

# Quantum-Enhanced Machine Learning: Algorithms and Challenges in the Noisy Intermediate-Scale Quantum Era

Anna Patricia Tan Torres\*

Department of Computer Science, Ateneo de Manila University, Quezon City, Philippines

\*Corresponding author: Anna Patricia Tan Torres, [annaptt@obf.ateneo.edu](mailto:annaptt@obf.ateneo.edu)

## Abstract

The rapid advancement of quantum computing hardware and software has ushered in a new era for computational science, promising to tackle problems that are intractable for classical computers. Parallely, machine learning (ML) has revolutionized data analysis and artificial intelligence. The confluence of these two fields, known as Quantum Machine Learning (QML), aims to harness the principles of quantum mechanics—such as superposition, entanglement, and interference—to enhance and redefine classical machine learning algorithms. This survey provides a comprehensive overview of the burgeoning field of QML. We begin by elucidating the fundamental quantum concepts that underpin QML approaches. We then present a detailed taxonomy and analysis of prominent QML algorithms, categorizing them into quantum-enhanced classical models and fully quantum models. Key families of algorithms discussed include Quantum Support Vector Machines (QSVM), Quantum Neural Networks (QNNs), and Quantum Generative Adversarial Networks (QGANs). For each, we explain the theoretical quantum advantage, the circuit implementation, and the current limitations. Furthermore, we delve into the critical challenges facing the practical realization of QML, including the pervasive issue of noise in Noisy Intermediate-Scale Quantum (NISQ) devices, the problem of barren plateaus in training quantum circuits, and the complexities of efficient data encoding (quantum feature maps). We also present a small-scale experimental simulation of a variational quantum classifier on a standardized dataset, demonstrating its principle and comparing its performance with a classical counterpart. Finally, the paper discusses the future trajectory of QML, highlighting the importance of error mitigation, hybrid quantum-classical architectures, and the development of quantum-specific learning inductive biases. We conclude that while fault-tolerant quantum computers remain on the horizon, QML represents a paradigm shift with the potential to deliver profound computational advantages in the coming decades.

## Keywords

Quantum Machine Learning, Quantum Computing, Quantum Algorithms, Variational Quantum Circuits, NISQ, Quantum Supremacy, Hybrid Algorithms

## 1. Introduction

The last decade has witnessed an unprecedented growth in both computational power and data generation. Machine learning, a subset of artificial intelligence, has been at the forefront of extracting valuable insights from this deluge of data, driving innovations in fields ranging from natural language processing to drug discovery. However, the computational demands of training increasingly complex models on massive datasets are pushing the limits of classical von Neumann architecture, leading to escalating energy costs and extended research timelines.

Simultaneously, quantum computing has emerged from the realm of theoretical physics into tangible, albeit nascent, hardware. Leveraging the counterintuitive principles of quantum mechanics, quantum computers process information in a fundamentally different way. While a classical bit is strictly a 0 or a 1, a quantum bit, or qubit, can exist in a superposition of both states simultaneously. Furthermore, qubits can be entangled, creating profound correlations that have no classical analogue. These properties grant quantum computers the potential for massive parallelism, offering exponential or polynomial speedups for specific classes of problems, such as integer factorization (Shor's algorithm) and unstructured database search (Grover's algorithm) [1].

The natural question arises: can the computational prowess of quantum computers be harnessed to accelerate and improve machine learning tasks? This intersection defines the field of Quantum Machine Learning (QML). The promise of QML is twofold: first, to achieve a computational speedup in executing existing ML algorithms, and second, to discover entirely new learning models that are native to the quantum world, potentially capable of recognizing patterns in quantum data that are inaccessible to classical methods.

This paper aims to provide a systematic survey of the current landscape of QML. We will explore the foundational principles, categorize and explain the major algorithmic families, critically assess the practical challenges—particularly in the NISQ era—and outline promising future research directions [2]. Our contribution is a structured synthesis of this

rapidly evolving field, designed to be accessible to computer scientists with a basic understanding of machine learning and quantum computing.

The remainder of this paper is organized as follows: Section 2 covers the essential background in quantum computing and machine learning. Section 3 presents a detailed survey of QML algorithms. Section 4 discusses the significant challenges and limitations. Section 5 describes a simple experimental setup and results. Section 6 explores future directions, and Section 7 concludes the work.

## 2. Background


To understand QML algorithms, a firm grasp of certain foundational concepts is necessary.

