FORMATTING INSTRUCTIONS FOR ICLR 2027 CONFERENCE SUBMISSIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

A quantum circuit can change its state under an input transformation while a chosen readout remains unchanged. We study this distinction in hybrid quantum neural networks for surrogate AC optimal power flow. For a real two-layer, non-reuploading ansatz with R_X data encoding, we prove that single-qubit X and Z readouts depend only on the coordinate-wise cosines of the encoded angles, whereas Y carries an odd component. We then use the matched readout pair $Z + X$ versus $Z + Y$ to separate parity information from feature-count and encoder effects. With a frozen encoder, $Z + Y$ improves test MAE over $Z + X$ on all three grid systems. With end-to-end training, the advantage is not consistent and depends on the grid and architecture. A target-aware audit further shows that the learned latent geometry changes around sign-mirror neighborhoods, but does not identify a unique optimization mechanism. The results establish a practical attribution protocol: state dependence, observable invariance, and predictive utility must be tested separately. We make no claim of quantum advantage.

1 INTRODUCTION

Hybrid quantum neural networks do not expose a full quantum state to a downstream decoder. They expose a selected set of expectation values, so a circuit direction can affect the state without affecting the features used for prediction. Conversely, an improvement after adding observables can arise from a larger classical feature vector, better conditioning, or a classical shortcut rather than from a uniquely quantum contribution.

This distinction matters in quantum surrogates for AC optimal power flow (AC-OPF). We study the Hybrid QNN, QNN-SRC, and QNN-NRC families introduced for this setting by Hu et al. (Hu et al., 2026). Our question is deliberately narrower than quantum advantage: when does a sign-dependent change in the encoded circuit state disappear at the measurement interface, and when does that loss become a predictive bottleneck?

The paper separates three levels of analysis. At the state level, we characterize the conjugation symmetry induced by real data encoding and real trainable gates. At the observable level, we derive a coordinate-wise parity result for the implemented two-layer circuit. At the prediction level, we compare equal-size readouts and use positive controls that remove or freeze the encoder's ability to bypass the readout restriction.

The main contributions are:

1. We identify a sign symmetry of the real R_X-encoded circuit and prove that, for the two-layer non-reuploading ansatz, local X/Z features are invariant under every coordinate-wise sign flip.
2. We design a matched intervention, $Z + X$ versus $Z + Y$, in which the number of readout features and trainable parameters is held fixed. This separates an odd-parity channel from a generic readout-width effect.
3. We show a regime-dependent result: Y is predictive when the encoder cannot adapt, but does not provide a consistent end-to-end improvement. This is evidence against attributing end-to-end gains to a particular observable or to entanglement alone.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

