Global Journal of Engineering Innovations & Interdisciplinary Research

Original Article

Quantum Machine Learning: A Review of Concepts, Algorithms, And Applications

Vanaparthi Kiranmai[1], Sruthi Thanugundala[2], Allvala Bhagyasree[3]

[1]Assistant Professor, Department of CSE(AI&ML), Guru Nanak Technical Campus, Ibrahimpatnam, Hyderabad
[2]Assistant Professor, Department of AIML, Keshav Memorial Engineering College, Hyderabad
[3]Assistant Professor, Department of Computer Science and Engineering, Kommuri Pratapreddy Institute of Technology, Hyderabad

Correspondence

Vanaparthi Kiranmai

Assistant Professor, Department of CSE(AI&ML), Guru Nanak Technical Campus, Ibrahimpatnam, Hyderabad

- Received Date: 24 Aug 2026
- Accepted Date: 20 Sep 2026
- Publication Date: 26 Sep 2026

Abstract

Quantum Machine Learning (QML) is an emerging interdisciplinary field that combines machine-learning principles with quantum computing to investigate new approaches for data representation, feature transformation, optimization, classification, and pattern recognition. QML exploits quantum phenomena such as superposition, entanglement, interference, and high-dimensional state representations to construct computational models that may provide advantages for selected problems. Major approaches include quantum support vector machines, quantum kernel methods, variational quantum classifiers, quantum neural networks, quantum convolutional neural networks, quantum generative models, and hybrid quantum-classical learning architectures. This review examines the fundamental concepts, algorithmic structures, data-encoding strategies, application areas, implementation procedures, and practical limitations of QML. Particular attention is given to near-term quantum devices, where limited qubit connectivity, noise, circuit depth, measurement overhead, and trainability strongly influence performance. The article also presents graph-ready normalized comparative datasets covering algorithm characteristics, application domains, challenges, and representative architectures. The review indicates that QML has considerable research potential, but practical deployment depends on improved hardware reliability, efficient encoding, robust optimization, standardized benchmarks, and meaningful comparisons with strong classical baselines.

Introduction

Machine learning has become a fundamental computational approach for analyzing large and complex datasets in areas such as healthcare, finance, cybersecurity, engineering, scientific computing, and intelligent automation. Classical machine-learning methods depend on conventional digital processors and mathematical representations of information using bits, vectors, matrices, and tensors. Although modern processors and accelerators have significantly improved computational capacity, several learning tasks continue to encounter challenges involving optimization complexity, high-dimensional feature spaces, large datasets, and computational resource requirements. Quantum computing introduces a fundamentally different information-processing model in which quantum states can represent combinations of computational states and quantum operations can manipulate them through interference and entanglement. These characteristics have motivated researchers to investigate how quantum computation can be integrated with machine learning. [1,2]

Quantum Machine Learning refers broadly to approaches that use quantum computing within one or more stages of a machine-learning workflow. The quantum component may process classical data encoded into quantum states, manipulate inherently quantum data, construct nonlinear feature maps, perform optimization, or generate probability distributions. QML therefore represents a family of methods rather than a single algorithm. Some approaches translate classical models such as support vector machines and nearest-neighbor methods into quantum settings, while others introduce parameterized quantum circuits that function as trainable models. Hybrid quantum-classical systems are especially important because present-day quantum processors generally require classical computers to perform data preparation, parameter optimization, result interpretation, and workflow management. [3,4]

A central motivation for QML is the possibility that quantum feature spaces may represent useful structures that are difficult to reproduce efficiently using classical computation. Quantum kernels, for example, encode classical data into quantum states and estimate similarities between those states, allowing a classical learner such as a support vector machine to operate on a quantum-derived kernel. Variational quantum circuits use trainable parameters to transform encoded data and produce measurement-based outputs.

Citation: Vanaparthi K, Thanugundala S, Allvala B. Quantum Machine Learning: A Review of Concepts, Algorithms, And Applications. GJEIIR. 2026;6(6):236.

