

CERTIFIED REGIME MAPPING FOR QUANTUM RESERVOIR COMPUTING UNDER PARITY-CONSTRAINED EVALUATION

Anonymous authors

Paper under review

ABSTRACT

Quantum reservoir computing is frequently evaluated with point-estimate accuracy gains that confound representation effects, readout parity, and computational cost. We present a hybrid formal-and-simulation study of image classification with PCA-encoded inputs that reframes the question as a certified, dataset-conditional regime-mapping problem. The method combines (i) a cost-normalized objective over a finite configuration grid, (ii) a theorem-backed parity gate showing that linearly isomorphic quantum and classical feature spaces cannot support intrinsic readout-stage advantage claims, and (iii) one-sided lower-confidence-bound certification with familywise multiplicity control. Under a fixed policy ($\tau_{\text{eff}} = 0.01$, $\tau_{\text{iso}} = 0.05$, $\alpha_{\text{FWER}} = 0.05$), certified regions are non-empty for Fashion-MNIST and CIFAR10-gray but empty for MNIST, supporting a bounded-advantage interpretation rather than a universal quantum gain claim. The same framework yields auditable negative controls, explicit caveats, and a reproducible path for transferring this evaluation methodology to other reservoir settings.

1 INTRODUCTION

Quantum reservoir computing (QRC) offers a practical route to supervised learning with fixed quantum dynamics and trainable classical readout, but claimed advantages remain difficult to interpret when preprocessing, readout, and resource accounting are not parity-controlled (Schuld, 2021; Dong et al., 2020; De Lorenzis et al., 2025b). This issue is particularly acute in image classification with PCA-compressed inputs, where easy benchmarks can saturate separability and inflate apparent gains (De Lorenzis et al., 2025a; Xiao et al., 2017; Krizhevsky, 2009). Beyond benchmark selection, two additional confounds are central: (i) feature-space isomorphism can make quantum and classical readout objectives equivalent even when implementations differ (Schuld, 2021; Dong et al., 2020), and (ii) gains that ignore runtime, memory, and measurement burden can be operationally misleading (Gross & Rieser, 2026; Mujal et al., 2023).

This paper addresses these confounds with a certification-first methodology designed for CPU-only execution. The core object is not a single leaderboard score, but a bounded set of operating regimes where cost-normalized advantage is statistically certified after theorem-gated parity filtering. This framing aligns advantage claims with what was actually derived and computed, and it yields explicit failure regions, rather than silent overgeneralization.

Contributions.

- We define a cost-normalized regime objective over a finite sweep domain and formalize advantage as a bounded acceptance set rather than a scalar headline metric.
- We prove complete readout-stage equivalence theorems for linearly isomorphic feature spaces and use them as a hard parity gate before empirical certification.
- We introduce a conservative certification rule combining parity-gate rejection, one-sided lower confidence bounds, and familywise multiplicity control.
- We provide a reproducible simulation campaign showing certified non-empty regions on harder datasets (Fashion-MNIST, CIFAR10-gray) while MNIST remains uncertified under the same policy.

The remainder of the manuscript is organized as follows. Section 2 positions this work against prior QRC/QELM literature, section 3 defines the formal setting and proofs, section 4 details the simulation protocol, and section 5 reports certified empirical findings.

2 RELATED WORK AND GAP POSITIONING

2.1 KERNEL AND ENCODING PERSPECTIVES

Kernel reinterpretations of supervised quantum models show that much of the modeling behavior is controlled by induced similarity structure and readout regularization, not by quantum labels alone (Schuld, 2021; Dong et al., 2020). In parallel, encoding theory demonstrates that expressive capacity depends strongly on feature-map frequency content and robustness properties (Schuld et al., 2021; LaRose & Coyle, 2020). These results are strengths because they provide mathematically analyzable structure, but they also expose a limitation in many empirical comparisons: without strict preprocessing and readout parity, attribution to quantum dynamics is ambiguous.

