

Assessing The Reality of Quantum Advantage in Machine Learning

Shridhar Jalindar Shelke¹, Purvaja Mahajan², Atharav Arvind Pujari³, Dr. Jyoti Chetan Bachhav⁴

^{1,2,3,4}*Bachelor of Computer Applications, MIT Arts, Commerce and Science College*

doi.org/10.64643/ijirt.209614-459

Abstract—Quantum computing has generated significant expectations for its potential to enhance machine learning (ML) tasks, with theoretical research proposing computational advantages ranging from polynomial to exponential speedups for specific subroutines. However, a critical disconnect remains: theoretical or asymptotic speedup claims do not automatically imply a practical, end-to-end advantage. Current evidence does not clearly establish whether these claimed advantages translate into measurable superiority over optimised classical approaches. This review critically assesses the relationship between theoretical claims of quantum advantage and the currently available practical and experimental evidence. We critically examine the evidence and assumptions underlying claimed speedups, including the overheads associated with data encoding and access, as well as the use of appropriate and competitive classical baselines for fair comparison. The scope of this assessment covers key limiting factors, such as scalability, trainability, hardware noise in the noisy intermediate-scale quantum (NISQ) era, and overall resource requirements. By distinguishing between theoretical potential and demonstrated practical performance, this work identifies where current evidence is convincing and where it remains insufficient. The review identifies the specific conditions required for the credible validation of quantum advantage and argues that future progress must be positioned around rigorous, realistic, and end-to-end evaluation rather than theoretical complexity claims alone.

Index Terms—quantum machine learning, quantum advantage, NISQ, classical baselines, computational complexity

I. INTRODUCTION

Quantum machine learning (QML) lies at the intersection of quantum computation and machine learning, with the central ambition of using quantum resources to improve how learning problems are represented, optimized, or solved. Quantum

computation can exploit superposition, interference, and, in appropriate settings, entanglement to process information differently from classical computation. This has motivated proposals for quantum learning models that may provide computational or statistical advantages for selected tasks. At the same time, the existence of a theoretical speedup for a particular subroutine does not by itself establish a practical advantage for an entire machine-learning workflow. Cerezo et al. (2022) emphasize both the opportunities for QML and the substantial challenges associated with trainability, data representation, and near-term hardware, while Ciliberto et al. (2018) similarly caution that the practical significance of quantum learning algorithms depends on assumptions about computation, data access, and comparison with classical alternatives.

This distinction is especially important because the term quantum advantage can refer to different forms of improvement. A quantum algorithm may exhibit better asymptotic query or computational complexity under a specified model, while a practical implementation must also account for preparing or encoding the data, executing the quantum circuit, estimating outputs through measurements, and performing any required classical optimization or post-processing.

Ciliberto et al. (2018) identify data loading and computational assumptions as central caveats when evaluating proposed QML speedups. Cerezo et al. (2022) further show that the prospects for advantage depend strongly on the type of data and task being considered, with quantum-generated data providing a particularly natural setting for quantum learning. These observations shift the question from whether quantum algorithms can be faster in principle to whether the complete learning pipeline remains advantageous once realistic costs are included.

The current hardware environment strengthens the need for this end-to-end perspective. Most contemporary quantum processors operate under noisy, resource-constrained conditions, which limits circuit depth, qubit counts, and the fidelity with which increasingly complex learning models can be executed. Consequently, many QML approaches use hybrid quantum-classical architectures in which a quantum circuit performs only part of the computation while classical hardware handles optimization, data preparation, or other components. Tychola et al. (2023) describe this hybrid setting as a practical response to present hardware limitations and report that performance can depend strongly on dataset complexity, feature mapping, and the available quantum resources. Zeguendry et al. (2023) likewise demonstrate several QML models through case studies while stressing that present implementations remain constrained by quantum hardware and encoding choices. Jiang et al. (2021) provide an end-to-end case study of mapping neural-network inference to quantum circuits, illustrating that successful execution on a quantum platform is technically distinct from demonstrating a net computational advantage over a classical implementation.

Accordingly, this review evaluates quantum advantage in machine learning using a deliberately critical boundary: theoretical or asymptotic improvement is treated as evidence of potential, whereas practical quantum advantage requires a credible end-to-end benefit over an appropriate classical alternative. The assessment therefore focuses on the assumptions behind claimed speedups, the competitiveness of classical baselines, the cost of data encoding and access, scalability and trainability, noise in the noisy intermediate-scale quantum regime, and total resource requirements. This framing is consistent with the central problem identified across the selected literature: QML contains important theoretical results and encouraging experiments, but the strength of those results varies substantially depending on the task, comparison method, and hardware setting (Cerezo et al., 2022; Ciliberto et al., 2018; Tychola et al., 2023; Zeguendry et al., 2023).

