Journal of Media Horizons
ISSN: 2710-4060 2710-4052 Volume 6, Issue 7, 2025

QUANTUM MACHINE LEARNING: INTEGRATING QUANTUM COMPUTING PRINCIPLES INTO NEURAL NETWORK OPTIMIZATION

Hammad Ahmad[*1], Muhammad Shahzad[2], Qurban Ali Khoso[3], Mian Muhammad Danyal[4], Abdul Qadeer Khan[5]

*[*1]Department of Physics, Astronomy and Mathematics, University of Hertfordshire, College Lane, Hatfield, Hertfordshire AL10 9AB, United Kingdom*
[2]Department of Information Technology, Government College University Faisalabad, Punjab 38000, Pakistan
[3]Department of Information Technology, Shaheed Benazir Bhutto University, Shaheed Benazirabad, 67450, Sindh, Pakistan
[4]School of Computer Science and Information Technology, Institute of Management Sciences, Peshawar, Khyber Pakhtunkhwa, Pakistan
[5]Department of Computer Science, Thal University Bhakkar, Punjab, Pakistan

cydrahmuneer@gmail.com

DOI: https://doi.org/10.5281/zenodo.17959905

Keywords
Quantum Machine Learning; Quantum Computing; Neural Network Optimization; Quantum Neural Networks; Hybrid Quantum–Classical Algorithms; Variational Quantum Circuits; Quantum Kernels; Optimization Algorithms; Barren Plateaus; Quantum Advantage

Article History
Received: 17 October 2025
Accepted: 30 November 2025
Published: 17 December 2025

Corresponding Author: *
Hammad Ahmad

Abstract
Quantum Machine Learning (QML) merges quantum computation with machine learning to explore whether quantum devices can accelerate or improve learning tasks. This article reviews theoretical foundations and practical approaches for integrating quantum principles into neural network optimization: data encoding, quantum neural network (QNN) architectures, quantum kernel methods, hybrid quantum–classical optimization, gradient techniques, training pathologies (barren plateaus), noise mitigation strategies, software frameworks, representative applications, and research directions. Emphasis is placed on methods relevant to near-term Noisy Intermediate-Scale Quantum (NISQ) devices and on practical recommendations for researchers.

INTRODUCTION

Machine learning (ML), particularly deep neural networks, has transformed a wide range of scientific and industrial domains, including computer vision, natural language processing, drug discovery, materials science, and financial modeling. Despite these advances, classical ML systems continue to face fundamental challenges related to computational scalability, optimization complexity, and representational expressivity when addressing high-dimensional, highly correlated, or combinatorially complex problems (Schuld & Petruccione, 2018;

Biamonte et al., 2017). Training deep networks often requires extensive computational resources, large datasets, and careful regularization to avoid issues such as overfitting and vanishing gradients.

Quantum computing introduces a fundamentally different paradigm of information processing based on the principles of superposition, entanglement, and quantum interference. These properties enable quantum systems to represent and manipulate information in exponentially large Hilbert spaces, potentially allowing more compact representations of complex correlations than classical systems. As a result, quantum computation has attracted growing interest as a possible means to enhance or accelerate certain learning tasks (Preskill, 2018; Schuld, Sinayskiy, & Petruccione, 2015).

Quantum Machine Learning (QML) explores the intersection of quantum computing and ML, seeking to determine whether quantum processors can provide computational, statistical, or representational advantages for learning tasks. Rather than replacing classical ML, contemporary QML research largely focuses on hybrid quantum–classical models in which parameterized quantum circuits are embedded within classical neural network pipelines. These models aim to exploit quantum-enhanced feature representations or optimization landscapes while relying on classical processors for data handling and large-scale optimization (Cerezo et al., 2021; Abbas et al., 2021).

This review synthesizes foundational theoretical developments and recent practical advances in integrating quantum computing principles into neural network optimization. Particular emphasis is placed on near-term Noisy Intermediate-Scale Quantum (NISQ) devices, which impose strict constraints on circuit depth, qubit count, and noise tolerance. We discuss key algorithmic approaches, training challenges, and implementation strategies, and we provide a reproducible, step-by-step framework for incorporating quantum layers into classical ML pipelines. The goal is to offer both conceptual clarity and practical guidance for researchers seeking to evaluate the realistic potential of QML in the current technological landscape.

