Quantum Advantage and Complexity

Audit advantage claims by matching access, output, accuracy and resource measures.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Translate a search scaling claim between N and $\log_2 N$.
- Separate query guarantees from physical execution claims.
- Identify assumptions in hardness and benchmark statements.

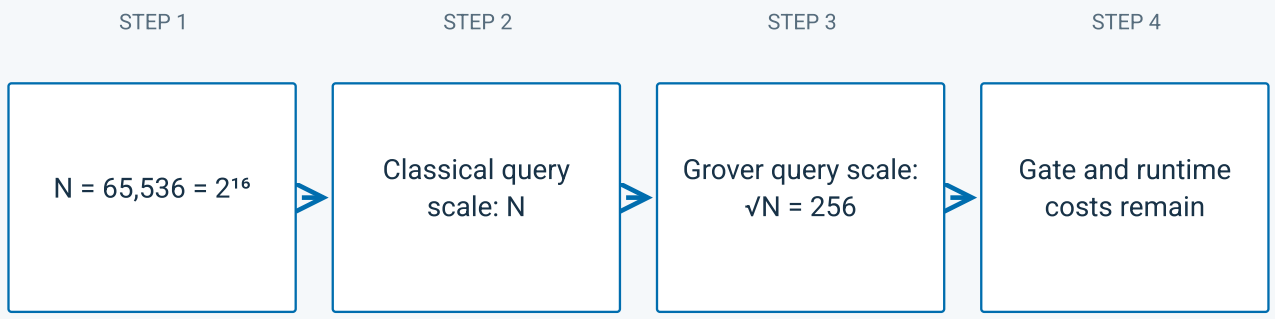
Ask what kind of advantage is being claimed

Quantum advantage is a comparison statement, not a property conferred by using qubits. It needs a task, input model, output requirement, accuracy target and resource measure. A theorem about oracle queries, an empirical sampling experiment and an end-to-end application benchmark can all be valuable, but they establish different things.

Prerequisites: algorithmic complexity, oracles and at least one complete algorithm. This lesson provides a vocabulary for reading performance claims without either dismissing all quantum results or accepting every large-state-space argument.

WORKED DIAGRAM · A35

Count the same resource



The improvement is quadratic in candidate count. Written in the 16-bit label length, both search scales remain exponential.

The comparison setup

Define the input-size variable. A graph may have n vertices and many edges; a search space may contain N candidates described by log-base-two N bits. State the allowed error probability and whether the task is exact, approximate, decision, search, estimation or sampling. Then identify the resource being counted: queries, gates, memory, depth or total elapsed time.

Match those choices across quantum and classical methods. Giving the quantum side coherent input access while making the classical side scan an ordinary file may be a legitimate oracle model, but it is not automatically a fair application comparison. Explain how both access interfaces arise in the intended setting.

Worked example: a square-root search claim

With one marked item among 65,536 candidates, the search labels have 16 bits. A classical black-box search has a query scale proportional to 65,536, while Grover's scale is proportional to 256. In the candidate-count variable N, that is quadratic improvement. In the label-length variable n, the scales are two to n and two to half n, both exponential.

Suppose the coherent quantum predicate costs one thousand elementary operations per query while the classical predicate costs one. The query improvement does not immediately imply a gate or time improvement at this size. Other overheads can further change the result. Conversely, at a different size or with a different predicate implementation, the trade-off can change. A resource model makes that discussion precise.

The algorithm also returns one candidate, not every marked value. If an application requests a large output list, output cost must be included. An advantage for a compact decision or observable cannot be silently promoted to an advantage for printing an exponentially large classical object.

Name the size variable before interpreting an asymptotic improvement.

$$N = 2^n, \qquad Q_{\text{classical}} = \Theta(N) = \Theta(2^n), \qquad Q_{\text{quantum}} = \Theta(\sqrt{N}) = \Theta(2^{n/2}), \qquad n = 16: \quad \sqrt{65\,536} = 256.$$

These are bounded-success unstructured-query scales, with n the label length. The improvement is quadratic in candidate count while both scales are exponential in label length; it is separate from the cost of implementing a query.