### 2.1 Fundamental Quantum Computing Concepts

- **Qubit and Superposition:** A qubit is the fundamental unit of quantum information. Its state  $|\psi\rangle$  is a vector in a two-dimensional complex Hilbert space and can be represented as  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , where  $|0\rangle$  and  $|1\rangle$  are the computational basis states, and  $\alpha$  and  $\beta$  are complex probability amplitudes satisfying  $|\alpha|^2 + |\beta|^2 = 1$ . Upon measurement, the qubit collapses to  $|0\rangle$  with probability  $|\alpha|^2$  or to  $|1\rangle$  with probability  $|\beta|^2$ . The ability to be in a combination of states is superposition [3].

- **Entanglement:** Entanglement is a strong correlation between qubits such that the state of one qubit cannot be described independently of the state of the others. A famous example is the Bell state:  $(|00\rangle + |11\rangle)/\sqrt{2}$ . Measuring the first qubit and finding it to be 0 instantly forces the second qubit into the 0 state, regardless of the physical distance between them.

- **Quantum Gates and Circuits:** Quantum operations are performed by unitary transformations called quantum gates. Examples include the Pauli-X (bit-flip), Hadamard (H) gate that creates superposition, and controlled-NOT (CNOT) gate that creates entanglement. A sequence of gates applied to a set of qubits forms a quantum circuit, the quantum analogue of a classical logic circuit. A simplified quantum circuit is shown in:



**Figure 1.** Quantum circuit for generating a bell state.

Figure 1 show a simple quantum circuit creating a Bell state. The Hadamard (H) gate puts the first qubit into superposition, and the CNOT gate entangles it with the second qubit. Measurement (M) collapses the state.

- **Measurement:** The process of observing a quantum system, which causes its wavefunction to collapse to a definite classical state. Measurement is probabilistic and is typically the final step in a quantum algorithm.

### 2.2 Machine Learning Preliminaries

Machine learning models can be broadly categorized into:

- **Supervised Learning:** Learning a mapping from inputs to outputs using a labeled dataset (e.g., classification, regression) [4].
- **Unsupervised Learning:** Finding hidden patterns or intrinsic structures in unlabeled data (e.g., clustering, dimensionality reduction).
- **Reinforcement Learning:** An agent learns to make decisions by performing actions and receiving rewards from an environment.

### 3. A Taxonomy of Quantum Machine Learning Algorithms

QML algorithms can be coarsely divided into two categories: those that are quantum versions of classical algorithms and those that are inherently quantum.

#### 3.1 Quantum-Enhanced Classical Models

These algorithms use quantum subroutines to speed up specific computational bottlenecks within a classical ML pipeline.

- Quantum Linear Algebra for ML: Many ML algorithms (e.g., Support Vector Machines, Principal Component Analysis) rely on linear algebra operations, particularly solving systems of linear equations or finding eigenvalues. The Harrow-Hassidim-Lloyd (HHL) algorithm solves a system of linear equations exponentially faster than the best-known classical algorithms, under certain conditions. This can be used to speed up the training of a Quantum Support Vector Machine (QSVM). The core idea is that the task of finding the hyperplane for an SVM can be framed as solving a linear system [5]. By using HHL, the QSVM can, in principle, achieve an exponential speedup in the feature space dimension. The process involves encoding the kernel matrix in a quantum state, using HHL to solve for the support vectors, and then using quantum state tomography to read out the result. However, the practical utility is currently limited by the resource requirements of HHL and the noise in current hardware.

- Quantum Annealing for Optimization: Quantum annealers, such as those built by D-Wave Systems, are specialized quantum computers designed to find the global minimum of complex optimization landscapes. This is directly applicable to unsupervised learning tasks like clustering (e.g., k-means) and Boltzmann machine training. The objective function of the ML problem is mapped to the energy landscape of the quantum annealer's Hamiltonian. The system naturally evolves towards its ground state, which corresponds to the optimal solution of the original problem [6].

#### 3.2 Fully Quantum Models

These models are natively designed for quantum processors and often employ a hybrid quantum-classical training loop.

- Variational Quantum Algorithms (VQAs) and Quantum Neural Networks (QNNs): VQAs are the leading paradigm for the NISQ era [7]. They use a parameterized quantum circuit (PQC), often called an ansatz, as their core. This PQC is analogous to a neural network, hence the term Quantum Neural Network. A typical VQA workflow, illustrated in figure as follows:




Figure 2. Hybrid quantum-classical workflow of a Variational Quantum Algorithm (VQA).

Figure 2 show the iterative hybrid workflow of a Variational Quantum Algorithm. Parameters are tuned on a classical computer based on the cost function evaluated from the results of a parameterized quantum circuit.

