

# Quantum Transfer Learning as a Noise Diagnostic Tool: Using Classical-to-Quantum Feature Handoff to Characterize Decoherence Sensitivity in Parameterized Quantum Circuits

**Dr. Zuhair Ahmed**

*Centre of Excellence for Technology Quantum and AI Canada (CETQAC)*

---

## ABSTRACT

Quantum machine learning algorithms hold promise for near-term quantum advantage, yet their performance under realistic hardware noise remains poorly characterized. In this work, we introduce a hybrid classical-to-quantum transfer learning architecture in which a fixed linear encoder maps two-dimensional classical input data to four quantum rotation angles, which are subsequently encoded into a parameterized quantum circuit (PQC) ansatz executed on the PKTRON v3.7.3 simulation framework using the PK NoisyLab 8Q virtual device. We systematically evaluate the degradation of binary classification accuracy and the quantum feature space under three distinct decoherence channels — amplitude damping, phase damping, and depolarizing noise — across five noise strength levels ( $p = 0.01$  to  $0.35$ ) and three entanglement depths ( $L = 1, 2, 3$  layers). Our results reveal that phase damping and depolarizing noise impose an immediate, strength-invariant accuracy reduction of 10% from the ideal baseline of 83.33%, while amplitude damping exhibits partial robustness at low noise levels. Frobenius distance analysis demonstrates that depolarizing noise induces the steepest degradation of the quantum feature space, reaching a saturation plateau at  $p = 0.20$ . Critically, deeper entanglement circuits achieve higher ideal accuracy but exhibit superlinear growth in Frobenius distance under noise, suggesting a fundamental accuracy-robustness tradeoff. We propose a Noise Sensitivity Score  $S(N)$  as a lightweight, application-level noise diagnostic metric, offering a tractable alternative to full quantum process tomography for NISQ-era devices.

**Keywords:** *quantum machine learning, transfer learning, noise characterization, decoherence, parameterized quantum circuits, NISQ, PKTRON, quantum noise diagnostics*

---

## I. Introduction

The rapid development of near-term quantum processors has catalyzed significant interest in quantum machine learning (QML), a field that seeks to leverage quantum mechanical phenomena — superposition, entanglement, and interference — to enhance classical learning tasks. Parameterized quantum circuits (PQCs) have emerged as the central computational primitive for QML, serving as the quantum analogue of neural network layers in hybrid classical-quantum architectures [1, 2]. Among the most promising of these hybrid approaches is quantum transfer learning, wherein a pre-trained or fixed classical network performs feature extraction, and a small quantum circuit executes the final classification in a high-dimensional Hilbert space [3].

However, a fundamental tension exists between the theoretical promise of quantum feature maps and the practical constraints of NISQ-era hardware. Contemporary quantum processors suffer from decoherence —

the progressive loss of quantum information due to interaction with the environment — manifesting as T1 energy relaxation (amplitude damping), T2 dephasing (phase damping), and stochastic gate errors (depolarizing noise) [4, 5]. While these noise sources are well-characterized at the gate level, their downstream effect on application-level metrics such as classification accuracy and learned feature geometry has received comparatively little systematic attention.

Existing QML literature largely evaluates algorithms on ideal, noise-free simulators [1, 6, 7] or studies noise in isolation without connecting hardware-level noise parameters to application-level performance degradation curves. This gap is significant: practitioners deploying QML algorithms on real hardware need quantitative guidance on which noise channels are most destructive to their specific circuit architecture and at what noise threshold their quantum advantage collapses.

In this work, we address this gap by introducing quantum transfer learning as a noise diagnostic instrument. We propose that the hybrid architecture — classical encoder followed by a quantum classification layer — provides an ideal testbed for isolating the effect of each noise channel on quantum feature extraction, since the classical component is noise-free by construction, leaving all observed performance degradation attributable solely to quantum decoherence. Our specific contributions are:

1. A hybrid classical-to-quantum transfer learning pipeline executed on the PKTRON v3.7.3 framework, using the PK NoisyLab 8Q virtual device with calibration-realistic noise parameters.
2. The first systematic, side-by-side comparison of amplitude damping, phase damping, and depolarizing noise on QML classification accuracy and feature space geometry under identical experimental conditions.
3. A Noise Sensitivity Score  $S(N)$  — a normalized, application-level metric for quantifying the destructiveness of a noise channel without requiring process tomography.
4. An empirical characterization of the accuracy-robustness tradeoff as a function of entanglement depth, with direct implications for circuit design on NISQ devices.

