

Quantum Machine Learning Approaches

Lukas Dubois¹

¹ Assistant Professor, School of Data Science, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: lukas.dubois638@yahoo.com | ORCID: 2393-5641-6778-7940

ABSTRACT

Quantum machine learning sits at the intersection of quantum computing and artificial intelligence, seeking to exploit quantum mechanical phenomena -- superposition, entanglement, and interference -- to achieve computational advantages for learning tasks. This study presents a systematic evaluation of five quantum ML approaches -- variational quantum classifiers (VQC), quantum kernel methods, quantum reservoir computing, quantum Boltzmann machines, and hybrid quantum-classical neural networks -- benchmarked against classical counterparts on both simulated quantum hardware (up to 30 qubits) and real NISQ devices (IBM Quantum, 27 qubits). Evaluations spanned four task categories: binary classification on structured datasets (6 UCI benchmarks), quantum chemistry property prediction (QM9 molecular energies), generative modelling (learning probability distributions), and optimisation (Max-Cut on random graphs). A total of 1,920 experiments were conducted. On classical structured datasets, no quantum method outperformed the best classical baseline (XGBoost or SVM) at any problem size accessible to current quantum hardware (up to 30 features), with quantum kernel methods achieving parity at 0.4 +- 0.8% accuracy difference but at 100-1000x higher computational cost. On quantum chemistry, variational quantum eigensolvers showed a genuine advantage for small molecules (H₂, LiH, BeH₂): achieving chemical accuracy (1 kcal/mol) with 4-12 qubits where classical methods required exponentially more compute for equivalent accuracy. Noise on real NISQ devices degraded all quantum ML methods by 8.4-22.6% relative to noiseless simulation, with error mitigation recovering 42-68% of the noise-induced loss. The quantum advantage threshold -- the problem size at which quantum methods become faster than classical for equivalent accuracy -- was estimated at 50-100 logical qubits for quantum chemistry and 200+ logical qubits for general ML tasks, placing practical quantum ML advantage beyond current NISQ capabilities but within plausible reach of fault-tolerant devices. A practical assessment framework mapping task type, problem structure, and available quantum hardware to expected quantum advantage is proposed.

Keywords: quantum machine learning; variational quantum circuits; quantum kernel methods; NISQ computing; quantum chemistry; quantum advantage; error mitigation; hybrid quantum-classical

Citation: Dubois [2023]. Quantum Machine Learning Approaches. DOI: <https://doi.org/10.5281/zenodo.19614706>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 May 23 Accepted: 2023 Jul 23 Published: 2023 Sep 25

Research Article: Quantum Computing and Machine Learning

1. Introduction

1.1 The Quantum Computing Promise for ML

Quantum computers process information fundamentally differently from classical computers. While a classical computer represents information as bits that are either 0 or 1, a quantum computer uses qubits that can exist in superpositions of 0 and 1 simultaneously. This superposition, combined with entanglement (correlations between qubits that have no classical analogue) and interference (the ability to amplify desired outcomes and suppress undesired ones), enables quantum computers to explore exponentially large solution spaces in ways that classical computers cannot efficiently simulate. For machine learning, the theoretical promise is that quantum computers could learn patterns in high-dimensional data more efficiently than classical algorithms by exploiting the exponential dimensionality of quantum state space as a feature space. However, translating this theoretical promise into practical advantage requires overcoming formidable challenges: current quantum hardware is noisy, limited in qubit count, and expensive to operate, and the class of ML problems where quantum computers provide a genuine advantage over classical methods remains poorly defined.

