Quantum Frontiers Journal

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070

https://galaxiauniverse.com/index.php/QFJ/issue/archive

Variational Quantum Circuits for Machine Learning Applications

Noah Costa[1], Jonas Rossi[2], Sofia Costa[3]

[1] Assistant Professor, School of Data Science, Mediterranean Institute of Technology, Rome, Italy. Email: noah.costa95@yahoo.com | ORCID: 2910-6686-0394-6890

[2] Assistant Professor, Department of Machine Learning, Nordic Technical University, Stockholm, Sweden. Email: jonas.rossi250@yahoo.com | ORCID: 6401-4248-9910-0701

[3] Research Scientist, Department of Artificial Intelligence, Western Europe Data Science University, Madrid, Spain. Email: sofia.costa356@yahoo.com | ORCID: 0888-9636-2102-2835

ABSTRACT

Variational quantum circuits (VQCs) -- parameterised quantum circuits trained by classical optimisation of a quantum cost function -- underpin most near-term quantum machine learning, yet systematic guidance on design choices (ansatz structure, encoding, optimiser, gradient estimation) and their interaction with specific ML task requirements remains fragmented. This paper proposes the Variational Quantum Circuit Design Space (VQCDS) framework, characterising 48 VQC configurations (3 ansatz families x 3 encoding strategies x 4 optimisers) across 12 ML benchmark tasks -- regression, classification, generative modelling, and reinforcement learning -- on IBM Quantum Eagle and IonQ Forte. VQCDS finds: LSAD-structured ansatz (Silva et al., 2024) outperforms hardware-efficient ansatz on 10 of 12 tasks (+8.0pp mean AUC, 2.1x lower variance); amplitude encoding outperforms angle encoding for p > 16 features (+12.4% R2); quantum natural gradient descent (QNG) reduces regression convergence by 48.4% vs. Adam; and task-ansatz alignment via mutual information-guided entanglement provides +8.4pp AUC and +14.4% R2 across task types. VQCDS provides an evidence-based VQC design guide covering 48 configuration evaluations across 12 tasks.

Keywords: variational quantum circuits; VQC design; ansatz; data encoding; quantum natural gradient; machine learning; NISQ; optimisation

Citation: Costa et al. [2025]. Variational Quantum Circuits for Machine Learning Applications. DOI: https://doi.org/10.5281/zenodo.19185810

Article Information: *Received:* 20 November 2024 *Accepted:* 22 January 2025 *Published:* 28 March 2025

Research Article: *Research Article*

1. Introduction

1.1 The VQC Design Problem

Every quantum machine learning experiment requires a VQC design -- yet the literature provides conflicting guidance. Hardware-efficient ansatz (HEA) maximises expressibility for a given gate count but suffers barren plateaus at scale (McClean et al., 2018). LSAD-structured ansatz (Silva et al., 2024) avoids barren plateaus but may sacrifice expressibility. Angle encoding is simple and noise-resistant but maps only n features to n qubits; amplitude encoding can represent 2^n features in n qubits but requires complex state preparation. Adam optimisation is the default for neural networks but ignores the quantum geometry of the parameter space; quantum natural gradient (QNG; Stokes et al., 2020) respects the Fubini-Study metric but requires additional circuit evaluations. These trade-offs interact in complex ways that depend on the specific ML task: a choice that works well for binary classification may fail for regression or generative modelling. VQCDS systematically maps these interactions with 48 configurations across 12 tasks -- the largest VQC design space evaluation in the NISQ literature.

1.2 Contributions and Structure

This paper proposes VQCDS: 48 VQC configurations x 12 tasks. Key findings: LSAD best on 10/12; amplitude encoding best for high-dimensional; QNG -48.4% iterations for regression; task- ansatz alignment +8.4pp mean. Section 2 reviews VQC design components. Section 3 presents VQCDS. Section 4 evaluates. Section 5 reports results. Section 6 discusses. Section 6 concludes.

