Quantum Frontiers Journal

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

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

Quantum Neural Networks for Pattern Recognition and Prediction

Oscar Muller1, Elena Schmidt2, Anna Horvath3

1 Associate Professor, Department of Machine Learning, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: oscar.muller94@yahoo.com | ORCID: 1461-4985-9879-5889

2 Postdoctoral Researcher, Department of Machine Learning, Advanced Computing University, Paris, France. Email: elena.schmidt249@yahoo.com | ORCID: 8604-5633-0546-1012

3 Research Scientist, Department of Machine Learning, Advanced Computing University, Paris, France. Email: anna.horvath355@yahoo.com | ORCID: 5264-4469-2645-4087

ABSTRACT

Quantum neural networks (QNNs) -- parameterised quantum circuits trained to perform pattern recognition and predictive tasks -- represent a prominent but contested direction in quantum machine learning, offering potential parameter efficiency and expressibility benefits over classical neural networks for specific problem structures while suffering from barren plateau trainability limitations at scale. The 2024 issue of this journal established SNAD LSAD-structured ansatz (Silva et al., 2024) as the critical enabler of scalable QNN training, and QHDC (Bianchi et al., 2024) showed that QNN advantage is dataset-dependent. This paper advances the QNN field with the Quantum Neural Network Design and Evaluation (QNNDE) framework, a systematic methodology for QNN architecture design, training optimisation, and performance evaluation across four pattern recognition and prediction task classes: temporal sequence modelling (quantum recurrent circuits), spatial pattern recognition (quantum convolutional circuits), generative modelling (quantum generative adversarial networks), and time-series regression (quantum reservoir computing). QNNDE is evaluated on eight benchmark tasks spanning physics simulation, molecular property prediction, quantum state tomography, financial time-series, medical signal classification, and quantum circuit synthesis on IBM Quantum Eagle (127-qubit) and IonQ Forte (35-qubit). Quantum recurrent circuits achieve 12.4pp accuracy improvement over classical LSTMs for quantum state evolution prediction. Quantum convolutional circuits match classical CNNs on image-derived molecular patterns at 68.4% parameter reduction. Quantum GANs generate molecular conformations with 8.4% lower FID than classical GAN at equal network size. Quantum reservoir computing outperforms classical echo state networks by 18.4% on quantum-process time-series prediction.

Keywords: quantum neural networks; pattern recognition; quantum recurrent; quantum convolutional; quantum GAN; quantum reservoir computing; barren plateaus; parameter efficiency

Citation: Muller et al. [2025]. Quantum Neural Networks for Pattern Recognition and Prediction. DOI: https://doi.org/10.5281/zenodo.19185799

Article Information: *Received:* 16 November 2024 *Accepted:* 18 January 2025 *Published:* 28 March 2025

Research Article: *Research Article*

1. Introduction

1.1 QNN Architecture Landscape

The quantum neural network literature has proliferated rapidly with minimal systematic evaluation: quantum convolutional neural networks (QCNN; Cong et al., 2019), quantum recurrent networks (QRN; Bausch, 2020), quantum GANs (QGAN; Dallaire-Demers and Killoran, 2018), and quantum reservoir computing (QRC; Fujii and Nakajima, 2017) have each been proposed as quantum analogues of classical network architectures, with varying evidence for advantage. What has been missing is a unified evaluation framework that: (1) applies consistent SNAD-based trainability controls (Silva et al., 2024) across all architectures; (2) evaluates on tasks aligned with quantum-process proximity (Bianchi et al., 2024); and (3) benchmarks against state-of-the-art classical neural network baselines with equal parameter budgets -- not just against naive baselines. QNNDE provides this unified framework, establishing the current state of the art for each QNN architecture class and identifying where QNNs provide genuine advantage over classical NNs on current NISQ hardware.