Quantum neural networks extend this concept by combining parameterized circuits, quantum gates, measurements, and classical optimization. However, theoretical possibility does not automatically translate into practical advantage. Current evidence emphasizes that hardware limitations, noise, encoding costs, training difficulties, and comparison methodology remain important considerations. [5,6]

Recent reviews demonstrate continued growth of research into QML algorithms and applications. A systematic literature review covering publications from 2017 to 2023 identified 94 relevant studies and reported extensive use of quantum implementations of classical algorithms, quantum neural models, and classification-oriented approaches. Image classification was identified as a particularly active application domain. More recent work has broadened the discussion toward near-term quantum devices, digital health, quantum kernels, trainability, and hybrid architectures. These studies also highlight a persistent gap between small-scale demonstrations and reliable large-scale practical deployment. Consequently, QML research increasingly focuses not only on algorithmic novelty but also on noise robustness, resource efficiency, benchmark design, and the conditions under which quantum methods can provide useful computational value. [7–10]

Literature Review

Early research in QML established the conceptual relationship between quantum computing and learning systems. Foundational work demonstrated that quantum processors could potentially support machine-learning tasks through quantum representations, quantum algorithms, and quantum-inspired learning mechanisms. A major theoretical development was the interpretation of quantum data encoding as a feature map into a Hilbert space. In this framework, a quantum circuit transforms classical input data into quantum states, and the resulting state relationships can be exploited by kernel-based learning techniques. This idea provided a systematic bridge between quantum circuits and established statistical learning theory and became an important foundation for quantum kernel methods and quantum support vector machines. [1,5,11]

Parameterized quantum circuits subsequently became one of the most important mechanisms for developing trainable QML models. In a typical variational architecture, classical input features are encoded through quantum gates, a sequence of parameterized gates transforms the quantum state, and selected measurements produce model outputs. Classical optimization routines iteratively update the circuit parameters using measurement results. This hybrid structure is attractive for near-term devices because it avoids requiring a fully fault-tolerant quantum computer. Variational quantum classifiers and quantum neural networks belong to this family. Their effectiveness depends strongly on ansatz design, parameter initialization, circuit depth, optimization strategy, measurement strategy, and device noise. [3,4,12]

Quantum convolutional neural networks represent another important research direction. Inspired by classical convolutional architectures, QCNNs use local quantum operations, entanglement, pooling-like procedures, and measurements to learn hierarchical representations. They have been investigated for classification of quantum states and selected classical-data problems. Quantum generative models provide another branch of QML research, where parameterized quantum circuits are used to approximate or sample complex probability distributions. Quantum Generative Adversarial Networks combine quantum generators with classical or quantum discriminators, creating a framework for generative modeling. Although these architectures are promising, their practical value depends on the availability of sufficiently controllable quantum hardware and efficient training procedures. [7,13]

A significant portion of contemporary QML research focuses on near-term devices. Current processors are affected by decoherence, gate imperfections, measurement errors, limited connectivity, finite coherence time, and restricted circuit depth. These effects become particularly important when QML models require repeated circuit execution during optimization. A recent review of QML on near-term quantum devices emphasized data encoding, ansatz design, error mitigation, gradient estimation, and comparisons with classical approaches as major practical issues. Consequently, QML algorithms must be designed with hardware characteristics in mind rather than treating the quantum processor as an ideal mathematical abstraction. [9]

Trainability is another major research concern. Variational models can suffer from barren plateaus, where gradients become extremely small and optimization becomes difficult as system size or circuit complexity increases. A comprehensive review published in 2025 identified barren plateaus as a major obstacle in variational quantum computing and showed that the phenomenon can be influenced by architecture, initialization, observables, cost functions, and hardware noise. Quantum kernels also face concentration effects in which kernel values may become insufficiently informative as circuit complexity increases. These findings indicate that simply increasing circuit depth or expressivity is not necessarily beneficial and may instead reduce the usefulness of the learning landscape. [8,14]

Recent systematic research in application areas further demonstrates both promise and limitations. In digital health, QML has been explored for classification, prediction, and analysis of medical data using quantum kernels, QNNs, QCNNs, and hybrid learning systems. However, published studies do not establish a universal quantum advantage over classical models. Similarly, applications in finance, high-energy physics, chemistry, cybersecurity, image analysis, and optimization remain active research areas. Contemporary surveys increasingly emphasize standardized evaluation, reproducibility, fair comparison with classical baselines, resource accounting, and practical deployment requirements. Thus, the direction of QML research is moving from demonstrating that quantum learning is possible toward determining where, when, and under what computational conditions it is useful. [7,10,15]

