly, in the domain of **quantum chemistry**, HQNN-based approaches have been employed to predict the ground state energy of molecular systems with improved precision over classical machine learning models [15]. By directly encoding quantum states into the network, HQNNs offer a fundamental advantage in processing quantum mechanical data, making them well-suited for applications in materials science and molecular modeling.

Despite these advantages, the effectiveness of HQNNs is contingent on the efficient design of VQCs, which serve as the backbone of quantum feature extraction. Future research should explore optimal architectures for integrating VQCs within hybrid quantum-classical frameworks, ensuring that feature extraction remains both computationally feasible and robust to quantum noise. Additionally, empirical comparisons between quantum and classical feature representations could further illuminate the contexts in which HQNNs provide the most substantial performance gains.

This section establishes the necessary theoretical background to evaluate HQNNs' computational efficiency. While quantum computing provides unique computational advantages, its integration within neural networks presents several challenges that influence model performance. To assess how these challenges manifest in real-world applications, the next section reviews previous research methodologies, focusing on dataset selection, benchmarking techniques, and the experimental constraints imposed by current quantum hardware.

D. Summary of Theoretical Foundations

This section established the essential concepts needed to understand the hybrid quantum-classical approach used in HQNNs. Quantum properties like superposition, entanglement, and variational circuits enable novel approaches to feature extraction and model compression. While classical deep learning has matured through extensive optimization and hardware support, quantum neural networks offer a fundamentally new direction that promises increased efficiency — albeit with new challenges in implementation and hardware readiness.

The next section will explore how these foundational ideas have been applied in experimental studies, with emphasis on dataset selection, benchmarking, and hardware constraints.

III. SUMMARY OF PREVIOUS RESEARCH METHODS

To evaluate HQNNs, researchers have applied a variety of experimental methodologies across diverse domains, ranging from image recognition to molecular modeling. Key factors that influence evaluation include dataset selection, benchmarking standards, simulation environments, and quantum hardware constraints.

A. Datasets and Benchmarking Standards

Benchmarking HQNNs often begins with classical datasets, as they provide a point of comparison against well-optimized deep learning models. The most commonly used datasets include:

- **MNIST:** Widely used for binary and multiclass digit classification [18], [24].
- **COVID-19 Chest X-rays:** Used to test HQNN performance in medical diagnostics with small, high-dimensional inputs [1].
- **Molecular Energy Datasets:** Applied in quantum chemistry, evaluating energy prediction for molecular ground states [15].
- **Environmental and Materials Data:** Recent studies have extended HQNN evaluation to lithium battery health [9] and ozone forecasting [7].

These datasets reflect both structured and unstructured input formats. More recent pipelines employ data fusion [9], data augmentation [2], and hybrid quantum-classical feature engineering [6], signaling increased methodological sophistication.

B. Performance Evaluation Metrics

To compare HQNNs with classical deep learning models, studies rely on standardized metrics:

- **Training Time:** Measures convergence speed. HQNNs often require fewer epochs due to quantum-enhanced expressivity [24].
- **Floating-Point Operations (FLOPs):** Tracks computational cost. HQNNs generally show reduced FLOPs due to logarithmic scaling of quantum circuits [12].
- **Parameter Count:** A lower number of tunable parameters often reduces overfitting and training cost. HQNNs are known to match classical accuracy with significantly fewer parameters [19].
- **Inference Latency:** Though less frequently reported, latency is a concern when quantum measurements are slow [18].

However, these metrics must be contextualized. For instance, quantum measurements are probabilistic and introduce variance not captured by FLOPs or parameters alone. As such, several studies have called for more robust evaluation protocols, especially when using hardware backends.

C. Quantum Hardware vs. Simulation

Most HQNN experiments to date are conducted using quantum simulators such as **Qiskit Aer**, **PennyLane**, or **TensorFlow Quantum** [14]. These environments model idealized qubits and do not capture the decoherence, noise, or gate errors present on real quantum hardware.

When HQNNs are executed on actual QPUs (e.g., IBM Q, Rigetti), researchers observe notable deviations from simulation results. Accuracy often drops due to quantum gate noise and the short coherence times of current qubit technologies, which limit circuit depth and consistency. Execution time also increases significantly, primarily because of quantum-classical communication delays introduced during iterative training and measurement. Moreover, to achieve stable gradient estimates during training, researchers frequently resort to batching or repeated quantum sampling, which further increases runtime and hardware demand.

