

Pauli Noise Annealing for Escaping Optimization Stagnation in Variational Quantum Eigensolvers

Jungyun Lee¹, Yujin Kim¹, and Daniel K. Park^{*1,2,3}

¹Department of Statistics and Data Science, Yonsei University, Seoul, Republic of Korea

²Department of Applied Statistics, Yonsei University, Seoul, Republic of Korea

³Department of Quantum Information, Yonsei University, Seoul, Republic of Korea

Abstract

Variational quantum eigensolvers (VQEs) provide a near-term approach to ground-state energy estimation, but their performance is hindered by nonconvex loss landscapes with local minima and saddle points. Although noise is usually viewed as detrimental, recent studies show that structured Pauli noise can instead regularize loss landscapes by suppressing high-frequency components. Motivated by this idea, we apply Pauli-noise annealing to VQE training. Numerical simulations show improved convergence reliability over noiseless and finite-shot baselines. We further study annealing schedules and extended noise-factor regimes, showing that stronger landscape deformation can enhance saddle-point escape. These results support adaptive noise annealing for reliable noise-assisted VQE.

1 Introduction

Variational quantum algorithms (VQAs), particularly the variational quantum eigensolver (VQE), are promising NISQ-era approaches for estimating ground-state energies via quantum-classical hybrid optimization [1, 2]. However, their performance can be limited by barren plateaus, local minima, and saddle points in nonconvex energy landscapes [3–6]. Noise has traditionally been viewed as an obstacle, as it can exacerbate these problems by inducing loss-function concentration and suppressing gradients [7]. However, recent theoretical studies suggest a more nuanced picture, showing that suitable noise channels and circuit structures can instead prevent such concentration and need not be detrimental to trainability [8, 9]. In particular, finite-shot fluctuations may perturb optimization trajectories away from saddle points [10], and controlled Pauli noise suppresses high-frequency Fourier components and thereby smooths unfavorable local structures [11].

Motivated by these studies, we apply the Pauli-noise annealing framework [11] to VQE training, an application that has not been explored in previous work. Moreover, the dependence on the annealing schedule and the potential role of extending the noise factor μ over its full physically allowed range remain open questions. To address these gaps, we conduct statistical experiments on 8-qubit random mixed-field Ising model (MFIM) instances and show that Pauli-noise annealing improves convergence reliability toward near-ground-state solutions compared with noiseless and finite-shot sampling-based optimization [10]. We also compare an exponential-decay schedule with a manually staged schedule and find that both yield similar improvements. In addition, using a saddle-point escape example on a 4-qubit local- Z Hamiltonian, we examine the effect of moving beyond the conventional smoothing-only regime. We show that the sign-alternating deformation regime can qualitatively reshape the loss landscape and enable escape from stagnation in cases where smoothing alone fails. These results suggest that extending the range of the noise factor μ can enhance the optimization capability of Pauli-noise annealing and motivate adaptive strategies that select μ according to optimization progress.

*dkd.park@yonsei.ac.kr




Figure 1: VQE optimization failure rates on ten independent 8-qubit random mixed-field Ising model (MFIM) instances.

2 Method

Pauli Noise Regularization For a parameterized circuit $U(\theta)$ and observable H , the VQA loss $L(\theta) = \langle 0|U^\dagger(\theta)HU(\theta)|0\rangle$ can be decomposed as a trigonometric polynomial $L(\theta) = \sum_m L_m(\theta)$, where L_m contains order- m terms. The Pauli noise annealing method introduced in [11] smooths the loss landscape by suppressing each term according to its order: $L(\theta, \mu) = \sum_m (1 - \mu)^m L_m(\theta)$, where $\mu \in [0, 2]$ denotes the noise factor. For $0 \leq \mu \leq 1$, $(1 - \mu)^m$ increasingly suppresses high-frequency components and smooths the landscape [11]. We further explore the sign-alternating regime $1 < \mu \leq 2$, where the factor $(1 - \mu)^m$ alternates in sign with the order m , deforming the landscape rather than merely smoothing it. This can be implemented by applying the axis-matched Pauli noise channels $\mathcal{E}_{\mathbb{P}_\mu}(\rho) = (1 - \mu/2)\rho + (\mu/2)\mathbb{P}\rho\mathbb{P}$, $\mathbb{P} \in \{X, Y, Z\}$, before each parameterized single-qubit rotation, using the same μ throughout the circuit.

Noise Scheduling We consider two Pauli noise annealing schedules: The fixed noise annealing scheme uses $\mu_{\text{fix}}(t) = \mu_0 e^{-at/T}$ [11], with t the current step, T the total number of steps, μ_0 the initial noise strength, and a the decay-rate parameter, thereby continuously approaching the noiseless objective. Manual annealing instead follows a prescribed sequence $[\mu^{(0)}, \mu^{(1)}, \dots, \mu^{(K)} = 0]$. At stage k , optimization continues at $\mu^{(k)}$ until convergence, and the resulting parameters initialize the next stage. Thus, manual stages can have different durations, whereas the fixed schedule is determined solely by the step index. The final stage $\mu^{(K)} = 0$ recovers the noiseless objective.