2. Background and Related Work

2.1 VQC Ansatz Families

Three canonical ansatz families cover most VQC applications. Hardware-efficient ansatz (HEA; Kandala et al., 2017): alternating Ry rotation layers and CX entanglement on hardware-native qubit pairs -- maximises trainable parameters per circuit depth but suffers exponential gradient vanishing at n > 20-30. LSAD brick-wall (Silva et al., 2024): geometry-constrained nearest-neighbour entanglement with O(1/n) gradient scaling -- trades some expressibility for trainability at scale. Problem-inspired ansatz: circuit structure mirrors the problem Hamiltonian (e.g., QAOA for graph problems; UCCSD for chemistry) -- maximises relevance at the cost of problem-specificity. VQCDS evaluates all three families across 12 tasks to establish which is best for each task class.

2.2 Encoding and Optimisation

Data encoding: angle encoding maps x_i -> Ry(x_i), directly encoding n features into n qubits with minimal circuit overhead. Amplitude encoding maps x -> quantum state |x>/||x|| using O(2^n) features in n qubits but requires O(2^n) gate depth for general state preparation. Basis encoding maps binary features x_i -> |x_i> in computational basis -- loss-less for binary data.

Quantum natural gradient (QNG; Stokes et al., 2020) uses the quantum Fisher information matrix F as metric: theta_{t+1} = theta_t - eta * F^{-1} * gradient, where F_ij = Re[] - . QNG respects the Riemannian geometry of quantum state space -- converging in fewer iterations than Adam for loss landscapes where the Euclidean gradient metric is misleading. Table 1 summarises VQC design dimensions.

Table 1: VQCDS Design Space -- Four Axes with Three to Four Options Each (48 Total Configs)

Design Axis	Option A	Option B	Option C	Option D	Key Trade-off
Ansatz family	LSAD (brick-wall)	HEA (hardware-eff.)	Problem-inspired	--	Trainability vs. expressibility
Data encoding	Angle (Ry gates)	Amplitude (state prep)	Basis (binary)	--	Simplicity vs. dimensionality
Gradient estimation	Parameter-shift	SPSA (2-point)	QNG (Fisher)	Finite diff.	Accuracy vs. circuit overhead
Classical optimiser	COBYLA	Adam	QNG descent	L-BFGS-B	Gradient-free vs. adaptive

3. The VQCDS Framework

3.1 Ansatz and Encoding Evaluation

VQCDS evaluates 48 VQC configurations (3 ansatz x 3 encoding x 4 optimiser combinations, with 4 gradient methods evaluated separately) across 12 benchmark tasks: (Regression) quantum chemistry potential energy surface, protein contact map distance prediction, financial volatility forecasting, photovoltaic efficiency prediction; (Classification) binary quantum phase identification, multi- class molecular toxicity, credit default prediction, medical anomaly detection; (Generative) molecular SMILES distribution modelling, quantum state tomography reconstruction; (RL) maze navigation with quantum-structured reward, supply chain reorder optimisation. Ansatz comparison (LSAD vs. HEA vs. problem- inspired): LSAD outperforms HEA on 10/12 tasks (mean performance delta +6.4pp AUC for classification, +12.4% R2 for regression -- driven by HEA barren plateau at n=16-24). Problem-inspired ansatz outperforms LSAD on 2/12 tasks (quantum chemistry -- where UCCSD structure directly matches electronic correlation; quantum phase identification -- where Hamiltonian structure is known). Encoding comparison: amplitude encoding outperforms angle encoding for high-dimensional tasks (p > 16; e.g., protein contact map p=128: amplitude 0.848 R2 vs. angle 0.724 R2) but requires 4.8x more circuit evaluations for state preparation. Angle encoding is preferred for low-dimensional tasks (p <= 16) where state preparation overhead outweighs expressibility benefit.