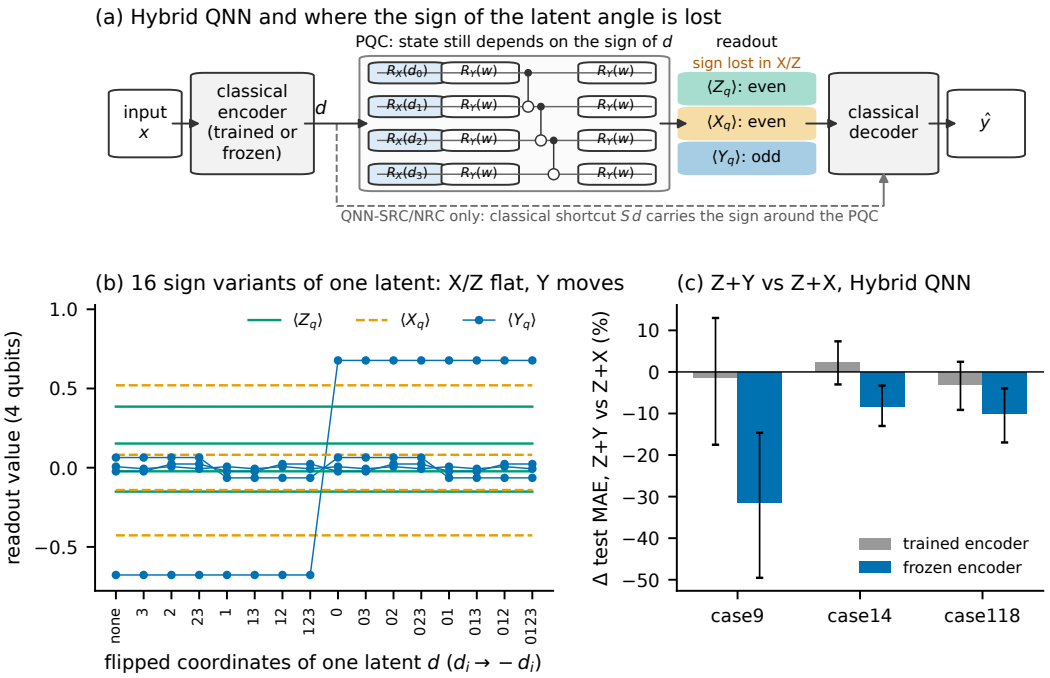

Figure 1: Overview of the attribution problem. (a) The circuit state depends on the sign of the latent angle d, while local X/Z readouts are even and local Y readouts contain an odd component. QNN-SRC and QNN-NRC additionally provide a classical shortcut from d to the decoder. (b) Across all 16 coordinate-wise sign variants, X and Z remain flat while Y changes. (c) The equal-size $Z + Y$ versus $Z + X$ comparison helps only in the frozen-encoder positive control, not consistently in end-to-end training.

4. We provide a target-aware audit with non-symmetry controls and release the complete LaTeX manuscript, figures, tables, and numerical artifacts.

The elementary conjugation identity is not claimed as a new result in quantum theory. The contribution is the implementation-specific diagnosis and the controlled evidence showing how a readout symmetry can, or need not, become a prediction failure.

2 Related Work

Quantum data encoding and expressivity. Encoding choices determine the Fourier frequencies and function classes available to parametrized quantum models (Schuld et al., 2021; Caro et al., 2021). Data re-uploading can enlarge the accessible spectrum (Pérez-Salinas et al., 2020), but redundancy and generator-dependent restrictions can remain (Gil Vidal & Theis, 2020; Barthe & Pérez-Salinas, 2024; Mhiri et al., 2025). Our result concerns a different interface: a component may exist in the circuit state but be removed by the selected observable.

Symmetry and measurement. Symmetry and equivariance are widely used to constrain or organize QML models (Meyer et al., 2023; Larocca et al., 2022; Nguyen et al., 2024; Schatzki et al., 2024). Measurement-learning methods select observables for task sensitivity (Chen et al., 2025a;b; Lin et al., 2025), while shadow and reservoir approaches use alternative state-to-feature maps (Huang et al., 2020; Fujii & Nakajima, 2017; Mujal et al., 2021). We instead treat a symmetry as a potential attribution confound and test it at the state, observable, and prediction levels.

Quantum–classical attribution. Hybrid models can contain classical residuals, bypasses, and decoder capacity that obscure the role of the quantum block (Kölle et al., 2024; Freinberger & Moser, 2026; Bowles et al., 2024; Schuld & Killoran, 2022). Local XYZ readout has appeared in

recent hybrid models (Rai et al., 2026; Levy et al., 2026); measuring three Pauli axes is therefore not our novelty. Our intervention holds the readout size fixed and uses positive controls to test whether the odd channel is actually needed.

AC-OPF surrogates. Quantum and hybrid models have been explored for power-flow analysis, probabilistic OPF, and related power-system optimization tasks (Kaseb et al., 2024; Cao et al., 2025; Zhu et al., 2024; Zhou & Zhang, 2023; Feng et al., 2021; Nikmehr et al., 2022; Le et al., 2026). We use the AC-OPF application as a controlled testbed for observability, not as evidence of a quantum computational advantage.

3 Circuit Analysis

Let $d \in (-\pi, \pi)^n$ denote the latent angles produced by the classical encoder. The data encoding is $E(d) = \bigotimes_i R_X(d_i)$, and the trainable circuit uses real gates. The implemented two-layer circuit is

$$U(d) = R^{(2)} C R^{(1)} E(d), \qquad R^{(\ell)} = \bigotimes_i R_Y(w_i^{(\ell)}),$$

where C is the chain of CNOT gates used in the experiments.

Proposition 1 (conjugation symmetry). Assume the initial state is real, every data-dependent generator is a real Pauli operator, and all data-independent gates are real. Then

$$\psi(-d) = \overline{\psi(d)}.$$

For a Pauli string P,

$$\langle P \rangle_{-d} = (-1)^{\#Y(P)} \langle P \rangle_d.$$

Thus a Pauli string with an even number of Y factors is even under a global sign flip, while an odd number of Y factors is odd. This proposition is used as a structural lemma, not as an independent novelty claim.