Methodology

This review adopts a structured analytical methodology for examining the concepts, algorithms, architectures, applications, and practical limitations of QML. Relevant scientific literature was organized according to algorithmic family, quantum data representation, training mechanism, application domain, hardware assumptions, and performance considerations. Particular attention was given to foundational publications, systematic reviews, near-term quantum-device studies, and recent research addressing trainability and quantum-kernel behavior. The methodological objective was not to claim a new experimental benchmark but to synthesize major research directions and construct a consistent comparative framework suitable for academic interpretation and graphical analysis. [3,7,8]

The review first identifies the core relationship between machine learning and quantum information processing. Classical inputs are assumed to undergo preprocessing followed by a quantum encoding operation. Depending on the model, the encoded data may pass through a fixed feature map, a variational circuit, a kernel estimation process, a quantum neural architecture, or a generative circuit. Measurements transform the quantum state into classical outputs, while a classical optimizer may update trainable parameters. This structure allows QML systems to combine quantum processing with well-established machine-learning workflows rather than replacing classical computing completely. For comparative evaluation, this review considers predictive performance, circuit requirements, training complexity, data-encoding overhead, scalability, noise sensitivity, and application readiness. Where numerical comparisons are needed for graphical presentation, normalized illustrative values are used to place different QML concepts on a common visualization scale. These values are not claimed to be results from a single dataset or quantum processor. They are intended to provide a mathematically consistent representation of comparative trends discussed across the literature.

The primary performance metric used in the comparative framework is:

Accuracy = (Number of Correct Predictions / Total Number of Predictions) × 100

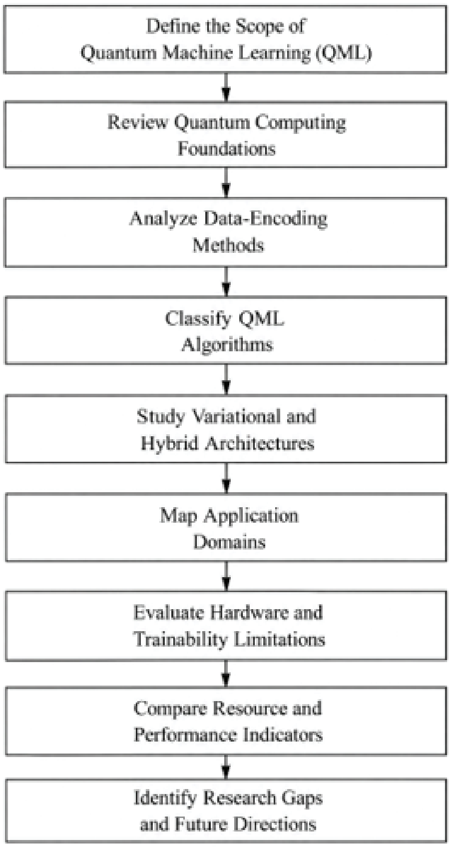

Methodology flow chart

The final synthesis groups QML into three broad implementation classes. The first class consists of quantum versions of classical learning methods, including quantum support vector machines and quantum nearest-neighbor methods. The second consists of trainable parameterized circuits, including variational quantum classifiers, quantum neural networks, and quantum convolutional structures. The third contains hybrid and generative architectures in which quantum and classical processing stages cooperate. This organization supports comparison while recognizing that many contemporary systems combine characteristics from more than one category.

Implementation and Results

Because this is a review article rather than a single experimental study, implementation refers to the construction of a comparative QML evaluation framework. The framework organizes algorithmic characteristics into numerical fields that can be visualized through bar graphs, radar charts, vertical comparisons, and bubble plots. The numerical values below are normalized illustrative values developed specifically for graphical representation. A higher value in a performance or readiness column represents a more favorable relative position, whereas a higher challenge value represents greater difficulty. Therefore, the values should not be interpreted as universal experimental measurements.

Table 1: Normalized comparison of representative QML algorithms

Algorithm	Qubits	Circuit Depth	Training Cost Index	Perfor-mance Index
QSVM	4	12	42	84
VQC	4	18	56	81
QNN	6	20	61	83
QCNN	8	24	68	87
QGAN	8	28	76	78

