Dequantisation, quantum-inspired classical algorithms and honest baselines

Audit quantum advantage claims with matched data access, output formats and stronger classical algorithms.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Construct a classical length-squared sampler.
- Calculate a preprocessing break-even point.
- Distinguish a specific dequantisation from universal classical simulation.

A quantum idea can improve a classical algorithm

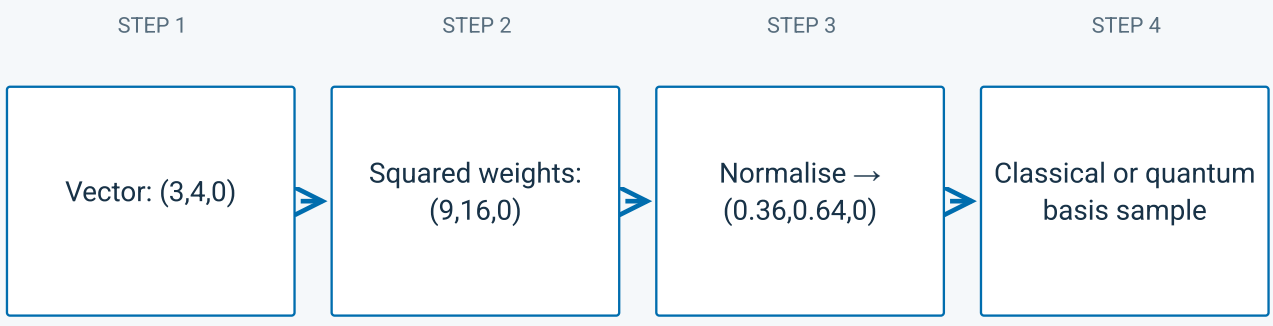
Dequantisation asks whether a quantum algorithm's useful behaviour can be reproduced classically under comparable structural and input-access assumptions. A quantum-inspired classical algorithm borrows ideas such as length-squared sampling or low-rank spectral processing without using a quantum processor. This is a productive outcome: understanding the proposed quantum speedup can reveal a better classical method.

The term does not mean every quantum computation is efficiently classically simulable. A successful dequantisation result applies to a specified task, data structure and approximation regime. It may remove an exponential separation while leaving polynomial differences, different constants or other quantum advantages. Avoid turning one important example into a universal conclusion about the field.

Review measurement probabilities, input-access costs and complexity comparisons. The central question is whether the classical baseline receives access comparable to what the quantum algorithm assumes.

WORKED DIAGRAM · AA30

The same basis-sampling distribution



Both samplers can return indices with probabilities 0.36 and 0.64 under the stated access. This matches one distribution, not arbitrary later interference. The hypothetical total-cost lines cross at 1000 queries and cost 11.

Work a length-squared sample

Take the real vector (3,4,0). Its squared norm is 25. A quantum amplitude encoding has entries three fifths, four fifths and zero, so basis measurement samples indices with probabilities nine twenty-fifths, sixteen twenty-fifths and zero. A classical sampler can produce exactly that index distribution if it has the corresponding cumulative squared weights.

For example, draw a uniform random number between zero and one. Return the first index if the number is below 0.36, otherwise return the second. The third has zero probability. This reproduces the measurement distribution without representing a coherent quantum state. It does not reproduce arbitrary later quantum interference; matching one output distribution is a smaller task.

That distinction explains both the usefulness and the limitation of sampling analogies. Some matrix algorithms only need approximate samples or inner-product estimates that can be assembled from such access. Others rely on coherent transformations with no known equally efficient classical counterpart. State which operations the dequantisation actually reproduces.

A classical sampler can reproduce this particular squared-amplitude distribution.

$$v = (3,4,0), \quad \|v\|^2 = 25, \qquad p_j = \frac{|v_j|^2}{\|v\|^2}, \qquad (p_0,p_1,p_2) = \left(\frac{9}{25}, \frac{16}{25}, 0\right).$$

The distribution agrees with measuring the amplitude-encoded vector. Reproducing this sampling interface does not also reproduce arbitrary subsequent quantum interference; a dequantisation claim must specify which operations are matched.

The recommendation example

Tang's classical recommendation result uses length-squared sampling access to obtain a low-rank recommendation sample with favourable dimension dependence, challenging an earlier claimed exponential quantum separation for that task. Its assumptions and output format are central to the comparison. Read the recommendation dequantisation paper.