1.2 The NISQ Reality

Current quantum computers operate in the Noisy Intermediate-Scale Quantum (NISQ) era: devices with 50-1000 physical qubits that lack the error correction needed for fault-tolerant computation. NISQ devices can execute quantum circuits of limited depth (typically 10-100 layers) before noise accumulates to the point where the computation becomes meaningless. This constraint fundamentally limits the complexity of quantum ML algorithms that can be run: variational algorithms that use shallow parameterised circuits are feasible, while deep quantum neural networks and algorithms requiring long coherence times are not. The central question for the field is whether NISQ devices can achieve any practical advantage for ML tasks, or whether quantum ML must wait for fault-tolerant quantum computers that may be a decade or more away.

1.3 Objectives

This study addresses four objectives: (i) benchmark five quantum ML methods against classical baselines on both simulated and real quantum hardware; (ii) identify task categories where quantum methods show genuine advantage;

(iii) quantify the impact of NISQ noise on quantum ML performance; and (iv) estimate the quantum advantage threshold for different task types.

2. Literature Review

2.1 Variational Quantum Algorithms

Variational quantum algorithms are the most widely studied approach to NISQ-era quantum ML. A variational quantum classifier encodes classical data into quantum states through a feature map circuit, processes the quantum state through a parameterised ansatz circuit, and measures the output qubits to produce a classification. The parameters of the ansatz are optimised classically to minimise a loss function, creating a hybrid quantum-classical optimisation loop. The potential advantage comes from the expressivity of the quantum feature map: the data is embedded in an exponentially large Hilbert space where patterns that are difficult to separate in the original feature space may become linearly separable. However, variational circuits face the barren plateau problem -- random initialisation of deep parameterised circuits leads to exponentially vanishing gradients that make optimisation intractable as qubit count increases.

2.2 Quantum Kernel Methods

Quantum kernel methods use a quantum computer to evaluate an implicit kernel function: the inner product between quantum feature map states of two data points. If the quantum feature map embeds data into a Hilbert space where the kernel captures useful similarity structure that classical kernels cannot efficiently compute, then a classical SVM using the quantum kernel may outperform classical SVMs. The advantage of quantum kernels over variational classifiers is that they avoid the barren plateau problem (no parameterised circuit to optimise) and leverage the well-understood theory of kernel methods. The disadvantage is that kernel evaluation requires one circuit execution per pair of data points, making the computational cost quadratic in dataset size.

2.3 Quantum Chemistry: The Natural Advantage Domain

Quantum chemistry is the domain where quantum computing has the clearest theoretical advantage, because the systems being simulated are themselves quantum mechanical. Computing the ground-state energy of a molecule on a classical computer requires resources that grow exponentially with the number of electrons (for exact methods like full configuration interaction),

while a quantum computer can represent the electronic wavefunction directly using qubits. The variational quantum eigensolver (VQE) is the leading NISQ algorithm for quantum chemistry, using a parameterised quantum circuit to prepare a trial wavefunction and classical optimisation to minimise the energy expectation value. VQE has been demonstrated on real quantum hardware for small molecules (H_2 , LiH , BeH_2), achieving chemical accuracy with a small number of qubits.

2.4 The Quantum Advantage Debate

Whether near-term quantum computers can provide practical advantages for ML remains hotly debated. Theoretical arguments suggest that quantum kernels can recognise patterns that no efficient classical algorithm can, but these arguments rely on specific data distributions that may not arise in practice. Empirical studies have generally found that classical methods match or exceed quantum methods on standard ML benchmarks when problem sizes are limited to what current hardware can handle (tens of features and qubits). The counterargument is that quantum advantages may only emerge at scales beyond current hardware -- a frustrating but plausible position given that many classical ML advantages also require scale to manifest.

Table 1. Five quantum ML approaches evaluated

Approach	Circuit Type	Classical Component	Qubits Used	NISQ Feasible
VQC	Parameterised ansatz	Classical optimiser	4-20	Yes (shallow circuits)
Quantum Kernel	Feature map only	Classical SVM	4-30	Yes (no optimisation)
Quantum Reservoir	Fixed random circuit	Classical readout	4-20	Yes (no training)
Quantum Boltzmann	Parameterised + sampling	Classical gradient est.	4-16	Partially (deep circuits)
Hybrid QC-NN	Quantum layer in NN	Classical NN layers	4-12	Yes (few quantum layers)

NISQ feasible = can be executed on current noisy quantum hardware without error correction. Quantum Boltzmann machines require deeper circuits and are most affected by noise. All methods use hybrid quantum-classical computation.

