Amplitude estimation variants and modern iterative approaches

Compare probability-estimation methods that use amplification experiments without a large phase register.

Undergraduate 21 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Calculate amplified probabilities from $p = \sin^2\theta$.
- Explain sensitivity and phase-branch ambiguity.
- Compare total queries with maximum coherent depth.

Estimate a probability without a large phase register

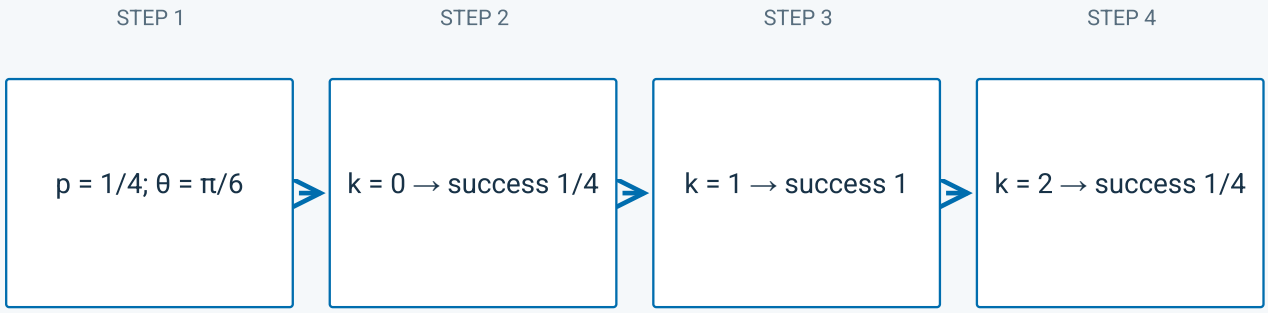
Suppose a circuit prepares a good outcome with unknown probability p. Ordinary sampling estimates p by running the circuit repeatedly and counting successes. Quantum amplitude estimation uses coherent repetitions of a preparation and its inverse to make p influence an interference pattern. Modern variants replace the large phase-estimation register with repeated smaller experiments and classical inference.

"Without phase estimation" does not mean "without coherent depth." Many such methods still apply powers of a Grover-like operator. They reduce some circuit components while retaining the need for reliable repeated coherent access. The right comparison counts total preparation-oracle uses, maximum circuit depth, shots, confidence and classical processing.

Review probability amplitudes, sampling uncertainty and Grover interference. The key mathematical substitution is p equal to sine squared of an angle theta between zero and pi over two.

WORKED DIAGRAM · AA03

One probability, several amplification depths



The measured success is sin²((2k+1)θ). At θ = π/6, depths 0, 1 and 2 give probabilities 1/4, 1 and 1/4. Inference must recover the original probability, not confuse it with an amplified one.

Specify the preparation setup

Assume a unitary A prepares a state whose good component has squared norm p, and a reversible predicate identifies the good subspace. The algorithm also needs A's inverse or equivalent reflection access. Its output is a classical estimate of p, or a confidence interval with chosen additive error and failure probability. Preparing the original distribution and computing the good predicate are part of the cost.

After k amplification iterations, a good measurement has probability

$$q_k = \sin^2((2k+1)\theta), \qquad p = \sin^2\theta.$$

Different values of k probe different oscillations of the same unknown angle. Large k makes a small angle change more visible, but also creates ambiguity because sine squared repeats. An estimation algorithm must manage that ambiguity instead of fitting any convenient oscillation branch.

Work an exact small example

Take p equal to one quarter, so theta is pi over six. With no amplification, success probability is one quarter. With k equal to one, the angle is pi over two and success is certain. With k equal to two, the angle is five pi over six and success returns to one quarter. Thus a high success frequency after amplification does not itself mean the original p was near one.

Now compare a nearby p, such as 0.26. The amplified probabilities move along the same oscillatory formula. A collection of shot counts at several k values can distinguish nearby hypotheses more effectively than only k equal to zero, provided the experiment is sufficiently accurate and the inference identifies the correct angle interval. Finite counts and physical noise limit how much information a long coherent circuit actually supplies.