Chia and colleagues develop a broader framework for sampling-based low-rank matrix arithmetic and dequantisation under suitable access assumptions. The result concerns structured matrices and a particular sampling model, not arbitrary dense matrix input or a general simulation of every QSVT circuit. Read the low-rank sampling framework.

Match input resources before comparing time

Suppose a quantum algorithm is granted a preprocessed tree that prepares row states efficiently. The same stored norms may also support a strong classical sampler. Comparing the quantum online query time with a classical full scan of raw data would then give the quantum side a more powerful input interface. Charge the preprocessing to both methods or explain why the application supplies one interface naturally.

A naturally quantum input can change the situation. If the data is an unknown physical state supplied as copies, a classical array description may be costly or impossible to obtain efficiently. Dequantisation based on classical sample-and-query access may not apply. Input format is part of the computational problem, not a nuisance to be standardised away without justification.

Work a preprocessing break-even example

Use deliberately hypothetical costs to illustrate accounting. Method C has setup time one unit and costs 0.01 units per query. Method Q has setup time ten units and costs 0.001 units per query. Their total costs are one plus 0.01 times q and ten plus 0.001 times q. They are equal at one thousand queries, when each costs eleven units.

For fewer queries, C is cheaper; for more, Q is cheaper under these invented assumptions. The example makes no hardware claim. It shows why an online speedup can be useful when setup is amortised, yet irrelevant for a one-off task. A benchmark should report both setup and repeated-use regimes.

Solve the total-cost equality before deciding which online rate is useful.

$$T_C(q) = 1 + 0.01q, \qquad T_Q(q) = 10 + 0.001q, \qquad T_C(q) = T_Q(q) \iff q = 1000, \qquad T_C(1000) = T_Q(1000) = 11.$$

q is the number of repeated queries and all costs here are hypothetical units. The setup cost is amortised only when the workload actually reuses the prepared interface.

Low rank is a substantive assumption

If a matrix is well approximated by a small number of singular directions, sampling can reveal useful structure without reading every entry. Complexity can then depend strongly on rank, conditioning and approximation tolerance rather than only dimension. If the effective rank grows with problem size, the favourable scaling may disappear.

The same warning applies to quantum algorithms that assume low rank. A fair comparison should not hold rank constant for one side while allowing it to grow for the other. Use the same data family and error metric, including how much spectral information is discarded. A low-rank approximation good in one norm may not preserve every downstream prediction.

Follow a baseline audit

Write the exact task and output. List the quantum access primitives and their preprocessing. Ask which classical sampling or query operations the same data structure supports. Identify structural promises such as low rank or bounded condition number. Compare algorithms at matched error, confidence and output format, including setup, memory and repeated use.

Then test strong classical implementations, not only a simple textbook baseline. Randomised numerical linear algebra, sketching, tensor networks and problem-specific heuristics can change the relevant comparison. A new classical result does not invalidate the quantum mathematics; it changes the advantage claim and may reveal which part of the quantum proposal supplied the real benefit.

Distinguish evidence levels

A query separation is a theorem under an access model. A dequantisation is a classical algorithm under matched assumptions. A benchmark measures implementations on selected instances. A hardware demonstration validates an experimental process. These can support one another, but none should be silently substituted for the others.

For students, the durable lesson is to ask “compared with what, given which access, producing which output?” Those three questions often reveal more than a dramatic exponential symbol. They also make quantum research constructive: a proposal can teach us a new algorithmic idea even when a stronger classical method later narrows its advantage.

Practice with answers

Exercise 1. Does a classical sampler matching the (3,4,0) state’s basis probabilities reproduce every later quantum circuit on that state?

> Check the worked answer
>
> No. It reproduces one distribution. Later coherent gates can depend on relative phases and amplitude interference that the sampled index alone does not retain. A dequantisation must specify the complete task it simulates.

Exercise 2. In the break-even example, what is each method’s total cost for one hundred queries?

> Check the worked answer
>
> C costs one plus one, or two units. Q costs ten plus 0.1, or 10.1 units. The smaller per-query cost does not compensate for the larger setup at that workload. The numbers are an accounting illustration, not a measured quantum benchmark.

Check your understanding

PAUSE & THINK

Does matching one state's basis probabilities simulate every later quantum circuit?

- [ ] Yes
- [ ] No; relative phase and interference can matter later
- [ ] Only if there are two nonzero entries

Answer: No; relative phase and interference can matter later. A sampled index does not retain the full coherent state. A dequantisation must specify the complete task and access assumptions it reproduces.

Sources

- Read the recommendation dequantisation paper
- Read the low-rank sampling framework