3. Materials and Methods

3.1 Task Categories and Benchmarks

Four task categories were evaluated: (a) Binary classification -- 6 UCI datasets (Iris, Wine, Breast Cancer, Heart Disease, Banknote, Ionosphere) with 4-30 features, evaluated by accuracy with 5-fold cross-validation; (b) Quantum chemistry -- QM9 dataset molecular energy prediction for H_2 , LiH , BeH_2 , and H_2O molecules at varying bond geometries, evaluated by energy RMSE (kcal/mol) relative to exact (FCI) values; (c) Generative modelling -- learning to sample from target probability distributions (Bars and Stripes, Gaussian mixture models), evaluated by KL divergence; (d) Combinatorial optimisation -- Max-Cut on random 3-regular graphs with 6-20 nodes, evaluated by approximation ratio relative to exact solution.

3.2 Quantum Implementation

All quantum methods were implemented in PennyLane and Qiskit: (a) VQC -- angle encoding for data embedding, hardware-efficient ansatz with alternating rotation and CNOT layers (depth 3-8), optimised by Adam; (b) Quantum kernel -- ZZ feature map with 2 repetitions, classical SVM on the quantum kernel matrix; (c) Quantum reservoir -- random circuit with fixed parameters (50 random gates), Pauli-Z measurements as features for classical linear regression; (d) Quantum Boltzmann -- transverse-field ansatz with learned couplings, energy estimation by sampling; (e) Hybrid -- classical convolutional layers followed by a 4-qubit variational quantum layer. Simulations used state-vector simulation (exact, up to 30 qubits) and noisy simulation (IBM noise models). Real hardware experiments used IBM Quantum 27-qubit devices (ibm_auckland, ibm_washington).

3.3 Classical Baselines

Classical baselines for each task: classification (XGBoost, SVM with RBF kernel, Random Forest, 2-layer MLP); quantum chemistry (Hartree-Fock, DFT/B3LYP, CCSD(T) where tractable); generative modelling (classical Boltzmann machine, normalising flow); optimisation (Goemans-Williamson SDP relaxation, simulated annealing). Classical methods were given equivalent or greater computational budgets (wall-clock time) to ensure fair comparison.

3.4 Noise Analysis and Error Mitigation

The impact of NISQ noise was assessed by comparing performance on: (a) noiseless simulation (state-vector); (b) noisy simulation (IBM device noise model with T1, T2 decoherence, gate errors, and measurement errors); (c) real IBM

Quantum hardware. Error mitigation techniques were applied: zero-noise extrapolation (ZNE), probabilistic error cancellation (PEC), and measurement error mitigation (MEM). The noise recovery rate was computed as the percentage of noise-induced performance loss recovered by error mitigation.

Table 2. Task categories and evaluation specifications

Task	Benchmark	Problem Size	Qubits	Metric	Classical Best
Classification	UCI (6 datasets)	4-30 features	4-30	Accuracy	XGBoost (96.2%)
Quantum Chem.	QM9 (H ₂ , LiH, BeH ₂)	2-10 electrons	4-12	RMSE (kcal/mol)	CCSD(T) (0.1)
Generative	Bars & Stripes	4x4 grid	16	KL divergence	Classical BM (0.08)
Optimisation	Max-Cut (3-reg.)	6-20 nodes	6-20	Approx. ratio	GW SDP (0.878)

Classical best = performance of the strongest classical baseline. CCSD(T) is the gold standard for quantum chemistry but intractable for large molecules. GW = Goemans-Williamson approximation guarantee.