3 Results

We conduct numerical VQE experiments comparing noiseless VQE, fixed Pauli-noise annealing, manual Pauli-noise annealing, and finite-shot VQE using a strongly entangling layer ansatz [12]. The Pauli noise channels can be implemented by randomized Pauli sampling with probability $\mu/2$.

Our statistical experiment uses ten 8-qubit random MFIM Hamiltonians, which correspond to Ising models with both transverse and longitudinal fields, $H = \sum_i J_{zz} Z_i Z_{i+1} + \sum_i (J_x X_i + J_z Z_i)$. For each instance, J_{zz} , J_x , and J_z are randomly sampled from $[-2, 2]$. Every method is evaluated from the same 1000 random seeds, and the optimized parameters are evaluated on the noiseless objective to obtain E_{test} . We measure the final solution quality by $\Delta_{\text{norm}} = (E_{\text{test}} - E_g)/(E_{\text{top}} - E_g)$, where E_g and E_{top} are the smallest and largest eigenvalues. Figure 1 reports failure rates defined by $\Delta_{\text{norm}} > 0.1$. Across these instances, both Pauli-noise schedules produce fewer failures than the noiseless baseline, whereas




Figure 2: Saddle-point escape experiment on local- Z Hamiltonian $\sum_i Z_i$.

finite-shot noise does not consistently avoid poor solutions. This result demonstrates that Pauli noise can improve convergence reliability by regularizing the optimization landscape [11].

We additionally examine a saddle-point stagnation example using the 4-qubit local- Z Hamiltonian $H = \sum_i Z_i$ considered in Ref. [10]. From 1000 random initial ansatz parameters sampled from $[0, 2\pi)$, we first identify initial points for which the noiseless optimization process converges to a suboptimal, nearly stationary point and select the highest-ranked trapped representative according to our saddle diagnostic, which prioritizes endpoints with small gradient norm and negative minimum Hessian eigenvalue. Figure 2 compares trajectories from this fixed initialization. At $\mu_0 = \mu^{(0)} = 0.9$, noiseless, fixed, and manual annealing remain trapped near $E \simeq -2$. Finite-shot optimization can occasionally escape, but the escape is not reliable: the 1000- and 100-shot cases fail in all 10 trials, while the 30-shot case escapes in 2 out of 10 trials. Together with Fig. 1, this indicates that stochastic shot noise alone does not consistently avoid poor local structures, and that the smoothing regime $0 \leq \mu \leq 1$ is insufficient for this selected initial condition. In contrast, for $\mu_0 = \mu^{(0)} = 1.4$, the manual Pauli-noise annealing schedule reaches the global minimum, whereas the fixed schedule and the noiseless trajectory remain trapped. This suggests that successful escape in this example requires a sufficiently strong initial landscape deformation together with an annealing schedule that maintains the deformation long enough to leave the stagnation region.

4 Conclusion and Future Work

We investigated Pauli-noise annealing as a noise-assisted VQE optimization strategy. Across ten 8-qubit MFIM instances, both annealing schedules reduce failures relative to noiseless and finite-shot baselines. In the 4-qubit saddle-point stagnation example, extending the noise regime enables a stronger landscape deformation that allows the manual schedule to leave the stagnation region when schedules in the smoothing regime remain trapped. These results suggest that Pauli-noise annealing can improve convergence reliability, while the sign-alternating deformation regime may help escape stagnation. Future work should develop adaptive schedules that select the noise strength and stage duration based on optimization progress, and test their scalability on larger VQE problems.

References

- [1] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):4213, 2014.
- [2] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
- [3] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.
- [4] S. S. Du, C. Jin, J. D. Lee, M. I. Jordan, A. Singh, and B. Poczos. Gradient descent can take exponential time to escape saddle points. *Advances in Neural Information Processing Systems*, 30, 2017.
- [5] L. Bittel and M. Kliesch. Training variational quantum algorithms is NP-hard. *Physical Review Letters*, 127(12):120502, 2021.
- [6] E. R. Anschuetz and B. T. Kiani. Quantum variational algorithms are swamped with traps. *Nature Communications*, 13(1):7760, 2022.
- [7] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1):6961, 2021.
- [8] G. Crognaletti, M. Grossi, and A. Bassi. Estimates of loss function concentration in noisy parametrized quantum circuits. *PRX Quantum*, 7(2):020336, 2026.
- [9] A. A. Mele, A. Angrisani, S. Ghosh, S. Khatrri, J. Eisert, D. Stilck França, and Y. Quek. Noise-induced shallow circuits and the absence of barren plateaus. *Nature Physics*, pages 1–6, 2026.
- [10] J. Liu, F. Wilde, A. A. Mele, X. Jin, L. Jiang, and J. Eisert. Stochastic noise can be helpful for variational quantum algorithms. *Physical Review A*, 111(5):052441, 2025.
- [11] D. S. Bagaev, M. A. Gavreev, A. S. Mastiukova, A. K. Fedorov, and N. A. Nemkov. Regularizing quantum loss landscapes by noise injection. *Physical Review A*, 112(3):032417, 2025.
- [12] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020.