Empirical Behavior Under Shot Noise

Variational Quantum Algorithms · Module 5 — Regularization

Node 5.9

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Regularization **Node:** 5.9

This is the **closing node** of Module 5. It asks the empirical question that nodes 5.5–5.8 set up: *under realistic shot noise, does the dynamical-modification view of regularization actually deliver more stable, more reliable variational optimization than the unregularized baseline?*

The answer is "yes, in specific regimes," and this node makes the regimes precise.

Why This Is the Right Question

Module 5's argument has been theoretical:

- Standard prior-injection (node 5.1) is the wrong frame for VQA because the optimizer never reaches MAP (node 5.5).
- Regularization should be understood as modifying the **trajectory**, not the **target** (5.5).
- This unlocks specific tools: parameter-space techniques (5.6), objective-space techniques (5.7), and the two-stage protocol with cosine annealing (5.8).

A theoretical argument is necessary but not sufficient. The cash value is whether running real VQE instances with the protocol gives better outcomes than running them without. "Better" here means: lower variance across seeds, more reliable convergence, comparable or better final energies, and improved measurement efficiency.

Experimental Setup

The setup mirrors the empirical chapters of the underlying thesis (see node 5.5 for the framing).

- **Problem class:** variational quantum eigensolver instances on fixed Hamiltonians.
- **Ansatz:** hardware-efficient `TwoLocal` circuits, varying depth and qubit count.
- **Noise model:** finite-shot sampling noise on every cost evaluation.
- **Optimizer:** HOPSO (see Module 6) with and without the two-stage protocol.
- **Replicates:** multiple independent random seeds per condition.

The independent variable is the regularization configuration (off / parameter-space only / objective-space only / two-stage protocol). The dependent variables are stability and efficiency metrics defined below.

Metrics

Convergence speed alone is the wrong yardstick under stochastic noise. The relevant metrics are:

1. **Best-energy variance across seeds.** How much does the final reported energy depend on initialization? Low variance means the protocol is reliable; high variance means the optimizer is finding different local minima or failing in seed-dependent ways.
2. **Worst-case run quality.** The 90th-percentile worst seed, not the median seed, is what determines whether you can trust a single production run.
3. **Iteration-vs-energy trajectory variance.** Even within a single seed, how jittery is the descent? A noisy trajectory is a sign of unstable feedback.
4. **Measurement-budget efficiency.** Energy as a function of total shots consumed, not iteration count. In NISQ, shots are the resource; iterations are an internal accounting unit.

Expected Results (Predictive)

Note: results pending the next vqe-lab sweep. The predictions below come from the dynamical analysis of nodes 5.5–5.8 and from preliminary observations on smaller instances.

Prediction 1. Under high shot noise, the two-stage protocol reduces best-energy variance across seeds compared to the unregularized baseline.

Prediction 2. Under low shot noise, the protocol confers little or no benefit and may slightly slow convergence due to the Stage A overhead.

Prediction 3. Parameter-space techniques (5.6) help most when instability comes from large parameter excursions; objective-space techniques (5.7) help most when instability comes from cost-signal volatility. The two-stage protocol blends both.

Prediction 4. The improvement in **worst-case** seeds is larger than the improvement in **median** seeds. This is the practically important effect: regularization tightens the distribution rather than shifting its center.

Prediction 5. Measurement-budget efficiency (metric 4) improves under the protocol whenever stability improves, because fewer wasted iterations chasing noise-induced bad updates means fewer wasted shots.

What a Negative Result Would Mean

If the protocol fails to improve any of the metrics, the dynamical-modification framing is not falsified — but its operational utility is. The framing would still be **correct** as a description of what regularization does in VQA; it would just turn out that under the specific noise levels and ansatz classes tested, the standard unregularized optimizer is already good enough.

In that case the recommendation would shift: skip the two-stage protocol in low-noise regimes, retain it as a tool for high-noise hardware-on-loop runs. This is a more nuanced position than "the protocol works" or "the protocol fails."

Structural Role

This is the empirical capstone of Module 5. It is also the **bridge** to Module 6 (Optimization Dynamics), where HOPSO is treated in full and the interaction between structured optimizer dynamics and regularization is analyzed. Many of the regimes in which regularization helps are regimes where HOPSO already does part of the job — the question of how the two interact is the natural next step.

Relations to Other Concepts

- **A/B testing under stochasticity** — how to compare two optimizers when their outputs are random variables.
- **Reproducibility in VQA literature** — many published VQE results are single-seed; this section argues for distributional reporting.
- **Measurement-optimizer co-design** — Module 10 picks this up.

Visualization

Pending. Diagram set: 1. *Box-and-whisker plot of final energies across seeds, four conditions side by side.* 2. *Energy-vs-shots curves with confidence bands, showing the protocol's narrower band.* 3. *Per-seed trajectories overlaid, illustrating the variance reduction visually.*

Exercises

1. Design a statistical test for whether two optimizers give significantly different best-energy distributions, given 20 seeds per condition. What is the appropriate test and why?
2. Suppose the protocol improves the median final energy by 1% but reduces the 90th-percentile-worst final energy by 10%. Which result is more important for production deployment, and why?
3. (VQA) Why is *iteration count* a misleading x-axis when comparing optimizers in the shot-noisy regime? What should the x-axis be instead?

Research Extensions

- **Joint shot-and-regularization scheduling** — co-adapting the shot budget and the regularization weight per iteration.
- **Hardware-on-loop validation** — running the protocol on real quantum hardware, where noise is correlated and time-varying rather than i.i.d.
- **Generalization to other ansatz families** — the protocol was tuned on `TwoLocal`; how transferable is it to UCCSD, ADAPT-VQE, or QAOA?

Cross-links to Other Courses

- **Statistics Module 8** — hypothesis testing for stochastic optimizers.
- **Optimization Dynamics (Module 6)** — the next chapter in the same story.
- **Hardware Noise Models (Module 9)** — the source of the noise this node fights against.

Variational Quantum Algorithms · Module 5 — Regularization · Node 5.9
Concepts: regularization, stochastic-optimization, variance-and-covariance, expectation-value
Prerequisites: 5.5, 5.6, 5.7, 5.8

hbar.university — own.your.cognition