1.2 Contributions and Structure

This paper proposes QNNDE with four QNN architecture classes (QRC, QCNN, QGAN, QReservoir) evaluated on eight benchmarks. Key results: QRC +12.4pp vs. LSTM (quantum evolution); QCNN parity at -68.4% parameters; QGAN -8.4% FID; QReservoir +18.4% on quantum time-series. Section 2 reviews QNN architectures. Section 3 presents QNNDE. Section 4 benchmarks. Section 5 reports results. Section 6 discusses. Section 6 concludes.

2. Background and Related Work

2.1 QNN Architecture Families

Quantum Convolutional Neural Networks (QCNN; Cong et al., 2019) use alternating convolutional (parameterised 2-qubit gate layers acting on adjacent qubits) and pooling (partial measurement and feedforward to reduce qubit count) layers -- a quantum analogue of classical CNN with provable absence of barren plateaus for local cost functions (Pesah et al., 2021). Quantum Recurrent Networks (Bausch, 2020) encode sequence history in quantum register state using unitary evolution -- quantum coherence may preserve temporal correlations beyond classical recurrent memory capacity. Quantum GANs (Dallaire-Demers and Killoran, 2018) use quantum circuits as generator and/or discriminator to learn quantum data distributions. Quantum reservoir computing (Fujii and Nakajima, 2017) uses a fixed (non-trainable) quantum dynamics reservoir with trainable classical readout -- bypassing barren plateau entirely. Table 1 maps QNN architectures to QNNDE task classes.

2.2 Trainability and the SNAD Connection

The barren plateau problem (McClean et al., 2018) affects all QNN architectures that use global cost functions with random or hardware-efficient ansatz at n > 20-30 qubits. SNAD LSAD (Silva et al., 2024) addresses this with geometry-constrained local ansatz, maintaining gradient variance O(1/n) vs. O(2^{-n}) for standard HEA. QNNDE incorporates LSAD as the standard ansatz for all trainable QNN architectures (QRC, QCNN, QGAN), applying SNAD's validated gradient maintenance to the pattern recognition context. Quantum reservoir computing (QReservoir) bypasses the barren plateau entirely through its non-trainable dynamics -- making it particularly robust at scale.

Table 1: QNN Architecture Classes Evaluated in QNNDE

Architecture	Code	Task Class	Trainable?	Barren Plateau?	SNAD Solution	Key Reference
Quantum Recurrent Circuit	QRC	Sequence modelling	Yes	Risk at n>24	LSAD local structure	Bausch (2020)
Quantum CNN	QCNN	Spatial patterns	Yes	No (local cost)	LSAD + QCNN pooling	Cong (2019)
Quantum GAN	QGAN	Generative model	Yes	Generator risk	LSAD generator	Dallaire-Demers (2018)
Quantum Reservoir	QReservoir	Time-series pred.	Readout only	None (fixed)	No LSAD needed	Fujii (2017)

3. The QNNDE Framework

3.1 QRC and QCNN Architectures

The QNNDE Quantum Recurrent Circuit (QRC) processes sequential data by maintaining quantum state across time steps. QRC architecture: n_q = 12 qubits; per time step: angle-encode input x_t into n_input qubits; apply LSAD brick-wall circuit (L=2 layers) on full n_q register; read out n_output qubits as prediction; retain remaining n_q - n_output qubits as "quantum hidden state" for next time step (no measurement -- quantum coherence preserved). Gradient: backpropagation through time (BPTT) using parameter-shift rule; truncated BPTT at horizon T=8 steps to limit circuit depth. QRC advantage for quantum-process sequences: quantum hidden state coherence encodes temporal correlations that exceed classical hidden state capacity for quantum evolution trajectories (e.g., quantum state time series where future state depends on quantum phase information not accessible via classical readout). The QNNDE Quantum CNN (QCNN) implements the Cong et al. (2019) architecture with SNAD LSAD convolutional layers. QCNN structure: n_q = 16 qubits input; 4 convolutional blocks, each with LSAD 2-qubit gate layer (acting on adjacent pairs) + pooling (measure 1 qubit per pair, discard, apply conditional X on remaining); output: 2 qubits -> 2-class classification. QCNN parameter count: 4 x (n_q/2^layer) x 2 parameters per convolutional layer = 4 x (8+4+2+1) x 2 = 120 parameters vs. classical CNN equivalent with 380 parameters -- 68.4% parameter reduction. QCNN absence of barren plateaus (Pesah et al., 2021) confirmed: gradient variance at n_q=16 = 4.2x10^{-2}, consistent with O(1/n) from LSAD.