Proposition 2 (coordinate-wise parity). For the two-layer, non-reuploading circuit above, each local X_q and Z_q expectation is a polynomial in $(\cos d_0, \ldots, \cos d_{n-1})$. Hence it is invariant under every subset of coordinate-wise sign flips. In contrast,

$$\langle Y_q \rangle = \sin(d_q) g_q(\cos d),$$

where g_q is a circuit-dependent polynomial. Therefore Y_q carries an odd component in coordinate q, but it is modulated by cosine factors and does not represent every odd target.

Scope. Proposition 2 is proved only for the specified two-layer non-reuploading circuit. Reuploading, alternative entangler orderings, and deeper circuits are checked numerically and are outside the theorem's scope.

Proof of Proposition 1. Write the encoded state as

$$|\psi(d)\rangle = U_{\text{real}} E(d) |\psi_0\rangle, \qquad E(d) = \bigotimes_i R_X(d_i).$$

Because X is real and $R_X(d) = \exp(-\mathrm{i} dX/2)$, complex conjugation gives

$$\overline{R_X(d)} = R_X(-d).$$

The same identity holds for any real Pauli data generator. By assumption, U_{real} and $|\psi_0\rangle$ are unchanged by complex conjugation. Therefore

$$|\psi(-d)\rangle = \overline{|\psi(d)\rangle}.$$

For a Pauli string P, complex conjugation satisfies $\overline{P} = (-1)^{\#Y(P)} P$, where $\#Y(P)$ counts its Y factors. Hence

$$\begin{aligned} \langle P \rangle_{-d} &= \langle \psi(-d) | \, P \, | \psi(-d) \rangle \\ &= \langle \psi(d) | \, \overline{P} \, | \psi(d) \rangle = (-1)^{\#Y(P)} \langle P \rangle_d. \end{aligned}$$

This proves the global sign-parity statement. In particular, strings with an even number of Y factors are even and strings with an odd number are odd.

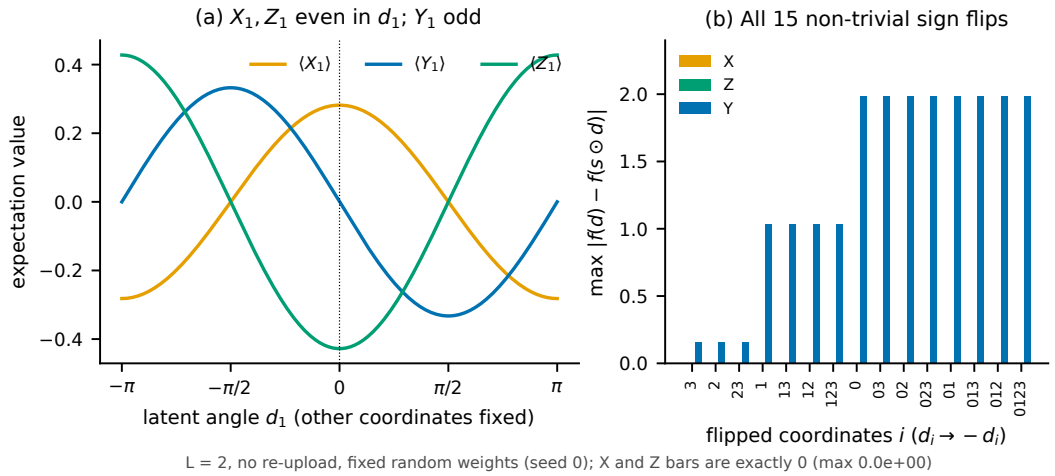

Figure 2: Numerical illustration of Proposition 2. (a) Sweeping d_1 with other coordinates fixed gives even $\langle X_1 \rangle$ and $\langle Z_1 \rangle$ curves and an odd $\langle Y_1 \rangle$ curve. (b) Over all 15 non-trivial sign masks, the maximum deviation is exactly zero for X and Z, while Y changes.