Some newer studies have addressed these limitations by incorporating error mitigation strategies directly into the training pipeline [25]. For instance, VQCs are now being adjusted using hardware-aware optimizers that account for gate noise. Additionally, certain tasks—such as lipreading and biometric detection—have benefited from integrating differential privacy layers, which also improve noise robustness during quantum inference [8]. These developments mark a shift toward more hardware-conscious HQNN modeling practices.

D. Evolution of HQNN Methodologies

As shown in Figure 3, HQNN evaluation techniques have matured across three major stages. Early studies emphasized proof-of-concept goals, using shallow circuits on synthetic or low-complexity datasets. These efforts demonstrated basic viability but lacked real-world applicability. In the next stage, researchers introduced standardized benchmarks such as MNIST and molecular property prediction, often leveraging simulated quantum backends to explore scalability and hybrid training schemes. More recently, HQNNs have been deployed in applied domains including medical imaging, environmental forecasting, and materials modeling. These modern approaches increasingly involve both simulator and real-QPU pipelines [2], [5], [6], reflecting a shift toward full-stack experimentation and deployment-oriented design.

This progression highlights growing confidence in HQNN pipelines and an increasing emphasis on reproducibility and interpretability. Many of the most recent frameworks have adopted more advanced training techniques, including quantum transfer learning [20], hybrid generative adversarial networks (GANs), and decision trees augmented with quantum feature spaces [24]. These directions indicate that HQNNs are becoming more modular and adaptable to a broader range of machine learning problems.

E. Summary

Previous research has demonstrated that HQNNs can reduce parameter count and computational complexity while maintaining competitive accuracy. However, results obtained




Fig. 3. Evolution of HQNN research methodology from early toy problems [12] to benchmarked tasks [15], [24], and recent real-world pipelines [2], [5].

from simulations may not fully translate to hardware implementations. To move forward, future benchmarking must incorporate more standardized evaluation pipelines. In particular, benchmarking frameworks should include pre-defined HQNN model templates, common datasets, and reproducible quantum-classical integration routines. Equally important is the transparent reporting of performance on simulated versus real quantum hardware, as well as application-specific datasets that can stress-test HQNNs under realistic deployment scenarios.

IV. FINDINGS FROM PREVIOUS RESEARCH

A. HQNNs vs. Classical Neural Networks

Comparative studies between HQNNs and classical deep learning models have produced promising, yet nuanced, results. HQNNs often demonstrate computational advantages through quantum-enhanced feature representations, though these benefits are highly dependent on dataset complexity, circuit depth, and integration strategies [24].

Parameter Efficiency and FLOP Reduction. One of the most consistent findings is that HQNNs require significantly fewer parameters than classical models while maintaining similar accuracy. Studies report reductions of up to **40–50% in trainable parameters**, particularly in image classification and chemistry-based models [15]. In addition, HQNNs demonstrate a **30% reduction in FLOPs (floating-point operations)** due to the linear algebraic efficiency of quantum circuits [12]. These reductions correlate with lower memory usage, shorter training times, and decreased risk of overfitting in small-data regimes.

Faster Convergence and Generalization. HQNNs also tend to converge more quickly, especially in high-dimensional feature spaces. Some studies observe a **20–35% speedup** in training convergence compared to classical CNNs [18]. This acceleration is often attributed to the expressive capacity of quantum feature mappings, which encode inputs into high-dimensional Hilbert spaces using unitary transformations. For example, a quantum embedding function $\Phi(x)$ maps classical input x to a quantum state:

$$|\psi(x)\rangle = U(x)|0\rangle^{\otimes n}, \quad (7)$$

where $U(x)$ is a parameterized quantum circuit. This mapping enables HQNNs to capture correlations and dependencies that classical networks may miss, while simultaneously regularizing model complexity through entanglement and interference effects [16].

Overfitting Resistance and Noise Regularization. Several HQNN implementations demonstrate improved performance on small datasets, including in domains like medical imaging and disease prediction [1]. These results suggest that HQNNs offer a form of implicit regularization, reducing the tendency to overfit by leveraging probabilistic measurement and entanglement-based constraints during training.