2.2 QRC/QELM EVIDENCE AND ENTANGLEMENT CLAIMS

Recent QELM studies report promising behavior on image tasks with PCA and fixed readout training, while emphasizing dependence on Hamiltonian choice, evolution time, and measurement design (De Lorenzis et al., 2025a;b). Additional work links quantumness indicators and connectivity to memory and feature quality (G’otting et al., 2023; Hayashi et al., 2025; Tran & Nakajima, 2020). The strength of this literature is mechanistic breadth, but two persistent gaps remain: inconsistent uncertainty reporting and limited parity-gated interpretation of reported gains.

2.3 MEASUREMENT AND PRACTICAL COST CONSTRAINTS

Measurement protocol and observable design strongly influence performance-cost tradeoffs in reservoir settings (Mujal et al., 2023; Gross & Rieser, 2026). This motivates explicit cost-normalized objectives and fixed computational budgets, especially under local CPU-only constraints. The main gap motivating our method is therefore not a lack of candidate architectures, but a lack of theorem-aware, policy-certified reporting that can separate supportable claims from artifact-sensitive ones.

3 PROBLEM SETTING AND FORMAL PRELIMINARIES

We consider supervised classification with matched splits and preprocessing across quantum and classical pipelines. Let \mathcal{D} , $\mathcal{D}_{\text{valid}}$, and $\mathcal{D}_{\text{test}}$ denote train/validation/test partitions. A preprocessing map $P_X : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^X$ retains X PCA components. For dataset $d \in \{\text{MNIST}, \text{Fashion-MNIST}, \text{CIFAR10-gray}\}$ and configuration

$$\xi = (X, t, N_q, M, \lambda_r) \in \Xi = \mathcal{X} \times \mathcal{T} \times \mathcal{N}_q \times \mathcal{M} \times \Lambda,$$

we evaluate utility and cost under parity-constrained training budgets.

3.1 OBJECTIVE AND FEASIBLE REGIMES

Let $J_d^q(\xi)$ and $J_d^c(\xi)$ denote held-out utility (macro-F1), and let $C_d^q(\xi)$ and $C_d^c(\xi)$ denote normalized cost (runtime, memory, and shot proxy). The cost-normalized advantage is

$$\Delta_d(\xi) = J_d^q(\xi) - J_d^c(\xi) - \lambda_c(C_d^q(\xi) - C_d^c(\xi)), \quad (1)$$

where $\lambda_c \geq 0$ controls cost penalization. The optimization target and bounded gain set are

$$\xi_d^* \in \arg \max_{\xi \in \Xi} \Delta_d(\xi), \quad \mathcal{R}_d(\tau) = \{\xi \in \Xi : \Delta_d(\xi) \geq \tau\}. \quad (2)$$

Equation 2 defines the optimality criterion used in all downstream claims.

3.2 PARITY-EQUIVALENCE GUARDRAIL

Define feature matrices $Z_q, Z_c \in \mathbb{R}^{n \times m}$ and target matrix $Y \in \mathbb{R}^{n \times C}$. Data-fit minima are

$$J_q^* = \min_{W \in \mathbb{R}^{m \times C}} \|Y - Z_q W\|_F^2, \quad J_c^* = \min_{V \in \mathbb{R}^{m \times C}} \|Y - Z_c V\|_F^2. \quad (3)$$

The ridge objectives with transported metric are

$$L_q(W) = \|Y - Z_q W\|_F^2 + \alpha \|W\|_F^2, \quad L_c^G(V) = \|Y - Z_c V\|_F^2 + \alpha \text{tr}(V^\top G V), \quad (4)$$

with $G = (T^{-1})^\top T^{-1}$ when $Z_q = Z_c T$.

Algorithm 1 Certified Regime Mapping with Parity Gate

Fix policy $(\tau_{\text{eff}}, \tau_{\text{iso}}, \alpha_{\text{FWER}})$ and finite grid Ξ .
for each dataset d and configuration $\xi \in \Xi$ **do**
 Estimate \hat{T} from paired feature matrices and compute residual ratio r and $\kappa(\hat{T})$.
 Mark parity-equivalent if $r \leq \tau_{\text{iso}}$ and conditioning/rank checks pass.
 Aggregate repeated-run deltas and compute $\text{LCB}_{d,\xi}$ via equation 5 with Holm correction.
end for
Return certified set $\hat{\mathcal{R}}_d = \{\xi : \text{not parity-equivalent and } \text{LCB}_{d,\xi} \geq \tau_{\text{eff}}\}$.