II. LITERATURE REVIEW

Early and foundational QML literature established the conceptual basis for using quantum computation in

learning problems, but it also introduced the distinction between formal speedups and realizable performance. Ciliberto et al. (2018) survey quantum learning through the lens of computational complexity and explicitly compare quantum methods with their classical counterparts. Their analysis highlights two recurring motivations for QML: the increasing computational demands of modern machine learning and theoretical results showing separations between classical and quantum learning under particular assumptions. However, the same review stresses that many quantum algorithms for machine learning contain caveats that limit immediate practical applicability, including assumptions about efficient access to data and the need to compare quantum routines with strong classical algorithms rather than with weak or structurally mismatched baselines.

Subsequent work broadened the focus from algorithmic speedup to the broader conditions under which QML could become useful. Cerezo et al. (2022) organize the field around quantum neural networks, quantum kernels, data encoding, generalization, trainability, and opportunities for quantum advantage. Their review makes clear that QML performance cannot be separated from the representation of the data or from the optimization landscape of the quantum model. In particular, trainability problems such as barren plateaus can make parameter optimization increasingly difficult as circuits scale, while noise and finite measurement resources create additional practical costs. The same work distinguishes settings involving classical data from settings involving quantum data, suggesting that learning directly from quantum processes may provide a more natural path to advantages that are difficult to reproduce classically. This perspective is important for evaluating broad claims that QML will automatically outperform classical machine learning on ordinary classical datasets.

Empirical and application-oriented reviews provide a more mixed picture. Tychola et al. (2023) compare support vector machines with quantum support vector machines and report that the quantum model performs better on some more complex datasets, while also acknowledging that present quantum hardware lacks sufficient resources for large-scale execution. Their study therefore illustrates why higher predictive accuracy on selected benchmarks should not automatically be interpreted as computational

quantum advantage: accuracy, runtime, memory requirements, circuit execution, and hardware scale are different dimensions of performance. Zegundry et al. (2023) reach a similarly qualified position through case studies involving quantum convolutional neural networks, quantum support vector machines, and variational quantum classifiers. These implementations demonstrate that QML models can be executed and meaningfully compared with classical counterparts, but they do not establish a general end-to-end superiority that transfers across datasets, algorithms, and hardware conditions.

The literature also shows why implementation evidence must be interpreted separately from theoretical complexity claims. Jiang et al. (2021) demonstrate the mapping of a trained neural-network inference workflow to quantum circuits using IBM-oriented tools and processors. Such work is valuable because it exposes practical design requirements that are often hidden in abstract algorithm descriptions, including the translation of classical model components into quantum operations and the constraints imposed by available processors. However, an end-to-end implementation is not itself proof of end-to-end advantage. To support a strong advantage claim, an implementation must demonstrate that the complete quantum or hybrid pipeline provides a meaningful benefit after the costs of encoding, circuit execution, measurement, optimization, and classical post-processing are included and after comparison with an appropriately optimized classical workflow.

Across the selected studies, three themes recur. First, the strongest theoretical claims are conditional: their significance depends on assumptions about the learning problem, data-access model, and complexity of the best available classical method (Ciliberto et al., 2018; Cerezo et al., 2022). Second, empirical results are heterogeneous. Specific QML models can match or exceed classical performance in particular experiments, but the results depend on dataset properties, feature maps, circuit design, and hardware or simulation conditions (Tychola et al., 2023; Zegundry et al., 2023). Third, the transition from small or simulated demonstrations to practically relevant quantum advantage remains constrained by trainability, noise, limited hardware resources, and the cost of interfacing classical data with quantum

processors (Cerezo et al., 2022; Tychola et al., 2023). These themes indicate that quantum advantage should be evaluated as a systems-level claim rather than inferred from a single accuracy score or an isolated asymptotic speedup.

The research gap addressed by this review follows directly from this fragmented evidence base. Existing literature provides theoretical arguments for quantum speedups, surveys of QML architectures, and proof-of-concept comparisons with classical models, yet these forms of evidence are often discussed using different definitions of advantage and different evaluation conditions. What remains insufficiently established is whether the claimed benefits survive when the full computational pipeline and a competitive classical baseline are considered together. Therefore, the present review synthesizes theoretical, simulation-based, and hardware-oriented evidence to assess where quantum advantage in machine learning is genuinely demonstrated, where it remains conditional on restrictive assumptions, and where current evidence is still inadequate. This focus is intended to separate credible progress from claims that are mathematically valid but not yet practically realizable, thereby identifying the conditions required for a defensible demonstration of practical quantum advantage.