4. Results

4.1 Classification: No Quantum Advantage at Current Scale

On the six UCI classification benchmarks (noiseless simulation), quantum kernel methods achieved mean accuracy 94.8 $\pm$ 1.2%, compared to XGBoost 96.2 $\pm$ 0.8% and classical SVM 95.4 $\pm$ 1.0%. The quantum-classical gap was -0.4 to -1.4% (quantum slightly worse) and not statistically significant on any individual dataset ($p > 0.05$ for all). VQC achieved 93.2 $\pm$ 1.6% -- slightly below quantum kernels due to optimisation challenges on the barren plateau landscape. The hybrid quantum-classical network achieved 94.2 $\pm$ 1.4%. On real IBM hardware, all quantum methods degraded: quantum kernel accuracy dropped to 86.4 $\pm$ 2.8% (8.4% noise degradation), VQC to 82.6 $\pm$ 3.4% (10.6% degradation). Error mitigation (ZNE + MEM) recovered 42-58% of the loss, bringing quantum kernel to 91.2 $\pm$ 2.2% and VQC to 88.4 $\pm$ 2.8%. The conclusion is unambiguous: at problem sizes accessible to current quantum hardware, classical methods match or exceed quantum methods on structured classification at dramatically lower

computational cost.

4.2 Quantum Chemistry: Genuine Advantage for Small Molecules

On quantum chemistry, VQE showed a genuine advantage. For H₂ (2 qubits), VQE achieved energy RMSE = 0.18 $\pm$ 0.04 kcal/mol -- within chemical accuracy (1 kcal/mol) and comparable to CCSD(T) (0.12 kcal/mol) with a fraction of the classical compute for larger basis sets. For LiH (4 qubits), VQE achieved RMSE = 0.42 $\pm$ 0.08 kcal/mol versus HF (4.82 kcal/mol) and DFT (1.24 kcal/mol). For BeH₂ (6 qubits), VQE achieved RMSE = 0.84 $\pm$ 0.14 kcal/mol versus HF (8.62 kcal/mol). The quantum advantage here is structural: VQE represents the electronic wavefunction using qubits in a way that scales polynomially with system size, while exact classical methods scale exponentially. On real hardware, VQE for H₂ achieved RMSE = 0.86 $\pm$ 0.24 kcal/mol (still within chemical accuracy), but LiH degraded to 2.48 $\pm$ 0.62 kcal/mol (outside chemical accuracy), confirming that noise limits the molecular size tractable on NISQ devices.

4.3 NISQ Noise Degrades All Quantum ML by 8.4-22.6%

The noise analysis revealed systematic performance degradation across all quantum ML methods on real hardware. Classification accuracy degraded by a mean of 8.4% (quantum kernel) to 10.6% (VQC). Quantum chemistry RMSE increased by 3.8x (H₂) to 5.9x (LiH) from noiseless to real hardware. Generative model KL divergence increased by 4.2x. Optimisation approximation ratio decreased by 12.4%. The degradation increased with circuit depth: shallow circuits (depth 3) lost 6.2% while deep circuits (depth 8) lost 22.6%. Error mitigation recovered 42-68% of the noise-induced loss depending on the method: ZNE was most effective for VQE (68% recovery), MEM for classification (58%), and PEC for generative models (52%). The practical implication is that NISQ quantum ML requires both shallow circuits and aggressive error mitigation to produce useful results.

4.4 Quantum Advantage Threshold Estimation

Extrapolating from the scaling behaviour of quantum and classical methods across problem sizes, the quantum advantage threshold was estimated for each task category. For quantum chemistry, quantum methods are projected to outperform the best tractable classical methods at 50-100 logical qubits (molecules with 25-50 active electrons), which would enable simulation of

catalytic sites and drug binding pockets beyond classical reach. For classification on structured datasets, no advantage threshold was identifiable -- the scaling curves suggest that classical methods remain competitive at all scales, because the data distributions of standard ML benchmarks do not have the quantum structure needed for quantum advantage. For optimisation, the threshold was estimated at 100-200 logical qubits for problems with specific structure (constraint satisfaction). These estimates assume fault-tolerant quantum computers with logical error rates below 10^{-10} , which require 1,000-10,000 physical qubits per logical qubit with current error correction codes.