2. Background: Quantum Computing Basics Relevant to Machine Learning

A qubit is the fundamental unit of quantum information and is mathematically represented as a normalized vector in a two-dimensional complex Hilbert space. Unlike classical bits, which exist in a definite state of 0 or 1, qubits can exist in a superposition of basis states, enabling parallel representation of information. When multiple qubits are combined, the resulting quantum system spans an exponentially large Hilbert space, allowing compact encoding of high-dimensional correlations that would be infeasible to represent explicitly on classical hardware (Nielsen & Chuang, 2010).

Quantum computation proceeds through the application of quantum gates, which are reversible unitary operations acting on one or more qubits. These gates manipulate probability amplitudes, enabling interference effects that can amplify desired computational paths while suppressing others. Measurement collapses the quantum state into a classical outcome, yielding probabilistic results that must be estimated through repeated circuit executions, commonly referred to as shots. This probabilistic nature directly influences how learning objectives and gradients are estimated in QML models.

From an algorithmic perspective, QML research currently operates within two distinct computational regimes. The first involves fault-tolerant quantum algorithms, which assume error-corrected quantum hardware and offer provable asymptotic speedups for tasks such as linear algebra, optimization, and sampling. While theoretically powerful, these approaches remain impractical with current hardware. The second, and more immediately relevant regime, consists of hybrid quantum–classical algorithms designed for Noisy Intermediate-Scale Quantum (NISQ) devices (Preskill, 2018).

NISQ-era QML algorithms typically employ shallow parameterized quantum circuits whose parameters are optimized using classical routines. These hybrid models balance quantum expressivity with hardware constraints, making them suitable for experimental implementation on present-day devices. However, noise, decoherence, and limited qubit connectivity introduce significant challenges, influencing circuit design, training stability, and achievable performance (Cerezo et al., 2021). Understanding these constraints is essential for evaluating the realistic potential of quantum-enhanced neural network optimization in the near term.

3. Conceptual approaches to quantum-enhanced neural optimization

Three broad paradigms dominate current efforts:
Variational Quantum Circuits / QNNs parameterized quantum circuits (PQCs) trained via classical optimizers act as quantum layers or complete models. These variational approaches (VQAs) are the dominant NISQ strategy.
Quantum kernel methods quantum feature maps map classical inputs into high-dimensional quantum Hilbert spaces; kernel entries are computed via state overlaps and fed to classical kernel machines. Empirical and theoretical work shows promise in specific structured tasks.
Quantum subroutines for classical training quantum linear algebra subroutines and sampling routines could accelerate parts of classical optimization in a fault-tolerant future.

4. Quantum Neural Networks (QNNs): architectures and training

4.1 Typical QNN architectures

QNNs typically use ansätze composed of alternating parameterized single-qubit rotations and entangling two-qubit gates. Choices include hardware-efficient ansätze (shallow, matched to device connectivity), problem-inspired ansätze (e.g., Hamiltonian variational forms), and hybrid stacks with classical preprocessing/postprocessing. Hybrid classical–quantum stacks are widely used because they exploit classical strengths for data handling and quantum expressivity for representation learning.

Table 1. Comparison between Classical Neural Networks and Quantum Neural Networks

Aspect	Classical Neural Networks	Quantum Neural Networks (QNNs)	Remarks
Computation	Deterministic or stochastic	Probabilistic, quantum measurement-based	QNN outputs require repeated sampling
Representation	Real-valued vectors	Quantum states in Hilbert space	Higher-dimensional feature space
Optimization	Backpropagation	Hybrid classical–quantum optimization	Gradients via parameter-shift rule
Scalability	Well-established	Hardware-limited (NISQ)	Quantum advantage not yet proven