III. METHODOLOGY

This section presents the methodological framework used to assess the current evidence for quantum advantage in machine learning (QML). The review adopts a comparative and critical perspective that distinguishes theoretical or asymptotic speedup claims from practical and experimental evidence. Consistent with the scope of the review, the assessment examines the assumptions underlying claimed speedups, the use of appropriate classical baselines, data-encoding and access overheads, scalability, trainability, NISQ-era hardware noise, and overall resource requirements.

Review Approach and Scope

The review distinguishes theoretical potential from practical implementation in QML. The evidence considered spans foundational theoretical work and rigorous speedup results, empirical studies conducted through simulation, early-stage implementations on physical quantum hardware, and review or survey

literature that synthesizes recurring challenges in the field (Liu et al., 2021; Schuld et al., 2015; Bowles et al., 2024; Zeguendry et al., 2023).

A central methodological concern is the critical examination of the quantum-advantage narrative. The review therefore evaluates whether reported improvements support a theoretical separation, an empirical performance benefit, or evidence that is sufficiently practical to contribute to an end-to-end claim of advantage. This distinction follows the broader concern that the goal of QML cannot be assessed solely by whether a quantum model appears to outperform a classical model under restricted conditions (Schuld & Killoran, 2022).

Evidence Classification

The literature is classified into distinct evidence types so that theoretical, simulation-based, hardware-based, and secondary-review evidence are not interpreted as equivalent demonstrations of quantum advantage.

Theoretical and Asymptotic Evidence

Theoretical evidence includes rigorous mathematical analyses of quantum speedups, learning separations, and methods based on quantum feature maps or kernel estimation.

The assessment considers the assumptions on which these results depend, including the structure of the learning problem and the classical hardness assumptions used to establish a separation. Rigorous supervised-learning speedup results are therefore examined as theoretical evidence while remaining distinct from practical implementation evidence (Liu et al., 2021; Olaitan et al., 2025; Zaman et al., 2025).

Simulation-Based Evidence

Simulation-based evidence includes studies in which quantum circuits or QML models are executed using classical quantum simulators. These studies are examined for comparative model performance, the influence of data encoding and entanglement, and the rigor of benchmarking against classical machine-learning approaches. Examples in the evidence base include multidimensional QML-versus-classical comparisons, financial prediction benchmarks, and benchmarking studies focused on how experimental design choices affect apparent QML performance (Ahmad et al., 2026; Bowles et al., 2024; Vhaduri et al., 2026).

Hardware-Based Evidence

Hardware-based evidence consists of experiments conducted on physical quantum processors. These studies are examined with attention to real-device constraints such as gate errors, decoherence, limited qubit capacity, and other NISQ-era effects. Hardware-based evaluations in the evidence base include quantum-enhanced machine-learning experiments and QML case studies implemented on physical devices (Bowles et al., 2024; Zeguendry et al., 2023).

Secondary Review Evidence

Review and survey literature is used as secondary evidence to synthesize recurring methodological issues and broader limitations across QML. These sources provide context on scalability, trainability, hardware limitations, error correction, algorithm development, and the progression from NISQ systems toward fault-tolerant quantum computing (Gokhale et al., 2026; Olaitan et al., 2025; Tychola et al., 2023; Ullah & Garcia-Zapirain, 2024; Zaman et al., 2025).

Evaluation Framework

A multidimensional evaluation framework is applied to assess quantum-advantage claims across the selected evidence. The framework focuses on the following parameters identified in the literature.

Speedups and Classical Baselines

The review distinguishes between theoretical claims of exponential or polynomial speedup and empirical performance improvements reported in QML studies. Central to this assessment is the rigor of the classical baseline. The methodology examines whether quantum models are compared against relevant and competitive classical methods and whether the comparison is appropriate to the task being evaluated (Ahmad et al., 2026; Bowles et al., 2024; Schuld & Killoran, 2022).

Data Encoding and Circuit Complexity

The methodology evaluates data-encoding methods that form the interface between classical data and quantum states. Angle encoding, amplitude encoding, quantum feature maps, and kernel-based representations are considered in relation to scalability and circuit requirements. Circuit depth and width are also considered where they affect embedding dimensionality, hardware capacity, or sampling-

related limitations (Ahmad et al., 2026; Bowles et al., 2024; Liu et al., 2021; Olaitan et al., 2025; Zaman et al., 2025).

Scalability, Trainability, and Noise

Scalability is assessed in relation to qubit numbers, feature dimensionality, sample size, and increasing circuit complexity. Trainability is examined through limitations such as barren plateaus in quantum neural networks where relevant. Noise impact is evaluated in relation to growing qubit and gate demands and the absence of full error correction in current NISQ systems (Bowles et al., 2024; Gokhale et al., 2026; Olaitan et al., 2025; Vhaduri et al., 2026; Zaman et al., 2025).