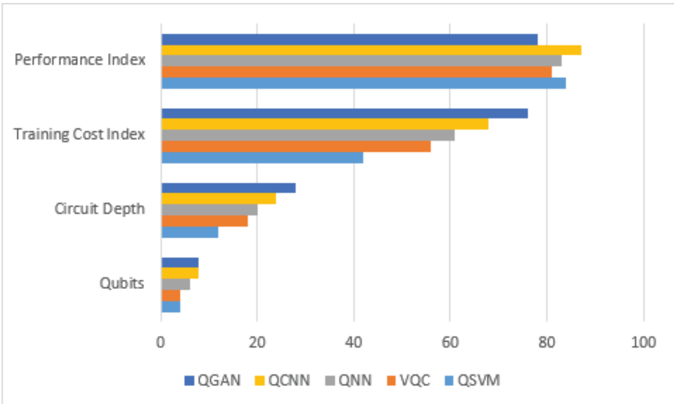

Fig 1: *Horizontal bar graph comparing algorithm performance indices*

The data in Table 1 indicate a general trade-off between architectural complexity and computational requirements. QCNN is assigned the highest illustrative performance index, but it also requires greater circuit depth and training cost than simpler models. QSVM exhibits a relatively low training-cost index and moderate qubit requirements, reflecting the attraction of quantum-kernel approaches for structured classification tasks. VQC and QNN provide flexible learning mechanisms but require repeated parameter optimization. QGAN receives a lower illustrative performance index because generative modeling generally introduces additional training complexity. These values are useful for producing a comparative graph, not for asserting superiority of one algorithm under all datasets or hardware conditions.

The horizontal bar graph represented by Fig 1 should place algorithm names on the vertical axis and the performance index on the horizontal axis. The five bars can therefore be plotted directly using values 84, 81, 83, 87, and 78. A secondary graphical analysis can be created from the qubit and depth columns to demonstrate how circuit complexity changes across model families. This visualization makes it clear that performance-oriented comparisons should not be separated from resource considerations. A model with a higher theoretical or illustrative performance index may still demand substantially more quantum resources, training iterations, measurements, or hardware stability.

Application-Domain Analysis

QML research extends across a wide range of application domains. Image classification is frequently studied because feature reduction can allow small image datasets to be mapped into manageable quantum representations. Healthcare applications include clinical classification, biomedical analysis, and digital-health prediction. Finance research addresses classification, portfolio-related optimization, fraud analysis, and risk modeling. Cybersecurity applications include anomaly detection, classification, and pattern recognition. Chemistry and materials research use quantum computation for problems closely related to quantum systems, while optimization applications investigate quantum-assisted approaches to combinatorial problems. [7,10,15]

To compare these areas systematically, a normalized five-factor readiness framework is used. Data Suitability represents how naturally the application can be expressed using structured quantum-compatible features. Model Maturity indicates the level of algorithmic development. Hardware Fit represents

Table 2: *Success-probability progression for a 64-state search*

Application Domain	Data Suitability	Model Maturity	Hardware Fit	Research Activity	Deployment Readiness
Image Classification	8.2	7.8	6.5	9.0	6.4
Healthcare Analytics	7.5	6.8	5.9	8.2	5.7
Finance	7.8	7.2	6.2	8.4	6.1
Cybersecurity	7.1	6.5	5.8	7.6	5.6
Chemistry and Materials	8.6	7.5	6.8	8.0	6.0

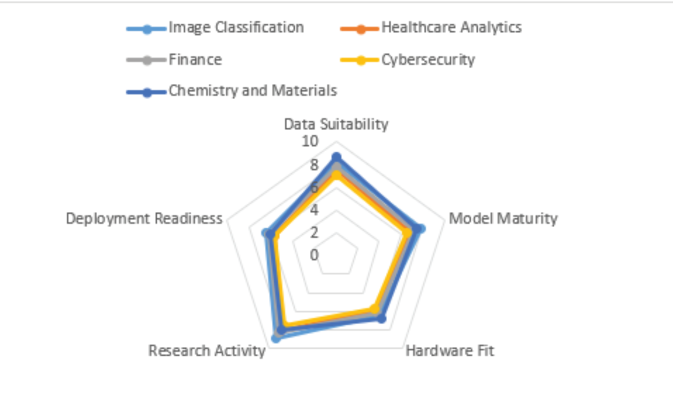

Fig 2: *Radar chart comparing QML application-domain indicators*

compatibility with near-term quantum resources. Research Activity reflects the relative intensity of published investigation. Deployment Readiness indicates how close the application is to practical implementation based on the constraints discussed in current research.

