Quantum Advantage Criteria

Variational Quantum Algorithms · Module 8 — Expressivity Vs Trainability

Node 8.7

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Expressivity vs Trainability **Node:** 8.7

This is the **closing node** of Module 8 — and arguably the most important one for anyone wondering whether all of this VQA machinery actually *matters*.

The question is simple: **when does a VQA actually beat a classical algorithm at the same task?**

The answer, as of 2026, is honest and uncomfortable: **almost never, with current hardware**. There are theoretical regimes where quantum advantage is provable. There are heuristic regimes where it's plausible. There are application-specific regimes where it's claimed. But for the vast majority of VQA experiments published in the past decade, the honest assessment is that a classical algorithm of comparable effort would do at least as well.

This node walks through the criteria for quantum advantage, the regimes where it's been claimed, the regimes where it's been *disproved* (a sobering category), and the open questions about where future advantage might come from. It is the essential reality check for the rest of Module 8 and indeed for the entire VQA literature.

Three Notions Of Quantum Advantage

The phrase "quantum advantage" means different things in different contexts. The three main notions:

1. Asymptotic advantage (provable). A quantum algorithm has an asymptotically better runtime (poly vs super-poly, or poly vs exp) than the best classical algorithm for the same task. Examples: Shor's algorithm, Grover's algorithm. Holds for all sufficiently large inputs.

2. Empirical advantage (heuristic). On a specific benchmark or instance, the quantum algorithm produces better results (lower cost, faster) than a specific classical baseline. Holds only for the specific benchmark; may not generalize.

3. Sample complexity advantage (practical). The quantum algorithm uses fewer samples (less data) to achieve a given accuracy than any classical algorithm. Especially relevant for QML, where data is expensive.

VQAs are usually assessed in the second sense (empirical) or, more recently, the third (sample). The first (asymptotic) is rarely on offer for VQAs, because the variational structure makes asymptotic analysis hard.

What Provable Asymptotic Advantage Requires

For a VQA to give provable asymptotic advantage, several conditions must hold:

1. The classical baseline is well-characterized. You must know the best classical algorithm for the same task, or have a complexity-theoretic lower bound. For most problems VQAs target (chemistry ground states, optimization), the best classical algorithms are heuristics — there's no clean lower bound.

2. The quantum algorithm has a provable runtime. VQAs have variational structure with classical optimization in the loop, and the convergence of the classical optimizer is itself unprovable in general. So the *whole* VQA runtime is unprovable.

3. The quantum part actually does something exponentially hard for classical methods. Many quantum-inspired classical algorithms (DMRG, neural network states, Monte Carlo) can do what shallow VQAs do, sometimes faster.

Net result: no VQA has provable asymptotic advantage as of 2026. Several have *plausible* asymptotic advantage based on conjectured complexity classes (e.g., assuming $P \neq NP$), but no rigorous proof.

Empirical Advantage And The Dequantization Problem

The empirical advantage story is complicated by **dequantization** — the discovery that many quantum algorithms thought to give exponential speedup actually have classical algorithms with similar performance.

Notable dequantizations:

- **Tang's quantum-inspired classical recommendation algorithm (2018):** showed that the quantum recommendation algorithm of Kerenidis and Prakash, which appeared to give exponential speedup, has a classical analogue with polynomial-time runtime.
- **Quantum-inspired classical PCA, regression, etc.** Many "quantum machine learning" speedups have classical analogues.
- **Tensor network simulations** of VQAs that claimed empirical advantage have repeatedly shown that DMRG or MPS methods reach the same accuracy with less compute.

The pattern: every time a VQA claims empirical advantage, the classical ML/numerical community responds with a quantum-inspired algorithm that closes the gap.

This is sobering. It does *not* mean quantum advantage is impossible — it means **the bar is high**. To claim a quantum advantage, you need to show your problem instance is hard for *all* known classical methods (DMRG, NN-states, Monte Carlo, etc.), not just one specific baseline.

When Quantum Advantage Is Plausible For VQAs

Despite the dequantization caution, there are regimes where VQA quantum advantage remains plausible:

1. Strongly correlated electron systems beyond DMRG range. DMRG works well for 1D and quasi-1D systems but degrades in 2D and 3D. For 2D Hubbard model at intermediate doping, no classical method scales well. A VQE on a 2D Hubbard ansatz may be the right tool — *if* the VQA can be made to converge.

2. Quantum chemistry with strong static correlation. Coupled-cluster methods (CCSD, CCSD(T)) fail for systems with strong static correlation (transition metals, bond-breaking). UCCSD via VQE might handle these — but the trainability and noise problems make it hard to demonstrate.

3. Sampling problems related to BosonSampling/IQP. Some sampling tasks are *provably* classically hard (Aaronson-Arkhipov for BosonSampling, similar for IQP circuits). These are not VQAs in the strict sense, but they share architectural features. A VQA whose output distribution is in this hardness class would inherit the advantage.