Resource Requirements

Beyond predictive performance, the methodology considers computational and implementation costs reported in the literature. These include parameter count, memory requirements, circuit resources, and data-access assumptions, including the practical limitations associated with quantum random-access memory (qRAM) where discussed (Olaitan et al., 2025; Vhaduri et al., 2026).

Comparative Assessment

The methodology incorporates a comparative assessment to identify agreement, discrepancies, and differences in evidential strength across theoretical, simulation-based, and hardware-based studies. Comparisons are interpreted in relation to the evaluation criteria above, particularly the form of the claimed speedup, the strength of classical baselines, resource costs, and the constraints of the implementation setting.

Evidence-Strength Assessment

The strength of the evidence is assessed according to its underlying assumptions and the level at which the claim is demonstrated.

- Theoretical rigor: examined through the robustness of the proof, the hardness assumptions used, and the conditions under which the claimed speedup is maintained (Liu et al., 2021).
- Simulation fidelity: examined in relation to experimental scale and the extent to which simulator-based results can support broader claims beyond small or controlled settings (Ahmad et al.,

2026; Bowles et al., 2024; Schuld & Killoran, 2022).

- Hardware resilience: examined through sensitivity to noise and performance on physical devices relative to the comparison conditions reported in the studies (Bowles et al., 2024; Zeguendry et al., 2023).

Methodological Limitations

The interpretation of quantum-advantage evidence is limited by heterogeneity in benchmarking practices. The studies use different datasets, problem settings, model designs, and experimental conditions, which makes direct cross-study comparison difficult (Ahmad et al., 2026; Tychola et al., 2023; Vhaduri et al., 2026; Zeguendry et al., 2023).

Variation in classical baselines also affects the comparability of reported advantages because standardized best-in-class classical comparisons are not consistently used (Bowles et al., 2024; Schuld & Killoran, 2022). In addition, the gap between theoretical requirements and current NISQ hardware remains relevant because of limitations in qubit counts, gate fidelities, noise, and error correction, while many empirical studies remain restricted in scale (Bowles et al., 2024; Eisert & Preskill, 2025; Gokhale et al., 2026; Olaitan et al., 2025; Vhaduri et al., 2026; Zaman et al., 2025).

IV. RESULTS

Theoretical evidence: advantage is real but conditional
The reviewed literature provides strong evidence that quantum machine learning can exhibit genuine computational advantages under carefully specified conditions, but the strength of these results depends on what is counted as the learning task and on the assumptions surrounding data access, computation, and output. Yamasaki et al. (2023, pp. 1, 6) establish a general complexity-theoretic route from quantum computational advantage to end-to-end advantage in supervised learning with classical data. Their construction shows that, for a broad class of functions that are efficiently computable on a quantum computer but hard for classical computation, a polynomial-time quantum learner can solve corresponding learning tasks for which polynomial-time classical learning is provably hard under the stated assumptions. Importantly, the authors themselves distinguish this theoretical result from practical demonstration:

realizing the advantage experimentally still requires a quantum computation whose performance exceeds the best feasible classical procedure at a meaningful problem size. Thus, the result strengthens the theoretical basis of QML advantage without by itself establishing a practical advantage on ordinary real-world data sets.

Cerezo et al. (2022, pp. 574–575) further narrow the conditions under which large advantages are most plausible. Their analysis argues that the clearest near-term opportunities arise when the data are quantum mechanical in origin, because a quantum learner can access information without first converting a large classical data set into quantum states. For such quantum data, advantages may occur in sample complexity as well as time complexity. In contrast, for data of purely classical origin, the review notes that no general exponential advantage is known and that the expected advantage is more likely to be polynomial or problem-specific. This distinction is central to the present review: a complexity separation for a quantum subroutine, or for a specially constructed learning problem, cannot automatically be interpreted as a practical advantage for conventional machine-learning pipelines.

Empirical comparisons: performance gains are inconsistent

The empirical evidence is substantially more mixed than the theoretical literature. Bowles et al. (2024, pp. 14–15) provide one of the most direct stress tests of claims that contemporary QML models outperform classical alternatives.

Across their benchmark suite, the authors report that prototypical out-of-the-box classical models systematically outperform the quantum models. For kernel methods, the classical support vector classifier performs best on four of seven benchmarks, while a quantum-kernel-inspired model is consistently among the weakest. The conventional convolutional neural network also outperforms the quantum convolutional models in the tested tasks. More broadly, the study does not observe a systematic improvement of quantum models relative to classical baselines as problem difficulty or dimensionality increases. These findings challenge the practice of inferring quantum advantage from isolated demonstrations on a small number of favorable data sets.