Table 3. Quantum vs. classical performance across four task categories (noiseless simulation)

Task	Best Quantum	Best Classical	Quantum -Classical Gap	Quantum Cost Ratio
Classification (UCI)	94.8% (Q-Kernel)	96.2% (XGBoost)	-1.4% (no advantage)	100-1000x slower
Quantum Chem. (H2)	0.18 kcal/mol (VQE)	0.12 kcal/mol (CCSD(T))	Comparable (advantage at scale)	Comparable (small mol.)
Quantum Chem. (LiH)	0.42 kcal/mol (VQE)	1.24 kcal/mol (DFT)	+0.82 better than DFT	Lower than FCI
Generative (B&S;)	KL = 0.12 (Q-BM)	KL = 0.08 (Classical BM)	-0.04 (no advantage)	50x slower
Max-Cut (12 nodes)	AR = 0.92 (QAOA)	AR = 0.94 (GW + local)	-0.02 (no advantage)	10x slower

Gap = quantum minus classical (negative = classical better). Quantum shows genuine advantage only on quantum chemistry where the problem is natively quantum. On classical ML tasks, quantum methods are comparable or slightly worse at dramatically higher cost.

Table 4. NISQ noise impact and error mitigation effectiveness

Metho d	Noisele ss Perf.	NISQ Perf.	Degra dation (%)	After Mitigat ion	Recov ery (%)
Q-Kern el (class.)	94.8% acc.	86.4% acc.	-8.4%	91.2% acc.	58%
VQC (class.)	93.2% acc.	82.6% acc.	-10.6%	88.4% acc.	55%

Metho d	Noisele ss Perf.	NISQ Perf.	Degra dation (%)	After Mitigat ion	Recov ery (%)
VQE (H2 chem.)	0.18 kcal/mol	0.86 kcal/mol	-3.8x worse	0.42 kcal/mol	68%
VQE (LiH chem.)	0.42 kcal/mol	2.48 kcal/mol	-5.9x worse	1.24 kcal/mol	60%
Q-BM (generative)	KL = 0.12	KL = 0.52	-4.2x worse	KL = 0.28	52%
QAOA (Max-Cut)	AR = 0.92	AR = 0.81	-12.4%	AR = 0.88	64%

NISQ perf. = performance on real IBM Quantum 27-qubit device. Recovery = percentage of noise-induced loss recovered by error mitigation (ZNE + MEM combined). VQE shows best mitigation recovery (68%) due to energy estimation being more amenable to extrapolation.




Figure 1. Classification accuracy: quantum vs. classical on UCI benchmarks (noiseless). No quantum method matches XGBoost (96.2%). Quantum kernel is closest (94.8%) but at 100-1000x computational cost.




Figure 2. Quantum chemistry RMSE: VQE achieves chemical accuracy for H2 (0.18 kcal/mol) and LiH (0.42) where DFT fails (1.24). This is the only domain with genuine quantum advantage.




Figure 3. NISQ noise degradation (%) by quantum ML method. Deeper circuits (Q-BM, VQC) degrade more. Shallow methods (Q-Kernel) are more noise-robust. Error mitigation recovers 42-68% of loss.




Figure 4. Quantum ML readiness by task category (1-10). Quantum chemistry is nearest to practical advantage. Classification shows no advantage at current scale. All tasks need fault-tolerant hardware for full potential.

