

# Next-Generation Quantum Neural Networks: Enhancing Efficiency, Security, and Privacy

Nouhaila Innan<sup>\*†</sup>, Muhammad Kashif<sup>\*†</sup>, Alberto Marchisio<sup>\*†</sup>, Mohamed Bennai<sup>‡</sup>, Muhammad Shafique<sup>\*†</sup>

<sup>\*</sup> eBrain Lab, Division of Engineering, New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

<sup>†</sup> Center for Quantum and Topological Systems, NYUAD Research Institute, New York University Abu Dhabi, UAE

<sup>‡</sup> Quantum Physics and Spintronics Team, LPMC, Faculty of Sciences Ben M'sick, Hassan II University of Casablanca, Morocco

Emails: {nouhaila.innan, muhammadkashif, alberto.marchisio, muhammad.shafique}@nyu.edu, mohamed.bennai@univh2c.ma

**Abstract**—This paper provides an integrated perspective on addressing key challenges in developing reliable and secure Quantum Neural Networks (QNNs) in the Noisy Intermediate-Scale Quantum (NISQ) era. In this paper, we present an integrated framework that leverages and combines existing approaches to enhance QNN efficiency, security, and privacy. Specifically, established optimization strategies, including efficient parameter initialization, residual quantum circuit connections, and systematic quantum architecture exploration, are integrated to mitigate issues such as barren plateaus and error propagation. Moreover, the methodology incorporates current defensive mechanisms against adversarial attacks. Finally, Quantum Federated Learning (QFL) is adopted within this framework to facilitate privacy-preserving collaborative training across distributed quantum systems. Collectively, this synthesized approach seeks to enhance the robustness and real-world applicability of QNNs, laying the foundation for reliable quantum-enhanced machine learning applications in finance, healthcare, and cybersecurity.

**Index Terms**—Quantum Neural Networks, Quantum Machine Learning, Quantum Federated Learning

## I. INTRODUCTION

Quantum Neural Networks (QNNs) have emerged as a promising paradigm at the intersection of quantum computing and machine learning, offering potential advantages in processing complex data structures and solving computationally intensive tasks [1]. In the current Noisy Intermediate-Scale Quantum (NISQ) era, characterized by quantum processors with limited qubit counts and susceptibility to noise [2], hybrid quantum-classical QNNs are envisioned to leverage quantum mechanics to achieve high performance in specific applications [3]–[5].

Despite their theoretical potential, the practical deployment of QNNs faces several significant challenges inherent to NISQ devices. One of the critical issues is the occurrence of barren plateaus in the optimization landscape, where gradients vanish exponentially with system size [6], which hinders the effective training of variational quantum circuits. Additionally, the limited number of qubits and their short coherence times constrain the scalability of QNNs, making it challenging to handle high-dimensional data. The presence of noise and decoherence further complicates the reliable execution of quantum circuits, affecting the expressibility and robustness of QNN models [7]. Moreover, concerns regarding the security and privacy of quantum machine learning systems, especially in adversarial settings, remain largely unexplored.

To address these challenges, we propose a comprehensive cross-layer methodology aimed at enhancing the efficiency, security, and privacy of QNNs in the NISQ era. Our contributions are as follows:

- **Trainability and Scalability Optimization:** We deploy techniques to mitigate barren plateaus, such as layer-wise training and parameter initialization strategies, and explore quantum circuit cutting methods to enable scalable QNN architectures on limited qubit devices.
- **Noise-Aware Design and Architecture Exploration:** We conduct an in-depth analysis of QNN robustness against various noise models, including phase flip, bit flip, and depolarizing channels, and guide the design of robust QNNs.
- **Security and Privacy Enhancements:** We investigate the vulnerability of QNNs to adversarial attacks and their respective defense mechanisms. Furthermore, we explore the integration of federated learning and encryption techniques to ensure data privacy in distributed QNN scenarios.
- **Application Outlook:** We provide an overview of emerging applications for next-generation QNNs, highlighting their potential impact in fields such as intelligent transportation, finance, and healthcare.

## II. METHODOLOGY

We present an end-to-end framework designed to support the development of next-generation QNNs, with a focus on enhancing their trainability, scalability, and ensuring security and privacy. An overview of our proposed methodology is illustrated in Fig. 1.

### A. Trainability and Scalability

**Trainability** is a key challenge in developing QNNs, primarily due to barren plateaus (BP), regions in the optimization landscape where the gradient variance vanishes exponentially with system size [6], rendering gradient-based training infeasible (see Fig. 2).