The same benchmark also weakens the assumption that explicitly quantum circuit ingredients automatically generate useful learning power. Bowles et al. (2024, pp. 14–15) find that models without entangling gates can perform surprisingly well and, in some cases, match or improve on related entangling designs. This does not show that entanglement is irrelevant to QML in general; rather, it shows that good predictive performance on small classification tasks is not sufficient evidence that a model is exploiting a uniquely quantum computational mechanism. For a claim of quantum advantage to be credible, the observed gain must therefore be tied to a resource or scaling property that cannot be reproduced by an efficient classical model.

Other experimental studies report more favorable results for quantum models, but the gains remain modest and strongly dependent on the evaluation setting. Simões et al. (2023, p. 6207) compare quantum support vector machines and quantum neural networks with classical counterparts on five relatively small data sets using both a simulator and a real quantum computer. Their best quantum neural-network configurations reach average accuracies of approximately 0.86 on the simulator and 0.85 on the quantum computer, compared with 0.78 for the best classical neural networks in their setup. The authors also report that QSVMs outperform their classical counterparts by roughly 3–4 percentage points on average and that quantum neural networks can improve on QSVM performance by up to about 5 percentage points. These results demonstrate that present-day quantum implementations can be competitive on constrained tasks. However, the study explicitly notes the small scale of the data sets, the absence of a consistently best circuit or optimizer, and the need to test larger, more error-prone hardware before drawing broader conclusions.

A similar tension appears in the SVM comparisons reported by Tychola et al. (2023, pp. 17–18). In their experiments, QSVM performs better than classical SVM on two of three data sets and wins nine of ten folds on the more complex Spam Base data set, although the differences are not large. At the same time, the quantum implementation is substantially more demanding in the simulated environment: the authors report at least 25 GB of memory and more than three hours of training per data set for QSVM, compared with less than 1 GB and less than one hour

for classical SVM. Their results therefore separate predictive performance from computational advantage. A small accuracy improvement can coexist with worse end-to-end resource use, so accuracy alone cannot validate a practical quantum advantage.

Small-scale case studies provide the same warning. Zeguendry et al. (2023, p. 30) report a variational quantum classifier that slightly exceeds classical classifiers on a reduced Iris classification task, but the data are dimensionally reduced to fit a two-qubit implementation. The study itself emphasizes the need to reduce features and circuit depth because of hardware limitations and calls for larger data sets and more capable quantum devices. Such demonstrations are useful proofs of implementation and algorithmic feasibility, but they remain insufficient evidence of a scalable advantage over optimized classical learning.

Practical overheads and NISQ constraints

The gap between theoretical and practical performance is explained in large part by costs that are often outside the asymptotic core of the quantum algorithm. Classical information must be prepared in a form accessible to a quantum circuit, and the output of a quantum computation must be estimated through measurement. These steps introduce state-preparation, encoding, sampling, and decoding costs that can dominate the useful quantum computation. The surveyed literature repeatedly identifies data encoding as a critical bottleneck for classical-data QML. In the large survey by Zaman et al. (2025, p. 23), current QML is characterized by limited qubit counts, short coherence times, non-negligible error rates, high implementation cost, and dependence on hybrid classical-quantum optimization. The same survey notes that barren plateaus can make gradients vanish rapidly with system size, creating a direct trainability barrier for variational QML models.

Hardware noise further changes the interpretation of nominal speedups. Cerezo et al. (2022, pp. 574–575) explain that noise accumulates through deeper circuits, affects data preparation and quantum-kernel estimation, and can induce noise-related trainability problems. Error mitigation may improve performance, but it also adds additional executions and resources. Consequently, a theoretical reduction in query or gate complexity can be offset by the number of circuit repetitions, calibration steps, measurements, classical optimization iterations, and error-control procedures

required to obtain a stable answer on physical hardware. This makes end-to-end resource accounting essential when moving from asymptotic claims to practical comparisons.

Synthesis of the results

Across the reviewed evidence, three patterns are consistent. First, quantum advantage is best established at the level of theory or in restricted learning settings with carefully defined computational assumptions. Second, current empirical QML can match or occasionally outperform classical models on particular small-scale tasks, but the direction and magnitude of the advantage vary with the data set, model family, hyperparameters, preprocessing choices, and classical baseline. Third, present hardware and workflow overheads prevent these task-level improvements from being interpreted as a general end-to-end advantage. The evidence therefore supports the existence of quantum advantage as a conditional phenomenon, while not supporting a broad claim that contemporary QML is practically superior to optimized classical machine learning on ordinary classical data.