The radar chart in Fig 2 can use the five columns after the application name as its five axes. Each application domain is represented by a separate polygon with five numerical coordinates. Such a graph provides a multidimensional view rather than reducing QML evaluation to one score. Image classification receives a relatively high research-activity value because it has been a recurring experimental application in the literature. Chemistry and materials receive a high data-suitability value because quantum computation is naturally connected to the simulation of quantum systems. Healthcare and cybersecurity show strong research activity but lower deployment-readiness values because privacy, dataset scale, validation requirements, hardware limitations, and reproducibility remain significant considerations.

The results also demonstrate that an application domain may be scientifically attractive without being immediately deployable. High research activity does not necessarily imply production readiness. In healthcare, for instance, model validation must account for data quality, patient diversity, reliability, privacy, and regulatory requirements. In finance, changing market behavior and the cost of quantum execution must be considered. Cybersecurity models require stable detection performance against evolving threats. Chemistry and materials applications may have strong conceptual alignment with quantum computing but can involve highly specialized encodings and simulation requirements. Therefore, application-level feasibility must be examined separately from algorithm-level novelty.

QML Challenges And Resource Limitations

The practical implementation of QML is constrained by multiple interacting challenges. Quantum noise affects gate operations and measurements, while finite coherence restricts how long a useful circuit can evolve before information quality deteriorates. Connectivity constraints may require additional operations when logical qubits are mapped to physical hardware. Data encoding can become a bottleneck when classical features greatly exceed the number of available qubits. Parameterized models may also encounter barren plateaus, where gradients become increasingly difficult to estimate and optimization slows

Table 3: *Normalized severity values for major QML challenges*

Challenge	**Error Sensitivity**	**Encoding Overhead**	**Optimization Difficulty**	**Scalability Loss**
Quantum Noise	91	58	64	82
Data Encoding	63	89	51	74
Barren Plateaus	72	42	94	88
Limited Connectivity	76	55	67	79
Measurement Overhead	71	61	73	84

significantly. These limitations are recognized as important obstacles to scalable variational quantum computing. [8,9]

Table 3 represents the major QML challenges on a normalized severity scale. The Error Sensitivity Index indicates the relative impact of imperfect quantum operations. Encoding Overhead reflects the burden of mapping classical data into quantum states. Optimization Difficulty measures the relative difficulty of training parameterized circuits. Scalability Loss represents the extent to which increasing problem size can reduce practical usability. A larger value indicates greater challenge.

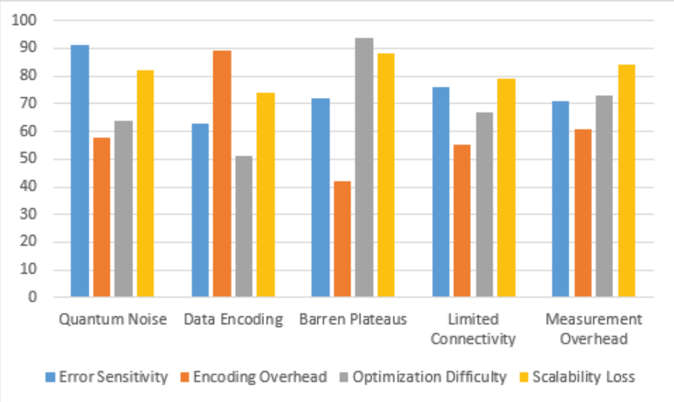

Fig 3: *Vertical bar graph comparing challenge severity levels*

Fig 3 should be generated as a grouped vertical bar graph using the four numerical columns of Table 3. Barren plateaus have the highest illustrative optimization-difficulty value of 94, while data encoding has the highest encoding-overhead value of 89. Quantum noise has the highest error-sensitivity value of 91, demonstrating that hardware imperfections can affect several stages of the QML workflow. Measurement overhead receives a high scalability-loss value because variational training and quantum-kernel evaluation may require repeated circuit execution and sampling. These patterns reinforce the importance of evaluating QML models using both predictive and computational metrics.

Barren plateaus require particular attention because they can make gradient-based optimization ineffective. The 2025 review by Larocca and colleagues describes barren plateaus as exponentially flat optimization landscapes that can emerge depending on circuit structure, cost functions, initial states, observables, and noise. This finding changes the interpretation of model expressivity: making a circuit more complex does not necessarily increase useful learning capacity if the training landscape becomes inaccessible. Strategies such as problem-informed ansatz design, layerwise training, improved initialization, local cost functions, and hardware-aware circuit construction are therefore important research directions. [8] Data encoding represents another fundamental issue. Classical datasets may contain hundreds or thousands of features, whereas near-term quantum processors can only directly manipulate a limited number of useful qubits. Feature selection, dimensionality reduction, amplitude encoding, angle encoding, basis encoding, and data re-uploading have therefore become central components of QML design. Encoding strategies introduce their own trade-offs. More compact encodings may require deeper circuits or additional preprocessing, while simple rotation-based encodings may be easier to execute but require repeated feature insertion for high-dimensional data. Effective QML design must consequently treat classical preprocessing and quantum encoding as a single integrated stage.