## II. Background and Theoretical Framework

### A. Parameterized Quantum Circuits

A parameterized quantum circuit is a unitary transformation  $U(\theta)$  applied to an initial quantum state  $|0\rangle^{\otimes n}$ , where  $\theta$  denotes a vector of trainable or data-encoding parameters. In the data encoding context, the circuit implements a feature map  $\varphi: \mathbb{R}^d \rightarrow \mathcal{H}$ , mapping classical input  $x$  to a quantum state  $|\varphi(x)\rangle$  in Hilbert space  $\mathcal{H}$  of dimension  $2^n$ , where  $n$  is the number of qubits. The measurement outcome of  $|\varphi(x)\rangle$  serves as the quantum-derived feature for downstream classification.

In this work, the PQC ansatz encodes a 4-dimensional angle vector  $\theta = [\theta_0, \theta_1, \theta_2, \theta_3]$  through alternating rotation and entanglement layers. For  $L$  repetitions, the circuit implements:

$$U(\theta, L) = \prod_{l=1}^L [CX_{ring} \cdot R_y(\theta_0/l) R_z(\theta_1/l) R_y(\theta_2/l) R_z(\theta_3/l)] \cdot R_y(\theta_0/2) R_y(\theta_1/2)$$

where  $CX_{ring}$  denotes a ring entanglement layer ( $q_0 \rightarrow q_1 \rightarrow q_2 \rightarrow q_3 \rightarrow q_0$ ) and the diminishing scale factor  $1/l$  prevents over-rotation in deeper circuits. The quantum feature is extracted as the Z-expectation value on qubit 0:

$$f(x) = \langle Z_0 \rangle = \sum_{\{i: bit_0(i)=0\}} \rho_{ii} - \sum_{\{i: bit_0(i)=1\}} \rho_{ii}$$

where  $\rho$  is the output density matrix (or, in the ideal case, the outer product of the output statevector).

## B. Classical-to-Quantum Transfer Learning

Transfer learning in classical deep learning refers to the practice of reusing representations learned by a large model on a related task. In the quantum context, transfer learning denotes a hybrid architecture wherein a classical network — typically pre-trained — performs feature extraction, and the resulting compressed representation is passed to a quantum circuit for the final prediction [3, 8].

In our architecture, the classical encoder is a fixed linear layer  $W \in \mathbb{R}^{4 \times 2}$  with bias  $b \in \mathbb{R}^4$ , followed by a tanh nonlinearity scaled by  $\pi$ :

$$\theta = \tanh(Wx + b) \cdot \pi, \quad \theta \in (-\pi, \pi)^4$$

This maps the 2D input  $x \in \mathbb{R}^2$  to 4 quantum rotation angles, which are bounded within the range required for meaningful qubit rotations. The weights  $W$  and  $b$  are fixed (not trained), simulating a frozen pre-trained encoder. This design choice is deliberate: it isolates all learnable expressivity in the quantum layer, ensuring that observed performance changes are attributable to quantum circuit behavior — and, under noise, to quantum decoherence.

## C. Noise Models

We evaluate three physically motivated noise channels, each corresponding to a distinct physical mechanism in superconducting qubit hardware:

Amplitude Damping (AD) models energy relaxation from  $|1\rangle$  to  $|0\rangle$  with characteristic time  $T_1$ . Its Kraus operators are  $K_0 = [[1, 0], [0, \sqrt{1-p}]]$  and  $K_1 = [[0, \sqrt{p}], [0, 0]]$ , where  $p = 1 - \exp(-t/T_1)$  for gate duration  $t$ .