BPs are linked to random initialization [8], global cost functions [9], hardware noise [10], and the expressibility of parameterized quantum circuits (PQCs) [11]. Mitigation strategies include careful parameter initialization, residual learning, and noise-assisted training. In particular, Xavier initialization has shown superior performance in reducing BP effects [8], as illustrated in Fig. 3, while narrower initialization ranges further enhance trainability [12]. Residual connections, inspired by

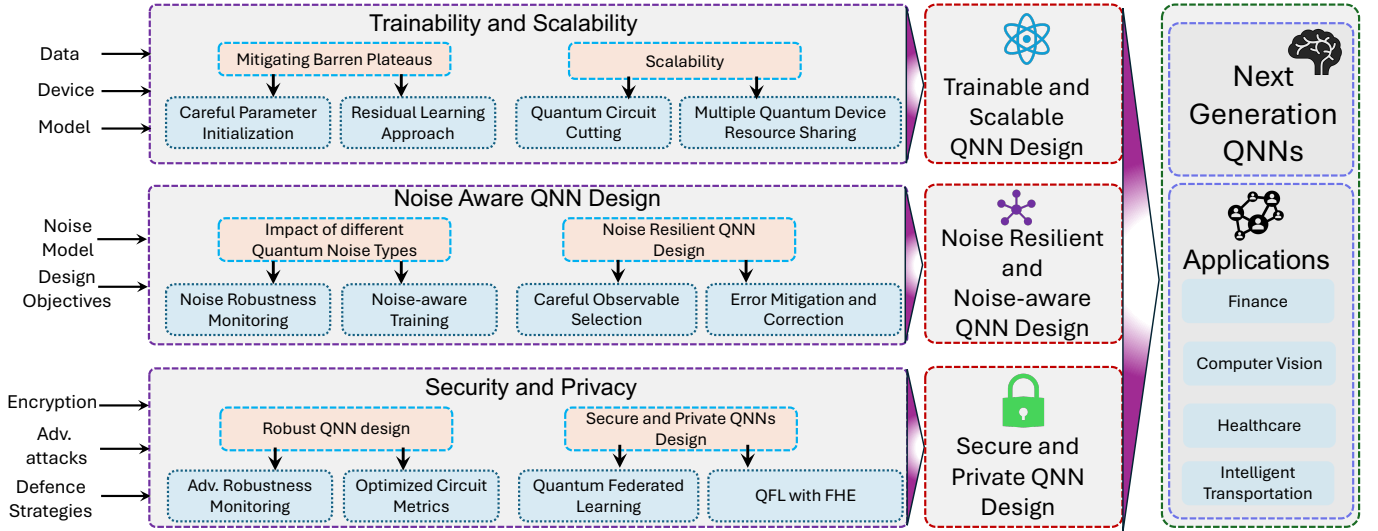


Fig. 1: Overview of our methodology to design efficient, robust, and secure QNNs.




Fig. 2: Demonstration of Barren Plateaus.




Fig. 3: Variance decay and training results with different initialization techniques [8].

classical deep networks, preserve gradient flow and enable deeper QNNs [13], [14].

**Scalability** is constrained by limited qubit counts and noise in NISQ hardware. Current QNN architectures are restricted to small-scale models that fit within the available quantum resources. To overcome this, quantum circuit cutting has been proposed [15], allowing large circuits to be decomposed into smaller subcircuits. As shown in Fig. 4, a 6-qubit circuit can be executed using 4-qubit subcircuits. Intermediate measurement results are stored and reused in subsequent subcircuits. Importantly, this approach preserves model accuracy while enabling deployment on limited hardware.

#### B. Noise-Aware Design and QNN Architecture Exploration

In the NISQ era, integrating hardware noise considerations into QNN design is vital for practical deployment. Noise-aware training and architecture exploration are essential to identify robust configurations that perform reliably under realistic noise conditions.

The work in [16] presents a detailed study on how different




Fig. 4: Quantum circuit representation when cutting a 6-qubit circuit into four 4-qubit subcircuits [15].

quantum noise types affect QNN architectures, specifically Quantumvolational Neural Networks (QuanNNs) and Quantum Convolutional Neural Networks (QCNNs), using real-world datasets. As shown in Fig. 5(a), QuanNNs demonstrate robustness against phase and bit flip noise but degrade under depolarizing noise. In contrast, Fig. 5(b) shows that QCNNs maintain high accuracy on MNIST under amplitude damping noise but struggle with the more complex Fashion-MNIST, illustrating dataset-dependent resilience. This analysis underscores the importance of selecting noise-resilient architectures and informs the design of QNNs better suited to noisy environments. It also highlights when quantum error mitigation or correction becomes crucial to maintain model reliability.