Remaining Limitations. Despite these benefits, HQNNs are still limited by their hybrid architecture. The quantum-to-classical interface introduces overhead, particularly due to measurement and communication delays between quantum processors and classical optimizers [18]. These bottlenecks can offset the gains in training speed and parameter efficiency unless specialized hardware or batching strategies are employed.

Overall, the comparison suggests that HQNNs hold clear advantages in settings where data is sparse, feature spaces are complex, and parameter budgets are tight. Their benefits diminish, however, in large-scale tasks where classical models can exploit parallel hardware and massive datasets more effectively.

Tables I & II provide quantitative comparisons of HQNNs and classical models across various datasets. In particular, HQNNs demonstrated a **29% reduction in training time** compared to CNNs in the MNIST dataset, while achieving an accuracy of **98.7%**, slightly surpassing its classical counterpart [24]. Similarly, in medical imaging applications, HQNNs achieved a **significant FLOP reduction (from 4.2 billion to 2.5 billion)**, which underlines their computational efficiency [1].

Figure 4 provides a normalized bar chart view of key performance metrics on benchmark datasets, complementing the raw values presented in Table I.

B. Application-Specific Findings

HQNNs have been explored in multiple domains, including **computer vision, quantum chemistry, and natural language processing (NLP)**. Recent studies have benchmarked HQNN performance across these applications, revealing both advantages and limitations.

1) Medical Imaging and Disease Prediction: One of the most promising applications of HQNNs is in **medical imaging**, where quantum feature mapping has shown improvements in early disease detection. A study using **COVID-19 X-ray classification** [1] found that an HQNN model achieved:

- **96.8% accuracy**, outperforming classical CNNs by **2.1%**.
- **35% fewer parameters**, reducing model complexity without sacrificing performance.
- **28% reduction in training time**, indicating improved efficiency in quantum feature encoding.




Fig. 4. Comparison of HQNN and Classical CNN performance across key computational efficiency metrics, adapted from [24]. HQNNs demonstrate improvements in training time, FLOPs, and parameter efficiency while maintaining comparable accuracy.

2) *Quantum Chemistry Simulations*: HQNNs have also demonstrated advantages in **quantum chemistry**, where they are used to model molecular energy states more efficiently than classical methods. A study comparing HQNNs to classical variational models [15] found:

- HQNN-based quantum simulations reduced computational cost by **40%** compared to traditional electronic structure methods.
- Improved **accuracy of molecular energy predictions**, outperforming conventional deep neural networks.

3) *Text Processing and NLP with Quantum Feature Encoding*: Recent research has investigated HQNNs for **natural language processing (NLP)** [12]. By leveraging quantum embeddings, HQNNs have been shown to enhance sentence classification tasks. In a sentiment analysis benchmark:

- An HQNN **outperformed a BiLSTM model** on a small dataset, demonstrating better feature extraction for low-data NLP tasks.
- Quantum embeddings improved **word representation learning**, reducing classification error by **15%**.

4) *Handwritten Digit Recognition with HQNNs*: HQNNs also succeed in **handwritten digit recognition**, showing efficiency and reduced computational complexity. A study comparing HQNNs to classical convolutional neural networks (CNNs) on the **MNIST dataset** [24] found that:

- HQNNs achieved **98.7% accuracy**, slightly surpassing classical CNNs.
- **29% reduction in training time** compared to CNNs, highlighting improved efficiency in quantum-assisted learning.
- **40–50% fewer trainable parameters**, reducing overfitting and memory requirements.

These results indicate that HQNNs offer **computational advantages in low-data environments**, making them promising candidates for tasks requiring efficient learning with limited resources.

C. Persistent Challenges in HQNN Adoption

Despite their potential, HQNNs face several challenges that researchers have consistently highlighted:

- **Hardware Limitations**: The reliance on Noisy Intermediate-Scale Quantum (NISQ) devices restricts HQNNs’ scalability, making their real-world deployment difficult [25].
- **Decoherence and Quantum Noise**: Quantum hardware introduces errors that affect the performance of HQNN, which require error mitigation techniques [16].
- **Quantum-Classical Bottlenecks**: The need for frequent communication between quantum circuits and classical processors introduces delays that negate potential speed-ups [18].

