

Gradient Variance Measurements in Parameterized Quantum Circuits

Usama Inam Paracha, Abrar Ul Haque, Muhammad Mustafa Umar Gondel, Uman Khalid and Hyundong Shin

Department of Electronics and Information Convergence Engineering, Kyung Hee University, Korea

Email: hshin@khu.ac.kr

Abstract—Variational Quantum Algorithms (VQAs) face the well-known *barren plateau* problem, where gradients vanish exponentially with qubit number or circuit depth, impeding optimization. Prior works show that *local cost functions* mitigate this issue in ideal, noiseless settings. However, near-term quantum devices suffer from sampling noise, gate noise, and decoherence, raising the question: how robust are local cost landscapes under realistic noise? We investigate this empirically using a hardware-efficient ansatz, parameter-shift gradients, and a custom noise model with depolarizing, thermal relaxation, and readout errors. By sweeping circuit depth and qubit count, we plot the local cost landscape in noiseless vs. noisy cases (10,000 shots). Our results reveal progressive flattening of local cost contours under noise, delineating a practical boundary where trainability collapses. This empirical map complements existing theory on noise-induced barren plateaus and informs ansatz design for noisy intermediate-scale quantum (NISQ) devices.

I. INTRODUCTION

Variational Quantum Algorithms (VQAs) combine parameterized quantum circuits with classical optimization to solve problems in chemistry, optimization, and machine learning [1]–[4]. Their training involves computing gradients of cost functions with respect to circuit parameters. However, these gradients can vanish significantly with increased system size or circuit depth, posing a major challenge to effective optimization [5].

We empirically study the trainability of local cost functions under realistic noise conditions, using parameter-shift gradient evaluations, a hardware-efficient ansatz, and a custom Qiskit noise model. Our main contributions include:

- Detailed contour maps of local cost gradients as a function of qubit number and circuit depth, contrasting the noiseless and noisy regimes.
- Quantitative characterization of the trainability boundary, highlighting the impact of noise-induced gradient suppression on variational optimization.

This work advances the understanding of VQA performance limitations in realistic noisy quantum hardware environments.

II. METHODOLOGY

We employ a hardware-efficient ansatz consisting of L layers acting on an n -qubit register. Each layer applies sequential single-qubit rotations about the X , Y , and Z axes on every qubit, with rotation angles parameterizing the circuit. These are followed by an entangling sublayer composed of alternating controlled-NOT gates arranged according to the device connectivity. The rotation parameters are independently

initialized from a uniform distribution over $[0, 2\pi]$. This ansatz balances expressibility and hardware compatibility by limiting circuit depth while creating entanglement across qubits.

Only the local cost function is considered in this work, as prior research has established that local cost functions exhibit significantly improved trainability compared to global cost functions, particularly in mitigating barren plateau phenomena [6].

$$C_{\text{local}}(\theta) = \sum_{i=1}^n \langle Z_i \rangle, \quad (1)$$

The local cost function is defined as the sum of expectation values of single-qubit Z operators, omitting the usual normalization factor for convenience. Although this alters the absolute scale of the cost, it preserves the qualitative features relevant for assessing trainability and noise effects, and aligns with the implementation used in the numerical experiments.

A. Noise Model

A custom noise model is constructed in Qiskit incorporating several key noise sources relevant for realistic quantum hardware. Single-qubit depolarizing errors are modeled with a probability $p_1 = 10^{-3}$, while two-qubit depolarizing errors occur with probability $p_2 = 10^{-2}$. Thermal relaxation is included with characteristic times $T_1 = 80 \mu\text{s}$ and $T_2 = 60 \mu\text{s}$, and gate durations are set to 35 ns for the $u1$ gate, 160 ns for the $u2$ gate, and 400 ns for the controlled-NOT gate. Measurement noise is captured through readout error channels with probability $p_{ro} = 0.02$. The noise processes are applied after every quantum gate operation, and measurement noise is modeled via bit-flip channels acting on the readout outcomes.

B. Gradient Estimation

Gradients are computed using the parameter-shift rule:

$$\frac{\partial C(\theta)}{\partial \theta_i} = \frac{1}{2} \left[C(\theta_i + \frac{\pi}{2}) - C(\theta_i - \frac{\pi}{2}) \right], \quad (2)$$

with 10,000 measurement shots per expectation to approximate hardware sampling.

For each (n, L) pair, 32 random parameter sets are sampled, their gradients are calculated, and the variance is recorded. The reported values correspond to the mean variance over 10 random initializations.