4. Quantum simulation of dynamics. Time evolution under a Hamiltonian is a natural quantum task. If the VQA prepares states involved in a hard simulation problem, the broader simulation pipeline can give advantage. The VQA part is typically state preparation, not the simulation itself.

5. Tasks with quantum-data inputs. If the data x is itself quantum (e.g., outputs of a quantum sensor), classical methods need to first measure and reconstruct, which is exponentially expensive. A quantum ML model can process quantum data directly, with provable sample complexity advantage (Huang et al. 2022).

6. Specific structured problems with hidden symmetries. Some VQAs exploit problem-specific structure (e.g., translation invariance, particle conservation) in ways classical algorithms don't naturally exploit. For these problems, a custom-designed VQA may beat generic classical methods.

These six are the **realistic** regimes for VQA advantage. None has been definitively demonstrated as of 2026, but all are active research areas.

The Sample Complexity Frontier

A more recent and more promising frontier: **sample complexity advantages**.

Setup. Suppose you want to learn some property of an unknown quantum state or process. A classical algorithm needs many samples (measurements) to reconstruct the state and then compute the property. A quantum algorithm can sometimes compute the property directly, with exponentially fewer samples.

Results (Huang-Kueng-Preskill 2020, Aharonov et al. 2022):

- **Shadow tomography:** $O(\log d)$ samples suffice for many properties of a d-dimensional state, vs $O(d)$ classically.
- **Quantum learning advantage:** for certain learning tasks (e.g., learning quantum process tomography), a quantum learner uses exponentially fewer samples than any classical learner, and this is *provable*.
- **Quantum data analysis:** the same task can require exponentially less data with quantum access.

Connection to VQAs: if a VQA is part of a sample-efficient protocol (e.g., learning a property of a state that the VQA prepared), the sample complexity advantage carries over. This is one of the cleanest theoretical justifications for VQA-based ML in 2026.

Caveat: sample complexity is not the same as runtime complexity. A protocol with low sample complexity may still have high computational complexity. The advantage is real but limited.

What Disproves Quantum Advantage

A useful counter-frame: what kinds of evidence *would* disprove a VQA quantum advantage claim?

1. **A classical algorithm matches the quantum algorithm's performance.** (Dequantization.)
2. **The classical baseline was incorrectly chosen.** Many "advantage" claims compare against weak baselines.
3. **The quantum algorithm doesn't actually scale.** Many VQAs are demonstrated at small n and don't extrapolate.
4. **The quantum algorithm is dominated by classical processing.** Some VQAs do most of their work classically (in the optimizer) and the quantum part is incidental.
5. **The benchmark doesn't reflect a real problem.** Synthetic benchmarks designed to be quantum-friendly don't transfer to real applications.

A *credible* quantum advantage claim must address all five of these objections. Most published VQA-advantage claims address one or two and ignore the rest.

This is the standard for the field in 2026, and it has gotten significantly stricter over the past five years as dequantization has matured.

The Honest 2026 Assessment

Putting it all together, here is the honest 2026 picture for VQA quantum advantage:

1. No proven asymptotic advantage for any VQA on any natural problem.

2. Several plausible empirical advantages in narrow regimes: - Strongly correlated chemistry / condensed matter beyond classical reach. - Sample complexity advantages for quantum data. - Specific sampling problems inherited from quantum supremacy circuits.

3. Many *disproven* advantages in broad regimes: - Most QML claims for classical data have been dequantized or matched by classical baselines. - Most VQA chemistry below 30 qubits has classical analogues that perform as well or better. - Most QAOA claims for combinatorial optimization don't hold up against modern classical heuristics.

4. The hardware bottleneck. Even if a VQA has theoretical advantage, current hardware (NISQ, ~100 qubits, ~1% noise per gate) cannot demonstrate it. The advantage *might* appear at fault-tolerant scales, but that is years away.

5. The most active research frontier: finding provable sample complexity advantages for VQAs on quantum data, where classical methods are exponentially worse by structural necessity.

This is uncomfortable but it is the field's current state. Anyone working on VQAs should know it.

What Would Change The Picture

For VQA quantum advantage to become *demonstrated* (not just plausible), several things would need to happen:

1. Better hardware. Lower noise ($p < 10^{-4}$ per gate), more qubits ($n > 200$ high-fidelity), and longer coherence times. This is the most important factor and the hardest to control.

2. Better ansätze. Ansätze designed using the expressivity-trainability tradeoff (Module 8.5), with explicit symmetry restrictions matched to the problem. This is a research area that this thesis contributes to.

3. Better optimizers. Hybrid optimizers (HOPSO, Module 6.8) and dynamics-side regularization (Module 5) that can navigate the difficult landscapes. Also a thesis contribution area.

4. Honest classical baselines. The community needs to stop comparing VQAs against weak classical baselines (random search, untuned SGD) and start comparing against state-of-the-art (DMRG, neural quantum states, professional optimization solvers).

5. Focus on quantum-data tasks. Where the input is naturally quantum, the comparison is more honest because classical methods lose information at the input boundary.