Representative Hybrid Architectures

Hybrid quantum-classical architectures are particularly relevant to current QML implementation because they allow quantum processors to perform specialized transformations while classical processors handle data preprocessing, optimization, control, and post-processing. A common workflow begins with classical feature preparation, followed by quantum encoding and parameterized circuit execution. Measurement outcomes are returned to the classical environment, which evaluates the loss and updates parameters. This iterative loop may continue for many optimization steps. Such an architecture provides flexibility but creates communication and execution overhead between the classical and quantum components.

Table 4 provides a graph-ready comparison of representative architectures. Qubits represent the number of active quantum variables in the illustrative model. Circuit Depth represents the approximate normalized depth category. Trainable Parameters represents the number of adjustable variables in the example configuration. The Performance Index provides a normalized comparison score.

Table 4: *Graph-ready comparison of representative QML architectures*

Architecture	Qubits	Circuit Depth	Trainable Parameters	Performance Index
Quantum Kernel Model	4	10	8	82
VQC Hybrid Model	4	16	12	80
QNN Hybrid Model	6	20	18	84
QCNN Hybrid Model	8	23	24	88
QGAN Hybrid Model	8	27	30	77

Fig 4: *Bubble chart comparing qubits, depth and parameters*

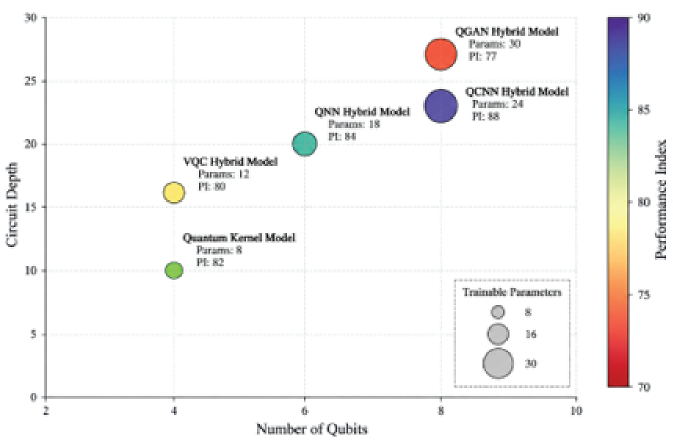

Fig 4 should be constructed as a bubble graph in which the horizontal axis represents qubit count, the vertical axis represents circuit depth, and the bubble size represents the number of trainable parameters. The table provides all required mathematical values. Quantum Kernel Model and VQC Hybrid Model occupy lower-qubit regions, while QNN, QCNN, and QGAN architectures occupy higher-complexity regions. QCNN has an illustrative performance index of 88 with eight qubits and 24 trainable parameters, while QGAN uses the same qubit count but a greater depth and parameter count. This graph visually demonstrates the relationship between model complexity and resource requirements.

Hybrid architectures also provide a practical bridge between classical machine learning and quantum computing. A complete QML workflow may perform normalization, principal-component analysis, feature selection, and data balancing classically, while quantum circuits perform nonlinear feature transformations or kernel evaluations. The optimizer can remain classical because current hardware does not generally eliminate the need for conventional numerical computation. This division of responsibilities is important for near-term QML and may remain relevant even as quantum hardware improves, because many preprocessing and decision-making operations are naturally suited to classical systems.