In [17], the impact of quantum noise on QNN trainability is investigated, revealing that BPs arise more readily in noisy settings. The study shows that selecting suitable measurement observables, particularly a custom Hermitian observable aligned with the learning objective, can substantially enhance QNN resilience and trainability across various noise types.

#### C. Security and Privacy

After achieving trainability, scalability, and noise resilience, ensuring the security and privacy of QNNs is crucial. Without this, even the most robust architectures remain vulnerable to adversarial threats and data leakage, a critical aspect in domains like healthcare, finance, and cybersecurity.

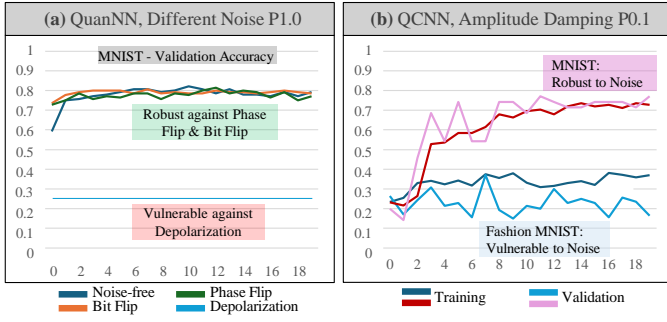


Fig. 5: Noise robustness of QNNs under different noise types and tasks. (a) QuanNN’s noise robustness under different noise channels with probability 1.0, for the MNIST dataset. (b) QCNN’s noise robustness under Amplitude Damping noise with probability 0.1, for the MNIST and Fashion MNIST datasets [16].

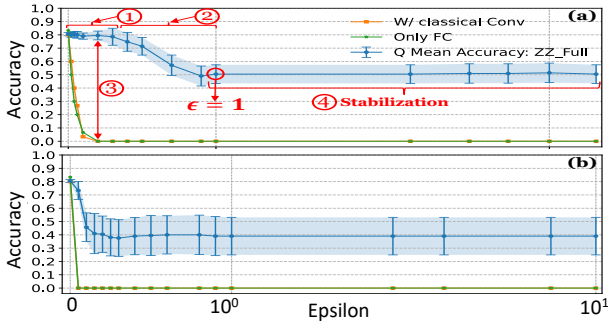


Fig. 6: Robustness evaluation of classical and quantum models using the ZZ full ansatz against adversarial attacks on the MNIST dataset: (a) FGSM and (b) PGD, tested across varying perturbation strengths [18].

### 1) Embedding Adversarial Robustness in QNN Design

Neural networks, including hybrid quantum-classical models, are susceptible to adversarial inputs. Thus, adversarial robustness should be embedded from the design phase. This includes analyzing circuit-level properties such as expressibility, entanglement, and gate configurations. Circuits with high expressibility and controlled entanglement, particularly those using Z-axis controlled rotations, demonstrate improved resistance to adversarial attacks [18]–[20].

Rather than relying on post-training defenses, a preemptive adversarial testing phase, simulating attacks like FGSM and PGD, should guide architectural choices. As shown in [21], this helps identify circuits that generalize well even under perturbations. Empirical results show that such circuits achieve up to 60% greater robustness on datasets like MNIST and Fashion-MNIST at low perturbation levels (see Fig. 6).

### 2) Ensuring Privacy through Quantum Federated Learning and Encryption

QNNs deployed in sensitive domains require decentralized, privacy-preserving training. We adopt a federated learning paradigm, where QNNs are trained locally and only encrypted model updates, not raw data, are shared.

Quantum Federated Learning (QFL) [22] enables distributed quantum nodes to train locally while maintaining data privacy. Each client uses quantum hardware or simulators to train QNNs and transmits only quantum-encoded parameters for global aggregation (see Fig. 7). This setup is ideal for data-




Fig. 7: Execution of QNNs on real quantum hardware within a federated learning setup, demonstrating stable performance across different IBM Quantum backends [22].

sensitive sectors, such as healthcare and finance.

To further protect updates, we integrate Fully Homomorphic Encryption (FHE) [23], enabling encrypted aggregation without exposing plaintext parameters. Though FHE can reduce performance, quantum-enhanced models, especially on multi-modal data, help retain accuracy. Across diverse datasets, this QFL-FHE pipeline maintains over 70% test accuracy while preserving privacy (see Table I).

TABLE I: QFL with FHE experiment results for various datasets.