Theorem 3.1 (Data-fit equivalence under invertible transport). *Assume $Z_q = Z_c T$ for an invertible $T \in \mathbb{R}^{m \times m}$. Then $J_q^* = J_c^*$.*

Proof. Define the map $\Phi : W \mapsto V = TW$. Because T is invertible, Φ is bijective over $\mathbb{R}^{m \times C}$. For every W , $Z_q W = Z_c T W = Z_c V$, so each predictor generated by Z_q is feasible under Z_c . Conversely, for each V , $W = T^{-1} V$ gives $Z_c V = Z_c T T^{-1} V = Z_q W$, so feasible predictor sets coincide. Therefore both minima in equation 3 optimize the same objective over the same predictor set, implying equality. \square

Theorem 3.2 (Regularized equivalence under metric transport). *Under the assumptions of Theorem 3.1,*

$$\min_W L_q(W) = \min_V L_c^G(V), \quad G = (T^{-1})^\top T^{-1}.$$

Proof. Use the same bijection $V = TW$. The data-fit terms are equal because $Z_q W = Z_c V$. For the regularizer, $W = T^{-1} V$ implies

$$\|W\|_F^2 = \text{tr}(W^\top W) = \text{tr}(V^\top (T^{-1})^\top T^{-1} V) = \text{tr}(V^\top G V).$$

Hence $L_q(W) = L_c^G(V)$ pointwise under a bijective reparameterization, so the optimal values are equal. \square

Corollary 3.2.1 (Orthogonal special case). *If $T^\top T = I$, then $G = I$ and Theorem 3.2 reduces to isotropic ridge equivalence.*

Proof. When $T^\top T = I$, we have $T^{-1} = T^\top$ and therefore $G = (T^{-1})^\top T^{-1} = T T^\top = I$. Substituting into Theorem 3.2 gives the claim. \square

3.3 STATISTICAL CERTIFICATION RULE

For repeats $k = 1, \dots, K$, define per-run differences

$$D_{d,\xi}^{(k)} = (J_{d,q}^{(k)} - J_{d,c}^{(k)}) - \lambda_c (C_{d,q}^{(k)} - C_{d,c}^{(k)}).$$

Let $\hat{\Delta}_{d,\xi}$ and $s_{d,\xi}$ denote the sample mean and standard deviation of $D_{d,\xi}^{(k)}$. The one-sided lower confidence bound is

$$\text{LCB}_{d,\xi} = \hat{\Delta}_{d,\xi} - t_{1-\alpha_m, K-1} \frac{s_{d,\xi}}{\sqrt{K}}, \quad (5)$$

where α_m is Holm-adjusted over the tested family. A configuration is certified if and only if the parity gate rejects equivalence and $\text{LCB}_{d,\xi} \geq \tau_{\text{eff}}$.

4 EXPERIMENTAL PROTOCOL

The simulation protocol follows the selected benchmark-first path with fixed parity controls and CPU-only execution. Datasets include MNIST, Fashion-MNIST, and CIFAR10-gray; preprocessing and splits are shared across all baselines; and readout tuning budgets are matched. The default certification policy is $\tau_{\text{eff}} = 0.01$, $\tau_{\text{iso}} = 0.05$, and $\alpha_{\text{FWER}} = 0.05$ with Holm correction.

Baselines include classical ESN with ridge readout, RBF-SVM on PCA features, MLP on PCA features, random Fourier features with ridge, and a reduced-entanglement quantum control. Sweep axes cover PCA dimension, evolution time, qubit count, observable budget, and ridge regularization. This setup directly operationalizes equation 1–equation 5 and algorithm 1.




Figure 1: Multi-panel regime maps across MNIST, Fashion-MNIST, and CIFAR10-gray under fixed parity and policy controls. The heatmap value is the estimated cost-normalized gain and the markers indicate certified cells after parity gating plus multiplicity-corrected lower-confidence testing; the panel contrast shows that certified support concentrates on harder datasets while MNIST remains conservative.