This example also explains why one measurement at a large k is not enough. Several different theta values can give the same q_k. The earlier low-depth observations or maintained interval rule are what disambiguate them. Precision comes from a designed sequence of experiments, not from reading many decimal places from one detector click.

Repeated amplification changes the observation probability in an oscillatory way.

$$p = \sin^2\theta, \qquad q_k(p) = \sin^2((2k+1)\theta), \qquad p = \frac{1}{4}: \quad (q_0, q_1, q_2) = \left(\frac{1}{4}, 1, \frac{1}{4}\right).$$

p is the original success probability, k the amplification depth and q_k the probability measured at that depth. Several depths or an interval rule are needed to disambiguate the original angle from repeated oscillations.

Compare common approaches

Canonical amplitude estimation uses phase estimation on the Grover operator. It introduces counting qubits and controlled powers, then converts an estimated eigenphase into p. Iterative methods maintain an angle interval, choose a safe amplification depth and update the interval from measured frequencies. Maximum-likelihood variants collect data across chosen depths and fit a likelihood model. These are related approaches with different circuit and inference requirements.

Grinko and colleagues' iterative method provides a rigorous amplitude-estimation procedure without a quantum phase-estimation register, using Grover experiments and classical interval updates. Its analysis gives an asymptotic query improvement over ordinary Monte Carlo under the stated ideal access model. That theorem is not an automatic hardware speedup for every probability-estimation task. Read Iterative Quantum Amplitude Estimation.

Walk through a disciplined experiment

Begin with low-depth measurements to locate the plausible theta interval. Choose a k that increases sensitivity without wrapping that interval across unresolved sine-squared branches. Run the selected circuit enough times to estimate q_k with a stated statistical bound. Map the permitted q_k range back into the selected angle branch, update theta and continue until the desired p precision is reached. Record every depth and shot count.

A likelihood implementation follows a different inference rule but still needs care with multiple maxima and uncertainty calibration. An apparently narrow fit can be misleading if the model ignores noise or excludes the true branch. Confidence statements must follow the method's assumptions, not merely the width of a plotted numerical optimiser output.

Costs and limits

Ordinary independent sampling typically needs order inverse-error-squared trials for additive error. Coherent amplitude-estimation methods can improve the query dependence toward inverse error, with confidence and method-dependent factors. That comparison assumes useful coherent access to A, its inverse and the predicate. It does not include arbitrary data-loading machinery for free.

Long amplification circuits can accumulate noise, flatten oscillations and bias inference. A shallower variant may sacrifice asymptotic benefit to fit a device's coherence budget. Compare total query count and maximum depth separately; reducing one does not always reduce the other. Also include the cost of reversible arithmetic used to turn an application's quantity into a success amplitude.

Practice with answers

Exercise 1. At p equal to one quarter, what does a success probability of one after one amplification iteration tell you about p before amplification?

Check the worked answer

Under the known branch and formula, it is consistent with p equal to one quarter. It does not mean p was one. Amplification changes the measured probability, and inversion requires the angle model and branch information.

Exercise 2. A method removes all phase-register qubits but uses deeper powers of A. Is it necessarily cheaper?

Check the worked answer

No. Compare ancilla count, controlled operations, total uses of A and its inverse, maximum depth, repetitions, noise sensitivity and inference cost. “Fewer qubits” is one resource improvement, not a complete end-to-end cost statement.

Check your understanding

PAUSE & THINK

If amplification makes success certain at k = 1 for this example, what was the original p?

- ☐ 1/4
- ☐ 1/2
- ☐ 1

Answer: 1/4. At $\theta = \pi/6$, the amplified angle is $\pi/2$. Amplification changes the observed probability; it does not imply the original p was one.

Sources

- Read Iterative Quantum Amplitude Estimation