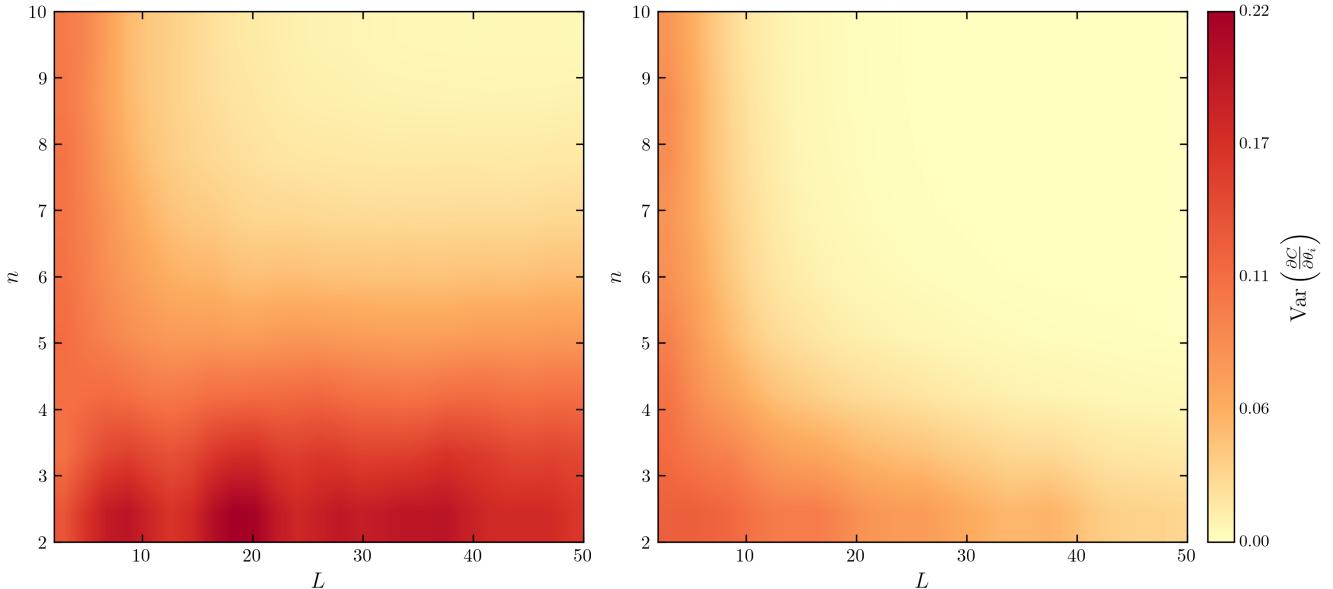


Figure 1. Mean gradient variance as a function of circuit depth (L) and number of qubits (n). Left panel shows results for noiseless simulations, while the right panel depicts noisy simulations with 10,000 shots. The color scale represents the magnitude of the mean gradient variance, with flattening regions indicating the onset of barren plateaus.

III. RESULTS

Figure 1 compares the mean gradient variance for the local cost function as a function of circuit depth (L) and qubit number (n) under both noiseless (left panel) and noisy (right panel) conditions. In the noiseless case, the gradient landscape retains contrast across moderate depths and system sizes, with vanishing gradients emerging only gradually as number of qubits and depth increases. This behavior indicates that local cost functions can delay the onset of barren plateaus, maintaining trainability even as circuit complexity grows under ideal conditions.

When noise is introduced, the gradient variance diminishes rapidly beyond moderate depths and qubit counts, producing flattening and barren plateaus at smaller system sizes compared to the noiseless case. Hardware noise therefore accelerates plateau formation and negates much of the robustness provided by local cost functions, highlighting that noise-induced barren plateaus remain a fundamental limitation on the practical scalability of variational quantum algorithms in current NISQ devices.

IV. CONCLUSION

The empirical analysis conducted in this work highlights the nuanced role of local cost functions in variational quantum algorithms under realistic noise conditions. While local costs demonstrate robust gradient behavior and mitigate barren plateau effects in noiseless simulations, the inclusion of hardware-relevant noise significantly accelerates gradient vanishing, limiting trainability at shallower circuit depths and smaller system sizes. These findings underscore the critical impact of noise on the scalability of variational algorithms

and the necessity for continued development of noise-resilient cost function designs and error mitigation strategies to advance practical quantum computing on near-term devices.

ACKNOWLEDGMENT

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) under RS-2025-00556064, by the MSIT (Ministry of Science and ICT), Korea, under the Convergence security core talent training business support program (IITP-2025-RS-2023-00266615) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation), and by a grant from Kyung Hee University in 2023 (KHU-20233663).

REFERENCES

- [1] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, and L. Cincio, “Variational Quantum Algorithms,” *Nat. Rev. Phys.*, vol. 3, no. 9, pp. 625–644, Sep 2021.
- [2] B. Narottama, A. U. Haq, J. A. Ansere, N. Simmons, B. Canberk, S. L. Cotton, H. Shin, and T. Q. Duong, “Quantum Deep Reinforcement Learning for Digital Twin-Enabled 6G Networks and Semantic Communications: Considerations for Adoption and Security,” *IEEE Trans. Netw. Sci. Eng.*, Sep 2025.
- [3] U. Khalid, U. I. Paracha, Z. Naveed, T. Q. Duong, M. Z. Win, and H. Shin, “Quantum fusion intelligence for integrated satellite-ground remote sensing,” *IEEE Wireless Commun.*, vol. 32, no. 3, pp. 46–55, Jun. 2025.
- [4] S. M. A. Rizvi, U. Khalid, S. Chatzinotas, T. Q. Duong, and H. Shin, “Controlled quantum semantic communication for industrial CPS networks,” *IEEE Trans. Netw. Sci. Eng.*, pp. 1–14, Jul. 2025.
- [5] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren Plateaus in Quantum Neural Network Training Landscapes,” *Nat. Commun.*, vol. 9, no. 1, p. 4812, Nov 2018.
- [6] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, “Cost Function Dependent Barren Plateaus in Shallow Parametrized Quantum Circuits,” *Nat. Commun.*, vol. 12, no. 1, p. 1791, Mar 2021.