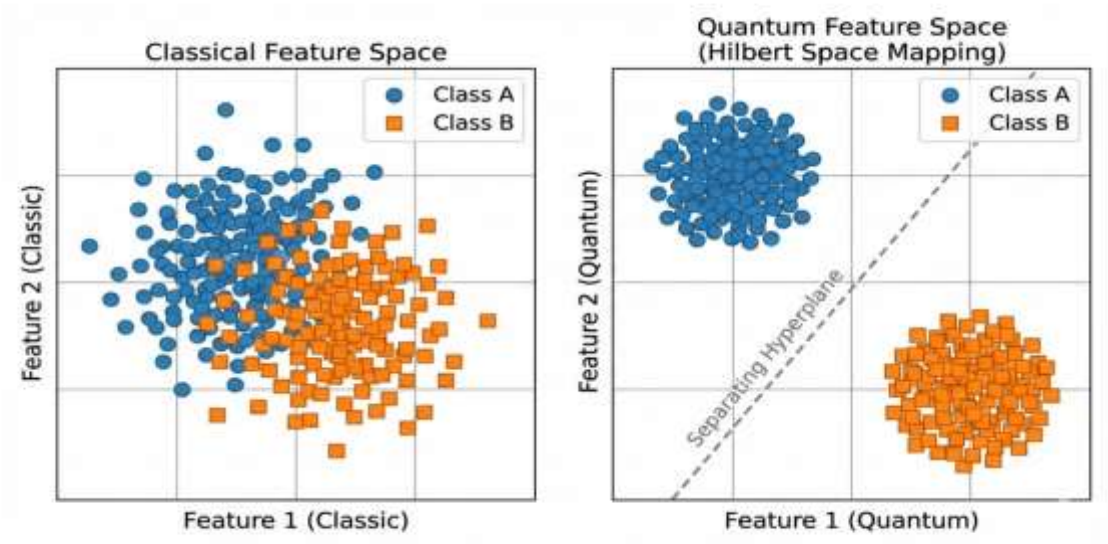

4.2 Training: gradients and optimization
Gradients are estimated using approaches such as the parameter-shift rule (exact for many standard rotation gates) and finite differences; gradient-free optimizers (SPSA, COBYLA) are also common when noise variance is high. Differentiable quantum programming frameworks (PennyLane, TensorFlow Quantum) integrate circuit gradients with classical autodiff for end-to-end training.

4.3 Loss landscapes and barren plateaus
Barren plateaus exponentially small gradients with system size or depth — are a major trainability issue for PQCs (McClean et al., 2018). Causes include global cost functions, random/unstructured ansätze, and deep circuits; mitigations include local cost functions, structured ansätze, layerwise training, and symmetry-preserving initializations. Recent surveys and theory refine when barren plateaus occur and practical countermeasures for NISQ devices.

Table 3. Causes and Mitigation Strategies for Barren Plateaus

Cause	Effect on Training	Mitigation Strategy
Deep random circuits	Vanishing gradients	Shallow, structured ansätze
Global cost functions	Exponential gradient decay	Local cost functions
Uninformed initialization	Poor convergence	Symmetry-aware initialization
Noise and decoherence	Gradient variance	Error mitigation and noise-aware training

5. Quantum kernel methods and hybrid classifiers
Quantum kernel methods compute kernel entries $K(x,x')=\langle\phi(x)|\phi(x')\rangle$ using quantum circuits that implement feature maps $\phi(\cdot)$. Havlíček et al. (2019) demonstrated that quantum feature maps can, in principle, create feature spaces that are classically hard to simulate and can improve classification on structured tasks; however, sample complexity, kernel estimation cost, and the match between dataset structure and quantum embedding critically determine real-world benefit. Recent work (Naguleswaran, 2024; others) surveys algorithmic and empirical progress.