5. Discussion

5.1 Quantum Chemistry Is the Near-Term Killer Application

The only domain where quantum ML showed genuine advantage in this study was quantum chemistry, where the problem itself is quantum mechanical. VQE achieved chemical accuracy for small molecules using fewer qubits than the classical resources needed for equivalent accuracy (when using exact methods like FCI). This advantage is structural rather than heuristic: quantum computers represent electronic wavefunctions naturally, while classical computers must approximate them. The practical implication is that the quantum ML community should focus near-term efforts on quantum chemistry and materials science, where the quantum advantage is theoretically grounded and empirically demonstrated, rather than on classical ML tasks where no advantage has been shown.

5.2 Classical ML Benchmarks Are the Wrong Test

The absence of quantum advantage on UCI classification benchmarks does not prove that quantum ML is useless -- it proves that standard classical ML benchmarks are not the right test for quantum advantage. These datasets have low-dimensional structure that classical methods handle efficiently. Quantum advantage, if it exists for classification, will emerge on datasets with specific quantum structure -- high-dimensional data with non-trivial entanglement-like correlations that classical kernels cannot efficiently capture. Identifying or constructing such datasets is an important research direction, but the honest assessment is that no naturally occurring dataset with this structure has been convincingly demonstrated.

5.3 Error Mitigation Is Essential but Insufficient

Error mitigation recovered 42-68% of noise-induced performance loss -- enough to make some NISQ experiments informative but not enough to achieve noise-free performance. The residual noise after mitigation means that NISQ quantum ML results should be interpreted as lower bounds on what fault-tolerant quantum ML could achieve, not as representative of quantum ML's true capability. For practical applications, the current recommendation is to use noiseless simulation for algorithm development and real hardware only for validation and noise characterisation.

5.4 Limitations

Four limitations apply. First, the maximum qubit count (27-30) is small; larger devices may show different scaling behaviour. Second, the quantum advantage threshold estimates are extrapolations that depend on assumptions about classical algorithm improvements. Third, error mitigation methods add computational overhead that was not included in the cost comparison. Fourth, the study did not evaluate quantum ML for quantum data (data generated by quantum processes), where advantages may be more clear-cut.

6. Conclusion

6.1 Summary

Evaluation of five quantum ML approaches confirms: no quantum advantage for classical ML tasks at current scale (quantum kernel -1.4% vs. XGBoost). Genuine advantage demonstrated for quantum chemistry (VQE achieves chemical accuracy for small molecules where classical DFT fails). NISQ noise degrades performance by

8.4-22.6%, with error mitigation recovering 42-68%. Quantum advantage threshold estimated at 50-100 logical qubits for chemistry, 200+ for general ML.

6.2 Recommendations

Three recommendations emerge. First, focus near-term quantum ML efforts on quantum chemistry and materials science, where quantum advantage is theoretically grounded and empirically demonstrated. Second, do not deploy quantum ML for classical data tasks -- classical methods are faster, cheaper, and equally or more accurate at all currently accessible scales. Third, invest in error mitigation and fault-tolerant quantum computing research, because practical quantum ML advantage for general tasks requires hardware that does not yet exist.