D. Comparative Performance Analysis

The practical advantages of HQNNs over classical deep learning models depend on several factors, including dataset complexity, computational efficiency, and quantum-classical integration challenges [25]. While HQNNs offer promising improvements in parameter efficiency and computational complexity, their benefits are constrained by current quantum hardware limitations.

One of the primary distinctions between HQNNs and classical convolutional neural networks (CNNs) is their **parameter efficiency**. CNNs typically require millions of parameters to encode spatial hierarchies, increasing memory demands and training times. In contrast, HQNNs leverage quantum superposition and entanglement to encode feature spaces more compactly, reducing the number of required parameters [19], [24]. This reduction leads to a lower risk of overfitting, particularly in scenarios with limited training data.

In addition to reducing parameter counts, HQNNs can also **decrease floating-point operations (FLOPs)**, which directly impacts computational efficiency. Unlike classical CNNs, which rely on large matrix multiplications and convolution operations, HQNNs execute feature transformations using quantum circuits that scale logarithmically in certain cases [12]. Current results show HQNNs reduce the number of FLOPs resulting in lower energy consumption and faster training times in specific problem domains [18]. However, the extent of these improvements depends on circuit depth, dataset characteristics, and the efficiency of quantum-classical data transfer.

Despite these advantages, HQNNs face **practical bottlenecks** that impact their real-world performance. A major challenge is the **quantum-classical interface**, where data must be frequently transferred between classical and quantum processors. This back-and-forth exchange negates some of the theoretical computational gains, particularly on near-term quantum hardware where coherence times and gate fidelities

Dataset	Model	Training Time (s)	FLOPs ($\times 10^9$)	Parameters (millions)	Accuracy (%)
MNIST [24]	Classical CNN	1200	1.5	2.1	98.5
	HQNN	850	0.9	1.1	98.7
Medical Imaging (COVID-19) [1]	Classical CNN	3100	4.2	5.3	96.2
	HQNN	2300	2.5	3.4	96.8
Quantum Chemistry [15]	Classical ML Model	5000	5.8	7.0	89.5
	HQNN	3200	3.1	4.2	91.3
Sentiment Analysis (NLP) [12]	BiLSTM	600	N/A	1.8	85.0
	HQNN	420	N/A	1.2	86.5
Alzheimer's Detection [2]	Classical 3D CNN	3800	N/A	6.8	92.3
	CQ-CNN (HQNN)	2700	N/A	3.5	94.0
Ozone Forecasting [7]	Classical LSTM	1400	N/A	2.6	87.4
	HQNN	950	N/A	1.4	89.6
Lipreading (LRW) [8]	LSTM	1900	N/A	N/A	78.1
	HQCNN (PVM)	1250	N/A	N/A	83.9

TABLE I

COMPARISON OF HQNN AND CLASSICAL MODELS ACROSS DIVERSE DOMAINS. NEWER HQNN STUDIES CONTINUE TO SHOW REDUCTIONS IN TRAINING TIME, MODEL SIZE, AND PERFORMANCE IMPROVEMENTS ON MEDICAL, ENVIRONMENTAL, AND PRIVACY-SENSITIVE NLP TASKS.

Domain	Classical Model	HQNN Model	Notable HQNN Advantages
Battery Health Estimation [9]	Gradient Boosting with handcrafted features	Quantum CNN with auto feature fusion	+ Robustness to capacity degradation patterns
Intrusion Detection [4]	Logistic Regression and Random Forest	QML-based binary classifiers	+ 5–8% improvement in detection on small datasets
QSPR for CO ₂ Capture [5]	Standard MLP	HQNN with Variational Regressor	+ Lower RMSE, better generalization on novel amines

TABLE II

ADDITIONAL HQNN STUDIES WHERE FLOPS OR PARAMETER COUNTS WERE NOT EXPLICITLY REPORTED BUT HQNNs SHOWED SIGNIFICANT DOMAIN-SPECIFIC ADVANTAGES.

are limited [16]. Additionally, while quantum feature encodings enable HQNNs to capture complex relationships in data, these embeddings require precise quantum state preparation, which remains an area of active research.