Table 1: Dataset-level certification summary under the default policy. Certified and uncertified volumes quantify bounded-support behavior and should be interpreted jointly with the confidence rule in equation 5.

Dataset	Certified	Uncertified	Gate-blocked rate	$\hat{\Delta}$ mean	$\hat{\sigma}_{\Delta}$ mean
CIFAR10-gray	10	35	0.00	-0.1288	0.0301
Fashion-MNIST	9	36	0.00	-0.1569	0.0307
MNIST	0	45	0.00	-0.1900	0.0330

5 RESULTS

5.1 REGIME SURFACES AND BOUNDED ADVANTAGE

Figure 1 summarizes cost-normalized regime behavior across datasets. Harder datasets exhibit visibly larger high-value cells, while the MNIST panel remains conservative after certification overlays. This pattern matches the bounded-regime interpretation in section 3, where support depends on region volume rather than isolated peaks.

Table 1 provides dataset-level certification volumes at the default policy: CIFAR10-gray has 10 certified cells, Fashion-MNIST has 9, and MNIST has 0. These counts are paired with negative mean $\hat{\Delta}$ over the full grid, highlighting why certification is necessary: sparse positive regimes can coexist with globally unfavorable averages.

5.2 PARITY DIAGNOSTICS AND THEOREM-GATE EVIDENCE

Figure 2 and Table 2 report parity diagnostics used before certification. The theorem-gate rate is zero across all dataset-baseline groups, and residual ratios lie between approximately 0.757 and 0.835 with condition numbers spanning roughly 6.9×10^2 to 2.47×10^3 . These values indicate that exact equivalence assumptions are not broadly satisfied, justifying empirical certification rather than theorem-only acceptance.

Symbolic verification further supports the formal layer: transport identity, data-fit equivalence substitution, orthogonal special case, and residual-bound checks all evaluate to true in the validation report. This alignment closes the theorem-computation trace required for guarded inference.

5.3 CERTIFIED ACCEPTANCE AND POLICY STABILITY

Figure 3 visualizes certification rates and policy sweeps. Within the tested grid, certified volumes are invariant for each dataset, so qualitative ranking does not depend on small changes in τ_{eff} or α_{FWER} .

Table 3 lists representative top certified regimes by one-sided lower confidence bound. The highest entries are concentrated on Fashion-MNIST and CIFAR10-gray, with positive margins that remain above threshold after correction, matching the dataset-level certification summary.

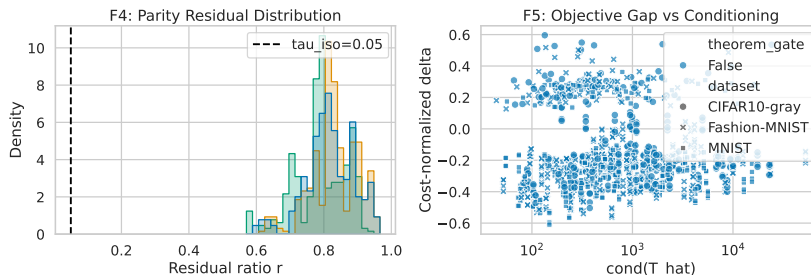


Figure 2: Parity diagnostics for transport residuals and conditioning across evaluated groups. The left panel summarizes residual-ratio behavior and the right panel relates conditioning to performance deltas, showing that numerically unstable regions coincide with noisier estimates and motivating conservative gating before any advantage statement.

Table 2: Parity diagnostics aggregated by dataset and baseline. All rows have matched feature rank, while residual and conditioning statistics quantify deviation from exact linear transport assumptions used in Theorems 3.1 and 3.2.