The implementation literature also emphasizes the importance of fair classical baselines. A QML model should not be evaluated only by whether it can classify or predict a target. A meaningful study should compare accuracy, precision, recall, computational cost, data-preparation requirements, circuit depth, number of measurements, robustness to noise, and scalability against competitive classical algorithms. This is particularly important because small datasets can sometimes allow classical methods to achieve strong performance with substantially lower operational complexity. Research therefore needs standardized benchmarks that prevent QML results from being interpreted without consideration of the complete computational pipeline. [6,10,16]

Conclusion

Quantum Machine Learning represents an important research intersection between quantum computing and artificial intelligence, offering new methods for feature mapping, classification, optimization, neural modeling, and generative learning. The reviewed literature shows that major QML approaches include quantum kernel methods, quantum support vector machines, variational quantum classifiers, quantum neural networks, quantum convolutional neural networks, and hybrid quantum-classical systems. Current research demonstrates promising results across image classification, healthcare, finance, cybersecurity, chemistry, and other application areas, but these results remain strongly dependent on problem structure and implementation conditions. Near-term hardware introduces significant limitations through noise, restricted connectivity, measurement overhead, data-encoding costs, and limited circuit depth. Barren plateaus further complicate the training of parameterized quantum models. Consequently, meaningful progress requires hardware-aware architecture design, efficient encoding, improved optimization, robust error mitigation, reproducible benchmarks, and fair classical comparisons. QML should therefore be evaluated through a complete computational perspective that considers accuracy, resource requirements, scalability, and practical deployment rather than predictive performance alone. The field is continuing to evolve toward application-specific and hybrid solutions that exploit quantum processing where it provides the greatest potential computational value.

References

1. J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum Machine Learning," Nature, vol. 549, pp. 195–202, 2017.
2. M. Schuld, I. Sinayskiy, and F. Petruccione, "An Introduction to Quantum Machine Learning," Contemporary Physics, vol. 56, pp. 172–185, 2015.
3. M. Cerezo, A. Arrasmith, R. B. M. B. et al., "Variational Quantum Algorithms," Nature Reviews Physics, vol. 3, pp. 625–644, 2021.
4. M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized Quantum Circuits as Machine Learning Models," Quantum Science and Technology, vol. 4, 043001, 2019.
5. M. Schuld and N. Killoran, "Quantum Machine Learning in Feature Hilbert Spaces," Physical Review Letters, vol. 122, 040504, 2019.
6. M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and Opportunities in Quantum Machine Learning," Nature Computational Science, 2022.
7. D. Peral-García, J. Cruz-Benito, and F. J. García-Peñalvo, "Systematic Literature Review: Quantum Machine Learning and Its Applications," Computer Science Review, vol. 51, article 100619, 2024.
8. M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, "Barren Plateaus in Variational Quantum Computing," Nature Reviews Physics, vol. 7, pp. 174–189, 2025.
9. Y. Gujju, A. Matsuo, and R. Raymond, "Quantum Machine Learning on Near-Term Quantum Devices: Current State of Supervised and Unsupervised Techniques for Real-World Applications," Physical Review Applied, vol. 21, 067001, 2024.
10. R. S. Gupta, C. E. Wood, T. Engstrom, J. D. Pole, and S. Shrapnel, "A Systematic Review of Quantum Machine Learning for Digital Health," npj Digital Medicine, vol. 8, article 237, 2025.
11. M. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. Chow, and J. M. Gambetta, "Supervised Learning with Quantum-Enhanced Feature Spaces," Nature, vol. 567, pp. 209–212, 2019.
12. V. Havlíček-related and variational learning approaches, together with foundational parameterized-circuit research, have established practical mechanisms for hybrid quantum-classical learning systems.
13. I. Cong, S. Choi, and M. D. Lukin, "Quantum Convolutional Neural Networks," Nature Physics, vol. 15, pp. 1273–1278, 2019.
14. S. Thanasilp, S. Wang, M. Cerezo, and Z. Holmes, "Exponential Concentration in Quantum Kernel Methods," Nature Communications, vol. 15, article 5200, 2024.
15. W. Lu, Y. Lu, J. Li, A. Sigov, L. Ratkin, and L. A. Ivanov, "Quantum Machine Learning: A Comprehensive Survey of Models, Challenges, and Future Directions," Theoretical Computer Science, vol. 1074, article 115923, 2026.
16. M. Schuld, R. Sweke, and J. J. Meyer, "Effect of Data Encoding on the Expressivity of Variational Quantum-Machine-Learning Models," Physical Review A, vol. 103, 032430, 2021.
17. S. Jerbi, C. Gyurik, S. C. Marshall, R. Molteni, and V. Dunjko, "Shadows of Quantum Machine Learning," Nature Communications, vol. 15, article 5676, 2024.
18. M. Cerezo, A. Poremba, L. Cincio, and P. J. Coles, "Limitations and Opportunities in Quantum Machine Learning," with subsequent developments in quantum kernels, variational models, and near-term quantum architectures.