If all five happen, VQAs may demonstrate genuine advantage in the late 2020s or 2030s. If only some happen, VQA advantage will remain a theoretical possibility for longer.

Closing Module 8

Module 8 has been a long arc:

- **8.1-8.3:** what expressivity means (definitions, metrics, universal approximation).
- **8.4:** what trainability means (gradient variance bounds, hierarchy of theorems).
- **8.5:** the explicit tradeoff curve relating them (Ragone Lie-algebraic theory).
- **8.6:** the data-encoding extension for quantum machine learning.
- **8.7 (this node):** the honest reality check on whether any of this gives quantum advantage.

The narrative thread: **the expressivity-trainability tradeoff is the central design constraint for VQAs**, and within that constraint, only narrow regimes plausibly give quantum advantage. The thesis's contribution is to provide tools (regularization, hybrid optimizers, dynamics-side framing) for navigating the constraint better — to find more of those narrow regimes, and to make them work in practice.

Module 9 (Hardware Noise Models) and Module 10 (Adaptive Control Systems) take this theoretical foundation and ask what it looks like in real hardware deployment.

Structural Role

This node is the **reality check** of Module 8 and of the VQA section of the course as a whole. It states the criteria for quantum advantage, walks through what's been claimed and disproven, and gives an honest 2026 assessment of the field's state.

It is the natural endpoint of the theoretical arc: having built up the machinery for understanding VQAs in detail, it asks the meta-question of whether VQAs are *worth* understanding in detail, and answers "yes, but with caveats."

Relations to Other Concepts

- **Tang (2018)** — quantum-inspired classical recommendation algorithm, the start of dequantization.
- **Aaronson-Arkhipov BosonSampling** — sampling-based quantum supremacy.
- **Huang-Kueng-Preskill (2020)** — classical shadows and learning quantum data.
- **Aharonov et al. (2022)** — provable sample complexity advantages.
- **Module 4 (Barren Plateaus)** — the trainability obstacle that constrains advantage.

Worked Example — Comparing VQE Against DMRG On A Spin Chain

Consider a 20-site Heisenberg spin chain — a benchmark VQA target.

Classical baseline: DMRG. - Finds the ground state to chemical accuracy in seconds on a laptop. - Scales to ~100 sites without difficulty. - Has been the gold standard for 1D spin systems for 30 years.

VQE attempt (2026 hardware, ~100 qubits, 1% noise): - Hardware-Efficient Ansatz, depth ~10. - 1000-shot gradient estimates, Adam optimizer. - Converges (after careful initialization) to ~1% accuracy. - Total runtime: hours of quantum compute time.

Comparison: - DMRG: seconds, much higher accuracy. - VQE: hours, lower accuracy. - DMRG wins decisively.

For 2D Heisenberg (20×20 lattice): - DMRG accuracy degrades sharply with width — limited to ~10 wide. - VQE on 400 qubits: not yet achievable on real hardware. - *In principle*, this is where VQE could win, if hardware were available.

The honest assessment: for benchmarks where we *can* compare, classical methods win. For benchmarks where we *can't* compare yet (very large 2D lattices), nobody knows. Hardware progress will determine the outcome.

Visualization

Pending. A 2D plot with axes "system size" and "accuracy" showing the regimes where DMRG, neural network states, Monte Carlo, and VQE are competitive. DMRG dominates 1D, NN-states dominate moderate *dimensions, Monte Carlo dominates fermionic systems, and a small region in the upper-right is "where VQE might eventually win." Caption: "Quantum advantage is a small slice of problem space, currently empty."*

Exercises

1. List three problems where, in principle, a VQA could beat the best known classical algorithm. For each, identify the classical baseline you're comparing against.
2. The dequantization argument says that many quantum ML speedups have classical analogues. Can amplitude encoding be dequantized? Explain.
3. (VQA) For a 30-qubit VQE on a chemistry molecule (e.g., LiH), estimate the total quantum compute time needed to reach chemical accuracy at current noise levels. Compare to DMRG's runtime on the same problem.

Research Extensions

- **Provable quantum advantage on quantum data** — extending the Huang et al. results to broader classes of problems.
- **VQA + classical preprocessing pipelines** — using VQAs as one component of a larger workflow, where the advantage may be in the integration.
- **Honest benchmarking standards** — methodology for VQA comparisons that resists motivated reasoning.
- **Quantum-classical hybrid advantage** — regimes where the joint quantum + classical algorithm beats either alone.

Cross-links to Other Courses

- **Computational Complexity** — BQP, NP, polynomial hierarchy, complexity-class separations.
- **Classical Numerical Methods** — DMRG, neural quantum states, Monte Carlo (the competitors).
- **Module 9 (Hardware Noise Models)** — the hardware limits on advantage.
- **Module 4 (Barren Plateaus)** — the trainability obstacle that constrains advantage.

Variational Quantum Algorithms · Module 8 — Expressivity Vs Trainability · Node 8.7
Concepts: computational-complexity, parameterized-circuits
Prerequisites: 8.5, 8.6

hbar.university — own.your.cognition