Noise-Induced Barren Plateaus

Variational Quantum Algorithms · Module 4 — Barren Plateaus

Node 4.5

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Barren Plateaus **Node:** 4.5

Everything in nodes 4.1-4.4 treated the ansatz as an ideal unitary and the noise as purely statistical (shot noise). Under that assumption, plateaus arise from the *expressivity* of the ansatz — deep, random, 2-design-like circuits concentrate.

This node introduces a **second mechanism** for barren plateaus, one that arises from **hardware noise** rather than ansatz expressivity. **Wang, Fontana, Cerezo, Sharma, Sone, Cincio, Coles (2021): "Noise-induced barren plateaus in variational quantum algorithms."**

The short version: when every layer of an ansatz is followed by incoherent noise (e.g., depolarizing noise), the effective state after L layers is exponentially close to the maximally mixed state $I/2^n$, for *any* initial state and *any* parameters. On the maximally mixed state, every observable has the same expectation value (its trace over dimension), so the cost function is a constant — zero gradient, perfect plateau.

This is a **hardware-origin** barren plateau: the noise eats the signal at a rate set by the gate fidelity, and past a critical depth no ansatz design can recover from it.

The Noise-Induced Mechanism

Suppose every gate in the ansatz is followed by depolarizing noise with strength p:

$$\mathcal{E}_p(\rho) = (1-p)\rho + p\frac{I}{2^n}.$$

After one noisy layer, the state is a convex combination of the ideal state and the maximally mixed state. After L layers, the "ideal component" has weight $(1-p)^L$, and the "noise component" has weight $1-(1-p)^L$.

For L such that $(1-p)^L \approx 2^{-n}$, we have

$$L \approx \frac{n \log 2}{p}.$$

At this depth, the state is exponentially close to maximally mixed, **for any input and any parameters**. Every local observable has expectation value $\mathrm{Tr}(O)/2^n$, which for traceless observables is zero. The gradient is zero, and the optimization is hopeless.

The Wang et al. Theorem

Formally (informal statement):

Theorem (Wang et al. 2021): For an ansatz $U(\boldsymbol{\theta})$ interleaved with depolarizing noise channels of strength p per layer, and for any observable O with bounded norm, the gradient variance satisfies

$$\mathrm{Var}\left[\frac{\partial C}{\partial \theta_i}\right] \leq \mathrm{const} \cdot (1-p)^{2L} \cdot \|O\|^2.$$

This is **exponential in depth**, not in qubit count — a critical difference from the standard 2-design plateau.

Implications:

1. **Depth-independent from the 2-design viewpoint.** You don't need the ansatz to be deep enough to form a 2-design; noise alone creates the plateau if depth exceeds the noise-dominated threshold $L \sim 1/p$.
2. **Qubit-count-independent.** The plateau depth threshold depends on the noise rate, not the qubit count. A 4-qubit noisy deep circuit has the same kind of plateau as a 40-qubit one, if the per-gate noise rate is the same.
3. **Unavoidable by ansatz redesign.** Restricting the ansatz to a low-entanglement subspace does NOT help, because the noise acts on every state, not just on highly entangled ones. You need better *hardware* (lower p) or shallower circuits.
4. **Local cost functions don't help.** Unlike the Cerezo 2021 result, local cost functions give the same noise-induced plateau scaling. The issue is that the full state is corrupted, so no local observation helps.

This makes noise-induced plateaus a qualitatively **harder** problem than 2-design plateaus: the usual ansatz-side fixes don't work.

Practical Depth Limit From Noise

For a device with gate error rate p, the depth at which noise-induced plateaus engage is $L \sim 1/p$.

On 2026 superconducting hardware: - Single-qubit gate error: $\sim 10^{-4}$. - Two-qubit gate error: $\sim 10^{-3}$.

The effective per-layer noise rate is dominated by two-qubit gates, so $p \sim 10^{-3}$ and the plateau threshold is at depth $L \sim 10^3$. This is *far* beyond the depths achievable in current VQA experiments (typically $L < 50$), so noise-induced plateaus are not the bottleneck on today's hardware.