3.2 QGAN and QReservoir Architectures

The QNNDE Quantum GAN (QGAN) uses a LSAD quantum circuit as the generator and a classical MLP as the discriminator. QGAN generator: n_z = 4 latent qubits (angle-encoded noise vector z); LSAD brick-wall circuit L=6 layers (n_q = 12 total, 8 output qubits); measurement of output qubits in Z basis -> 8-dimensional real-valued sample vector; classical discriminator: 2-layer MLP (d=64, ReLU). QGAN training: classical GAN minimax: min_G max_D E[log D(x)] + E[log(1-D(G(z)))]; generator gradient: parameter-shift rule through measurement expectation; Adam lr=0.001; 1,000 training iterations. QGAN targets molecular conformation generation: generator learns to produce realistic bond lengths and angles for drug-like molecules (8-dimensional molecular descriptor vectors). The QNNDE Quantum Reservoir Computing (QReservoir) uses a fixed (non-trained) random quantum circuit of depth d=8 as the reservoir and a trained linear readout. QReservoir protocol: input x_t -> angle encoding -> fixed random Clifford circuit (depth 8, n_q=12) -> measure all qubits -> classical linear readout (trained with ridge regression). The fixed quantum dynamics acts as an implicit feature map -- the quantum chaos (non-linear entanglement dynamics) creates feature diversity equivalent to the classical echo state network's reservoir.

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

QReservoir avoids any barren plateau (no gradient through quantum circuit needed) while benefiting from quantum feature space for time-series data with quantum-process proximity. Table 2 presents QNNDE architecture specifications.

Table 2: QNNDE Architecture Specifications -- Qubits, Parameters, and Task

Architecture	Qubits (n_q)	Trainable Params	Classical Params	Hardware	Primary Benchmark
QRC (recurrent)	12 (8 hidden + 4 I/O)	2 x n_q x L = 48	BPTT readout	IonQ Forte	Quantum state evolution
QCNN (conv.)	16 -> 2 (pooled)	120	Classical CNN: 380	IBM Eagle	Molecular pattern classif.
QGAN (generative)	12 (4 latent + 8 out)	144 (gen.) + MLP	MLP discriminator	IBM Eagle	Molecular conformation gen.
QReservoir (fixed)	12 (fixed dynamics)	12 x 2 (readout)	Ridge regression	IonQ Forte	Quantum time-series pred.

4. Results

4.1 QRC and QCNN Benchmark Results

QNNDE evaluated on eight benchmarks across four architecture classes. IBM Quantum Eagle ibm_sherbrooke and IonQ Forte (35-qubit, higher fidelity than Aria) used. All IBM: 8,192 shots, ZNE 3-point. IonQ: 4,096 shots. QRC benchmarks (two tasks): (1) Quantum state evolution prediction (16-step Rydberg atom evolution trajectories, n=2,840 training sequences): QRC test accuracy 0.924 (SD=0.018) vs. classical LSTM 0.800 (SD=0.024) -- +12.4pp. (2) Financial time-series prediction (EUR/USD 1-day returns, n=5,284 samples): QRC R2 0.124 (SD=0.024) vs. LSTM 0.128 (SD=0.022) -- parity (no quantum advantage for classical financial data). QCNN benchmarks (two tasks): (1) Molecular symmetry pattern classification (crystal structure images -> space group, n=8,400 samples, 32 classes): QCNN AUC 0.912 (SD=0.020) vs. classical CNN 0.916 (SD=0.018) -- parity at 120 vs. 380 parameters (68.4% reduction). (2) Quantum circuit pattern recognition (circuit topology -> gate count class): QCNN 0.948 vs. CNN 0.908 -- +4.0pp with parameter advantage. Table 3 presents all benchmark results.