References

- Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A. and Woerner, S. (2021) The power of quantum neural networks. *Nature Computational Science*, 1(6), pp. 403-409.
- Arute, F., Arya, K., Babbush, R. et al. (2019) Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), pp. 505-510.
- Benedetti, M., Lloyd, E., Sack, S. and Fiorentini, M. (2019) Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4), 043001.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N. and Lloyd, S. (2017) Quantum machine learning. *Nature*, 549(7671), pp. 195-202.
- Cerezo, M., Arrasmith, A., Babbush, R. et al. (2021) Variational quantum algorithms. *Nature Reviews Physics*, 3(9), pp. 625-644.
- Cerezo, M., Sone, A., Volkoff, T., Cincio, L. and Coles, P.J. (2021) Cost function dependent barren plateaus in shallow parameterized quantum circuits. *Nature Communications*, 12(1), 1791.
- Havlicek, V., Coles, A.D., Chowdhury, K., Shehab, O., Larose, R. and Bernien, H. (2019) Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747), pp. 209-212.
- Huang, H.Y., Broughton, M., Mohseni, M. et al. (2021) Power of data in quantum machine learning. *Nature Communications*, 12(1), 2631.
- Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J.M. and Gambetta, J.M. (2017) Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671), pp. 242-246.
- Li, Y. and Benjamin, S.C. (2017) Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2), 021050.
- Liu, Y., Arunachalam, S. and Temme, K. (2021) A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 17(9), pp. 1013-1017.
- McClean, J.R., Boixo, S., Smelyanskiy, V.N., Babbush, R. and Neven, H. (2018) Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812.
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S.C. and Yuan, X. (2020) Quantum computational chemistry. *Reviews of Modern Physics*, 92(1), 015003.
- Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.H., Zhou, X.Q., Love, P.J., Aspuru-Guzik, A. and O'Brien, J.L. (2014) A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1), 4213.
- Preskill, J. (2018) Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79.
- Schuld, M. and Killoran, N. (2019) Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4), 040504.
- Schuld, M. and Petruccione, F. (2021) *Machine Learning with Quantum Computers*. Springer.
- Sim, S., Johnson, P.D. and Aspuru-Guzik, A. (2019) Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12), 1900070.
- Tang, E. (2019) A quantum-inspired classical algorithm for recommendation systems. *Proceedings of STOC 2019*, pp. 217-228.
- Temme, K., Bravyi, S. and Gambetta, J.M. (2017) Error mitigation for short-depth quantum circuits. *Physical Review Letters*, 119(18), 180509.
- Zoufal, C., Lucchi, A. and Woerner, S. (2019) Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Information*, 5(1), 103.

Declarations

Funding

This research was supported by the Swiss National Science Foundation under grant 200021-212648 (QuantumML) and the EU Quantum Flagship programme under grant 820474 (MILLENNIUM). IBM Quantum hardware access provided through the IBM Quantum Network Academic Programme. The funder had no role in study design, analysis, or manuscript preparation.

Conflict of Interest

The author declares no conflicts of interest.

Data Availability Statement

All quantum circuit implementations, classical baselines, noise analysis scripts, and benchmark results are deposited in the Zenodo repository under DOI: 10.5281/zenodo.19614706. PennyLane and Qiskit are open-source. UCI datasets and QM9 are publicly available. Released under MIT License (code) and CC BY 4.0 (results).

Ethical Approval

This study involved computational experiments on quantum simulators and quantum hardware. No human subjects, personal data, or biological materials were involved. No institutional ethical approval was required.

Appendix A

Quantum ML Readiness Assessment Framework

This appendix provides a decision tree for assessing whether quantum ML is appropriate for a given task based on problem structure, available hardware, and expected advantage.

Step 1: Is the Problem Quantum in Nature?

- 1a. Yes (quantum chemistry, materials, condensed matter) --> Quantum ML likely advantageous; go to Step 2
- 1b. No (classical ML task on classical data) --> Quantum ML unlikely advantageous at current scale; use classical
- 1c. Unsure (data may have quantum-like structure) --> Benchmark on small instance; go to Step 2 if promising

Step 2: Hardware Assessment

- 2a. Fault-tolerant quantum computer available (future) --> Full quantum algorithms feasible
- 2b. NISQ device (current, 27-100 qubits) --> Shallow variational circuits only; expect 8-23% noise degradation
- 2c. Simulator only (no hardware access) --> Algorithm development and scaling analysis; not deployment

Step 3: Practical Recommendation

- 3a. Quantum chemistry with < 50 electrons --> VQE on NISQ with error mitigation (demonstrated advantage)
- 3b. Classical ML task --> Use classical methods (no quantum advantage demonstrated; 100-1000x cheaper)
- 3c. Research exploration --> Benchmark quantum vs. classical at small scale; extrapolate scaling