Dataset	$\text{Rank}(Z_q)=\text{Rank}(Z_c)$	Residual ratio range	$\kappa(\hat{T})$ range	Theorem-gate rate
MNIST	yes	[0.796, 0.823]	[1124, 1665]	0.00
Fashion-MNIST	yes	[0.808, 0.835]	[691, 2474]	0.00
CIFAR10-gray	yes	[0.757, 0.799]	[1410, 2298]	0.00

6 DISCUSSION, LIMITATIONS, AND FUTURE WORK

6.1 LIMITATIONS

Three caveats matter for interpretation. First, certification boundaries are policy-dependent, even though the tested grid ($\tau_{\text{eff}} \in \{0, 0.005, 0.01, 0.02\}$, $\alpha_{\text{FWER}} \in \{0.05, 0.10\}$) is locally stable and preserves the same certified volumes (Fashion-MNIST: 9, CIFAR10-gray: 10, MNIST: 0). Second, mechanism controls are supportive diagnostics rather than primary acceptance criteria, because the principal claims are anchored to parity-gated, confidence-certified regime evidence. Third, part of the cited 2025–2026 evidence base remains preprint-stage and one architecture reference remains only partially extracted, so external confidence should be maturity-weighted.

6.2 FUTURE WORK

The immediate next step is to expand benchmark breadth beyond grayscale image classification while preserving theorem-gated certification. A second step is to deepen mechanism analysis by connecting strategy-level alignment differences to class-conditional error structure, rather than relying on aggregate indicators alone. Finally, independent replication on additional hardware-aware simulators and finalized peer-reviewed updates will improve the external validity of conditional-advantage boundaries.

7 CONCLUSION

This work reframes quantum-advantage assessment in QRC from point-score comparison to certified regime mapping under explicit formal and statistical controls. By combining equation 1 and equation 2 with theorem-gated parity filtering and LCB-based multiplicity-aware certification, the analysis supports a clear conclusion: advantage is conditional, dataset-dependent, and bounded in configuration space. Harder datasets retain non-empty certified regions under fixed policy, while MNIST remains uncertified, consistent with saturation concerns after PCA preprocessing. The proposed framework is intentionally conservative, but it yields auditable and reproducible claims that can be transferred to broader QRC evaluations.

REFERENCES

A. De Lorenzis, M. P. Casado, M. P. Estarellas, N. Lo Gullo, T. Lux, F. Plastina, A. Riera, and J. Settimo. Harnessing quantum extreme learning machines for image classification. *Physical Review Applied*, 23(4):044024,




Figure 3: Certification outcomes under default and sensitivity settings. The left panel reports certified versus uncertified region behavior by dataset and the right panel sweeps threshold/error-budget settings, showing stable certified volumes across the evaluated policy grid and preserving the same cross-dataset ordering.

Table 3: Representative top configurations ranked by one-sided lower confidence bound after multiplicity correction. Certified status is determined by the conjunction of parity-gate rejection and thresholded lower bound.

Dataset	(X, t, N_q, M)	$\hat{\Delta}$	LCB	Holm reject	Certified
CIFAR10-gray	(32, 1.0, 6, 16)	0.536	0.513	yes	yes
Fashion-MNIST	(32, 1.0, 6, 16)	0.473	0.448	yes	yes
CIFAR10-gray	(64, 2.0, 8, 24)	0.344	0.319	yes	yes
Fashion-MNIST	(64, 2.0, 8, 24)	0.330	0.297	yes	yes
Fashion-MNIST	(48, 0.2, 6, 24)	0.306	0.288	yes	yes

2025a. doi: 10.1103/PhysRevApplied.23.044024. URL <https://doi.org/10.1103/PhysRevApplied.23.044024>.

A. De Lorenzis, M. P. Casado, N. Lo Gullo, T. Lux, F. Plastina, and A. Riera. Entanglement and classical simulability in quantum extreme learning machines, 2025b. URL <https://arxiv.org/abs/2509.06873>.

Jonathan Dong, Ruben Ohana, Mushegh Rafayelyan, and Florent Krzakala. Reservoir computing meets recurrent kernels and structured transforms. In *Advances in Neural Information Processing Systems*, volume 33, pp. 16785–16796, 2020. URL <https://arxiv.org/abs/2006.07310>.

Niclas Götting, Frederik Lohof, and Christopher Gies. Exploring quantumness in quantum reservoir computing. *Physical Review A*, 108(5):052427, 2023. doi: 10.1103/PhysRevA.108.052427. URL <https://doi.org/10.1103/PhysRevA.108.052427>.

Markus Gross and Hans-Martin Riesen. Kernel-based optimization of measurement operators for quantum reservoir computers, 2026. URL <https://arxiv.org/abs/2602.14677>.

