Quantum Amplitude Estimation

Estimate a prepared success probability through coherent phase information and compare its precision costs with sampling.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Write the coherent preparation and good-flag task description.
- Recover p from either amplification eigenphase.
- Distinguish additive error, relative error and access costs.

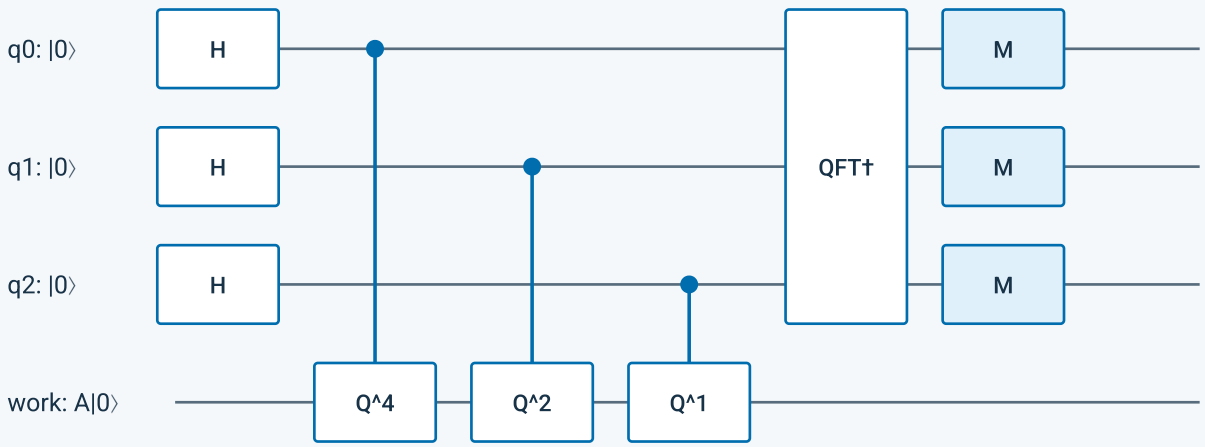
Estimate a success probability coherently

Quantum amplitude estimation asks for the probability that a coherent preparation lands in a chosen good subspace. Ordinary sampling estimates that probability by repeatedly preparing and measuring. Quantum methods can instead learn it through the phase of an amplification operator. Under their access assumptions, they improve the dependence on the requested additive accuracy.

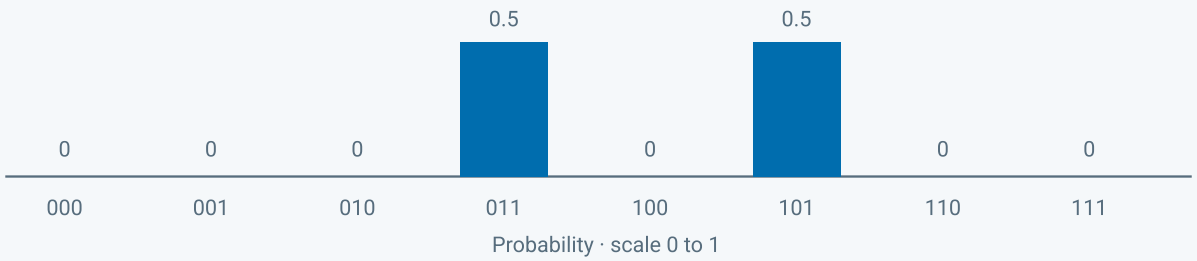
Begin with amplitude amplification and phase estimation. If phase estimation is unfamiliar, first read the worked example here as a geometric story: probability determines an angle, repeated coherent operations reveal the angle, and a classical conversion recovers the probability.

WORKED DIAGRAM · A09

Amplitude estimation: two phases, one probability



Counting-register output



Phase 3/8 → success probability 0.853553

Phase 5/8 → success probability 0.853553

The prepared good probability is (2+√2)/4. Eigenphase labels 011 and 101 each have probability 1/2 and both recover that same probability.

What goes in, what comes out, what we assume

A unitary A prepares a state whose good component has squared weight p. A flag or predicate identifies that component coherently. The output is a classical estimate of p within a stated additive error and with a stated success probability. The algorithm does not return the complete prepared state or the entire distribution over its basis labels.

