

Machine Learning-based Quantum Error Mitigation for Variational Algorithms

Nikita Korolev^{a,1}, Kirill Lakhmanskiy¹ and Daniil Rabinovich^{1,2}

¹*Russian Quantum Center, Skolkovo, Moscow, Russian Federation*

²*Skolkovo Institute of Science and Technology, Moscow, Russian Federation*

Variational Quantum Algorithms (VQAs) are promising candidates for near-term quantum computing but remain limited by hardware noise. While Zero-Noise Extrapolation (ZNE) relies on explicit noise scaling and hardware control, accurate noise models and calibration data are often unavailable. We therefore propose a practical machine learning-based quantum error mitigation (ML-QEM) protocol that learns a mitigation map directly from noisy measurements without requiring an explicit noise model.

The mitigation map is obtained by training a supervised machine learning model to reconstruct noiseless Pauli expectation values from noisy measurements, $f_\phi : \mathbf{P}_{\text{noisy}} \rightarrow \mathbf{P}_{\text{ideal}}$, where $\mathbf{P}$ denotes the vector of measured Pauli expectation values. The mitigated energy is then computed as $\hat{E} = \sum_i h_i \langle \hat{P}_i \rangle$, where $\hat{\mathbf{P}} = f_\phi(\mathbf{P}_{\text{noisy}})$. Training data are generated from efficiently classically simulable Clifford and near-Clifford circuits by replacing parametrized rotations in the ansatz with randomly sampled Clifford gates and, probabilistically, Haar-random rotations. The resulting noisy and noiseless expectation values define the supervised learning dataset, allowing the trained model to generalize to circuits with arbitrary variational parameters without requiring additional quantum resources.

We benchmark the protocol on the Variational Quantum Eigensolver (VQE) for 12-qubit Sherrington–Kirkpatrick Hamiltonians under depolarizing, Pauli, and composite noise models. Among the considered regression models, Ridge Regression trained on near-Clifford datasets achieved the best overall performance. As shown in Fig. 1, ML-QEM consistently shifts the error distributions toward lower values, reducing energy estimation errors by several-fold across all noise models. While ZNE performs well at low noise levels, its accuracy deteriorates as noise increases, whereas ML-QEM remains robust and becomes comparable to or outperforms ZNE in the high-noise regime. Moreover, post-optimization mitigation achieves essentially the same accuracy as in-loop mitigation at a substantially lower computational cost.




FIG. 1. **Absolute error $|E^{\text{ideal}} - \hat{E}|$ distribution for each considered noise model with corresponding noise strength 0.05** Here ML-QEM depicts the results of best acquired models for each case, i.e. Ridge regression trained on near-Clifford data for depolarization (a), Pauli noise (b), and Ridge regression trained on Clifford data for the composite noise model (c).

^a n.korolev@rqc.ru