Dataset	Test Loss	Train Accuracy	Test Accuracy	Time (sec)
CIFAR-10	0.0937	97.90%	71.12%	9747.32 ± 2.23
DNA Sequence	0.782	100.00%	94.32%	7123.91 ± 2.91
MRI Scan	0.360	100.00%	88.75%	7851.86 ± 3.54
PCOS	1.090	100.00%	70.15%	3942.60 ± 1.65
RAVDESS	0.83	94.53%	76.43%	1140.76 ± 1.69
DNA+MRI Multimodal	DNA: 0.174 MRI: 0.713	DNA: 99.64% MRI: 100%	DNA: 95.31% MRI: 87.26%	10314.34 ± 6.28

Together, these methods form a robust and secure QNN development pipeline. With adversarial defense and encrypted federated learning, QNNs are now capable of trustworthy deployment in sensitive real-world applications.

## III. APPLICATIONS

QNNs have shown significant promise across various domains, demonstrating practical advantages in tasks that demand high accuracy, security, and computational efficiency. Following the development of scalable, noise-aware, and secure QNN architectures, recent research has begun translating these models into real-world applications. In the financial sector, QNNs have been successfully deployed for fraud detection and loan eligibility prediction [24], [25]. Privacy-preserving frameworks that integrate federated learning with quantum layers have achieved precision rates exceeding 95%, even under noisy and distributed settings [26]. QNNs trained on structured financial data have also achieved up to 98% accuracy in predicting loan approvals, aided by dropout mechanisms and robust quantum circuit design [27]. In quantum information processing, QNN-based models have enhanced the efficiency of quantum state tomography by reducing the number of required measurements without compromising reconstruction fidelity. These advances are particularly beneficial for scaling to larger quantum systems [28]. While in healthcare, QNNs have enabled precise multi-omics integration for lung cancer classification, uncovering key biomarkers with exceptional diagnostic accuracy [29]. In intelligent transporta-

tion systems, QNNs have been used to process large-scale traffic data, achieving classification accuracies above 97% and demonstrating strong robustness under noise, highlighting their potential for deployment in urban mobility infrastructures [30]. Similarly, hybrid quantum-classical models with attention mechanisms have been employed for image super-resolution, offering competitive quality while reducing parameter counts, thus aligning with current hardware limitations in the NISQ era [31]. These applications underscore the versatility and practical relevance of QNNs, reinforcing their potential as a foundational technology across sectors where performance, privacy, and resilience are paramount.

#### IV. CONCLUSION AND OUTLOOK

In this paper, we presented an integrated, cross-layer methodology for advancing QNNs in the NISQ era. By synthesizing optimization techniques, architecture exploration, noise-aware design, adversarial robustness, and federated privacy-preserving learning, our framework systematically addresses the core limitations of QNNs related to trainability, scalability, and trustworthiness. We demonstrated how approaches such as parameter initialization, residual connections, quantum circuit cutting, and noise-informed design choices can significantly improve model performance on current quantum hardware. Furthermore, we established a foundation for secure and privacy-aware QNN training via adversarial testing and QFL augmented with encryption techniques.

Looking forward, as quantum hardware continues to evolve and mature, our approach offers a scalable path toward deploying robust and secure quantum machine learning models in real-world applications. Future work will involve deeper integration of quantum error correction techniques, exploration of hardware-efficient QNN architectures, and automated quantum architecture search to further streamline development. Additionally, expanding federated and privacy-enhanced quantum learning to edge and cloud environments will be key to enabling widespread adoption. Ultimately, this work serves as a step toward realizing the next-generation QNN systems capable of transforming critical sectors such as healthcare, finance, and intelligent infrastructure.

#### ACKNOWLEDGMENT

This work was supported in part by the NYUAD Center for Quantum and Topological Systems (CQTS), funded by Tamkeen under the NYUAD Research Institute grant CG008, and the Center for Cyber Security (CCS), funded by Tamkeen under the NYUAD Research Institute Award G1104.

