Quantum Sampling Algorithms

Define sampling accuracy and choose tests sensitive to joint distributions.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Distinguish sampling from full probability evaluation.
- Calculate total variation distance in a four-outcome example.
- Explain the limited scope of verification statistics.

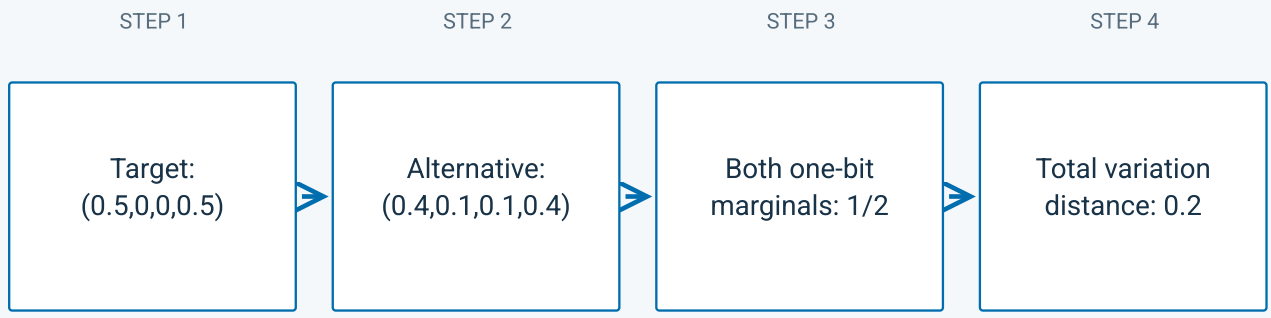
Sometimes the requested output is a random sample

A sampling algorithm produces outcomes drawn from a specified distribution. The task is not necessarily to calculate one exact number or find an optimum. Quantum experiments naturally produce samples after state preparation and measurement, so sampling tasks are central to both useful algorithms and investigations of classical simulation difficulty.

Prerequisites: measurement, entanglement and simulation evidence. Keep the target distribution, approximation metric and access model visible throughout the discussion. A sample from an unknown device is not automatically evidence that its intended distribution was realised.

WORKED DIAGRAM · A32

Same marginals, different samples



The target always gives agreeing bits; the alternative agrees with probability 0.8. Identical local histograms do not establish identical joint distributions.

The sampling setup

Specify a family of circuits or physical preparations, the output labels, the number of requested samples and an allowed distributional error. State whether preprocessing is allowed and whether the input circuit changes between tasks. Sampling one bitstring is different from outputting a table of every probability, which may have exponential size.

An exact sampler reproduces the distribution precisely in the mathematical model. An approximate sampler is judged by a metric such as total variation distance, not merely by whether a few prominent outcomes look similar. Physical devices also introduce noise and drift, requiring a connection between the mathematical target and observed records.

Worked example: joint distributions matter

The ideal Bell preparation has probabilities one half on 00 and 11, and zero on 01 and 10. Consider a different distribution with probabilities 0.4, 0.1, 0.1 and 0.4 in that order. Each individual bit is still equally likely to be zero or one in both distributions. Marginal histograms alone cannot distinguish them.

The total variation distance is half the sum of the absolute differences, giving 0.2. The probability of agreeing bits is one for the Bell target but 0.8 for the alternative. A simple parity statistic therefore detects a joint difference that local frequencies hide. This is a small example of choosing a validation statistic sensitive to the intended property.

Neither distribution is classically difficult to sample. For the ideal target, toss one classical fair coin and copy its result into both output positions. Quantum entanglement in the preparation does not imply that its particular Z-output distribution is hard to reproduce classically. The computational task concerns the specified outputs and access model, not just the state’s name.

A distributional metric detects a joint difference that single-bit marginal histograms miss.

$$P = (0.5, 0, 0, 0.5), \quad Q = (0.4, 0.1, 0.1, 0.4), \qquad d_{\mathrm{TV}}(P, Q) = \frac{1}{2} \sum_{x} |P(x) - Q(x)| = 0.2.$$

Entries use order $00, 01, 10, 11$. Both distributions have uniform individual-bit marginals, but their agreement probabilities are one and 0.8 respectively. Total variation compares the full requested output distribution.