V. DISCUSSION

Why theoretical speedup does not imply end-to-end advantage

The central finding of this review is that the phrase quantum advantage refers to several distinct claims that should not be merged. A proof that a quantum subroutine has better asymptotic complexity establishes an algorithmic separation under a specified computational model. A learning-theory result such as that of Yamasaki et al. (2023, pp. 1, 6) goes further by embedding a quantum computational separation into an end-to-end supervised learning task, but it still depends on a constructed function class and complexity assumptions. An experimental comparison, in contrast, asks whether a complete implemented pipeline performs better on finite hardware after state preparation, optimization, measurement, and classical post-processing are included. Evidence at one level does not automatically validate the next.

This distinction explains why strong theory can coexist with ambiguous experiments. Theoretical analyses often focus on the part of a problem where

quantum computation is most favorable, whereas practical QML must also move information into and out of the quantum processor. If those interfaces scale poorly, or if many circuit repetitions are required to overcome noise and sampling uncertainty, the nominal speedup can be reduced or eliminated. The empirical resource measurements reported by Tychola et al. (2023, pp. 17–18) illustrate this directly: better predictive scores on some data sets do not translate into lower memory use or shorter training time in the evaluated implementation.

Classical baselines are part of the claim, not a secondary detail

A credible quantum-advantage claim is meaningful only relative to an appropriate classical comparator. The Bowles et al. (2024, pp. 14–15) benchmark shows that conclusions can change when simple but strong classical models are included and tuned under comparable conditions. Their classical SVC, MLP, and CNN baselines frequently outperform quantum alternatives, despite the latter being drawn from prominent QML proposals. The result does not prove that all QML models lack value; it shows that a weak or poorly optimized classical baseline can create an artificial impression of quantum superiority. Accordingly, future studies should report the classical model class, optimization budget, preprocessing, hyperparameter search, compute budget, and statistical uncertainty with the same care used to describe the quantum circuit.

The disagreement between Bowles et al. (2024, pp. 14–15), Simões et al. (2023, p. 6207), and Tychola et al. (2023, pp. 17–18) is therefore informative rather than contradictory. Simões et al. obtain small average advantages for quantum classifiers on five data sets, whereas Bowles et al. find systematic superiority of out-of-the-box classical models in a broader benchmark. Tychola et al. observe data-set-dependent QSVM accuracy gains but worse computational efficiency. Together, these studies indicate that present QML performance is highly sensitive to benchmark design. A robust advantage should persist across reasonable choices of data splits, preprocessing, classical baselines, circuit architectures, and optimization procedures, and it should become clearer rather than disappear as the comparison becomes more rigorous.

Where quantum advantage appears most credible

The source of the data is an important boundary condition. Cerezo et al. (2022, pp. 574–575) argue that quantum-native data are a particularly promising domain because the information already exists in a quantum representation and may contain correlations that are expensive to reproduce classically. This setting avoids part of the classical-to-quantum loading problem and makes sample-complexity advantages more plausible. By contrast, most conventional machine-learning applications begin with classical data. In that regime, encoding overhead and strong classical algorithms narrow the space in which a practically relevant quantum advantage can appear. This does not rule out such an advantage; Yamasaki et al. (2023, pp. 1, 6) provide an important theoretical construction using classical data. It does, however, make the burden of experimental validation substantially higher.

What would count as convincing practical evidence

Based on the combined results, a convincing demonstration of practical quantum advantage in machine learning would require more than higher accuracy on a single small data set. At minimum, the quantum and classical systems should solve the same end-to-end task under comparable data access and output requirements; the classical baselines should be competitive and adequately optimized; and the comparison should include wall-clock time or another clearly defined computational cost together with memory, circuit executions, measurement shots, and preprocessing overhead where relevant. The quantum result should also be demonstrated on physical hardware at a problem size large enough that classical simulation is not itself the execution method being compared. Finally, the advantage should exhibit a stable scaling trend or a reproducible resource separation rather than a one-off difference within statistical variation.

Current studies satisfy pieces of this standard but not the full set simultaneously. Simões et al. (2023, p. 6207) show real-hardware feasibility and small accuracy gains, yet on limited data sizes. Bowles et al. (2024, pp. 14–15) provide unusually rigorous comparative benchmarking, but the quantum circuits are evaluated in a regime that remains classically simulable and do not exhibit systematic superiority. Tychola et al. (2023, pp. 17–18) show task-dependent

predictive gains but also substantial simulated resource costs. Theoretical work such as Yamasaki et al. (2023, pp. 1, 6) supplies the complexity foundation needed to know that end-to-end advantage is possible in principle, but an experimental realization that closes the full theory-to-practice loop remains outstanding within the reviewed evidence.