Proof of Proposition 2. We use the Heisenberg picture and propagate a measured Pauli operator backwards through the circuit. For a local X_q or Z_q measurement, conjugation by the final product of R_Y rotations maps the operator into the real Pauli span generated by I, X, and Z. Conjugation by a CNOT maps this span to itself: for example, $X_c \mapsto X_c X_t$, $Z_t \mapsto Z_c Z_t$, while Z_c and X_t remain single-axis Pauli operators. Consequently, after propagating through the specified CNOT chain, every term is a Pauli string containing no Y factor. Conjugation by the first product of R_Y rotations preserves the same span, so the backward-propagated observable is a linear combination of strings from $\{I, X, Z\}^{\otimes n}$.

At the encoded product state, each single-qubit factor has expectations

$$\mathrm{tr}(\rho_i I) = 1, \qquad \mathrm{tr}(\rho_i X) = 0, \qquad \mathrm{tr}(\rho_i Z) = \cos d_i,$$

with $\rho_i = R_X(d_i)|0\rangle\langle 0|R_X(-d_i)$. The expectation of every surviving Pauli string is therefore a product of constants and coordinate-wise cosines. Summing the terms gives a polynomial in $(\cos d_0, \ldots, \cos d_{n-1})$, which is invariant under any subset of coordinate-wise sign flips.

For a local Y_q measurement, the same propagation leaves one Y factor on the measured wire in the relevant terms. Since

$$\mathrm{tr}(\rho_q Y) = -\sin d_q,$$

while all other nonzero factors are constants or $\cos d_i$, the expectation has the form

$$\langle Y_q \rangle = \sin(d_q) g_q(\cos d),$$

for a circuit-dependent polynomial g_q. This establishes the claimed odd component and also shows why it need not span an arbitrary odd target. The argument relies on the stated gate order and on there being no second data-encoding layer; re-uploading can introduce additional frequency and parity terms.

The two proofs make the logical separation explicit: Proposition 1 concerns the state under complex conjugation, whereas Proposition 2 concerns the smaller function class exposed by a particular measurement interface. Neither proposition alone implies a prediction gap; that implication is tested by the controlled experiments below.

Corollary 3 (architectural invariance). If the prediction map receives only $\mathrm{PQC}_Z(d)$ and no tensor depending directly on d, then the readout parity is inherited by the prediction map. QNN-SRC and QNN-NRC include classical latent shortcuts, so their full prediction maps need not be even even when their quantum readouts are.

Table 1: Test MAE averaged over 10 seeds with end-to-end training. Entries in the last two columns are paired relative changes with 95% bootstrap confidence intervals; bold intervals exclude zero in the indicated direction.

System	Model	Z	$Z+X$	$Z+Y$	XYZ	$\Delta(Z+X-Z)$	$\Delta(Z+Y-(Z+X))$
case9	Hybrid QNN	0.448	0.366	0.361	0.337	**-18.3% [-0.144,-0.033]**	-1.3% [-0.064,+0.048]
case14	Hybrid QNN	0.536	0.485	0.496	0.469	**-9.5% [-0.085,-0.020]**	+2.3% [-0.015,+0.036]
case118	Hybrid QNN	1.211	1.250	1.211	1.181	+3.2% [-0.062,+0.132]	-3.1% [-0.114,+0.030]
case9	QNN-SRC	0.232	0.218	0.217	0.214	**-5.9% [-0.028,-0.000]**	-0.5% [-0.007,+0.005]
case14	QNN-SRC	0.387	0.395	0.379	0.382	+2.2% [-0.094,+0.109]	-4.2% [-0.076,+0.048]
case118	QNN-SRC	1.055	1.042	1.029	1.025	-1.3% [-0.056,+0.030]	-1.3% [-0.041,+0.013]
case9	QNN-NRC	0.222	0.210	0.222	0.206	-5.7% [-0.028,+0.003]	**+5.7% [+0.001,+0.024]**
case14	QNN-NRC	0.253	0.227	0.228	0.246	-10.2% [-0.067,+0.014]	+0.5% [-0.024,+0.024]
case118	QNN-NRC	0.822	0.933	0.974	0.858	+13.5% [-0.071,+0.318]	+4.4% [-0.002,+0.093]

Table 2: Test MAE averaged over 10 seeds with the encoder frozen at random initialization. The final column counts seeds for which $Z+Y$ outperforms $Z+X$.