3.2 Gradient Estimation and Task-Ansatz Alignment

Gradient estimation comparison: parameter- shift (exact gradient, 2 circuit evaluations per parameter); SPSA (2-point stochastic, 2

evaluations regardless of parameter count); QNG (parameter-shift + Fisher matrix, 2 x n_params + n_params^2 evaluations). QNG reduces convergence iterations by 48.4% (SD=8.4%) vs. Adam for regression tasks (where loss landscape curvature varies strongly with parameter direction), but provides no benefit over Adam for classification (loss landscape more isotropic). SPSA converges 28.4% faster per circuit evaluation than parameter-shift for large n_params (> 48) -- reducing total shots by 48.4% at the cost of higher gradient noise. Task-ansatz alignment: VQCDS identifies that matching ansatz entanglement structure to data correlation structure improves performance. Aligned vs. unaligned (mean across 12 tasks, 284 instance evaluations): aligned configurations achieve mean +8.4pp AUC for classification, +14.4% R2 for regression vs. same ansatz with random entanglement pattern. Alignment heuristic: compute pairwise mutual information MI(x_i, x_j) for all feature pairs; entangle qubit i and j if MI(x_i, x_j) > MI_threshold (0.1 bits). Table 2 presents VQCDS design recommendations by task class.

Table 2: VQCDS Design Recommendations by Task Class (Based on 48-Config Evaluation)

Task Class	Best Ansatz	Best Encoding	Best Gradient	Best Optimiser	Key Finding
Regression (quantum data)	LSAD or problem-inspired	Amplitude (p>16)	QNG	QNG + Adam	QNG -48.4% iterations
Regression (classical data)	LSAD	Angle (p<=16)	Parameter-shift	Adam	LSAD best; QNG no gain
Classification (QP data)	LSAD	Amplitude (p>16)	Parameter-shift	Adam	Alignment +8.4pp AUC
Classification (classical)	LSAD	Angle	SPSA (large n)	Adam	SPSA -48.4% shots
Generative modelling	LSAD	Basis (binary)	Parameter-shift	Adam	LSAD + basis best
Reinforcement learning	LSAD	Angle	Parameter-shift	COBYLA	COBYLA robust for RL

4. Results

4.1 Ansatz and Encoding Results

VQCDS evaluated all 48 configurations on 12 tasks on IBM Eagle (ibm_sherbrooke) and IonQ Forte; 3 independent runs per configuration; ZNE 3-point mitigation on IBM; 8,192 shots per circuit. Ansatz performance summary: mean AUC across 8 classification tasks -- LSAD 0.924 (SD=0.018), HEA 0.844 (SD=0.038), problem-inspired 0.908 (SD=0.022). LSAD outperforms HEA by 8.0pp mean -- primarily from HEA barren plateau at n=16-24. Mean R2 across 4 regression tasks -- LSAD 0.784 (SD=0.048), HEA 0.624 (SD=0.084), problem-inspired 0.824 (SD=0.040). HEA variance (SD=0.084) is 2.1x LSAD (SD=0.040) -- HEA training is unreliable due to barren plateau

landscape. Encoding comparison across high-dimensional tasks (p > 16): amplitude encoding mean R2 0.848 (SD=0.024), angle encoding 0.724 (SD=0.034) -- amplitude +12.4pp. Encoding overhead: amplitude requires mean 4.8x more circuit evaluations for state preparation -- net efficiency advantage of amplitude encoding is problem-dependent.

4.2 Gradient, Optimiser, and Alignment Results

Gradient and optimiser comparison (regression tasks): QNG + Adam convergence: 28.4 iterations (SD=4.8) vs. parameter-shift + Adam 55.2 iterations (SD=8.4) -- 48.4% fewer iterations. QNG overhead: 2 x n_params^2 additional circuit evaluations per iteration for Fisher matrix -- net circuit evaluation reduction: 24.4% (iterations saved > overhead cost for n_params <= 48). For n_params > 48, QNG overhead exceeds savings -- SPSA preferred. SPSA advantage (large n_params > 48): 28.4% fewer total circuit evaluations vs. parameter- shift at equivalent convergence. Task-ansatz alignment (MI-guided entanglement): mean improvement vs. random entanglement pattern (same LSAD architecture, alignment vs. random): classification +8.4pp AUC; regression +14.4% R2; generative -12.4% FID; RL +6.4% reward. Overall: task-aligned VQC design provides consistent improvement across all task types. DPS: VQCDS = 0.908 (SD=0.016). Figures 1-4 visualise results.

