

Practical Quantum Error Mitigation for Noisy Quantum Hardware

Daniil Rabinovich^{a,1,2} Zahar Sayapin,¹ Nikita Korolev,² and Kirill Lakhmanskiy²

¹*Skolkovo Institute of Science and Technology, Moscow, Russian Federation*

²*Russian Quantum Center, Skolkovo, Moscow, Russian Federation*

Abstract. Practical deployment of quantum algorithms on noisy intermediate-scale quantum (NISQ) devices requires error mitigation techniques that remain effective under realistic experimental constraints. Existing approaches often rely on accurate noise characterization, simplified assumptions about hardware errors, or substantial computational overhead, limiting their scalability in practical settings. In this work, we investigate two complementary strategies toward practical quantum error mitigation (QEM): exploiting experimentally accessible information about the hardware noise and learning mitigation transformations directly from measurement data.

The first approach develops a hardware-aware extension of zero-noise extrapolation (ZNE) based on structured circuit layout permutations, providing rigorous guarantees while reducing experimental overhead. The second approach introduces a data-driven QEM framework based on supervised machine learning, enabling the reconstruction of noiseless observables without requiring an explicit noise model. Together, these methods demonstrate two complementary directions for adaptive error mitigation: physics-informed protocols that exploit device characteristics and machine-learning approaches that infer mitigation strategies from data.

Hardware-aware Zero-Noise Extrapolation. When information about hardware noise is available, it can be directly incorporated into the mitigation procedure. Recently proposed layout permutation based zero-noise extrapolation (LP-ZNE) exploits the spatial inhomogeneity of error in quantum processors by modifying the assignment between abstract qubits in a circuit and physical qubits of the device [1]. Since different qubits and couplings exhibit different error rates, different circuit layouts naturally correspond to different effective noise strengths. Combining measurements obtained from these layouts enables estimation of the zero-noise expectation value. However, achieving first-order error cancellation with LP-ZNE requires averaging over all possible layouts, resulting in factorial overhead $n!$ for an n -qubit circuit.

We introduce Cyclic Layout Permutation Zero-Noise Extrapolation (CLP-ZNE), which replaces arbitrary layout averaging by structured cyclic permutations. The main steps of CLP-ZNE protocol are sketched in Figure 1. The theoretical guaranties for CLP-ZNE are summarized in two theorems [2]. For an n -qubit circuit with gates subject to a d -parameter noise model with a characteristic error rate q , we prove that CLP-ZNE reconstructs the noiseless expectation value with quadratic bias,

$$E_{\text{mitigated}} = E_0 + O(q^2), \quad (1)$$

using only $n(d + 1)$ circuit layouts, for circuits of linear connectivity. The construction can be extended to arbitrary connectivity while increasing the number of layouts to $O(dn^2)$.

We numerically benchmark CLP-ZNE using expectation-value estimation for random Sherrington-Kirkpatrick Hamiltonians under realistic device calibrations. The results, summarized in Table , demonstrate up to an order-of-magnitude reduction in estimation error while outperforming LP-ZNE and conventional circuit-folding ZNE across realistic hardware noise models.

Data-driven Quantum Error Mitigation. In realistic experimental scenarios, detailed calibration information may be incomplete or unavailable. This motivates a complementary strategy in which the mitigation transformation is learned directly from data rather than constructed from an

^a daniilrabinovich.quant@gmail.com




FIG. 1. **Schematic illustration of the proposed CLP-ZNE protocol applied to a four-qubit circuit under single-parameter noise ($d = 1$).** Circles represent physical qubits, while edges denote available two-qubit couplings, color-coded by their individual error rates q . Numbers on qubits denote assignment of abstract qubits from the circuit into the qubits of the device.

Noise type	Error suppression factor $ E_{\text{noisy}} - E_0 / E_{\text{mitigated}} - E_0 $			
	CLP-ZNE (our result)	LP-ZNE (random layouts)	ZNE (circuit-folding)	ZNE (gate-folding)
FakeTorino, 12-qubit circuits with cyclic connectivity				
Original noise	8.4 [5.1–12.4]	2.0 [0.7–4.9]	2.3 [1.0–5.0]	13.4 [8.0–18.3]
T_1/T_2 noise $\times 10$	12.8 [6–27]	1.6 [0.9–3.1]	1.4 [0.6–2.8]	5.0 [3.1–9.2]
Pauli noise	40.5 [25–60]	3.6 [1.5–8.7]	1.3 [0.6–2.9]	12.5 [5.4–25.8]
NClusterBackend, 9-qubit circuits with square-grid connectivity (3 by 3)				
Original noise	9.6 [4.6–16.6]	4.2 [2.1–11.3]	1.2 [0.5–2.5]	18.0 [11–29]
T_1/T_2 noise $\times 10$	7.4 [5.3–10.6]	1.6 [1.0–2.7]	0.9 [0.4–1.9]	5.8 [3.3–9.6]
Pauli noise	17.6 [12–25]	3.9 [1.8–7.5]	0.8 [0.4–1.9]	17.6 [8.2–37.6]

TABLE I. Comparison of ZNE based error mitigation methods across backends, noise settings, and circuit connectivities. Values are reported as median [Q1–Q3], where higher error suppression factors $|E_{\text{noisy}} - E_0|/|E_{\text{mitigated}} - E_0|$ indicate better performance.

explicit noise model. We introduce a machine-learning-based quantum error mitigation (ML-QEM) framework designed for variational quantum algorithms [3].