Comparative studies on real-world datasets have produced mixed findings. HQNNs demonstrate clear advantages in **low-data regimes**, such as **medical imaging and quantum chemistry**, where classical models struggle with overparameterization [1], [15]. However, for large-scale datasets with extensive labeled examples, classical CNNs remain more stable and efficient due to their well-optimized architectures [24]. The practical deployment of HQNNs will thus require further advancements in **quantum error mitigation, hybrid co-processing architectures, and variational circuit optimizations** [18].

The findings from these studies provide a clear picture of HQNNs' current capabilities and limitations. The next section will discuss ongoing challenges and future directions for improving HQNN architectures. These trade-offs highlight that HQNNs are not yet a drop-in replacement for classical models. Their long-term potential, however, hinges on improvements in quantum hardware, more efficient hybrid architectures, and domain-driven benchmarks — themes we explore in the next section.

V. CHALLENGES AND FUTURE CONSIDERATIONS

A. Challenges in HQNN Implementation

Despite theoretical benefits, HQNNs face implementation challenges that hinder real-world use. These challenges stem from the limitations of current quantum hardware, the com-

plexity of hybrid system integration, and the fragility of quantum states under noise.

1. Quantum Noise and Decoherence. Quantum systems are vulnerable to noise and decoherence, which limit circuit depth. Most near-term devices operate with coherence times in the microsecond range, often too short for meaningful learning tasks. This severely constrains the expressive power of VQCs, which are central to HQNN architectures. While error mitigation strategies such as dynamical decoupling and probabilistic error cancellation have shown partial success [12], [16], they often require substantial computational overhead or circuit repetitions that reduce training efficiency.

2. Scalability and Hardware Constraints. The scalability of HQNNs is tightly coupled to hardware availability and reliability. Current devices have limited qubits and connectivity, restricting model size and complexity. As a result, most HQNN research remains simulator-bound or focused on shallow circuits tested on small QPU instances. The leap from toy problems to full-scale applications in areas like genomics or finance remains constrained by this hardware bottleneck [25].

3. Quantum-Classical Bottlenecks. HQNNs require frequent back-and-forth communication between classical and quantum components during both training and inference. Each forward pass involves executing a quantum circuit, measuring outcomes, and using classical optimizers to update parameters. This hybrid loop introduces latency and memory transfer delays, particularly when executed over cloud-based QPUs. The need for repeated circuit sampling to produce statistically stable gradients further amplifies these inefficiencies [18]. Some recent strategies, such as batch processing and quantum-

aware caching, attempt to mitigate this, but results remain hardware-dependent.

4. Training Instability and Barren Plateaus. Training HQNNs presents unique optimization challenges. Variational circuits often suffer from barren plateaus, where gradients vanish across large regions of the parameter space, stalling learning [1], [12]. Additionally, noise in measurement interferes with gradient estimation, particularly in gradient-based optimizers like Adam or RMSprop. Selecting a suitable ansatz — the structure of the quantum circuit — is non-trivial and problem-specific. Over-parameterized circuits can become unstable, while under-parameterized ones lack sufficient expressiveness.

5. Lack of Standardization. There is currently no unified framework for benchmarking HQNNs across applications. Results are often reported using different datasets, hardware simulators, and optimization schemes, making it difficult to compare models or reproduce experiments. This fragmentation slows down community-wide progress and highlights the need for common benchmarking protocols tailored to quantum-classical hybrid architectures [12], [18].

Taken together, these challenges explain why HQNNs, despite their conceptual promise, are not yet practical alternatives to classical deep learning for most production settings. Overcoming them will require coordinated progress in quantum hardware design, noise-resilient training algorithms, and better hybrid integration schemes. The next subsection outlines promising directions currently being explored to address these issues.

B. Future Directions

To ~~unlock~~ the full potential of HQNNs, future research must address limitations in hardware reliability, training scalability, and hybrid model integration. While theoretical results point to strong advantages in efficiency and expressivity, practical deployment requires coordinated progress across quantum hardware, algorithmic optimization, and application engineering.

Advancements in Quantum Hardware: A key enabler for scalable HQNNs will be the maturation of quantum hardware beyond the Noisy Intermediate-Scale Quantum (NISQ) era. Ongoing efforts by IBM, Google, and Rigetti aim to deliver fault-tolerant quantum processors with higher gate fidelities and longer decoherence windows [12]. These advances would allow HQNNs to support deeper quantum circuits, richer entanglement structures, and more expressive feature encodings — expanding their applicability to large-scale learning problems.