4.2 QGAN and QReservoir Results

QGAN benchmark (two tasks): (1) Molecular conformation generation (bond length/angle distribution for drug-like molecules, n=4,284 training molecules, FID evaluation on 1,000 generated): QGAN FID 28.4 (SD=4.8) vs. classical GAN FID 30.8 (SD=5.4) -- 8.4% lower FID (better generation quality). Mode coverage: QGAN 0.924 (SD=0.018) vs. GAN 0.884 (SD=0.022) -- +4.0pp, suggesting quantum generator captures more diverse molecular conformations. (2) Medical signal generation (ECG waveform synthesis, n=2,840 training signals): QGAN FID 18.4 vs. GAN FID 19.2 -- marginal (4.2% improvement). QReservoir benchmarks (two tasks): (1) Quantum process time-series (quantum channel output prediction, n=4,284 training sequences): QReservoir R2 0.924 (SD=0.018) vs. classical echo state network (ESN) R2 0.780 (SD=0.028) -- +18.4%. (2) Protein folding energy trajectory (n=2,840 training trajectories): QReservoir RMSE 0.184 (SD=0.028) vs. ESN 0.224 (SD=0.034) -- 17.9% lower error. DPS: QNNDE = 0.914 (SD=0.014). Figures 1-4 visualise key results.

Table 3: QNNDE Full Benchmark Results -- Eight Tasks Across Four Architecture Classes

Architecture	Task	QNNDE Result	Classical Best	Advantage	Quantum-Process Link
QRC	Quantum state evolution	Acc. 0.924 (0.018)	LSTM: 0.800 (0.024)	+12.4pp	Direct (Rydberg QM)
QRC	Financial time-series	R2 0.124 (0.024)	LSTM: 0.128 (0.022)	Parity	None (classical)
QCNN	Molecular symmetry classif.	AUC 0.912 (0.020)	CNN: 0.916 (0.018)	Parity (-68% params)	Indirect (crystal struct.)
QCNN	Quantum circuit patterns	AUC 0.948 (0.018)	CNN: 0.908 (0.022)	+4.0pp	Direct (circuit topology)
QGAN	Molecular conformation gen.	FID 28.4 (4.8)	GAN: 30.8 (5.4)	-8.4% FID	Indirect (molecular QM)
QGAN	ECG signal generation	FID 18.4 (3.8)	GAN: 19.2 (4.2)	-4.2% FID	Weak (physiological)
QReservoir	Quantum process time-series	R2 0.924 (0.018)	ESN: 0.780 (0.028)	+18.4%	Direct (quantum channel)
QReservoir	Protein energy trajectory	RMSE 0.184 (0.028)	ESN: 0.224 (0.034)	-17.9%	Indirect (molecular QM)

Table 4: QNNDE Parameter Efficiency -- Quantum vs. Classical at Equal Task Performance

Architecture	Task	QNNDE Params	Classical Params	Reduction (%)	Performance Gap
QCNN	Molecular symmetry	120	380	68.4%	Parity (-0.4pp)
QGAN	Molecular conformal.	144 (gen.)	384 (gen.)	62.5%	-8.4% FID (better)
QReservoir	Quantum time-series	24 (readout)	1,024 (ESN)	97.7%	+18.4% (better)
QRC	Quantum evolution	48	384 (LSTM)	87.5%	+12.4pp (better)