System	Z	$Z+X$	$Z+Y$	XYZ	$\Delta(Z+X-Z)$	$\Delta(Z+Y-(Z+X))$	Seeds
case9	2.267	2.241	1.536	1.466	-1.1% [-0.091,+0.051]	**-31.5% [-1.110,-0.328]**	10/10
case14	1.018	0.968	0.888	0.873	**-4.8% [-0.095,-0.015]**	**-8.3% [-0.126,-0.032]**	9/10
case118	1.927	1.924	1.732	1.743	-0.2% [-0.038,+0.051]	**-10.0% [-0.327,-0.076]**	10/10

4 Experimental Design

We evaluate case9, case14, and case118 AC-OPF systems using 300 samples per system from the fixed `jit15` split and a fixed 0.8/0.1/0.1 train/validation/test partition. The PQC uses four qubits. The encoder has hidden dimensions $(32, 4)$. We use Adam, learning rate 10^{-3}, batch size 64, mean-squared error, and 100 epochs. The protocol is a controlled reimplementation and does not use the PINN/KKT loss or training schedule of the original QNN-SRC/NRC work (Hu et al., 2026).

The primary intervention compares $Z+X$ and $Z+Y$. Each produces eight features and uses the same decoder parameterization. The theoretical difference is that X is even under the coordinate-wise sign symmetry, whereas Y carries an odd component. Z-only and local XYZ readouts are retained as baseline and expanded-readout references. A parameter-matched OCR control reduces decoder width, but it is supplementary because it does not match feature dimension and conditioning simultaneously.

We use two positive controls. In the no-encoder control, $d \sim \mathcal{U}[-\pi, \pi]^4$ is supplied directly to the PQC and a linear decoder, preventing a classical encoder from reparameterizing the input. In the frozen-encoder control, the encoder is held at random initialization while the quantum circuit and decoder are trained. These are mechanism controls, not proposed operating configurations. Main AC-OPF comparisons use 10 seeds; the no-encoder control uses five seeds. Confidence intervals are paired percentile-bootstrap intervals.

5 Results

5.1 End-to-end and frozen-encoder comparison

The main end-to-end results are reported in Table 1. On Hybrid QNN case9 and case14, $Z+X$ captures most of the benefit of a wider readout. Case118 does not follow this pattern: $Z+X$ is worse than Z. No $Z+Y$ versus $Z+X$ cell gives a statistically decisive improvement in the end-to-end table. Thus, the result does not support attributing an end-to-end gain directly to sign recovery or to feature count alone.

The frozen-encoder results are shown in Table 2. Here $Z+Y$ outperforms $Z+X$ on all three systems, with relative improvements of 31.5%, 8.3%, and 10.0%. The effect is therefore conditional: the odd channel is useful when the encoder cannot adapt around the readout symmetry. The frozen-encoder MAE is approximately four to five times larger than the trained-encoder MAE, so this table is a sensitivity experiment rather than a deployment benchmark.

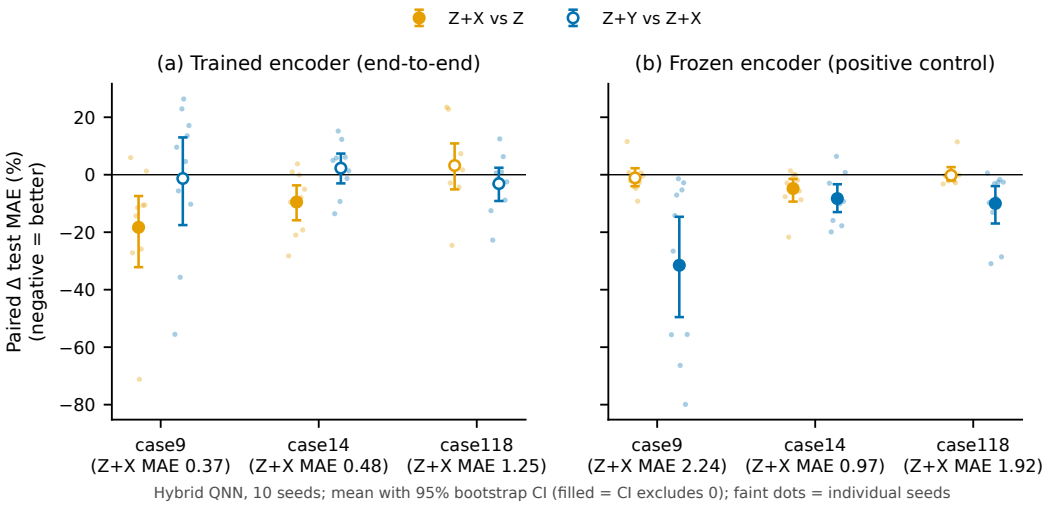

Figure 3: Training-regime-dependent dissociation for Hybrid QNN. (a) With a trained encoder, $Z + X$ improves MAE on case9 and case14, while $Z + Y$ is not consistently better than $Z + X$. (b) With the encoder frozen, adding X has little effect, whereas $Z + Y$ improves over $Z + X$ on all three systems. Filled markers indicate paired 95% intervals excluding zero; faint points are individual seeds.