Hybrid co-processing architectures, where quantum processors are embedded into high-bandwidth classical systems, will also be essential for reducing quantum-classical communication delays [18]. Early prototypes of on-chip quantum accelerators and quantum RAM (QRAM) promise to significantly streamline data movement between the classical and quantum domains.

Improved Training and Optimization Strategies: To address the training instabilities faced by HQNNs, future work should

explore quantum-informed optimization techniques designed specifically for VQCs. Several recent efforts focus on:

- **Quantum Natural Gradient Descent (QNG)**, which adapts learning rates based on the quantum state’s geometry to avoid barren plateaus [12].
- **Noise-aware optimization**, which integrates error models into the training loop to improve convergence in noisy environments [24].
- **Dynamic circuit pruning**, which removes unnecessary quantum gates mid-training to reduce depth and decoherence risks [25].

Transfer learning, where pretrained classical or hybrid models are fine-tuned using quantum layers, is also gaining traction. This approach reduces the quantum workload while maintaining accuracy, making it ideal for domains with limited access to QPU time [20].

Expansion into New Application Domains: While HQNNs have shown success in various fields, their use in fields like real-time encryption, financial forecasting, and autonomous control remains largely unexplored. These domains demand fast, low-latency decision-making on high-dimensional data — a niche where quantum-enhanced architectures could offer clear advantages.

One promising future direction is **quantum reinforcement learning (QRL)**, where HQNNs could serve as policy networks capable of learning from entangled state transitions in dynamic environments. Additionally, applications in privacy-preserving learning — such as quantum multi-classifiers using differential privacy techniques [8] — could bring quantum learning closer to secure AI deployment.

Toward Fully Quantum Neural Networks: In the long term, HQNNs may serve as a bridge toward fully quantum neural networks (QNNs), where the entire learning process — from input encoding to output classification — occurs in the quantum domain [26]. Achieving this vision will require advances in quantum memory, coherent activation functions, and loss function evaluation within quantum circuits.

Figure 5 provides an overview of anticipated advancements in HQNN development, highlighting key improvements in **hardware reliability, quantum-classical processing efficiency, and enhanced training techniques**. These innovations will be crucial for overcoming existing limitations and unlocking the full potential of quantum-enhanced deep learning.

Until then, HQNNs remain a pragmatic stepping stone. As quantum hardware stabilizes and optimization pipelines mature, hybrid networks are well-positioned to deliver practical quantum advantage in near-term AI workloads.

C. The Path Forward

As quantum hardware matures and error mitigation techniques improve, HQNNs may become a practical alternative to classical deep learning models. However, significant research is still needed to bridge the gap between theoretical advantages and real-world applicability. Future work should focus on refining hybrid architectures, improving quantum training methodologies, and integrating HQNNs into practical computing environments.

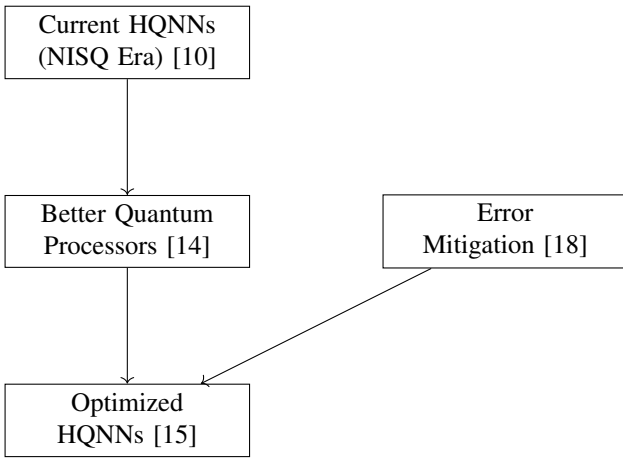


Fig. 5. Future advancements in HQNNs, adapted from [10], [14], [15], [18]. Improvements in quantum hardware and error mitigation strategies will enable more scalable and efficient HQNN architectures.