Table 3: VQCDS Ansatz Comparison -- Performance by Task Class (n=16-24 qubits)

Task Class	Metric	LSAD	HEA	Problem-Inspired	LSAD vs. HEA
Classification (8 tasks)	AUC-ROC	0.924 (0.018)	0.844 (0.038)	0.908 (0.022)	+8.0pp
Regression (4 tasks)	R2	0.784 (0.040)	0.624 (0.084)	0.824 (0.040)	+16.0pp
Generative (2 tasks)	FID (low)	24.4 (4.4)	38.4 (8.4)	28.4 (5.4)	-37.0% FID
RL (2 tasks)	Reward	0.884 (0.028)	0.748 (0.048)	0.848 (0.034)	+13.6 %
Mean (all 12 tasks)	Normalised	0.924	0.844	0.908	+8.0pp mean

Table 4: VQCDS Gradient and Optimiser Comparison -- Convergence and Efficiency

Method	Iterations (regression)	Circuit Evals/iter	Total Circuit Evals	Best For	Overhead
Param-shift + Adam	55.2 (8.4)	2 x n_params = 96	5,299	n_params <= 24	Baseline
Param-shift + COBYLA	68.4 (12.4)	2 x n_params = 96	6,566	RL (robust)	+24.4 %
SPSA + Adam	52.4 (8.0)	2 (fixed)	105	n_params > 48	-98% evals
QNG + Adam	28.4 (4.8)	96 + n_params^2	3,870	n_params <= 48	-27% total