The central challenge in supervised QEM lies in generating representative training data without requiring costly classical simulations of large parametrized circuits. We address this challenge by exploiting the efficient classical simulability of Clifford and near-Clifford circuits. Training datasets are generated by replacing parametrized rotations of the target ansatz with randomly sampled Clifford operations and, in the near-Clifford setting, by probabilistically introducing Haar-random rotations. The overall scheme is depicted in Figure 2. The measurement data from noisy and noiseless circuits is then used to learn the relation between noisy and noiseless expectations. The trained regression model can then be applied to the chosen ansatz structure at arbitrary parameters, providing error mitigation without additional quantum resources.

We benchmark the proposed framework using VQE for random Sherrington–Kirkpatrick Hamiltonians under depolarizing, Pauli, and composite noise models. Table II summarizes representative distributions of the absolute energy error after mitigation. Across all considered noise models, ML-



FIG. 2. **Workflow of the proposed machine learning quantum error mitigation protocol.** Training data are generated from efficiently simulable Clifford and near-Clifford circuits and used to train a regression model that predicts noiseless expectation values from noisy measurements.

TABLE II. Comparison of error mitigation methods across different noise settings in terms of the error suppression factors $|E^{\text{noisy}} - E^{\text{ideal}}|/|\hat{E} - E^{\text{ideal}}|$. Values are reported as median and [Q1–Q3] of the distributions across considered instances. Higher error suppression factors indicate better performance. Bold font indicates the best performing model.

p	Ridge		XGBoost		ZNE	Noisy RMSE
	near-Clifford	Clifford	near-Clifford	Clifford		
Depolarizing Noise						
0.01	8.2 [5.1–14.8]	6.6 [3.1–15.7]	0.8 [0.6–0.9]	1.2 [0.8–2.1]	44.4 [36.7–56.4]	1.849
0.05	8.0 [3.9–20.1]	3.6 [2.1–7.8]	1.9 [1.6–2.1]	5.0 [3.4–10.3]	4.0 [3.5–5.1]	6.507
0.1	3.5 [1.8–7.5]	2.3 [1.3–4.8]	2.2 [1.9–2.5]	5.3 [3.4–12.5]	2.6 [2.1–3.2]	9.289
Pauli Noise						
0.01	7.8 [4.7–13.0]	7.9 [4.6–19.9]	0.6 [0.5–0.8]	1.3 [0.8–3.0]	70.0 [43.2–148.5]	1.588
0.05	7.4 [3.6–14.6]	5.8 [3.6–12.4]	1.9 [1.5–2.1]	5.4 [3.0–14.7]	4.4 [3.7–5.6]	5.754
0.1	4.4 [2.6–8.7]	3.3 [2.1–7.7]	2.0 [1.8–2.3]	5.7 [3.5–9.7]	2.8 [2.3–3.5]	8.439
Composite						
0.01	7.8 [4.5–14.0]	5.9 [2.7–13.3]	1.1 [0.9–1.3]	0.9 [0.7–1.4]	19.0 [15.1–26.4]	2.886
0.05	5.5 [3.3–9.8]	4.2 [2.8–9.4]	1.8 [1.7–2.1]	6.1 [3.4–10.6]	2.8 [2.4–3.6]	8.446
0.1	3.4 [2.2–6.3]	2.5 [1.9–3.7]	1.6 [1.4–1.7]	2.4 [1.7–3.0]	1.0 [1.0–1.0]	10.825

QEM consistently suppresses the estimation error by several-fold, with Ridge regression trained on near-Clifford datasets providing the best overall performance. As the noise strength increases, the mitigation accuracy remains more stable than in conventional unitary-folding ZNE, which swiftly deteriorates. In the high-noise regime, the proposed ML-QEM protocol becomes comparable to, and can outperform standard ZNE.

Conclusion. The presented approaches highlight two complementary routes toward practical and adaptive quantum error mitigation. CLP-ZNE demonstrates how knowledge of hardware noise structure can be converted into efficient mitigation protocols with rigorous guarantees, while ML-QEM shows how mitigation strategies can be learned directly from experimental data when explicit noise characterization is unavailable. Together, these results suggest a broader framework in which physical understanding of quantum devices and data-driven inference can be combined to improve the reliability of near-term quantum algorithms.

-
- [1] Alexey Uvarov, Daniil Rabinovich, Olga Lakhmanskaya, Kirill Lakhmanskiy, Jacob Biamonte, and Soumik Adhikary. Mitigating quantum gate errors for variational eigensolvers using hardware-inspired zero-noise extrapolation. *Phys. Rev. A*, 110:012404, Jul 2024.
 - [2] Zahar Sayapin, Daniil Rabinovich, Nikita Korolev, and Kirill Lakhmanskiy. Zero-noise extrapolation via cyclic permutations of quantum circuit layouts. *arXiv preprint arXiv:2511.02901*, 2025.
 - [3] Nikita Korolev, Kirill Lakhmanskiy, and Daniil Rabinovich. Machine learning-based quantum error mitigation for variational algorithms. *arXiv preprint arXiv:2606.02697*, 2026.