- 1.Data Encoding (Embedding): Classical input data  $x$  is encoded into a quantum state using a feature map circuit  $U(x)$ . Common techniques include basis encoding, amplitude encoding, and angle encoding.
- 2.Parameterized Quantum Circuit (Ansatz): A sequence of quantum gates  $U(\theta)$ , with tunable parameters  $\theta$ , is applied. The structure of this circuit is critical and can be hardware-efficient or problem-inspired [8].
- 3.Measurement: The quantum state is measured, yielding a classical output (e.g., an expectation value of an observable).
- 4.Classical Optimization: The measurement output is used to compute a cost function  $C(\theta)$ . A classical optimizer (e.g., gradient descent, Adam) is used to adjust the parameters  $\theta$  to minimize this cost.
- 5.The loop (steps 2-4) is repeated until convergence.

QNNs have been applied to both classification and regression tasks. Their power is theorized to come from their ability to create highly complex, entangled states in a high-dimensional feature space that is difficult for classical computers to represent.

- Quantum Generative Adversarial Networks (QGANs): GANs are a class of generative models where a generator and a discriminator are trained in an adversarial game. In a QGAN, either the generator, the discriminator, or both can be replaced by quantum circuits [9]. A quantum generator may have an advantage in learning the underlying distribution of quantum data or certain complex classical distributions. The training follows a similar min-max optimization as classical GANs, but the gradients are computed with respect to the parameters of the quantum circuits.
- Quantum Boltzmann Machines (QBMs): Boltzmann Machines are generative models based on statistical mechanics. QBMs extend this concept by defining the energy of the model using a quantum Hamiltonian, which includes non-commuting terms, allowing it to represent more complex distributions than its classical counterpart. Training QBMs, however, remains challenging due to the difficulty of approximating the gradients of the quantum log-likelihood.

**Table 1.** Summary of key QML algorithms and their promised advantages.

Algorithm Category	Key Example	Promised Advantage	Key Challenge
Quantum-Enhanced	QSVM (via HHL)	Exponential speedup in feature space dimension	Resource-intensive; sensitive to noise
Quantum-Enhanced	Quantum Annealing for Clustering	Better solutions for non-convex optimization	Limited connectivity; precision issues
Fully Quantum	QNN/VQA	Efficient representation; suitable for NISQ	Barren plateaus; parameter training
Fully Quantum	QGAN	Efficient learning of quantum data	Training instability; mode collapse

Table 1 illustrates how different types of QML algorithms have potential advantages in speed, optimization, data representation, and learning capabilities, but also face challenges such as noise, hardware limitations, training difficulty, and instability. It highlights the potential and practical constraints of quantum machine learning.

### 3.3 The Power of Quantum Feature Spaces and Representation Learning

A fundamental concept underpinning many QML algorithms, particularly QNNs, is the exploitation of high-dimensional quantum feature spaces. This concept warrants a deeper discussion, as it is central to understanding the potential source of quantum advantage in machine learning [10].

In classical machine learning, the "kernel trick" allows algorithms like SVMs to operate in a high-dimensional (or even infinite-dimensional) feature space without explicitly computing the coordinates of the data in that space. Instead, they rely solely on the inner products between the images of all pairs of data in the feature space, computed via a kernel function. Quantum computing offers a natural and potentially more powerful analogue to this [11].

When classical data  $x$  is encoded into a quantum state using a feature map circuit  $U(\varphi(x))|0\rangle^{\hat{n}}$ , it is mapped to a quantum state  $|\varphi(x)\rangle$  in a  $2\hat{n}$ -dimensional Hilbert space. The inherent quantum parallelism allows for the computation of inner products and other properties in this vast space. The key insight is that the choice of the feature map  $U_\varphi$  defines the kernel:  $K(x_i, x_j) = |\langle \varphi(x_i) | \varphi(x_j) \rangle|^2$ .

This is the quantum kernel. A quantum computer can estimate this kernel directly by preparing the state  $|\varphi(x_i)\rangle$ , applying the inverse of the preparation circuit for  $|\varphi(x_j)\rangle$ , and measuring the probability of the all-zero state. For complex feature maps that are difficult to simulate classically, this quantum kernel could capture classically intractable similarities between data points.

The power of this approach lies in the expressibility of the feature map. Highly expressive feature maps can create highly entangled and complex states, effectively positioning data points in a feature space where they become more

easily separable, as illustrated in Figure 3. However, this power comes with a caveat. If the feature space is too large or the kernel is too specific, it can lead to overfitting, where the model learns the noise in the training data rather than the underlying pattern. This is linked to the problem of barren plateaus; highly expressive ansätze often exhibit gradients that vanish exponentially with the number of qubits. Therefore, a critical research direction is the design of problem-inspired feature maps that are sufficiently expressive to capture the relevant data structure without inducing barren plateaus or overfitting. This represents a form of quantum representation learning, where the goal is to find the optimal quantum embedding of classical data that facilitates a specific learning task.

