Quantum approximate counting and mean estimation

Estimate marked counts and bounded means by encoding them as a success probability.

Undergraduate 21 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Translate fraction error into count error.
- Construct the flag encoding for a bounded mean.
- Distinguish additive accuracy from rare-event relative accuracy.

Estimate how many, rather than find one

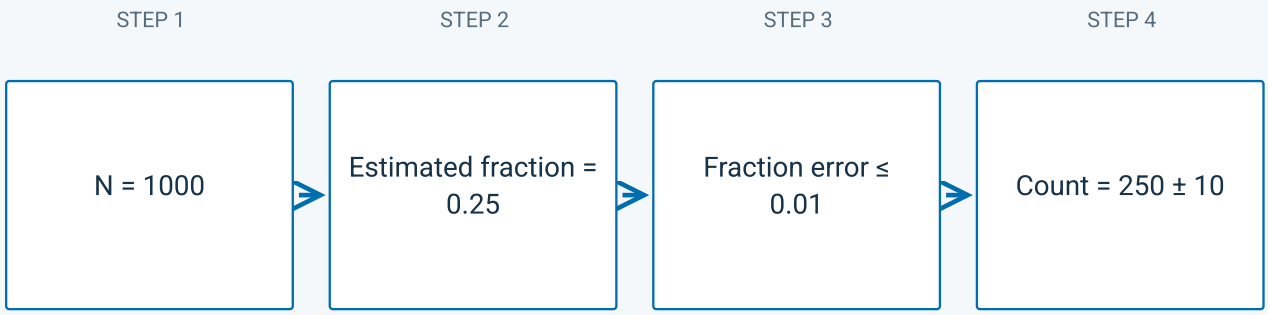
Search returns a marked item. Counting asks how many items satisfy the predicate. Exact counting may require high precision, while approximate counting can aim for an additive or relative error. Quantum counting connects the marked fraction to the rotation angle of a Grover operator, then estimates that angle. Mean estimation extends the idea from yes-or-no values to bounded numerical values.

Begin with the output you need. Estimating that roughly a quarter of a million candidates are acceptable is different from identifying every acceptable candidate. A small classical estimate can be easier to extract than a long list. Conversely, resolving whether the count is zero or one may require much finer relative precision than distinguishing a quarter from a half.

Review Grover interference, sampling uncertainty and oracles. The core relation is that the marked fraction is the squared sine of the initial Grover angle.

WORKED DIAGRAM · AA11

A fraction estimate becomes a count



Multiplying the fraction and its additive error by N gives the count estimate. A bounded-value flag must have squared amplitude v(x), so its overall success probability equals the mean of v.

State a counting setup

Let a predicate mark M of N labels and assume coherent access to it. The desired output is an estimate of M, with a stated error and failure probability. Additive error in the fraction translates to N times that error in the count. Relative error is more demanding when M is small. State whether N is known and whether the number of marked items is promised to be positive.

The algorithm needs the preparation and reflection operations used by amplitude estimation. A uniform preparation makes the good probability equal to M divided by N. A nonuniform preparation estimates a weighted probability instead. Counting and weighted mass estimation are different tasks unless the distribution is explicitly part of the task description.

Work a four-item count

Suppose one of four labels is marked. The initial good probability is one quarter and the Grover angle is pi over six. One amplification iteration makes success certain. If two labels are marked, the angle is pi over four; every odd multiple has squared sine one half, so success remains one half at each integer iteration. The oscillation pattern distinguishes these cases under the ideal model.

This does not mean one observed success tells you the count. Measurement is random in many settings, and different angles can produce the same amplified probability at one depth. A counting algorithm combines suitable coherent phase information or several amplification experiments with an inference rule. It must report uncertainty rather than selecting a count that happens to fit a single observation.

If an estimated fraction is 0.25 with additive error 0.01 and N is one thousand, the implied count estimate is 250 with additive error ten. To distinguish count 250 from 251, much finer fraction precision is needed. Precision should match the actual question, not the number of decimal places the interface can print.

The ideal amplification curves differ for one and two marked labels.

$$p = \frac{M}{N} = \sin^2\theta, \qquad q_k = \sin^2((2k+1)\theta), \qquad N = 4: \quad \begin{cases} M = 1: & q_0 = 1/4,\ q_1 = 1, \\ M = 2: & q_k = 1/2 \text{ for all integer } k \geq 0. \end{cases}$$

q_k is the measured success probability after k amplification steps. A single success is random evidence; a counting method needs several observations or coherent phase information and an uncertainty rule.

Extend counting to means

For values v(x) between zero and one, coherently compute v(x) and rotate a flag qubit so its probability of one is v(x). If x is prepared with probability distribution p(x), the flag's overall success probability is the weighted mean of v. Amplitude estimation can then target that mean. The rotation uses square roots of v and its complement, not v itself as an amplitude.

For example, a uniform distribution over values zero, one quarter, three quarters and one has mean one half. The good-flag probabilities average to one half, even though the individual branches differ. After estimating the flag probability, unscale or shift the result if the original quantity was mapped from another bounded interval.

Brassard and colleagues' amplitude-estimation framework includes approximate counting as an application. Its query guarantees assume coherent preparation and predicate access and distinguish relevant accuracy regimes. Read the amplitude-estimation and counting paper. The examples here explain the encoding mechanism rather than reproduce every optimal bound.

Follow the full procedure

Define the domain, distribution and numerical range. Construct a reversible predicate or bounded-value computation. Prepare the input distribution and flag rotation, then uncompute temporary arithmetic as required. Apply a chosen amplitude-estimation method, convert the probability estimate into a count or mean and propagate the statistical and arithmetic error bounds.

If the values are unbounded, a bounded rotation cannot be applied without additional handling. Truncation, tail bounds or bounded-variance methods introduce their own assumptions and errors. A rare but very large outcome can dominate a mean while being difficult to estimate. Do not transfer a theorem for bounded values to arbitrary heavy-tailed data without analysis.

Compare resources honestly

Ordinary sampling estimates a bounded mean with typical inverse-error-squared sample scaling. Coherent amplitude-estimation methods can improve the query dependence toward inverse error under suitable access. The cost of producing each coherent sample, evaluating v and applying inverse operations remains. A cheap classical sampler and a very expensive reversible sampler can reverse the practical comparison.

Relative error on a rare event is particularly sensitive to its probability. A fixed additive-error guarantee may be useless when the true mean is much smaller than the error tolerance. Specify the confidence target and whether the application can tolerate an interval including zero. This is especially important in reliability, finance and security examples involving rare failures.

Practice with answers

Exercise 1. An estimate of the marked fraction is 0.12 with error at most 0.005 for N equal to two thousand. What count estimate follows?

Check the worked answer

Multiply by N: the estimate is 240, with additive error at most ten. This is not an exact count and does not guarantee the same relative error for a much rarer marked set.

Exercise 2. Why must a value v be encoded using a flag amplitude whose squared magnitude is v?

Check the worked answer

Measurement probabilities are squared magnitudes. Using v itself as the amplitude would encode v squared and estimate a different mean. The conditional rotation must preserve normalisation by including the complementary flag amplitude too.

Check your understanding

PAUSE & THINK

For N=2000 and fraction estimate 0.12 ± 0.005, what count interval scale follows?

- ☐ 240 ± 10
- ☐ 240 ± 0.005
- ☐ 120 ± 5

Answer: 240 ± 10. Both the estimate and absolute error are multiplied by 2000, giving 240 and 10 respectively.

Sources

- Read the amplitude-estimation and counting paper