Phase Damping (PD) models pure dephasing with characteristic time  $T_2$ , destroying off-diagonal coherences without energy exchange. Its Kraus operators are  $K_0 = [[1, 0], [0, \sqrt{1-p}]]$  and  $K_1 = [[0, 0], [0, \sqrt{p}]]$ .

Depolarizing Noise (DP) models stochastic Pauli errors with equal probability, corresponding to gate infidelity. It applies Pauli operators  $\{I, X, Y, Z\}$  with probabilities  $\{1-p, p/4, p/4, p/4\}$  respectively.

## D. Noise Sensitivity Score

We propose the Noise Sensitivity Score  $S(N)$  as a normalized, application-level metric for quantifying the impact of noise channel  $N$  on classification performance:

$$S(N) = \max(0, (Acc_{ideal} - Acc_{noisy}) / Acc_{ideal}) \in [0, 1]$$

$S(N) = 0$  indicates no degradation;  $S(N) = 1$  indicates complete collapse to chance-level performance. This score provides a single-number summary of noise destructiveness for a given circuit-dataset combination, analogous in spirit to quantum volume [9] but applicable at the algorithm level rather than the hardware level.

# III. Methodology

## A. Dataset

We employ the scikit-learn two-moons synthetic dataset [10] with  $n = 80$  samples, noise parameter  $\sigma = 0.15$ , and random seed 42. The dataset comprises two interleaved half-circles in  $\mathbb{R}^2$ , providing a non-linearly separable binary classification task that cannot be solved by a linear classifier. Features are scaled to  $[0, \pi]$  via `MinMaxScaler` to ensure compatibility with quantum rotation angle encoding. The dataset is split into 56 training and 24 test samples.

## B. Experimental Framework

All experiments are executed using PKTRON v3.7.3, a production-grade Python quantum simulation framework developed by CETQAP. The PK NoisyLab 8Q virtual device is used for device-level execution reports. Noiseless (ideal) simulations use PKTRON's StatevectorSimulator; noisy simulations use the DensityMatrixSimulator with Lindblad-channel noise injection. Classification is performed using a Support Vector Machine (SVM) with RBF kernel ( $C = 1.0$ ,  $\gamma = \text{'scale'}$ ) from scikit-learn, trained on the quantum-extracted features.

## C. Experiment Design

Three experiments are conducted:

5. Ideal Baseline: Quantum features extracted under noise-free simulation ( $L = 2$  layers,  $n = 4$  qubits). SVM trained and evaluated on test set.
6. Noise Channel Sweep: Each of three noise types (amplitude damping, phase damping, depolarizing) evaluated at five strength levels  $p \in \{0.01, 0.05, 0.10, 0.20, 0.35\}$  with  $L = 2$  layers fixed.
7. Entanglement Depth Sensitivity: Circuit depth  $L \in \{1, 2, 3\}$  evaluated under fixed depolarizing noise  $p = 0.10$ , comparing ideal vs. noisy accuracy and Frobenius distance.

## IV. Experimental Setup

The PKTRON PK NoisyLab 8Q virtual device provides a calibration-realistic noise environment with the following per-qubit parameters as reported by the hardware execution report:

Qubit	Idle Dec. Loss	Gate Error	Readout Error	T1 ( $\mu\text{s}$ )	T2 ( $\mu\text{s}$ )
Q0	2.489%	4.300%	3.000%	~100	~80
Q1	2.642%	4.300%	3.000%	~100	~80
Q2	2.565%	4.200%	3.000%	~100	~80
Q3	2.719%	4.200%	3.000%	~100	~80

Table 1. PK NoisyLab 8Q per-qubit calibration parameters as reported by PKTRON hardware execution report. Circuit duration: 2.35  $\mu\text{s}$ , 18 gates, 0 inserted SWAPs. Estimated circuit fidelity: 66.96%.

The circuit consists of 18 gates ( $L = 2$  ansatz layers on 4 qubits), with a total estimated execution duration of 2.35  $\mu\text{s}$ . The dominant error sources identified by the PKTRON hardware execution report are T1/T2 idle decoherence on Q3 (2.72%) and gate error on Q0 (4.30%). No SWAP insertions were required, confirming that the ring entanglement topology is natively compatible with the device coupling map. The estimated end-to-end circuit fidelity of 66.96% is consistent with expectations for a 4-qubit, 18-gate circuit on a noisy near-term device.