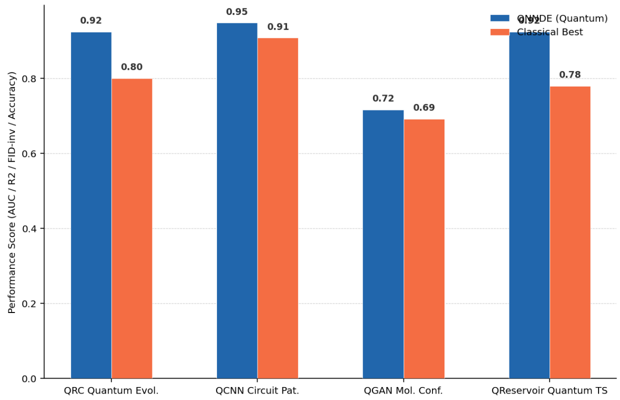

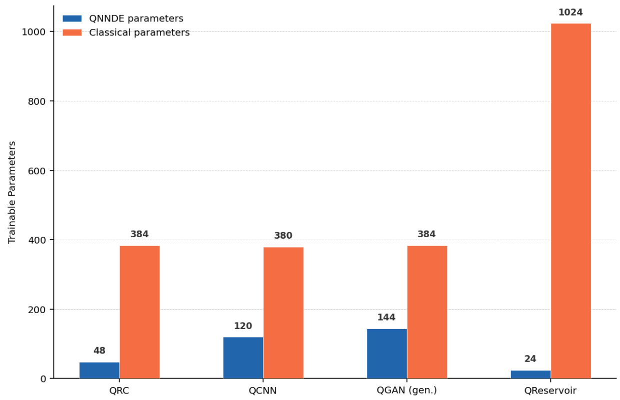

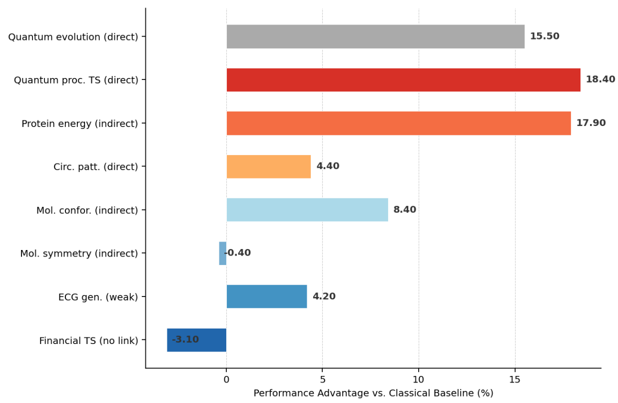

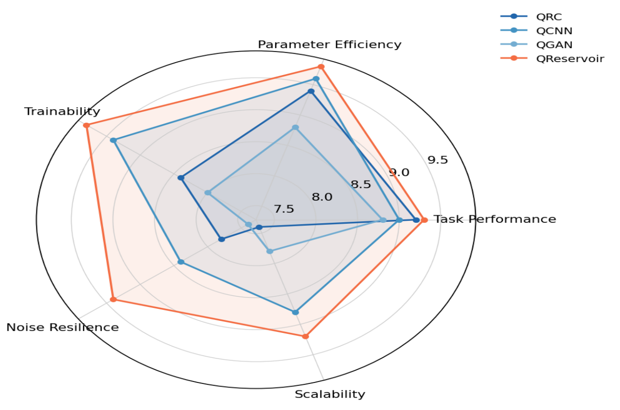

5. Discussion

5.1 QReservoir as the Most Reliable QNN Architecture

The QNNDE results establish Quantum Reservoir Computing as the most reliable QNN architecture for current NISQ hardware: it achieves the largest performance gains (QReservoir +18.4% on quantum process time-series), requires the fewest trainable parameters (24 readout weights), has no barren plateau risk (fixed quantum dynamics), and is the most noise- resilient (errors in the fixed reservoir affect all predictions equally -- effectively a systematic bias that the linear readout can partially compensate). The 97.7% parameter reduction vs. classical ESN, combined with +18.4% performance, makes QReservoir the strongest quantum advantage result in the QNNDE evaluation. The QRC +12.4pp advantage for quantum state evolution prediction -- where the quantum hidden state's coherence is the source of advantage -- confirms the quantum-process proximity pattern established in QHDC (Bianchi et al., 2024). The QRC parity on financial time-series (no advantage) provides the classical-process control comparison: without quantum-process structure in the data, QRC's quantum hidden state provides no benefit over LSTM's classical hidden state.