Proven separations and empirical evidence

Some query models have rigorous quantum lower and upper bounds. They compare abstract access to the same oracle under specified success criteria. General computational-class separations are often harder to prove. A polynomial-time quantum algorithm for factoring is a major result, but it is not a proof that no polynomial-time classical factoring algorithm exists.

Empirical advantage compares concrete implementations on a stated dataset and hardware. It can be practically important even without an asymptotic separation, but it depends on the tested alternatives and accounting rule. Future algorithms or hardware can change the comparison. Date the result and preserve its scope rather than describing it as an eternal verdict on all classical computation.

Hardness assumptions need their qualifiers

Sampling proposals often rely on complexity-theoretic assumptions, average-case arguments and approximation conditions. A result for exact sampling may not apply directly to a noisy approximate device. A validation statistic may test only part of the desired distribution. Primary research makes these distinctions explicit. Aaronson and Arkhipov, Bouland and colleagues.

These qualifications do not make the results unimportant. They tell you what evidence is needed to connect a mathematical claim with an experiment. The scientific task is to preserve that chain of reasoning rather than replacing it with a slogan about exponentially many amplitudes.

Classical structure changes the baseline

Product states and stabiliser circuits can have efficient classical descriptions despite containing many qubits or some entanglement. Low-depth or geometrically structured circuits may admit tensor-network methods. Problem-specific classical algorithms can outperform unstructured search. Choose a comparison that exploits relevant structure rather than deliberately disabling it on one side.

Likewise, quantum algorithms have useful structure and should not be judged only by a naive statevector simulation. A fair theoretical comparison uses the appropriate algorithmic models. A fair engineering comparison includes actual implementation costs. Mixing those levels produces claims that are difficult to assess or reproduce.

A step-by-step claim audit

Write the exact computational task in one sentence. List access assumptions and output requirements. Identify every parameter in the complexity expression, including precision and conditioning. Separate mathematical guarantees from heuristics or empirical observations. Add preparation, compilation, error handling, measurement and classical processing where relevant. Finally ask whether the best credible classical method solves the same task under the same rules.

For a beginner exercise, audit a sentence claiming that a QFT outputs all Fourier coefficients exponentially faster. The circuit efficiently transforms a quantum state under a specified access model, but measurement does not print every coefficient. The corrected statement must identify the downstream task that uses the transformed state and the cost of preparing its input.

Compare equal error criteria

An exact algorithm and a bounded-error algorithm solve different task descriptions unless the error criterion is explicitly aligned. A quantum procedure succeeding with high probability may need repetition or verification to meet an application's target. A classical randomised method may also gain substantially when a small error is allowed.

For a toy comparison, specify whether failure means a wrong label, an inconclusive outcome or a candidate rejected by a classical verifier. Those events can have different consequences and repetition strategies. Counting every attempt as one successful answer hides the success mechanism, while requiring certainty on only one side creates an unfair comparison.

Fair benchmarking as a scientific method

Benchmark studies should define instance distributions, targets, tuning budgets and uncertainty before selecting favourable results. A primary methodology paper discusses how speedup definitions and comparison choices affect conclusions. Rønnow and colleagues. This is a framework for asking better questions, not a rule that every useful result needs one universal benchmark score.

Exercises and worked answers

Exercise 1. Does an efficient quantum factoring algorithm prove factoring is outside classical polynomial time?

> Check the worked answer
>
> No. It gives a quantum upper bound. Proving a classical lower bound of that strength is a different unresolved task. State the best known algorithms and the mathematical result separately.

Exercise 2. A browser simulator takes longer on a Grover circuit than a classical search loop. Has Grover's query theorem been refuted?

<details>
<summary>Check the worked answer</summary>

No. The browser is executing a classical simulation with its own overhead and exponential state representation. The theorem counts coherent oracle queries in a quantum model. The timing may describe that simulator implementation, but it does not test the theorem's hardware resource model.
</details>

Check your understanding

PAUSE & THINK

Does polynomial-time quantum factoring prove no polynomial-time classical factoring algorithm exists?

- [ ] Yes
- [ ] No; an upper bound is not that lower bound
- [ ] Only for large moduli

Answer: No; an upper bound is not that lower bound. A quantum algorithm establishes a quantum upper bound. Proving a matching classical impossibility is a different mathematical task.

Sources

- Aaronson and Arkhipov
- Bouland and colleagues
- Rønnow and colleagues