On older hardware with $p \sim 10^{-2}$, the threshold is at depth $L \sim 100$, which *is* reachable and was part of why early VQAs (2016-2020) struggled.

The relationship is simple: **as hardware improves, noise-induced plateaus recede**. This is different from 2-design plateaus, which get *worse* as you try to scale to more qubits. The two plateaus have opposite trends with respect to "more hardware."

Noise-Induced vs 2-Design Plateau: Comparison

Feature	2-Design Plateau	Noise-Induced Plateau
Origin	Ansatz expressivity (Haar distribution)	Hardware noise (depolarizing)
Scaling	$O(2^{-2n})$ variance, exponential in n	$O((1-p)^{2L})$ variance, exponential in L

Feature	2-Design Plateau	Noise-Induced Plateau
Triggered by	Deep ansatz, random init	Any depth, bad hardware
Shallow circuits?	Can avoid	Still affected, just less
Local cost help?	Yes (Cerezo 2021)	No
Ansatz redesign help?	Yes (structure, restriction)	No
Fix	Restrict ansatz, warm start, local cost	Lower gate error, shorter depth
Trend with hardware progress	Gets worse as n grows	Gets better as p shrinks

This table is why the two mechanisms need separate treatment. They look the same operationally (flat training curves) but require different remedies.

Diagnosing Which Plateau You're In

If the training curve is flat, how do you tell which mechanism is responsible?

Test 1: Noiseless simulation. Simulate the same ansatz and cost function without hardware noise, using exact statevector. If the training still fails, it's a 2-design plateau (or some other ansatz/cost problem). If the training succeeds in simulation but fails on hardware, noise is contributing.

Test 2: Vary the depth. Reduce the circuit depth. If the plateau improves, noise is a significant factor (shorter circuits have less accumulated noise). If the plateau stays the same or gets worse as depth decreases, noise is not the dominant mechanism.

Test 3: Vary the qubit count. Reduce the qubit count at fixed depth. If the plateau improves significantly, you're in the 2-design regime (the plateau depends on n). If it stays the same, you're in the noise regime.

Test 4: Measure the effective quantum state. Tomography on the final state, or purity measurement. If the state is close to maximally mixed (purity $\sim 1/2^n$), you're in the noise regime. If the state is pure but the cost is flat, you're in the 2-design regime.

These diagnostics matter because they tell you where to invest effort: on the ansatz side, on the hardware side, or on both.

Coherent vs Incoherent Noise

The Wang et al. 2021 result assumes **incoherent** (depolarizing, dephasing) noise. **Coherent** noise — small systematic rotations of unitaries due to miscalibration — behaves differently. Coherent noise is a perturbation of the ideal unitary and doesn't directly concentrate the state. It can introduce optimization problems (shifted minima, accumulated phase errors) but doesn't create the noise-induced plateau per se.

In practice, real hardware has both: coherent errors from miscalibration (which you can mitigate via calibration or randomized compiling) and incoherent errors from decoherence (which are the dominant contributor to noise-induced plateaus). Randomized compiling (Wallman-Emerson) can convert coherent noise into incoherent noise, which makes the noise profile *simpler* but can also shift you closer to the noise-induced plateau regime.

The Noise-Plateau vs Noise-Optimization Distinction

There is a subtlety worth flagging: "noise makes optimization hard" is not the same as "noise-induced barren plateaus."

Shot noise (Node 3.7) corrupts the estimate of the cost function but doesn't touch the underlying state. More shots fixes it. It is a statistical problem.

Hardware noise (this node) corrupts the quantum state itself, not just the estimate. More shots do *not* fix it — the mean of the estimator is biased away from the true cost. It is a physical problem.

The confusion between the two has led to a few misunderstandings in the literature. When someone says "noise makes optimization slow," they might mean shot noise (fixable) or hardware noise (not fixable via statistics). These are genuinely different.

Module 9 (Hardware Noise Models) develops this distinction in more detail. For this node, the key point is: **noise-induced plateaus are a physical phenomenon where the state is wrong, not a statistical phenomenon where the estimator is noisy**.