## V. Results

### A. Ideal Baseline Performance

Under ideal (noiseless) simulation with  $L = 2$  entanglement layers, the hybrid quantum transfer learning model achieves a test accuracy of 83.33% (20/24 correct classifications) on the two-moons dataset. This establishes the performance ceiling against which all noisy results are compared. The result confirms that the quantum feature map, combined with the fixed classical encoder, extracts sufficient information from the two-moons geometry to achieve above-chance classification, validating the architecture design.

### B. Noise Channel Sweep Results

Table 2 presents the complete results of the noise channel sweep experiment across all three noise types and five strength levels.

Noise Type	Strength (p)	Accuracy	Sensitivity S(N)	Frob. Dist.	$\Delta$ Acc
Amplitude Damping	0.01	0.7917	0.0500	1.3528	- 4.16 %
	0.05	0.8333	0.0000	1.2062	0.00 %
	0.10	0.8333	0.0000	1.7129	0.00 %
	0.20	0.8333	0.0000	3.1223	0.00 %
	0.35	0.8333	0.0000	4.9680	0.00 %
Phase Damping	0.01	0.7500	0.1000	1.5150	- 8.33 %
	0.05	0.7500	0.1000	1.5150	- 8.33 %
	0.10	0.7500	0.1000	1.5150	- 8.33 %
	0.20	0.7500	0.1000	1.5150	- 8.33 %
	0.35	0.7500	0.1000	1.5150	- 8.33 %
Depolarizing	0.01	0.7500	0.1000	1.7466	- 8.33 %
	0.05	0.7500	0.1000	2.4602	- 8.33 %
	0.10	0.7500	0.1000	2.8447	-

Noise Type	Strength (p)	Accuracy	Sensitivity S(N)	Frob. Dist.	$\Delta$ Acc
					8.33 %
	0.20	0.7500	0.1000	3.0063	- 8.33 %
	0.35	0.7500	0.1000	3.0183	- 8.33 %

Table 2. Full noise channel sweep results. Ideal baseline accuracy = 0.8333. Frobenius distance measures the L2 norm between noisy and ideal quantum feature vectors across all 56 training samples.

Three key findings emerge from Table 2:

8. Amplitude Damping exhibits partial robustness. At  $p = 0.01$ , AD reduces accuracy by 4.16% ( $S = 0.05$ ). However, for  $p \geq 0.05$ , accuracy recovers fully to the ideal baseline ( $S = 0.00$ ), despite monotonically increasing Frobenius distances ( $1.21 \rightarrow 4.97$ ). This counterintuitive result indicates that AD at moderate-to-high strength levels shifts the quantum feature distribution in a direction that is still separable by the RBF-SVM, suggesting that energy relaxation noise may partially preserve the topology of the decision boundary for this dataset.
9. Phase Damping imposes an immediate, strength-invariant accuracy penalty. PD reduces accuracy to 75.00% ( $S = 0.10$ ) at all tested strength levels, with a constant Frobenius distance of 1.5150 across all  $p$  values. This striking saturation behavior indicates that phase damping induces a qualitative change in the quantum feature map at its onset ( $p = 0.01$ ), destroying the coherences responsible for the additional 8.33% of classification capacity, with no further degradation as noise strength increases. This suggests a threshold effect: once coherences are partially destroyed, the relevant information is lost regardless of noise magnitude.
10. Depolarizing noise combines the worst properties of both. It immediately reduces accuracy to 75.00% ( $S = 0.10$ , matching PD) while simultaneously inducing the largest Frobenius distances of all three noise types ( $1.75 \rightarrow 3.02$ ), indicating progressive and severe corruption of the feature space geometry. Notably, Frobenius distance saturates between  $p = 0.20$  and  $p = 0.35$  ( $3.006 \rightarrow 3.018$ ), suggesting that beyond a saturation threshold, the quantum state collapses to the maximally mixed state and further noise has no additional measurable effect.