#### REFERENCES

- [1] K. Zaman *et al.*, “A survey on quantum machine learning: Current trends, challenges, opportunities, and the road ahead,” *arXiv preprint arXiv:2310.10315*, 2023.
- [2] J. Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, 2018.
- [3] K. Zaman *et al.*, “A comparative analysis of hybrid-quantum classical neural networks,” in *World Congress in Computer Science, Computer Engineering & Applied Computing*, 2024.
- [4] M. Kashif, A. Marchisio, and M. Shafique, “Computational advantage in hybrid quantum neural networks: Myth or reality?,” in *DAC*, 2025.
- [5] M. Kashif and S. Al-Kuwari, “Design space exploration of hybrid quantum-classical neural networks,” *Electronics*, 2021.
- [6] J. R. McClean *et al.*, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications*, vol. 9, nov 2018.
- [7] M. Kashif *et al.*, “Investigating the effect of noise on the training performance of hybrid quantum neural networks,” in *IJCNN*, 2024.
- [8] M. Kashif *et al.*, “Alleviating barren plateaus in parameterized quantum machine learning circuits: Investigating advanced parameter initialization strategies,” in *DATE*, 2024.
- [9] M. Kashif and S. Al-Kuwari, “The impact of cost function globality and locality in hybrid quantum neural networks on nisq devices,” *Machine Learning: Science and Technology*, vol. 4, no. 1, p. 015004, 2023.
- [10] M. Kashif and M. Shafique, “Hqnet: Harnessing quantum noise for effective training of quantum neural networks in nisq era,” *arXiv:2402.08475*, 2024.
- [11] M. Kashif and S. Al-Kuwari, “The unified effect of data encoding, ansatz expressibility and entanglement on the trainability of hqnn,” *International Journal of Parallel, Emergent and Distributed Systems*, vol. 38, no. 5, pp. 362–400, 2023.
- [12] M. Kashif and M. Shafique, “The dilemma of random parameter initialization and barren plateaus in variational quantum algorithms,” 2024.
- [13] M. Kashif and S. Al-Kuwari, “Resqnets: a residual approach for mitigating barren plateaus in quantum neural networks,” *EPJ Quantum Technology*, vol. 11, no. 1, p. 4, 2024.
- [14] M. Kashif and M. Shafique, “Resqnns: towards enabling deep learning in quantum convolution neural networks,” *arXiv:2402.09146*, 2024.
- [15] A. Marchisio *et al.*, “Cutting is all you need: Execution of large-scale quantum neural networks on limited-qubit devices,” *arXiv preprint arXiv:2412.04844*, 2024.
- [16] T. Ahmed *et al.*, “Noisy hqnn: A comprehensive analysis of noise robustness in hybrid quantum neural networks,” in *IJCNN*, 2025.
- [17] M. Kashif and M. Shafique, “Nrqn: The role of observable selection in noise-resilient quantum neural networks,” in *World Congress in Computer Science, Computer Engineering & Applied Computing*, 2024.
- [18] W. E. Maouaki *et al.*, “Advqnn: A methodology for analyzing the adversarial robustness of quantum neural networks,” in *QSW*, 2024.
- [19] W. E. Maouaki *et al.*, “Robqnns: A methodology for robust quantum neural networks against adversarial attacks,” in *ICIPCW*, 2024.
- [20] W. E. Maouaki *et al.*, “Designing robust quantum neural networks: Exploring expressibility, entanglement, and control rotation gate selection for enhanced quantum models,” *arXiv preprint arXiv:2411.11870*, 2024.
- [21] W. E. Maouaki *et al.*, “Designing robust quantum neural networks via optimized circuit metrics,” *Advanced Quantum Technologies*, 2025.
- [22] N. Innan *et al.*, “FedQNN: Federated learning using quantum neural networks,” in *IJCNN*, 2024.
- [23] S. Dutta *et al.*, “MQFL-FHE: Multimodal quantum federated learning framework with fully homomorphic encryption,” *arXiv preprint arXiv:2412.01858*, 2024.
- [24] N. Innan *et al.*, “Financial fraud detection: a comparative study of quantum machine learning models,” *International Journal of Quantum Information*, 2024.
- [25] M. El Alami *et al.*, “Comparative performance analysis of quantum machine learning architectures for credit card fraud detection,” *arXiv preprint arXiv:2412.19441*, 2024.
- [26] N. Innan *et al.*, “QFNN-FFD: Quantum federated neural network for financial fraud detection,” *arXiv preprint arXiv:2404.02595*, 2024.
- [27] N. Innan *et al.*, “LEP-QNN: Loan eligibility prediction using quantum neural networks,” *arXiv preprint arXiv:2412.03158*, 2024.
- [28] N. Innan *et al.*, “Quantum state tomography using quantum machine learning,” *Quantum Machine Intelligence*, 2024.
- [29] M. K. Saggi and S. Kais, “Mqml: Multi-omic quantum machine learning based cancer classification, biomarker identification in human lung adenocarcinoma,” in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1, pp. 1713–1720, IEEE, 2024.
- [30] N. Innan *et al.*, “QNN-VRCS: A quantum neural network for vehicle road cooperation systems,” *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [31] S. Dutta *et al.*, “QUIET-SR: Quantum image enhancement transformer for single image super-resolution,” *arXiv preprint arXiv:2503.08759*, 2025.