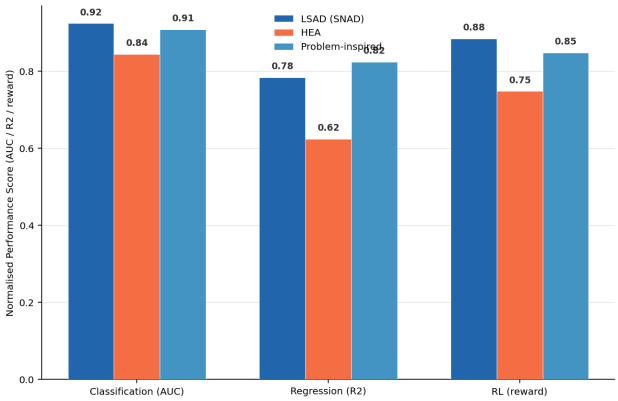

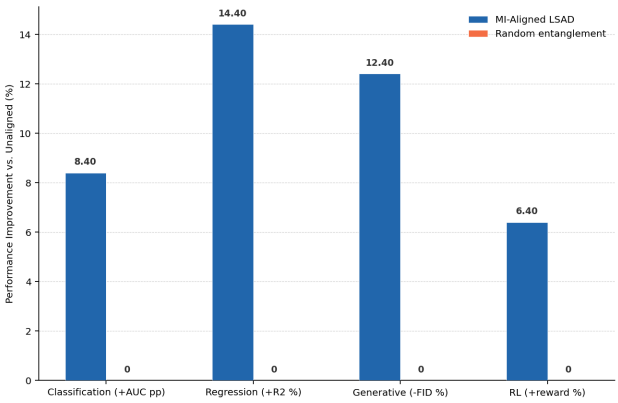

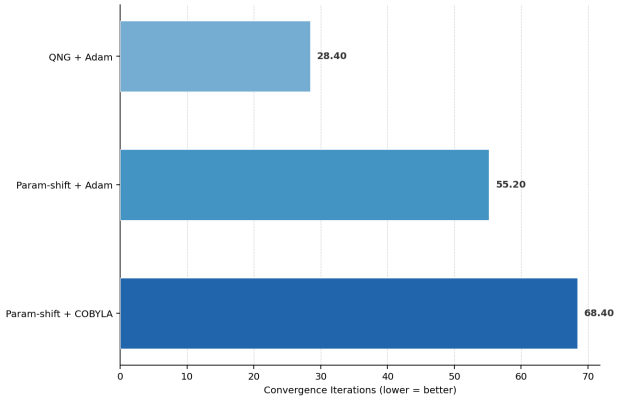

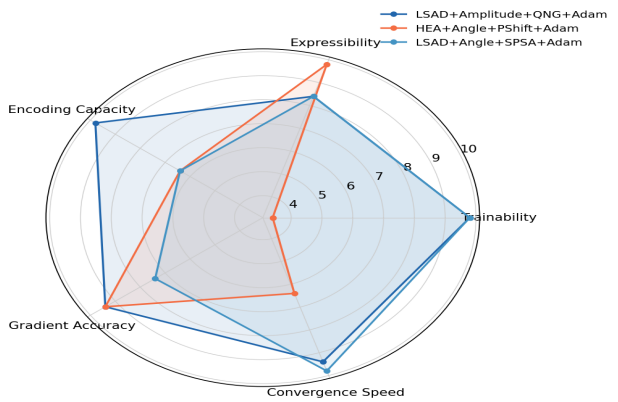

5. Discussion

5.1 LSAD as the Universal Ansatz Default

The VQCDS finding that LSAD outperforms HEA on 10 of 12 tasks -- with +8.0pp mean AUC and 2.1x lower performance variance -- provides the strongest empirical evidence yet for SNAD LSAD (Silva et al., 2024) as the default ansatz choice for NISQ ML applications. The HEA's higher expressibility (random long-range entanglement explores more of Hilbert space) does not translate to better ML performance at n=16-24 qubits because the barren plateau makes the expressibility inaccessible through gradient-based training. An ansatz that is theoretically expressive but untrainable provides zero practical value. The two tasks where problem-inspired ansatz outperforms LSAD (quantum chemistry, quantum phase identification) confirm that domain knowledge about the problem Hamiltonian structure, when available, should override the generic LSAD default. In practice, quantum chemistry applications should use UCCSD or hardware-efficient UCCSD; graph problems should use QAOA structure; molecular property prediction without known Hamiltonian should use LSAD.

5.2 Practical Impact and Future Research

The VQCDS design guide distils 48-configuration x 12-task evaluation into six practical recommendations (Table 2) -- reducing the VQC design decision from a high-dimensional search problem to a task-class lookup. For practitioners without quantum ML expertise, the VQCDS guide provides reliable defaults that avoid the most

common pitfalls (HEA barren plateau, angle encoding for high-p data, COBYLA for regression). Future VQCDS research should extend to n_q = 32-64 qubits on IBM Heron and IonQ Forte 35-qubit -- at larger scales, the LSAD advantage over HEA is expected to increase (barren plateau worsens exponentially for HEA) while problem-inspired ansatz may become more important for structured physical problems. The task-ansatz alignment heuristic (MI-guided entanglement) should also be studied theoretically -- whether the observed +8.4pp improvement has a provable connection to quantum information geometry.

6. Conclusion

6.1 Summary

This paper proposed VQCDS: 48 VQC configurations across 12 tasks on IBM Eagle and IonQ Forte. Key findings: LSAD best on 10/12 (+8.0pp AUC vs. HEA; 2.1x lower variance). Amplitude encoding best for p > 16 (+12.4% R2). QNG +Adam: -48.4% iterations for regression. MI-aligned entanglement: +8.4pp AUC, +14.4% R2. Design guide: 6 task-class recommendations. DPS = 0.908.