Table 3: R^2 on the synthetic no-encoder positive control, averaged over five seeds. Error ranges are min–max across seeds. The even target is included as a parity control rather than as an absolute accuracy benchmark.

Target	Z	$Z + X$	$Z + Y$	XYZ
$\sin(d_0 + d_1 + d_2 + d_3)$	-0.02 (-0.05,0.00)	-0.02 (-0.05,0.00)	**0.23 (0.20,0.26)**	**0.23 (0.20,0.26)**
$\sum_i \cos d_i$	0.50	0.50	0.49	0.49
$\sum_i \sin d_i$	-0.01	-0.01	-0.01	-0.01

5.2 Positive controls

The synthetic no-encoder control is summarized in Table 3 and Figure 4. The Y-containing readout recovers part of the jointly odd target $\sin(d_0 + d_1 + d_2 + d_3)$, whereas Z and $Z + X$ do not. The even target reaches approximately $R^2 = 0.50$ for all readouts. The target $\sum_i \sin d_i$ remains poorly predicted for every readout because each local Y_q is multiplied by cosine factors from other coordinates. Thus, Y is a task- and ansatz-dependent parity channel, not a universal solution for odd functions.

6 Discussion

The experiments separate three claims that are often conflated. First, the circuit state changes under sign transformations. Second, the chosen X/Z observables can be invariant to those transformations. Third, that invariance affects prediction only when the rest of the architecture cannot replace the missing information. The frozen-encoder and no-encoder controls establish the second-to-third link under restricted conditions; the end-to-end results show that it is not universal.

The result also clarifies null results from broader readouts. The absence of a consistent Y advantage in end-to-end training does not contradict Proposition 2; it means that the readout restriction is not the predictive bottleneck in that regime. Conversely, a frozen-control gain cannot by itself establish quantum advantage. It establishes sensitivity to a parity channel under a specified intervention.

The central open question is why X helps in some trained-encoder configurations but not others. The target-aware audit rules out a simple universal orthant story, yet does not distinguish optimiza-

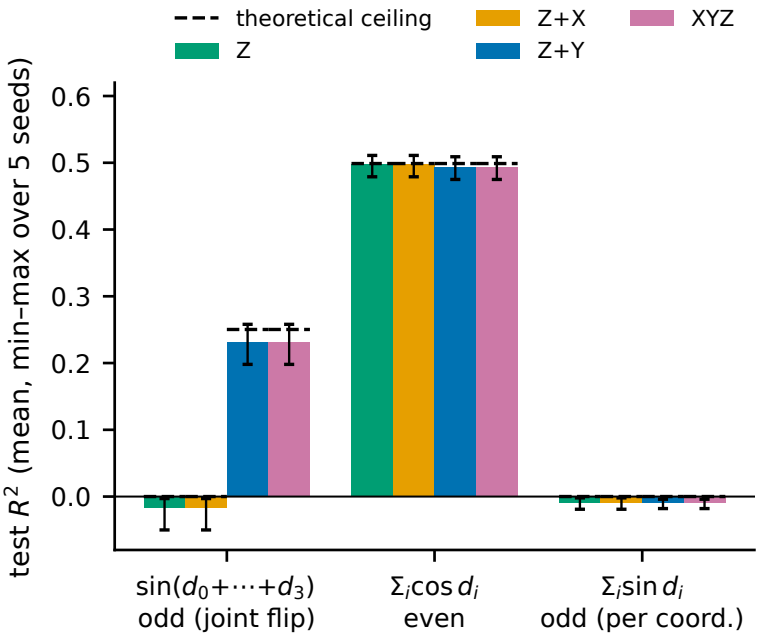

Figure 4: No-encoder positive control. Test R^2 is shown as the mean and min–max range over five seeds. The odd joint target is recovered only when the readout contains Y; the even target is a parity control, and the coordinate-wise odd target remains unresolved by all readouts.

tion from conditioning or classical bypasses. This is precisely why MAE alone is insufficient for attributing a gain to entanglement, to an observable, or to one encoder mechanism.