### C. Entanglement Depth Sensitivity

Table 3 presents the results of the entanglement depth experiment under fixed depolarizing noise  $p = 0.10$ .

Depth (L)	Ideal Acc.	Noisy Acc.	S(N)	Frobenius Dist.
1	0.7083	0.5417	0.2353	2.3023
2	0.8333	0.7500	0.1000	2.8391
3	0.8750	0.7917	0.0952	4.0114

Table 3. Entanglement depth sensitivity results under fixed depolarizing noise  $p = 0.10$ .  $S(N)$  = Noise Sensitivity Score.

The depth experiment reveals a fundamental accuracy-robustness tradeoff in parameterized quantum circuits. Deeper circuits ( $L = 3$ ) achieve higher ideal accuracy (87.50%) and exhibit a lower noise sensitivity score ( $S = 0.0952$ ) compared to shallower circuits ( $L = 1$ ,  $\text{Acc\_ideal} = 70.83\%$ ,  $S = 0.2353$ ). This is consistent with the theoretical understanding that deeper circuits explore a larger portion of Hilbert space, enabling more expressive feature maps.

However, deeper circuits also exhibit superlinear growth in Frobenius distance ( $2.30 \rightarrow 2.84 \rightarrow 4.01$  for  $L = 1, 2, 3$ ). This indicates that while deeper circuits are better at extracting the task-relevant features — and those features remain more recoverable after noise — the absolute corruption of the feature space is larger. The tradeoff is therefore between noise sensitivity per unit of ideal accuracy gained (lower for deeper circuits) and absolute feature space corruption (higher for deeper circuits).

Critically, the  $L = 1$  circuit suffers the most severe practical consequence of noise, with accuracy collapsing to 54.17% under  $p = 0.10$  depolarizing noise — barely above random chance (50%) for a balanced binary dataset. This represents a complete breakdown of the quantum classification advantage at a noise level that only reduces the  $L = 3$  circuit by 8.33%. This finding has direct practical implications: shallow quantum circuits should be avoided on NISQ devices, as they offer the worst noise resilience despite their reduced gate count.

## VI. Discussion

### A. Noise Channel Hierarchy for QML

Our results establish a clear hierarchy of noise destructiveness for quantum machine learning on the tested architecture:

Depolarizing > Phase Damping > Amplitude Damping (for feature space corruption)

Phase Damping = Depolarizing > Amplitude Damping (for classification accuracy impact)

This hierarchy has practical implications for hardware selection and error mitigation strategy. Phase damping, which arises primarily from low-frequency noise ( $1/f$  noise, flux noise in superconducting qubits), induces a hard accuracy ceiling that cannot be overcome by increasing noise strength. Practitioners should prioritize T2 extension (via dynamical decoupling or qubit design improvements) when deploying QML classifiers, as T2 degradation imposes the most robust and irreversible accuracy penalty.

Amplitude damping, by contrast, exhibits a remarkable self-correcting property in this experiment: at  $p \geq 0.05$ , accuracy returns to the ideal baseline despite substantially corrupted feature vectors. This suggests that the RBF-SVM classifier is robust to the specific type of feature space distortion induced by AD, possibly because AD preserves the relative ordering of feature values even as it shifts their absolute magnitude. This finding warrants further investigation with other classifiers and datasets.

### B. The Noise Sensitivity Score as a Diagnostic Tool

The proposed Noise Sensitivity Score  $S(N)$  provides a compact, interpretable summary of noise impact that complements existing hardware benchmarks. Unlike quantum volume [9], which characterizes hardware capability,  $S(N)$  characterizes algorithm-hardware co-performance for a specific task. Unlike randomized benchmarking,  $S(N)$  requires no specialized calibration circuits — it is computed directly from application-level accuracy measurements.

The  $S(N)$  values obtained in this work (0.00 for AD at  $p \geq 0.05$ , 0.10 for PD and DP at all tested levels, 0.24 for  $L=1$  under DP  $p=0.10$ ) suggest that the proposed metric is sensitive enough to distinguish between noise types and circuit depths while remaining interpretable. A score of  $S = 0.24$  for the  $L = 1$  circuit under depolarizing noise clearly flags this as an unacceptable operating regime, while  $S = 0.10$  for the  $L = 2$  circuit may be acceptable depending on application requirements.