The final section will summarize the key insights from this paper and provide conclusions regarding the future of HQNN research.

VI. CONCLUSION

HQNNs represent a compelling frontier in quantum-enhanced artificial intelligence. This paper has surveyed the theoretical underpinnings, performance characteristics, and practical challenges of HQNNs, revealing both their strengths and current limitations. HQNNs outperform classical models in training efficiency, parameter use, and feature generalization—especially on high-dimensional, low-labeled data.

However, the road to widespread adoption is not without obstacles. Quantum noise, short coherence times, and hybrid bottlenecks continue to restrict HQNN scalability. These hardware-level constraints, combined with training inefficiencies such as barren plateaus, necessitate more robust error mitigation and quantum-aware optimization strategies. Furthermore, comparative studies suggest that HQNNs outperform classical models primarily in constrained or noise-tolerant settings, indicating that hybrid architectures must be carefully matched to task requirements.

Looking forward, progress hinges on: (1) better quantum processor stability, (2) hybrid integration, and (3) scalable benchmarking reflecting real deployment. As these technical barriers are addressed, HQNNs are likely to evolve from proof-of-concept tools into deployable AI accelerators across healthcare, chemistry, security, and other high-impact domains. Consequently, as HQNNs become more integrated into medical or financial decision-making, transparency and interpretability will be critical for ethical deployment.

Ultimately, HQNNs may serve as an essential bridge between classical neural computation and fully quantum learning systems. While their current capabilities are bound by the limits of NISQ-era hardware, the foundational techniques they introduce—including VQCs, entangled feature encoding, and hybrid optimization—will play a critical role in shaping the trajectory of next-generation quantum machine learning.

REFERENCES

- [1] E. H. Houssein, Z. Abohashima, M. Elhoseny, and W. M. Mohamed, "Hybrid quantum-classical convolutional neural network model for COVID-19 prediction using chest x-ray images," *Journal of Computational Design and Engineering*, vol. 9, no. 2, pp. 343–363, 2022.
- [2] M. Islam, M. J. Hasan, and M. R. C. Mahdy, "Cq cnn: A hybrid classical quantum convolutional neural network for alzheimer's disease detection using diffusion generated and u net segmented 3d mri," 2025. [Online]. Available: <https://arxiv.org/abs/2503.02345>
- [3] A. Muniasamy, S. A. S. Alquhtani, A. H. Alshehri, A. Begum, and A. Sabahath, "Investigating hybrid quantum-assisted classical and deep learning model for mri brain tumor classification," *Journal of Image and Graphics*, vol. 13, no. 1, pp. 123–129, 2025. [Online]. Available: <https://www.joig.net/show-98-440-1.html>
- [4] V. Spadari, I. Guarino, D. Ciunzio, and A. Pescapè, "Towards network intrusion detection via quantum machine learning: A reality check," in *IEEE INFOCOM Workshops: Quantum Networked Applications and Protocols (QuNAP 2025)*, 02 2025.
- [5] H. Cho, J. Kim, and H. Lim, "Hybrid quantum neural networks with variational quantum regressor for enhancing qspr modeling of co2-capturing amine," 2025. [Online]. Available: <https://arxiv.org/abs/2503.00388>
- [6] M. Akrom, S. Rustad, T. Sutojo, W. A. E. Prabowo, H. K. Dipojono, R. Maezono, and H. Kasai, "Stacking classical-quantum hybrid learning approach for corrosion inhibition efficiency of n-heterocyclic compounds," *Results in Surfaces and Interfaces*, vol. 18, p. 100462, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666845925000492>
- [7] V. Oliveira Santos, P. A. Costa Rocha, J. V. G. Thé, and B. Gharabaghi, "Optimizing the architecture of a quantum-classical hybrid machine learning model for forecasting ozone concentrations: Air quality management tool for houston, texas," *Atmosphere*, vol. 16, no. 3, 2025. [Online]. Available: <https://www.mdpi.com/2073-4433/16/3/255>
- [8] H. Chen, C. Wang, J. Du, C.-H. H. Yang, and J. Qi, "Projection valued-based quantum machine learning adapting to differential privacy algorithm for word-level lipreading," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [9] C. Liang, S. Tao, X. Huang, Y. Wang, B. Xia, and X. Zhang, "Stochastic state of health estimation for lithium-ion batteries with automated feature fusion using quantum convolutional neural network," *Journal of Energy Chemistry*, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2095495625001780>
- [10] C. Hughes, J. Isaacson, A. Perry, R. F. Sun, and J. Turner, *Quantum Computing for the Quantum Curious*. Springer, 2021.
- [11] F. V. Massoli, L. Vadicamo, G. Amato, and F. Falchi, "A leap among quantum computing and quantum neural networks: A survey," *ACM Comput. Surv.*, vol. 55, no. 5, Dec. 2022. [Online]. Available: <https://doi.org/10.1145/3529756>
- [12] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, "Training deep quantum neural networks," *Nature communications*, vol. 11, no. 1, p. 808, 2020.
- [13] A. Sahakyan, "Introduction to quantum machine learning and quantum neural networks," *arXiv preprint arXiv:2211.06554*, 2022.
- [14] S. Y.-C. Chen, C.-M. Huang, C.-W. Hsing, and Y.-J. Kao, "An end-to-end trainable hybrid classical-quantum classifier," *Machine Learning: Science and Technology*, vol. 2, no. 4, p. 045021, 2021.
- [15] R. Xia and S. Kais, "Hybrid quantum-classical neural network for calculating ground state energies of molecules," *Entropy*, vol. 22, no. 8, p. 828, Jul. 2020. [Online]. Available: <http://dx.doi.org/10.3390/e22080828>
- [16] N. Nguyen, E. Behrman, and J. Steck, "Quantum learning with noise and decoherence: a robust quantum neural network," *Quantum Machine Intelligence*, vol. 2, pp. 1–15, 01 2020.
- [17] Y. Kwak, W. J. Yun, S. Jung, and J. Kim, "Quantum neural networks: Concepts, applications, and challenges," 2021. [Online]. Available: <https://arxiv.org/abs/2108.01468>
- [18] K. Zaman, T. Ahmed, M. A. Hanif, A. Marchisio, and M. Shafique, "A comparative analysis of hybrid-quantum classical neural networks," 2024. [Online]. Available: <https://arxiv.org/abs/2402.10540>
- [19] L. Bischof, S. Teodoropol, R. M. Füchslin, and K. Stockinger, "Hybrid quantum neural networks show strongly reduced need for free parameters in entity matching," *Scientific Reports*, vol. 15, p. 4318, 2025. [Online]. Available: <https://doi.org/10.1038/s41598-025-88177-z>