Structural Role

This node is the **noise-origin** node of Module 4. Where 4.1-4.4 discussed plateaus as a consequence of ansatz expressivity, this node introduces plateaus as a consequence of hardware noise. It is a bridge to Module 9 (Hardware Noise Models), which treats noise in full.

Relations to Other Concepts

- **Depolarizing channel** — the canonical incoherent noise model.
- **Wang et al. 2021** — the original noise-induced barren plateau paper.
- **Randomized compiling (Wallman-Emerson)** — the technique for converting coherent to incoherent noise.
- **Error mitigation** — techniques to estimate the noise-free cost from noisy measurements.
- **Module 9 (Hardware Noise)** — the broader treatment of noise channels.

Worked Example — Depth Threshold For Noise-Induced Plateau

Consider a device with two-qubit gate error rate $p = 10^{-3}$, single-qubit error rate 10^{-4}. An HEA layer on 10 qubits has roughly 10 two-qubit gates (nearest-neighbor CNOTs), so the effective per-layer noise rate is $p_{\text{layer}} \approx 10 \cdot 10^{-3} = 10^{-2}$.

The fidelity after L layers is $(1 - p_{\text{layer}})^L \approx e^{-0.01L}$.

At $L = 10$: fidelity ≈ 0.9. Mild noise, trainable. At $L = 50$: fidelity ≈ 0.6. Significant but usable. At $L = 100$: fidelity ≈ 0.37. Borderline. At $L = 200$: fidelity ≈ 0.13. Noise is dominant. At $L = 500$: fidelity ≈ 0.007. State is essentially maximally mixed. Noise-induced plateau engaged.

For scaling to $n = 20$, the per-layer noise rate roughly doubles (20 two-qubit gates per layer), so the thresholds shift to shorter depths. **The noise-induced plateau depth decreases as n grows**, because more gates per layer means more noise per layer. This is opposite to the 2-design plateau, which requires *more* depth to reach at larger n.

So at large n, the operationally-reachable depths may be in a window bounded *above* by noise and *above* (also) by expressivity-driven plateaus. This compresses the trainable depth window further.

Visualization

Pending. A 2D plot with depth L on the x-axis and qubit count n on the y-axis. Shaded regions: "2-design plateau" (upper-right), "noise-induced plateau" (lower-right, curving), "trainable region" (lower-left). As n grows, the trainable region shrinks from both sides. Caption: "Two plateaus, two directions. The trainable window shrinks as scale grows."

Exercises

1. For a device with single-qubit error 10^{-3} and two-qubit error 10^{-2}, compute the depth at which a 10-qubit ansatz crosses into the noise-induced plateau regime (state purity $< 1/2$).
2. Show that randomized compiling preserves the expected cost function but modifies the variance of the estimator. How does this interact with noise-induced plateaus?
3. (VQA) For a fixed device with noise rate p, what is the maximum number of qubits that can be meaningfully trained before noise-induced plateaus block progress? How does this depend on the required circuit depth for the problem?

Research Extensions

- **Error-mitigated barren plateaus** — whether ZNE or CDR can extend the trainable depth.
- **Noise-aware ansatz design** — ansätze that are robust to realistic noise profiles.
- **Bias-corrected gradient estimators** — using noise models to remove the bias that noise introduces.
- **Plateau-noise phase diagrams** — precise characterization of which (n, L, p) regions are trainable.

Cross-links to Other Courses

- **Open Quantum Systems** — master equations, depolarizing channel, dephasing.
- **Error Mitigation** — zero-noise extrapolation, Clifford data regression, probabilistic error cancellation.
- **Module 9 (Hardware Noise Models)** — the full treatment of noise in VQAs.
- **Module 3 (Shot Noise)** — the *other* noise that affects optimization.

Variational Quantum Algorithms · Module 4 — Barren Plateaus · Node 4.5
Concepts: barren-plateaus, quantum-noise, decoherence, density-matrix
Prerequisites: 4.2, 4.4
Enables: 4.6

hbar.university — own.your.cognition