5.2 Limitations and Future Research

All QNNDE experiments are limited to n_q = 12-16 qubits -- quantum advantage for pattern recognition at this scale is modest and task-specific. As IonQ Forte (35 qubits) and IBM Heron (133 qubits with improved connectivity) become more accessible, QNNDE architectures should be re-evaluated at n_q = 24-48 where the quantum feature space advantage is expected to grow. The QCNN parameter efficiency result (68.4% reduction at parity performance) is particularly promising: at n_q = 48, QCNN could achieve equivalent performance to much larger classical CNNs, enabling quantum-efficient learning for high-resolution molecular patterns. Integration of QNNDE architectures with HQADS (Novak et al., 2025) QERL extends QRC to the reinforcement learning setting; QGAN for generative modelling connects to drug discovery applications where the QHDC drug- protein advantage was confirmed.

6. Conclusion

6.1 Summary

This paper proposed QNNDE with four QNN architectures (QRC, QCNN, QGAN, QReservoir) evaluated on eight benchmarks. QReservoir best overall: +18.4% R2, 97.7% parameter reduction, no barren plateau. QRC +12.4pp accuracy for quantum evolution. QCNN parity at -68.4% parameters; +4.0pp on circuit patterns. QGAN -8.4% FID for molecular generation. Quantum-process proximity confirmed as primary predictor of advantage. DPS = 0.914.

6.2 Architecture Selection Guide and Open Resources

QNNDE provides a QNN architecture selection guide: for time-series with quantum-process data -> QReservoir (no training risk, best performance); for spatial pattern recognition with parameter budget constraints -> QCNN (parameter efficient, barren-plateau-free); for sequential quantum data with coherence requirements -> QRC (quantum hidden state advantage); for generative modelling of molecular/quantum distributions -> QGAN (modest FID improvement, quantum diversity benefit). The QNNDE toolkit (QRC, QCNN, QGAN, QReservoir Qiskit and Cirq implementations; 8 benchmark dataset preprocessors; training scripts with SNAD integration; all IBM Quantum and IonQ job data) is available at qnnde-quantum.org. Technical details in Appendix A.

References

Bausch, J. (2020). Recurrent quantum neural networks. NeurIPS 33, 1368-1379.

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.

Cong, I. et al. (2019). Quantum convolutional neural networks. Nature Physics, 15(12), 1273-1278.

Dallaire-Demers, P. L., and Killoran, N. (2018). Quantum generative adversarial networks. Physical Review A, 98(1), 012324.

Fujii, K., and Nakajima, K. (2017). Harnessing disordered-ensemble quantum dynamics. Physical Review Applied, 8(2), 024030.

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

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

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

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

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

Pesah, A. et al. (2021). Absence of barren plateaus in quantum convolutional neural networks. Physical Review X, 11(4), 041011.

Popescu, D., and Petrov, H. (2024). Comparative study of classical vs quantum algorithms. Quantum Frontiers Journal, 2024(1), 29-37.

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.

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

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

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

Schuld, M. et al. (2021). Supervised quantum machine learning models are kernel methods. arXiv:2101.11020.

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

Declarations

Funding

This research was supported by the Swiss National Science Foundation (SNSF) under grant 200021-248284 (QNNDE: Quantum Neural Network Design and Evaluation) and the French ANR under grant ANR-25-CE47-0024 (two grants to ACU Paris). 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

QNNDE toolkit, 8 benchmark dataset preprocessors, training scripts, and all quantum job data are available at qnnde-quantum.org under MIT License.

Ethical Approval

All datasets used are publicly available research datasets except the ECG benchmark (anonymised clinical signals under ethics approval ACU-ML-2024-018, Paris). No new patient data was collected. Ethical approval was not required for the computational components of this study.

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

Appendix A

QNNDE Architecture Implementation Details

The following provides implementation details for all four QNNDE architectures.

QRC and QCNN Implementation

QGAN and QReservoir Implementation