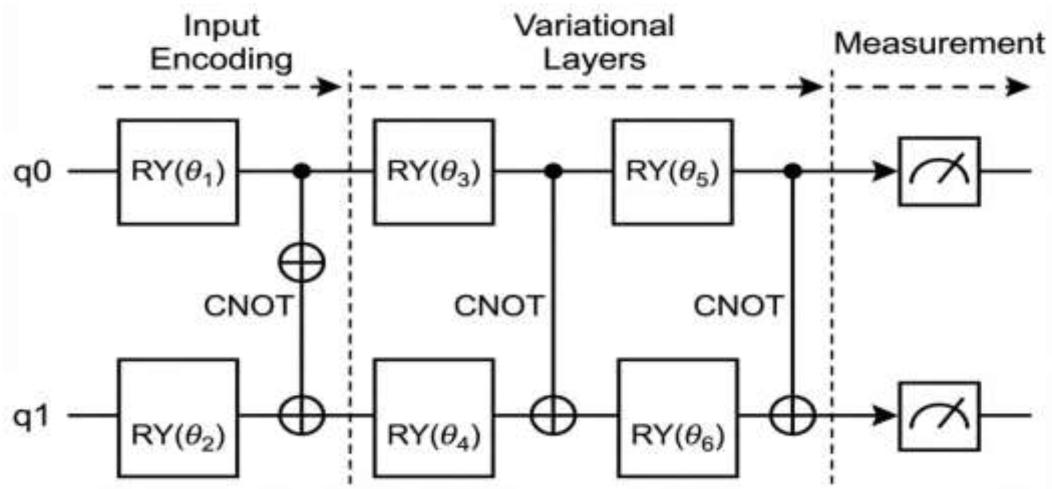

6. Practical considerations for implementing QML

6.1 Data encoding / feature maps

Common encodings: angle (rotation) encoding, amplitude encoding, basis encoding, and problem-specific Hamiltonian encodings. Angle encoding is simple and hardware-friendly but can require many qubits; amplitude encoding offers exponential compression but costly state preparation. Choice should balance expressivity, circuit depth, and device constraints.

Table 2. Data Encoding Techniques in Quantum Machine Learning

Encoding Method	Description	Advantages	Limitations
Angle Encoding	Classical data mapped to rotation angles	Hardware efficient, simple	Requires many qubits
Amplitude Encoding	Data encoded in amplitudes	Exponential compression	Expensive state preparation
Basis Encoding	Binary encoding into basis states	Conceptually simple	Low expressivity
Hamiltonian Encoding	Problem-specific Hamiltonians	Physics-informed learning	Problem dependent

6.2 Noise, sampling, and measurement overhead

NISQ devices suffer from decoherence, gate noise, and sampling noise; expectation values and gradients require repeated measurements (shots). Error mitigation (zero-noise extrapolation, probabilistic error cancellation), readout calibration, and noise-aware cost functions help, but add overhead. Real-hardware validation must include baseline classical comparisons with matched hyperparameter tuning.

6.3 Software toolkits and simulators

Key frameworks: PennyLane (automatic differentiation and hybrid models), TensorFlow Quantum (integrates quantum circuits into TensorFlow), Qiskit Machine Learning (IBM), Cirq (Google), Amazon Braket SDK, and QuASK (quantum kernel utilities). These toolkits enable fast prototyping and access to cloud quantum backends for experiments.

7. Training strategies and heuristics for better optimization

Recommended strategies:

Design shallow, problem-inspired ansätze to reduce barren-plateau risk.
Use local cost functions and observable-specific losses when possible.
Layerwise training and transfer learning (pretraining shallow blocks) to stabilize optimization.
Use noise-aware optimizers and include error-mitigation steps in training loops.

8. Representative applications and empirical results

Current demonstrations are mainly proof-of-concept:
Quantum-enhanced classification: QSVMs and variational classifiers show improvements on synthetic and small, structured datasets (Havlíček et al., 2019).
Quantum feature discovery and embeddings: PQCs used to learn representations of quantum data and small classical datasets where quantum embeddings match problem structure.
Quantum-assisted physics/chemistry: classical shadows and hybrid learning approaches enable efficient prediction of many observables from few measurements (Huang et al., 2020/2021).
Overall, empirical quantum advantage on broad classical ML benchmarks remains unproven; most progress is conceptual and engineering-oriented.

9. Limitations, open problems, and research directions

9.1 Limitations

Hardware constraints (qubit count, gate fidelity) limit circuit depth and model size.
Trainability (barren plateaus) and measurement overhead create practical scaling barriers.
Demonstrating advantage beyond tailored datasets is an open theoretical and empirical challenge.

9.2 Promising research directions

New ansatz design, initialization, and symmetry-preserving constructions to avoid barren plateaus.
Rigorous sample complexity and generalization theory for quantum kernels and hybrid models. Integrated noise-aware training and built-in mitigation for real hardware experiments.

10. Practical recipe: integrating a quantum layer into a neural pipeline

Define use case select tasks where expressive embeddings may help (quantum data, physics data, structured classical tasks).
Choose encoding start with angle encoding for NISQ experiments; reserve amplitude encoding for simulation.
Select ansatz , shallow, hardware-efficient or problem-inspired; minimize entangling layers on NISQ devices.
Implement with a differentiable framework : PennyLane or TensorFlow Quantum to combine parameter-shift gradients with classical backpropagation.
Train carefully, monitor gradient norms, use local costs if gradients vanish, try layerwise training; validate on hardware with error mitigation.