Sampling and probability evaluation are different tasks

A classical routine can sample some distributions without calculating every probability explicitly. Conversely, evaluating an isolated probability does not necessarily provide an efficient sampler for a huge outcome space. Hardness arguments must connect the relevant tasks carefully. “There are exponentially many probabilities” is not by itself a proof that sampling is hard.

Specialised classical simulation methods exploit structure such as stabilisers, low entanglement, tensor-network contraction patterns or limited circuit depth. A circuit family must be specified before assessing difficulty. Changing noise levels or approximation requirements can also change the strength of a classical method.

Two important research families

Boson sampling sends indistinguishable photons through a linear-optical network and samples detection patterns. Its complexity arguments connect optical amplitudes with matrix permanents and make explicit computational assumptions, particularly for approximate sampling. Aaronson and Arkhipov.

Random circuit sampling studies bitstrings produced by selected random quantum circuits. Primary complexity research examines hardness and verification under stated assumptions. Bouland and colleagues. These tasks are

useful for studying computational capability, but their existence does not automatically solve a practical optimisation or machine-learning problem.

Avoid treating a hardness conjecture as an unconditional theorem about every noisy implementation. Exact, approximate, average-case and experimentally sampled versions can require different arguments. A beginner can still follow the central question: can a classical method reproduce the requested distribution to the agreed accuracy using comparable resources?

How to validate a small sampler

For a small register, calculate the exact statevector and probabilities independently. Collect a fixed sample budget, compare frequencies with uncertainty and examine relevant joint statistics. Use fresh samples for any test selected after exploration. Check gate conventions and bit ordering before interpreting an apparent distribution mismatch.

For large systems, exhaustive probability comparison is unavailable. Special verification statistics may test limited aspects of the target. Report exactly what they certify and what alternatives they might fail to exclude. A high score under one statistic need not imply a small global distribution distance without an additional theorem or model assumption.

A second distinguishability check

For the four-outcome alternative, assign a signed parity of positive one to agreeing bits and negative one to disagreeing bits. Its expected parity is 0.8 minus 0.2, equal to 0.6. The ideal Bell-output target has parity expectation one. This statistic is directly estimable from bitstrings without reconstructing amplitudes.

A matching parity would still not identify every probability uniquely in a larger output space. Design tests around the claim and list remaining alternatives. Small counterexamples are a useful teaching tool: construct two distributions that agree on one statistic but differ on another before trusting that statistic as complete verification.

Resources and comparisons

Count circuit preparation, execution, readout and the number of samples needed for the application or verification statistic. Classical preprocessing and reuse of a fixed circuit can matter. Compare with the best relevant classical sampler under the same approximation metric, hardware budget and input family.

Sampling difficulty is one ingredient in a broader advantage argument. Application usefulness is another. A distribution may be challenging to reproduce yet have no established practical use beyond benchmarking. Conversely, a useful physical simulation may request only a few observables rather than a hard sampling distribution. These are distinct success criteria.

Exercises and worked answers

Exercise 1. Uniformly sample all four two-bit strings. What are the one-bit marginals and the probability of agreeing bits?

> Check the worked answer
>
> Each bit is fair, while agreement has probability one half. The marginals match the Bell target but the correlation does not. A local-frequency test would miss the difference.

Exercise 2. A device passes a test of one selected output probability. Has it been shown to sample the entire target distribution accurately?

Check the worked answer

Not without further assumptions or a theorem relating that test to the full metric. Many different distributions can share one probability or one summary statistic. State the limited property tested rather than promoting it to complete validation.

For an extension, design a small set of parity tests on three-bit samples. Work out which distinct distributions they fail to distinguish. This exercise builds the habit of asking what an observed statistic leaves uncertain before using it as evidence of a computational claim.

Check your understanding

PAUSE & THINK

What is the total variation distance between these two distributions?

- [ ] 0
- [ ] 0.2
- [ ] 0.4

Answer: 0.2. Half the sum of absolute differences is 1/2 × (0.1+0.1+0.1+0.1) = 0.2.

Sources

- Aaronson and Arkhipov
- Bouland and colleagues