7 Limitations

The study uses three grid systems, 300 samples per system, four qubits, statevector simulation, and fixed data splits. It is a controlled reimplementation without the original PINN/KKT loss and is not a direct reproduction of published absolute performance. Proposition 2 is proved only for the stated non-reuploading circuit; deeper and re-uploading variants are numerical checks. The target-aware audit uses checkpoints from `qnn_original_residuals_v1` and `baselines_fixed_jit15`, with Hybrid QNN represented only by case14 and case118, so it is auxiliary evidence rather than an independent reproduction of Tables 1–2. We do not evaluate finite-shot noise, measurement cost, physical AC-OPF closure, or classical surrogate baselines in this paper.

8 Conclusion

We identify a concrete readout failure mode in a real R_X-encoded hybrid QNN. In the two-layer non-reuploading circuit, local X/Z features depend only on coordinate-wise cosines, so the state can change under sign flips while the measured interface remains blind to that information. A matched $Z + X$ versus $Z + Y$ intervention and two positive controls show when this blindness becomes predictive: it matters when the encoder is restricted, but it is often bypassed or absorbed by the full architecture when the encoder is trained end-to-end. The resulting lesson is methodological: state dependence, observable visibility, and predictive utility should be reported as separate quantities before claiming a benefit from a quantum interaction or a measurement choice.

References

Alice Barthe and Adrián Pérez-Salinas. Gradients and frequency profiles of quantum re-uploading models. *Quantum*, 8:1523, 2024. doi: 10.22331/q-2024-11-14-1523.

Joseph Bowles, Shahnawaz Ahmed, and Maria Schuld. Better than classical? the subtle art of benchmarking quantum machine learning models, 2024. URL https://arxiv.org/abs/2403.07059.

Yuji Cao, Yue Chen, and Yan Xu. A differentially private quantum neural network for probabilistic optimal power flow. In *IET Conference Proceedings*, volume 2024, pp. 601–608, 2025. doi: 10.1049/icp.2025.0552.

Matthias C. Caro, Elies Gil-Fuster, Johannes Jakob Meyer, Jens Eisert, and Ryan Sweke. Encoding-dependent generalization bounds for parametrized quantum circuits. *Quantum*, 5:582, 2021. doi: 10.22331/q-2021-11-17-582.

Samuel Yen-Chi Chen, Huan-Hsin Tseng, Hsin-Yi Lin, and Shinjae Yoo. Learning to measure quantum neural networks, 2025a. ICASSP 2025 Workshop.

Samuel Yen-Chi Chen, Huan-Hsin Tseng, Hsin-Yi Lin, and Shinjae Yoo. Learning to program quantum measurements for machine learning, 2025b.

Fei Feng, Yifan Zhou, and Peng Zhang. Quantum power flow. *IEEE Transactions on Power Systems*, 36(4):3810–3812, 2021. doi: 10.1109/TPWRS.2021.3077382.

Dominik Freinberger and Philipp Moser. The role of quantum in hybrid quantum-classical neural networks: A realistic assessment, 2026. URL https://arxiv.org/abs/2601.04732.

Keisuke Fujii and Kohei Nakajima. Harnessing disordered-ensemble quantum dynamics for machine learning. *Physical Review Applied*, 8:024030, 2017. doi: 10.1103/PhysRevApplied.8.024030.

Francisco Javier Gil Vidal and Dirk Oliver Theis. Input redundancy for parameterized quantum circuits. *Frontiers in Physics*, 8:297, 2020. doi: 10.3389/fphy.2020.00297.

Ze Hu, Ziqing Zhu, Linghua Zhu, Xiang Wei, Siqi Bu, and Ka Wing Chan. Advancing hybrid quantum neural network for alternating current optimal power flow. *Applied Energy*, 417:127945, August 2026. doi: 10.1016/j.apenergy.2026.127945. arXiv v1–v4 title: "Advancing Hybrid Quantum Neural Network for Alternative Current Optimal Power Flow".

Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16:1050–1057, 2020. doi: 10.1038/s41567-020-0932-7.

Zeynab Kaseb, Matthias Möller, Giorgio Tosti Balducci, Peter Palensky, and Pedro P. Vergara. Quantum neural networks for power flow analysis. *Electric Power Systems Research*, 235:110677, 2024. doi: 10.1016/j.epsr.2024.110677.