### C. Implications for NISQ Device Selection

The PKTRON hardware execution report for the PK NoisyLab 8Q device reveals that the dominant error sources are gate errors ( $\sim 4.3\%$  per qubit) and idle decoherence ( $\sim 2.7\%$  per qubit per circuit execution). With an estimated circuit fidelity of 66.96% for the 18-gate, 4-qubit ansatz, the device operates well within the regime where noise effects are measurable but the circuit has not completely decohered.

Extrapolating from our noise sweep results, a device with gate errors corresponding to effective depolarizing probability  $p \approx 0.10$  per gate would be expected to reduce QML accuracy by approximately 10% from the ideal baseline for  $L = 2$  circuits. This provides a quantitative threshold for hardware selection: devices with gate error rates below the  $p = 0.05$  boundary should be preferred for QML classification tasks using this architecture.

## VII. Conclusion

We have presented a systematic experimental characterization of decoherence effects on a hybrid quantum transfer learning architecture, executed on the PKTRON v3.7.3 simulation framework with the PK NoisyLab 8Q virtual device. Our findings demonstrate that the three primary noise channels in superconducting quantum hardware — amplitude damping, phase damping, and depolarizing noise — produce qualitatively distinct degradation signatures in both classification accuracy and quantum feature space geometry.

Phase damping imposes an immediate, strength-invariant accuracy penalty of 10%, driven by the destruction of quantum coherences at the onset of noise. Depolarizing noise produces the largest Frobenius distance growth (saturation plateau at 3.02 for  $p \geq 0.20$ ) while matching phase damping in accuracy impact. Amplitude damping exhibits surprising resilience at moderate-to-high strength levels, with accuracy recovering to the ideal baseline despite substantial feature space corruption.

The entanglement depth analysis reveals a fundamental accuracy-robustness tradeoff: deeper circuits achieve higher ideal accuracy and lower noise sensitivity per unit accuracy, but incur larger absolute feature space corruption. Shallow circuits ( $L = 1$ ) are strongly contraindicated on NISQ devices, as they suffer catastrophic accuracy collapse (54.17%) under moderate depolarizing noise ( $p = 0.10$ ).

The proposed Noise Sensitivity Score  $S(N)$  provides a lightweight, application-level diagnostic metric that can guide hardware selection, error mitigation strategy, and circuit design decisions without requiring the exponential overhead of full process tomography. Future work will extend this diagnostic framework to larger qubit registers, more complex datasets, and trainable quantum transfer learning architectures on real superconducting hardware.

---

## References

- [1] M. Cerezo et al., "Variational quantum algorithms," *Nature Reviews Physics* 3, 625 (2021).



- [2] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum* 2, 79 (2018).
- [3] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," *Quantum* 4, 340 (2020).
- [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press (2000).
- [5] J. Preskill, *Lecture Notes on Quantum Computation*, Chapter 3: Quantum Noise and Open Quantum Systems, Caltech (2015).
- [6] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature* 567, 209–212 (2019).
- [7] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Physical Review Letters* 122, 040504 (2019).
- [8] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nature Physics* 15, 1273–1278 (2019).
- [9] A. W. Cross et al., "Validating quantum computers using randomized model circuits," *Physical Review A* 100, 032328 (2019).
- [10] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research* 12, 2825–2830 (2011).
- [11] CETQAP & PKTron Development Team, PKTRON v3.7.3: Pakistan's AI-Powered Quantum Simulation Framework, PyPI (2026). <https://pypi.org/project/pktron/>
- [12] L. Viola, E. Knill, and S. Lloyd, "Dynamical decoupling of open quantum systems," *Physical Review Letters* 82, 2417 (1999).
- [13] P. Aumann et al., "Demonstrating Record Fidelity for the Quantum Fourier Transform," arXiv:2604.12465 [quant-ph] (2026).