Overall interpretation

The evidence supports a cautious but non-dismissive conclusion. Quantum advantage in machine learning is not merely hypothetical: rigorous theory identifies learning tasks with genuine computational separations, and present experiments show that quantum models can be competitive and sometimes more accurate in restricted settings. At the same time, the current literature does not establish a general practical advantage for QML over optimized classical machine learning on real-world classical data. The dominant uncertainty is not whether quantum computation can ever aid learning, but whether the advantage survives data access, training, measurement, noise, resource costs, and strong classical competition at useful scales. On the evidence reviewed here, the most defensible position is therefore that quantum advantage in ML is presently conditional and task-specific, with practical end-to-end superiority still requiring more rigorous and scalable experimental validation.

VI. CONCLUSION

This review set out to determine whether the theoretical promise of quantum advantage in machine learning survives contact with practical, end-to-end evaluation. The evidence supports a conditional answer. At the level of theory, quantum advantage is genuine: Yamasaki et al. (2023) establish a rigorous complexity-theoretic pathway from quantum computational advantage to end-to-end supervised learning advantage, and Cerezo et al. (2022) identify quantum-native data as a setting where such advantages are especially plausible. These results confirm that quantum learning is not merely aspirational — under specified assumptions, provable separations from classical learning exist.

However, theoretical potential does not yet translate into demonstrated practical superiority. Empirical comparisons are heterogeneous rather than

confirmatory: Bowles et al. (2024) show that well-tuned classical baselines routinely outperform prominent QML models across a broad benchmark suite, while Simões et al. (2023) and Tychola et al. (2023) report modest quantum advantages that are either confined to small datasets or offset by substantially higher computational cost. Taken together, these findings indicate that predictive accuracy alone cannot establish quantum advantage; resource accounting, classical baseline strength, and data provenance must all be considered jointly. Compounding this, NISQ-era constraints — encoding overhead, barren plateaus, and hardware noise (Cerezo et al., 2022; Zaman et al., 2025) — continue to erode the margin between theoretical speedup and realized performance.

The most defensible conclusion, then, is that quantum advantage in machine learning is real but presently conditional and task-specific, strongest for quantum-native data and rigorously defined learning problems, and not yet established as a general practical benefit over optimized classical machine learning on ordinary classical data. Future progress will depend less on new asymptotic complexity results and more on rigorous, end-to-end, hardware-validated comparisons against competitive classical baselines — the standard this review has argued is necessary to distinguish credible advantage from mathematically valid but practically unrealized potential.

REFERENCES

- [1] R. Ahmad, M. Kashif, N. Innan, and M. Shafique, “Quantum vs. classical machine learning: A benchmark study for financial prediction,” *arXiv*, version 1, 2026, doi: 10.48550/ARXIV.2601.03802.
- [2] J. Bowles, S. Ahmed, and M. Schuld, “Better than classical? The subtle art of benchmarking quantum machine learning models,” *arXiv preprint arXiv:2403.07059*, 2024, doi: 10.48550/arXiv.2403.07059.
- [3] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, “Challenges and opportunities in quantum machine learning,” *Nature Computational Science*, vol. 2, no. 9, pp. 567–576, 2022, doi: 10.1038/s43588-022-00311-3.
- [4] A. M. R. Chy and M. F. Shahriyar, “Why classical-quantum machine learning works better