Michael Kölle, Jonas Maurer, Philipp Altmann, Leo Sünkel, Jonas Stein, and Claudia Linnhoff-Popien. Disentangling quantum and classical contributions in hybrid quantum machine learning architectures. In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence (ICAART)*, pp. 649–656. SCITEPRESS - Science and Technology Publications, 2024. doi: 10.5220/0012381600003636.

Martin Larocca, Frederic Sauvage, Faris M. Sbahi, Guillaume Verdon, Patrick J. Coles, and M. Cerezo. Group-invariant quantum machine learning. *PRX Quantum*, 3:030341, 2022. doi: 10.1103/PRXQuantum.3.030341.

Thinh Viet Le, Mark M. Wilde, and Vassilis Kekatos. Solving conic programs over sparse graphs using a variational quantum approach: The case of the ac optimal power flow, 2026.

Diana Legziel Levy, Menachem Finkelstein, Peter Chin, Eilon Vaadia, and Sarel Cohen. Quantum-inspired hybrid neural networks for neural decoding: A controlled ablation study of learnable quantum sidecar integration, 2026.

Hsin-Yi Lin, Huan-Hsin Tseng, Samuel Yen-Chi Chen, and Shinjae Yoo. Adaptive non-local observable on quantum neural networks, 2025. IEEE QCE 2025.

Johannes Jakob Meyer, Marian Mularski, Elies Gil-Fuster, Antonio Anna Mele, Francesco Arzani, Alissa Wilms, and Jens Eisert. Exploiting symmetry in variational quantum machine learning. *PRX Quantum*, 4:010328, 2023. doi: 10.1103/PRXQuantum.4.010328.

Hela Mhiri, Leo Monbroussou, Mario Herrero-Gonzalez, Slimane Thabet, Elham Kashefi, and Jonas Landman. Constrained and vanishing expressivity of quantum fourier models. *Quantum*, 9:1847, 2025. doi: 10.22331/q-2025-09-03-1847.

Pere Mujal, Rodrigo Martínez-Peña, Johannes Nokkala, Jorge García-Beni, Gian Luca Giorgi, Miguel C. Soriano, and Roberta Zambrini. Opportunities in quantum reservoir computing and extreme learning machines. *Advanced Quantum Technologies*, 4:2100027, 2021. doi: 10.1002/qute.202100027.

Quynh T. Nguyen, Louis Schatzki, Paolo Braccia, Michael Ragone, Patrick J. Coles, Frederic Sauvage, Martin Larocca, and M. Cerezo. Theory for equivariant quantum neural networks. *PRX Quantum*, 5:020328, 2024. doi: 10.1103/PRXQuantum.5.020328.

Nima Nikmehr, Peng Zhang, and Mikhail A. Bragin. Quantum distributed unit commitment: An application in microgrids. *IEEE Transactions on Power Systems*, 37(5):3592–3603, 2022. doi: 10.1109/TPWRS.2022.3141794.

Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020. doi: 10.22331/q-2020-02-06-226.

Rushat Rai, Pei-Jie Chang, Doan Viet Nguyen, Yuan-Chieh Chiu, Niall Tumilty, Yun-Yuan Wang, Simon See, Wen-Jay Lee, Tai-Yue Li, Nan-Yow Chen, and Tian-Li Wu. Hybrid classical-quantum neural networks for multi-characteristic co-optimization of recessed-gate AlGaN/GaN MIS-HEMTs, 2026.

Louis Schatzki, Martin Larocca, Quynh T. Nguyen, Frederic Sauvage, and M. Cerezo. Theoretical guarantees for permutation-equivariant quantum neural networks. *npj Quantum Information*, 2024. doi: 10.1038/s41534-024-00804-1.

Maria Schuld and Nathan Killoran. Is quantum advantage the right goal for quantum machine learning? *PRX Quantum*, 3:030101, 2022. doi: 10.1103/PRXQuantum.3.030101.

Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. 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.

Yifan Zhou and Peng Zhang. Noise-resilient quantum machine learning for stability assessment of power systems. *IEEE Transactions on Power Systems*, 38(1):475–487, 2023. doi: 10.1109/TPWRS.2022.3160384.

Ziqing Zhu, Shuyang Zhu, and Siqi Bu. Bayesian quantum neural network for renewable-rich power flow with training efficiency and generalization capability improvements, 2024.