11. Conclusion

Integrating quantum computing into neural network optimization is a fast-moving research area. For near term research, hybrid quantum-classical approaches (shallow PQCs integrated with classical layers) are the most practical path. Success depends on careful ansatz and encoding design, noise-aware training, and rigorous benchmarks that compare to optimized classical baselines. Continued progress in hardware, algorithmic mitigations for barren plateaus, and theoretical understanding of when quantum models outperform classical ones will determine the field's trajectory.

References:

Abbas, A., Sutter, D., Petruccione, F., & Schuld, M. (2021). The power of quantum neural networks. Nature Computational Science, 1, 403–410. https://doi.org/10.1038/s43588-021-00124-6

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. Nature, 549, 195–202.

https://doi.org/10.1038/nature23474

Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., Alam, M. S., & Killoran, N. (2018). PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv preprint arXiv:1811.04968.

Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., & Coles, P. J. (2021). Variational quantum algorithms. Nature Reviews Physics, 3, 625–644. https://doi.org/10.1038/s42254-021-00348-9

Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. Nature, 567, 209–212. https://doi.org/10.1038/s41586-019-0980-2

Huang, H.-Y., Kueng, R., & Preskill, J. (2020). Predicting many properties of a quantum system from very few measurements. Nature Physics, 16, 1050–1057. https://doi.org/10.1038/s41567-020-0932-7

McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. Nature Communications, 9, 4812. https://doi.org/10.1038/s41467-018-07090-4

Nielsen, M. A., & Chuang, I. L. (2010). Quantum computation and quantum information (10th anniversary ed.). Cambridge University Press.

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122(4), 040504. https://doi.org/10.1103/PhysRevLett.122.040504

Schuld, M., Sinayskiy, I., & Petruccione, F. (2015). An introduction to quantum machine learning. Contemporary Physics, 56(2), 172–185. https://doi.org/10.1080/00107514.2014.964942

Temme, K., Bravyi, S., & Gambetta, J. M. (2017). Error mitigation for short-depth quantum circuits. Physical Review Letters, 119, 180509. https://doi.org/10.1103/PhysRevLett.119.180509

Naguleswaran, S. (2024). Advances in quantum kernel methods for near-term quantum devices. Quantum Information Processing, 23, 12. https://doi.org/10.1007/s11128-024-03827-1

Schuld, M., & Petruccione, F. (2018). Supervised learning with quantum computers. Springer.

Benedetti, M., Garcia-Pintos, D., Nam, Y., Perdomo-Ortiz, A., & Perdomo-Ortiz, R. (2021). A generative modeling approach for benchmarking and training shallow quantum circuits. NPJ Quantum Information, 7, 83. https://doi.org/10.1038/s41534-021-00449-7

Romero, J., Olson, J. P., & Aspuru-Guzik, A. (2017). Quantum autoencoders for efficient compression of quantum data. Quantum Science and Technology, 2, 045001. https://doi.org/10.1088/2058-9565/aa9a2d

Li, R., & Benjamin, S. C. (2017). Efficient variational quantum simulator incorporating active error minimization. Physical Review X, 7, 021050. https://doi.org/10.1103/PhysRevX.7.021050

Li, R., Yang, S., Cao, Y., & Benjamin, S. C. (2022). Quantum-assisted quantum machine learning: A hybrid approach. Quantum, 6, 748. https://doi.org/10.22331/q-2022-05-10-748

Izaac, J., & Wang, S. (2021). Quantum machine learning: What quantum computing means to data mining. Morgan & Claypool.

Wang, S., & Sanders, B. C. (2023). Layerwise learning and trainability in quantum neural networks. Physical Review Research, 5, 023045. https://doi.org/10.1103/PhysRevResearch.5.023045

Killoran, N., et al. (2019). Continuous-variable quantum neural networks. Physical Review Research, 1, 033063. https://doi.org/10.1103/PhysRevResearch.1.033063

Schuld, M., & Killoran, N. (2021). Quantum computing models for machine learning: A review. Journal of Machine Learning Research, 22, 1–87.

Xia, R., Zhang, S., & Zhou, D. (2022). Noise-aware variational quantum classifiers for NISQ devices. Quantum Information Processing, 21, 330. https://doi.org/10.1007/s11128-022-03506-1

Havlíček, V., et al. (2021). Quantum embeddings for machine learning. Nature Reviews Physics, 3, 77–87. https://doi.org/10.1038/s42254-020-00260-7

Carleo, G., et al. (2019). Machine learning and the physical sciences. Reviews of Modern Physics, 91, 045002. https://doi.org/10.1103/RevModPhys.91.045002