- than classical machine learning,” in Proc. 2026 IEEE 2nd Int. Conf. Quantum Photonics, Artificial Intelligence & Networking (QPAIN), 2026, pp. 1–6, doi: 10.1109/QPAIN69676.2026.11546286.
- [5] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, “Quantum machine learning: A classical perspective,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 474, no. 2209, Art. no. 20170551, 2018, doi: 10.1098/rspa.2017.0551.
- [6] J. Eisert and J. Preskill, “Mind the gaps: The fraught road to quantum advantage,” *arXiv*, version 3, 2025, doi: 10.48550/ARXIV.2510.19928.
- [7] S. S. Gokhale, R. R. Dhote, and R. Delhibabu, “A review of quantum machine learning algorithms, applications, and emerging advantages,” *Discover Computing*, vol. 29, no. 1, Art. no. 226, 2026, doi: 10.1007/s10791-026-10085-1.
- [8] A. Kiani, “Circuit-optimized quantum kernels for battery degradation prediction across temperature regimes,” *Journal of Energy Storage*, vol. 160, Art. no. 121987, 2026, doi: 10.1016/j.est.2026.121987.
- [9] P. Lamichhane and D. B. Rawat, “Quantum machine learning: Recent advances, challenges, and perspectives,” *IEEE Access*, vol. 13, pp. 94057–94105, 2025, doi: 10.1109/ACCESS.2025.3573244.
- [10] Y. Liu, S. Arunachalam, and K. Temme, “A rigorous and robust quantum speed-up in supervised machine learning,” *Nature Physics*, vol. 17, no. 9, pp. 1013–1017, 2021, doi: 10.1038/s41567-021-01287-z.
- [11] N. Meyer, C. Ufrecht, G. Yammine, G. Kontes, C. Mutschler, and D. D. Scherer, “Benchmarking quantum reinforcement learning,” *arXiv*, version 2, 2025, doi: 10.48550/ARXIV.2501.15893.
- [12] B. Mohammadi, M. Chen, M. R. Nikoo, and A. Al-Maktoumi, “Quantum machine learning-driven surrogate modeling for efficient multi-objective optimization of CO₂ storage and geothermal energy extraction,” *Advances in Geo-Energy Research*, vol. 18, no. 2, pp. 137–152, 2025, doi: 10.46690/ager.2025.11.04.
- [13] O. F. Olaitan, S. O. Ayeni, A. Olosunde, F. C. Okeke, U. U. Okonkwo, C. G. Ochieze, O. V. Chukwujama, and O. N. Akatakpo, “Quantum computing in artificial intelligence: A review of quantum machine learning algorithms,” *Path of Science*, vol. 11, no. 5, Art. no. 7001, 2025, doi: 10.22178/pos.117-25.
- [14] L. Pira and P. Rebertost, “Fundamentals of quantum machine learning and robustness,” in *Quantum Robustness in Artificial Intelligence*, M. Usman, Ed. Cham, Switzerland: Springer Nature Switzerland, 2026, pp. 1–19, doi: 10.1007/978-3-032-11153-1_1.
- [15] M. Revythi and G. Koukiou, “Quantum machine learning and deep learning: Fundamentals, algorithms, techniques, and real-world applications,” *Machine Learning and Knowledge Extraction*, vol. 7, no. 3, Art. no. 75, 2025, doi: 10.3390/make7030075.
- [16] D. Ristè, M. P. Da Silva, C. A. Ryan, A. W. Cross, A. D. Córcoles, J. A. Smolin, J. M. Gambetta, J. M. Chow, and B. R. Johnson, “Demonstration of quantum advantage in machine learning,” *npj Quantum Information*, vol. 3, no. 1, Art. no. 16, 2017, doi: 10.1038/s41534-017-0017-3.
- [17] M. Schuld and N. Killoran, “Is quantum advantage the right goal for quantum machine learning?” *PRX Quantum*, vol. 3, no. 3, Art. no. 030101, 2022, doi: 10.1103/PRXQuantum.3.030101.
- [18] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015, doi: 10.1080/00107514.2014.964942.
- [19] R. D. M. Simoes, P. Huber, N. Meier, N. Smailov, R. M. Fuchsli, and K. Stockinger, “Experimental evaluation of quantum machine learning algorithms,” *IEEE Access*, vol. 11, pp. 6197–6208, 2023, doi: 10.1109/ACCESS.2023.3236409.
- [20] K. A. Tychola, T. Kalampokas, and G. A. Papakostas, “Quantum machine learning—An overview,” *Electronics*, vol. 12, no. 11, Art. no. 2379, 2023, doi: 10.3390/electronics12112379.
- [21] U. Ullah and B. Garcia-Zapirain, “Quantum machine learning revolution in healthcare: A systematic review of emerging perspectives and applications,” *IEEE Access*, vol. 12, pp. 11423–11450, 2024, doi: 10.1109/ACCESS.2024.3353461.

- [22] S. Vhaduri, R. Gammon, and S. Dibbo, “Do we really need quantum machine learning? A multidimensional empirical study,” arXiv, version 2, 2026, doi: 10.48550/ARXIV.2605.27923.
- [23] M. Wendlinger, K. Tschärke, and P. Debus, “A comparative analysis of adversarial robustness for quantum and classical machine learning models,” arXiv, version 1, 2024, doi: 10.48550/ARXIV.2404.16154.
- [24] M. T. West, S. M. Erfani, C. Leckie, M. Sevier, L. C. L. Hollenberg, and M. Usman, “Benchmarking adversarially robust quantum machine learning at scale,” *Physical Review Research*, vol. 5, no. 2, Art. no. 023186, 2023, doi: 10.1103/PhysRevResearch.5.023186.
- [25] H. Yamasaki, N. Isogai, and M. Murao, “Advantage of quantum machine learning from general computational advantages,” arXiv, version 1, 2023, doi: 10.48550/ARXIV.2312.03057.
- [26] K. Zaman, A. Marchisio, M. A. Hanif, and M. Shafique, “A survey on quantum machine learning: Current trends, challenges, opportunities, and the road ahead,” arXiv, version 4, doi: 10.48550/ARXIV.2310.10315.
- [27] A. Zeguendry, Z. Jarir, and M. Quafafou, “Quantum machine learning: A review and case studies,” *Entropy*, vol. 25, no. 2, Art. no. 287, 2023, doi: 10.3390/e25020287.