Canonical amplitude estimation assumes access to A, its inverse, the good-state reflection and controlled powers of the corresponding amplification iterate. Those are stronger resources than an ordinary source of independent samples. An application must explain how the preparation encodes the quantity it wants to estimate. For example, a Monte Carlo payoff may need arithmetic that converts a bounded payoff into a flag amplitude.

Why probability becomes phase

Write p as sine squared theta. The amplification operator has two relevant eigenvalues with phases plus and minus twice theta. Phase estimation can learn either branch. If the measured phase fraction is phi, the conversion p equals sine squared of pi phi gives the same probability for phi and one minus phi.

This relation is useful because eigenphases accumulate coherently under powers of the operator. Applying the iterate many times increases sensitivity to a small angle difference. It also increases circuit depth and demands reliable coherent operations. Precision is not obtained simply by adding output bits while keeping the physical work fixed.

The desired probability determines the rotation angle and hence the eigenphase.

$$A|0\rangle = \sqrt{1-p}|B\rangle + \sqrt{p}|G\rangle, \quad p = \sin^2\theta, \quad Q|v_\pm\rangle = e^{\pm 2i\theta}|v_\pm\rangle.$$

The good and bad states are normalised and orthogonal. Fix $Q = (2A|0\rangle\langle 0|A^\dagger - I)(I - 2\Pi_G)$, using the reflection convention in A07.

Both conjugate phase branches convert to the same probability estimate.

$$\phi \in \left\{ \frac{\theta}{\pi}, 1 - \frac{\theta}{\pi} \right\}, \qquad \widehat{\phi} = \frac{j}{2^t}, \qquad \hat{p} = \sin^2(\pi\widehat{\phi}).$$

t is the counting-register width and j the measured label. This is the canonical phase-estimation form of amplitude estimation.

Worked exactly representable case

Suppose the preparation has good probability sine squared of three pi over eight, equal to approximately 0.853553. Its relevant phase fractions are three eighths and five eighths. Three counting qubits can represent both fractions exactly, so an ideal canonical phase-estimation measurement returns numeric label three or five.

For label three, divide by eight to get three eighths and compute sine squared of three pi over eight. For label five, sine squared of five pi over eight gives the same result. Both branches therefore reconstruct approximately 85.36 percent. The example is exact at the phase grid level, apart from decimal rounding in the written probability.

Now suppose p is one quarter. Theta is pi over six and the phase fractions are one sixth and five sixths. A three-bit grid cannot represent those fractions exactly. The output spreads across several labels, and a single converted estimate is not guaranteed to equal one quarter. Additional precision or a statistically justified repetition strategy is needed.

The three-bit example has an exact radical expression, so its reconstructed value can be checked without rounding the phases.

$$p = \sin^2\frac{3\pi}{8} = \frac{2+\sqrt{2}}{4} \approx 0.853553, \qquad j \in \{3, 5\}, \qquad \sin^2\frac{3\pi}{8} = \sin^2\frac{5\pi}{8}.$$

For the prepared input, the two exactly resolved labels each have probability $1/2$ in this ideal canonical example.

Checkpoint 1

PAUSE & THINK

The three-bit phase register returns label 5. Which estimate follows?

- [ ] 5/8
- [ ] (2+√2)/4
- [ ] 1−(2+√2)/4

Answer: (2+√2)/4. The phase and probability are different quantities. Convert phase 5/8 with sine squared to recover approximately 0.853553.

Compare with ordinary sampling

For independent good/bad samples, the variance of the sample proportion is p times one minus p divided by the number of shots. Achieving a small additive error epsilon at constant confidence therefore generally needs a sample scale proportional to one over epsilon squared. This is the familiar square-root convergence of Monte Carlo sampling.

Canonical amplitude estimation achieves a coherent query scale proportional to one over epsilon for constant success probability, with confidence amplification adding further factors. That quadratic precision improvement is the main theoretical attraction. It is not a statement that every finance, integration or risk application becomes faster: preparation, reversible arithmetic and physical execution must also be paid for.

Relative error is a different requirement. Estimating a very small p to a fixed fraction of its value can be much more demanding than estimating it to a fixed additive tolerance. A result within 0.01 is useful for p near one half but may say little about a probability near one in a million. Choose an accuracy definition that matches the application.