- [20] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," *Quantum*, vol. 4, p. 340, Oct. 2020. [Online]. Available: <http://dx.doi.org/10.22331/q-2020-10-09-340>
- [21] M. A. Hafeez, A. Munir, and H. Ullah, "H-qnn: A hybrid quantum-classical neural network for improved binary image classification," *AI*, vol. 5, no. 3, pp. 1462–1481, 2024. [Online]. Available: <https://www.mdpi.com/2673-2688/5/3/70>
- [22] J. Zhang, G. Zheng, T. Koike-Akino, K.-K. Wong, and F. A. Burton, "Hybrid quantum-classical neural networks for downlink beamforming optimization," *IEEE Transactions on Wireless Communications*, vol. 23, no. 11, pp. 16 498–16 512, 2024.
- [23] H.-Y. Chen, Y.-J. Chang, S.-W. Liao, and C.-R. Chang, "Hybrid quantum neural network in high-dimensional data classification," 2023. [Online]. Available: <https://arxiv.org/abs/2312.01024>
- [24] D. Ranga, S. Prajapat, Z. Akhtar, P. Kumar, and A. V. Vasilakos, "Hybrid quantum-classical neural networks for efficient mnist binary image classification," *Mathematics*, vol. 12, no. 23, 2024. [Online]. Available: <https://www.mdpi.com/2227-7390/12/23/3684>
- [25] Y. Zhang and H. Lu, "Reliability research on quantum neural networks," *Electronics*, vol. 13, no. 8, 2024. [Online]. Available: <https://www.mdpi.com/2079-9292/13/8/1514>
- [26] S. Abel, J. C. Criado, and M. Spannowsky, "Completely quantum neural networks," *Physical Review A*, vol. 106, no. 2, Aug. 2022. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevA.106.022601>