#### 4. Critical Challenges in Quantum Machine Learning

The theoretical promise of QML is tempered by significant practical hurdles.

##### 4.1 The NISQ Era and Hardware Limitations

We are currently in the Noisy Intermediate-Scale Quantum era. NISQ devices typically have 50-1000 qubits but lack full error correction. They are plagued by decoherence (loss of quantum information), gate infidelities, and readout errors [12]. These noise sources severely limit the depth and complexity of quantum circuits that can be reliably run, often drowning out the desired signal. Error mitigation techniques, such as zero-noise extrapolation, are active areas of research but are not a panacea.

##### 4.2 The Barren Plateau Problem

A major obstacle in training VQAs is the barren plateau phenomenon. As the number of qubits increases, the gradient of the cost function vanishes exponentially for a large class of random parameterized quantum circuits. This makes it incredibly difficult for classical optimizers to find a direction for improvement, effectively stalling the training. Designing problem-specific ansätze with clever initializations is a key strategy to mitigate this [13].

##### 4.3 Data Encoding (Quantum Feature Maps)

Efficiently loading classical data into a quantum computer is a non-trivial problem, often referred to as the "input problem". The choice of encoding (basis, amplitude, angle) has profound implications for the algorithm's performance and resource requirements. Amplitude encoding can be very efficient (encoding  $2^n$  numbers in  $n$  qubits) but is difficult to prepare and can lead to exponentially small gradients (barren plateaus). Angle encoding is more common in NISQ algorithms but requires more qubits.

##### 4.4 The Practical Quantum Advantage Debate

A central and unresolved question is when, or if, a QML algorithm will demonstrate a clear, practical advantage over the best classical methods for a real-world, classically challenging problem [14]. While theoretical speedups exist, they often rely on assumptions that are not yet met by hardware, such as fault-tolerance or access to specific forms of quantum RAM (QRAM). Demonstrating a practical quantum advantage in machine learning remains the "holy grail" of the field.

#### 5. Experimental Illustration: A Variational Quantum Classifier


To ground the theoretical discussion, we implemented a simple variational quantum classifier (VQC) on a classical simulator to demonstrate the principle.

##### 5.1 Setup

- Dataset: We used the Iris dataset, a standard benchmark in ML, but restricted it to two classes (Setosa vs. Versicolor) and two features (sepal length and petal length) for simplicity and visualization.
- Quantum Simulator: We used the PennyLane library with a default qubit simulator to mimic an ideal, noiseless quantum device.
- Circuit Architecture:
  - Qubits: 2 qubits.
  - Encoding: Angle encoding. The two features ( $x_1, x_2$ ) were normalized and used as rotation angles:  $U(x) = RY(x_1) \times RY(x_2)$ . Ansatz: A simple, hardware-efficient ansatz consisting of layers of single-qubit rotations (RY gates) and entangling gates (CNOT). The number of layers was a tunable hyperparameter.
  - Measurement: The expectation value of the Pauli-Z observable on the first qubit was measured. This single value was interpreted as the model's output.
- Training: The cost function was the mean squared error between the quantum model's output and the binary labels. We used the Adam optimizer for classical parameter updates [15].

### 5.2 Results and Discussion


The results of our simulation are shown in:



**Figure 3.** Learned decision boundary of a variational quantum classifier on the iris dataset.

Figure 3 show the learned decision boundary of the variational quantum classifier on the two-class Iris dataset. The model successfully separates the two species (red and blue points) with a non-linear boundary.

The variational quantum classifier was able to successfully learn a decision boundary that separates the two classes of Iris flowers with high accuracy (>95%) on the training set. Figure 4 shows the convergence of the cost function over training iterations, demonstrating the effectiveness of the hybrid training loop:



**Figure 4.** Convergence behavior of the cost function in the variational quantum classifier.

Figure 4 mention convergence plot of the cost function for the variational quantum classifier. The steady decrease demonstrates the effectiveness of the hybrid training loop in minimizing the error.

This experiment serves as a proof-of-concept. However, it is crucial to note that this is a trivial problem for classical ML models (a simple linear classifier would perform perfectly). The quantum circuit here is shallow and was run on a perfect simulator. On real NISQ hardware, noise would likely degrade performance, and for such a simple problem, the

classical method would be vastly more efficient. The value of this experiment is pedagogical, illustrating the workflow and components of a VQC.

## 6. Future Directions

The path forward for QML is rich with research opportunities.