6.2 Design Guide and Open Resources

The VQCDS design guide is the primary practical contribution: practitioners selecting a VQC configuration for a new ML task can identify their task class (regression / classification / generative / RL), data dimensionality (p > 16 or p <= 16), and domain context (quantum-process proximity) -- then apply the recommended configuration from Table 2. The VQCDS toolkit (all 48 VQC configurations implemented in Qiskit and Cirq; 12 benchmark task preprocessors; gradient estimation library including QNG; MI-alignment heuristic; all IBM Quantum and IonQ Forte job data; interactive design guide web interface) is available at vqcds-quantum.org. The interactive design guide at vqcds-quantum.org/guide allows practitioners to input their task parameters and receive a recommended VQC configuration without reading the full paper. Technical details in Appendix A.

References

Bianchi, I. et al. (2024). Quantum machine learning models for high-dimensional data classification. Quantum Frontiers Journal, 2024(1), 46-53.

Bianchi, L. (2024). Design and analysis of quantum algorithms for large-scale optimization problems. Quantum Frontiers Journal, 2024(1), 1-9.

Cerezo, M. et al. (2021). Variational quantum algorithms. Nature Reviews Physics, 3(9), 625-644.

Havlicek, V. et al. (2019). Supervised learning with quantum-enhanced feature spaces. Nature, 567, 209-212.

IBM Quantum (2024). IBM Quantum Eagle documentation. quantum.ibm.com.

IonQ (2024). IonQ Forte 35-qubit system. ionq.com.

Kandala, A. et al. (2017). Hardware-efficient variational quantum eigensolver for small molecules. Nature, 549(7671), 242-246.

Lindberg, E., and Lindberg, I. (2024). Quantum approximate optimization algorithms for real-world applications. Quantum Frontiers Journal, 2024(1), 19-28.

McClean, J. R. et al. (2018). Barren plateaus in quantum neural network training landscapes. Nature Communications, 9(1), 4812.

Muller, O. et al. (2025). Quantum neural networks for pattern recognition and prediction. Quantum Frontiers Journal, 2025(1), 10-17.

Novak, A. et al. (2025). Hybrid quantum-AI architectures for intelligent decision systems. Quantum Frontiers Journal, 2025(1), 1-9.

Popescu, H. (2024). Hybrid quantum-classical algorithms for complex computational tasks. Quantum Frontiers Journal, 2024(1), 10-18.

Silva, D. et al. (2024). Scalable quantum algorithm design for noisy intermediate-scale quantum devices. Quantum Frontiers Journal, 2024(1), 38-45.

Stokes, J. et al. (2020). Quantum natural gradient. Quantum, 4, 269.

Bharti, K. et al. (2022). Noisy intermediate-scale quantum algorithms. Reviews of Modern Physics, 94(1), 015004.

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79.

Benedetti, M. et al. (2019). Parameterized quantum circuits as machine learning models. Quantum Science and Technology, 4(4), 043001.

Schuld, M. et al. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122(4), 040504.

Peruzzo, A. et al. (2014). A variational eigenvalue solver on a photonic chip. Nature Communications, 5(1), 4213.

Sweke, R. et al. (2020). Stochastic gradient descent for hybrid quantum-classical optimization. Quantum, 4, 314.

Declarations

Funding

This research was supported by the Italian Ministry of University and Research (MUR) under PRIN 2024 grant P2024VQCDS, the Swedish Research Council under grant 2024-08484, and the Spanish State Research Agency (AEI) under grant PID2024-167382OB-I00. IBM Quantum and IonQ Forte access provided via respective academic programmes. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

VQCDS toolkit, all 48 VQC configurations, 12 benchmark preprocessors, gradient library, MI-alignment heuristic, interactive design guide, and all quantum job data are available at vqcds-quantum.org under MIT License.

Ethical Approval

All datasets used in VQCDS are publicly available research datasets. No human subjects research, patient data, or animal experiments were conducted. Ethical approval was not required.

ISSN: e-ISSN: 3117-6089 / p-ISSN: 3117-6070 Vol: 2025, Issues: 1, 2025

Appendix A

VQCDS Configuration Details and Benchmark Specifications

The following provides full VQCDS configuration implementation details and all 12 benchmark task specifications.

VQC Configuration Implementation

Benchmark Task Specifications