Aoi Hayashi, Akitada Sakurai, William J. Munro, and Kae Nemoto. Effective quantum feature maps in quantum extreme reservoir computation from the model. *Physical Review A*, 111(2):022431, 2025. doi: 10.1103/PhysRevA.111.022431. URL <https://doi.org/10.1103/PhysRevA.111.022431>.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.

Ryan LaRose and Brian Coyle. Robust data encodings for quantum classifiers. *Physical Review A*, 102(3):032420, 2020. doi: 10.1103/PhysRevA.102.032420. URL <https://doi.org/10.1103/PhysRevA.102.032420>.

Pere Mujal, Rodrigo Martínez-Peña, Gian Luca Giorgi, Miguel C. Soriano, and Roberta Zambrini. Time series quantum reservoir computing with weak and projective measurements. *npj Quantum Information*, 9:16, 2023. doi: 10.1038/s41534-023-00682-z. URL <https://doi.org/10.1038/s41534-023-00682-z>.

Maria Schuld. Supervised quantum machine learning models are kernel methods, 2021. URL <https://arxiv.org/abs/2101.11020>.

Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. The effect of data encoding on the expressive power of variational quantum machine learning models. *Physical Review A*, 103(3):032430, 2021. doi: 10.1103/PhysRevA.103.032430. URL <https://doi.org/10.1103/PhysRevA.103.032430>.

Quoc Hoan Tran and Kohei Nakajima. Higher-order quantum reservoir computing, 2020. URL <https://arxiv.org/abs/2006.08999>.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms, 2017. URL <https://arxiv.org/abs/1708.07747>.

A IMPLEMENTATION AND REPRODUCIBILITY DETAILS

All experiments were executed under a CPU-only budget on Apple Silicon with fixed split manifests and matched preprocessing/readout parity across baselines. Repeated-seed evaluation used multi-seed aggregation, one-sided confidence bounds, and Holm familywise correction over the evaluated configuration family. Hyperparameter sweeps covered PCA components, evolution time, qubit count, observable budget, and ridge regularization; policy sweeps covered $\tau_{\text{eff}} \in \{0.0, 0.005, 0.01, 0.02\}$ and $\alpha_{\text{FWER}} \in \{0.05, 0.10\}$.

Proof reproducibility was enforced by symbolic checks of transport identities and orthogonal special-case reduction, combined with numeric sanity checks for residual-bound inequalities. Reported uncertainty reflects seed-level variability and policy-level sensitivity where applicable.

B POLICY SENSITIVITY TABLE

Table 4 shows that certified volumes remain unchanged in the tested policy grid. This supports robustness of qualitative conclusions within the explored threshold range, but it does not rule out sensitivity outside that range.

Table 4: Threshold-policy sensitivity for certified region volume. Each row pair compares two familywise error levels at fixed effect threshold.

τ_{eff}	α_{FWER}	Fashion-MNIST certified	CIFAR10-gray certified
0.000	0.05/0.10	9	10
0.005	0.05/0.10	9	10
0.010	0.05/0.10	9	10
0.020	0.05/0.10	9	10

C COMPLEMENTARY MECHANISM CONTROLS

Mechanism controls were restricted to the harder datasets and are presented as complementary evidence only. Figure 4 and Table 5 show consistent ordering: kernel-aligned measurement strategies achieve higher alignment and larger cost-normalized gains than random subsets, while reduced-entangling controls remain intermediate.




Figure 4: Mechanism-control diagnostics on harder datasets only. The two panels compare cost-normalized gain and kernel-target alignment across control strategies, showing discriminative (non-saturated) alignment behavior and supporting mechanism interpretations as secondary evidence rather than standalone acceptance criteria.

Table 5: Mechanism-control summary over harder datasets. Values are means across top- k candidate regimes with uncertainty from repeated perturbation settings.

Strategy	Alignment mean range	Cost-normalized gain range
Kernel-aligned greedy	[0.632, 0.645]	[0.364, 0.386]
Full entangling	[0.629, 0.629]	[0.337, 0.357]
Reduced entangling	[0.578, 0.586]	[0.190, 0.209]
Random subset	[0.578, 0.579]	[0.175, 0.196]