- **Error Mitigation and Correction:** Developing more sophisticated error mitigation techniques is essential for extracting meaningful results from NISQ devices. In the long term, the implementation of fault-tolerant quantum computing via quantum error correction will unlock the full potential of QML algorithms [16].
- **Algorithm-Hardware Co-Design:** Future QML algorithms must be designed with the constraints and strengths of specific quantum hardware architectures in mind [17]. This includes tailoring ansatzes to the native gate set and connectivity graph of a device.
- **Quantum-Centric Supercomputing:** The future likely lies in tight integration, where quantum processors act as accelerators within a larger classical HPC ecosystem. Developing software and middleware for these hybrid systems is a critical challenge [18].
- **Quantum Data:** As quantum sensors and networks advance, we will see an influx of native quantum data. QML models are uniquely positioned to analyze this data, a task for which classical computers are fundamentally ill-suited.
- **New Learning Theory:** A fundamental theory of learning for quantum data and quantum models is still under development. Understanding the generalization bounds, model capacity, and expressibility of QNNs is an active area of research [19].

## 7. Conclusion

Quantum Machine Learning stands at a fascinating crossroads between theoretical computer science, physics, and artificial intelligence. In this survey, we have outlined the foundational principles of quantum computing that empower QML, provided a detailed taxonomy of its major algorithmic families—from quantum-enhanced linear algebra to variational quantum circuits—and highlighted the profound challenges that currently hinder its widespread application, most notably the limitations of NISQ hardware and the barren plateau problem.

Our small-scale experimental demonstration confirms the conceptual validity of a variational quantum classifier, while also underscoring the vast gap between simulation-based proofs-of-concept and demonstrable quantum advantage on real-world problems. The future of QML is intrinsically tied to the advancement of quantum hardware. However, algorithmic innovation, particularly in error mitigation, ansatz design, and the development of a robust quantum learning theory, will be equally critical. While a general-purpose, fault-tolerant quantum computer may be years away, the research conducted today is laying the essential groundwork for a future where quantum computation fundamentally transforms the landscape of machine learning and artificial intelligence.

## References

- [1] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195-202. <https://doi.org/10.1038/nature23474>
- [2] Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- [3] 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(7747), 209-212. <https://doi.org/10.1038/s41586-019-0980-2>
- [4] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., ... & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625-644. <https://doi.org/10.1038/s42254-021-00348-9>
- [5] Schuld, M., & Petruccione, F. (2021). *Machine learning with quantum computers*. Springer. <https://doi.org/10.1007/978-3-030-83098-4>
- [6] Harrow, A. W., Hassidim, A., & Lloyd, S. (2009). Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), 150502. <https://doi.org/10.1103/PhysRevLett.103.150502>
- [7] Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13), 130503. <https://doi.org/10.1103/PhysRevLett.113.130503>
- [8] Albash, T., & Lidar, D. A. (2018). Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1), 015002. <https://doi.org/10.1103/RevModPhys.90.015002>
- [9] McClean, J. R., Romero, J., Babbush, R., & Aspuru-Guzik, A. (2016). The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2), 023023. <https://doi.org/10.1088/1367-2630/18/2/023023>
- [10] Lloyd, S., & Weedbrook, C. (2018). Quantum generative adversarial learning. *Physical Review Letters*, 121(4), 040502. <https://doi.org/10.1103/PhysRevLett.121.040502>
- [11] 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>
- [12] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27. <https://doi.org/10.48550/arXiv.1406.2661>
- [13] Amin, M. H., Andriyash, E., Rolfe, J., Kulchitsky, B., & Melko, R. (2018). Quantum Boltzmann machine. *Physical Review X*, 8(2), 021050. <https://doi.org/10.1103/PhysRevX.8.021050>

- [14] Endo, S., Cai, Z., Benjamin, S. C., & Yuan, X. (2021). Hybrid quantum-classical algorithms and quantum error mitigation. *Journal of the Physical Society of Japan*, 90(3), 032001. <https://doi.org/10.7566/JPSJ.90.032001>
- [15] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812. <https://doi.org/10.1038/s41467-018-07090-4>
- [16] Aaronson, S. (2015). Read the fine print. *Nature Physics*, 11(4), 291-293. <https://doi.org/10.1038/nphys3272>
- [17] Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., Alam, M. S., Ahmed, S., ... & Killoran, N. (2018). PennyLane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint*. <https://doi.org/10.48550/arXiv.1811.04968>
- [18] Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4), R2493. <https://doi.org/10.1103/PhysRevA.52.R2493>
- [19] Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J. M., & Gambetta, J. M. (2017). Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671), 242-246. <https://doi.org/10.1038/nature23879>