Ordinary sampling and canonical amplitude estimation use different precision resources.

$$\mathrm{Var}(\hat{p}_{\mathrm{sample}}) = \frac{p(1-p)}{S}, \qquad S = O(\epsilon^{-2}), \qquad Q_{\mathrm{coherent}} = O(\epsilon^{-1}).$$

S counts independent Bernoulli samples and ϵ is an additive probability tolerance. The scalings describe fixed-confidence general estimation, with coherent preparation, inverse and reflections available to the quantum method. They do not by themselves compare full application runtimes.

Variants change the implementation

The canonical construction uses a Fourier transform and controlled powers. Iterative, likelihood-based and other variants can reduce counting-register requirements or avoid the full Fourier circuit. They usually still rely on repeated coherent uses of preparation and amplification-related operations. Their confidence guarantees, depth limits and sample schedules differ.

Removing a Fourier transform does not automatically remove the cost of long coherent evolutions. Conversely, a method designed for shallower circuits may trade away some asymptotic precision advantage. Evaluate the actual schedule of powers, shot counts, estimator and confidence interval. An attractive fitted curve without a validated uncertainty analysis is not a complete estimation algorithm.

A small application model

Imagine a preparation that samples a toy scenario and rotates a flag so its one probability equals a bounded score. The joint good probability is then the mean score under that prepared scenario distribution. Amplitude estimation targets that mean, provided the distribution preparation and flag rotation are coherent and reversible.

If the scenarios are supplied only as a large ordinary data file, loading them into the required preparation may dominate. If the score calculation is expensive, each controlled amplification step repeats related arithmetic. The right comparison includes a strong classical estimator using the same information, variance-reduction opportunities and output accuracy. The oracle model is a way to separate assumptions, not an exemption from accounting for them.

Choose a meaningful error budget

Encoding a real quantity into a flag probability can introduce approximation error before estimation begins. If a score lies outside zero to one, rescale and shift it using known bounds, then undo that conversion in the final classical estimate. Uncertainty in those bounds or in the reversible arithmetic contributes to the application error. A more precise phase estimate cannot repair a biased input model. Allocate tolerances among distribution preparation, score approximation, coherent operation and statistical inference, and explain how their combined error meets the requested output accuracy.

Checkpoint 2

PAUSE & THINK

A flag-rotation circuit encodes a biased score. What does increasing QPE precision do?

- ☐ It automatically removes the encoding bias
- ☐ It estimates the encoded probability more precisely
- ☐ It makes the original data exact

Answer: It estimates the encoded probability more precisely. The algorithm targets the probability that the preparation actually encodes. Preparation, arithmetic and inference errors all contribute to the application result.

Chapter summary

- A coherent preparation encodes the target quantity as a good-state probability.
- Canonical amplitude estimation learns that probability through an amplification eigenphase.
- Its quadratic precision improvement assumes a stronger coherent interface than ordinary independent sampling.

A common misconception

An additive error of 0.01 is a one-percent relative error for every probability. Additive error bounds an absolute difference. For a probability of 10^{-6}, error 0.01 is far larger than the quantity being estimated.

Exercises

1. A phase-estimation run returns phi equal to one quarter. What probability estimate follows? Does phi equal to three quarters give a different result?
2. A report replaces one million ordinary samples with one thousand coherent calls but ignores a preparation requiring a million data accesses per call. What is missing from its runtime claim?

Worked answers

1. Sine squared of pi over four is one half. Three quarters gives the same sine squared, so both phases imply a probability estimate of one half.
2. Query count is only one resource layer. The analysis must include state preparation, inverse access, reversible score arithmetic, reflections, controlled implementation, error tolerance and readout. A coherent query advantage can be outweighed by expensive access, and the comparison must use matching accuracy and confidence.

Check your understanding

PAUSE & THINK

What is missing from a claim of O(1/ε) estimation based only on ordinary samples?

- [ ] The stronger coherent preparation, inverse and reflection access assumptions
- [ ] A larger probability than one
- [ ] A requirement that p be exactly one half

Answer: The stronger coherent preparation, inverse and reflection access assumptions. Canonical amplitude estimation uses coherent operations that ordinary independent sampling does not provide automatically. Their implementation cost must be counted.

Sources

Quantum Amplitude Amplification and Estimation is the primary reference for the canonical construction. The exact-grid example above is a teaching calculation that makes the two phase branches explicit.