Classical shadows

Reuse a record of random measurements to estimate many selected properties of a quantum state.

Undergraduate 18 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Construct a single-qubit classical snapshot.
- Estimate observables from recorded settings and outcomes.
- Explain the roles of precision, confidence and measurement choice.

Before you begin: Pauli measurements, expectation values, density matrices and sampling uncertainty.

Learn properties from reusable records

A quantum device repeatedly prepares the same state. We may want several properties: a spin expectation, correlations between qubits, or terms in a Hamiltonian. Reconstructing every density-matrix entry can be much more work than the task requires. Classical shadows use random measurement choices and classical post-processing to estimate selected properties from one reusable dataset.

The word "shadow" describes a classical record supporting predictions. The record contains measurement settings and outcomes. Fresh preparations are still required for new measurements; the procedure does not preserve an extra physical copy of the unknown state.

We will study the local Pauli version. On each qubit, choose X, Y or Z independently and uniformly, then measure in that basis. Store both the chosen basis and the resulting bit. Use eigenvalue $s = +1$ for bit 0 and $s = -1$ for bit 1.

Start with one qubit

If the selected basis is $P \in \{X, Y, Z\}$ and the measured eigenvalue is s, form the estimator

$$\hat{\rho} = \frac{I + 3sP}{2}.$$

It need not be a physical density matrix. For example, a Z measurement with outcome zero gives $\hat{\rho} = \mathrm{diag}(2, -1)$. Negative eigenvalues are allowed in this intermediate statistical estimator. Averaging over random settings and outcomes gives

$$\mathbb{E}[\hat{\rho}] = \rho.$$

To estimate $\langle Z \rangle$, take $\mathrm{Tr}(Z\hat{\rho})$. This equals $3s$ if the selected basis was Z, and zero if it was X or Y. The factor three compensates for using the matching basis only one third of the time.

Checkpoint. Does a snapshot value of 3 mean the physical spin exceeded its allowed range? No. Individual estimators can lie outside the observable's physical range; their mean estimates the expectation value.

Worked example: Work a six-shot record

Suppose the device prepares $|0\rangle$. A possible record is:

Shot	Basis	Outcome bit	Z estimator	X estimator
1	X	0	0	3
2	X	1	0	-3
3	Y	0	0	0
4	Y	1	0	0
5	Z	0	3	0
6	Z	0	3	0

The means are $\widehat{\langle Z\rangle} = 1$ and $\widehat{\langle X\rangle} = 0$, matching the ideal values for this illustrative record. Y outcomes contribute to the Y estimate in the same way. Six shots are enough to demonstrate the arithmetic; reliable precision needs a substantially larger dataset and uncertainty estimates.

For one qubit, the mean Z estimator has second moment 3 under uniform basis selection. Since its mean is $\langle Z\rangle$, its variance is $3 - \langle Z\rangle^2$. For $|0\rangle$, this is 2, giving standard error $\sqrt{2/N}$ for N independent snapshots. Directly measuring Z every time is more efficient when Z is the only target. Shadows become attractive when a shared record supports many targets.

Extend to correlations

For multiple qubits, tensor together the single-qubit estimators. To estimate a weight-k Pauli string, use the score 3^k times the product of measured signs if every nonidentity factor's basis matches; otherwise use zero. For $Z_1 Z_2$, a matching ZZ shot contributes $9 s_1 s_2$. Such matches occur with probability $1/9$ under uniform independent choices.

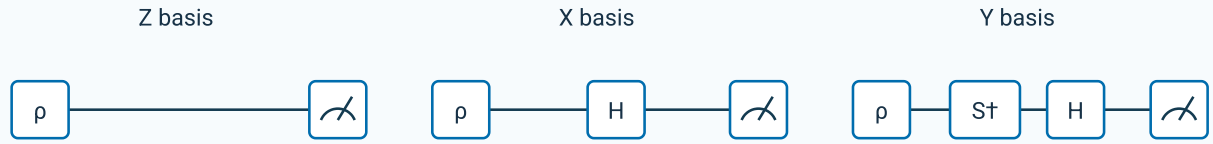
This gives an important cost trade-off. For M weight-at-most-k Pauli observables, suitable aggregation can achieve additive error ϵ and failure probability δ using a sample scale

$$N = O\left(\frac{3^k}{\epsilon^2} \log \frac{M}{\delta}\right).$$

The favourable logarithmic dependence on the number of targets comes with dependence on locality, precision and confidence. It should not be read as efficient prediction of every possible global property. Other shadow ensembles have other guarantees.

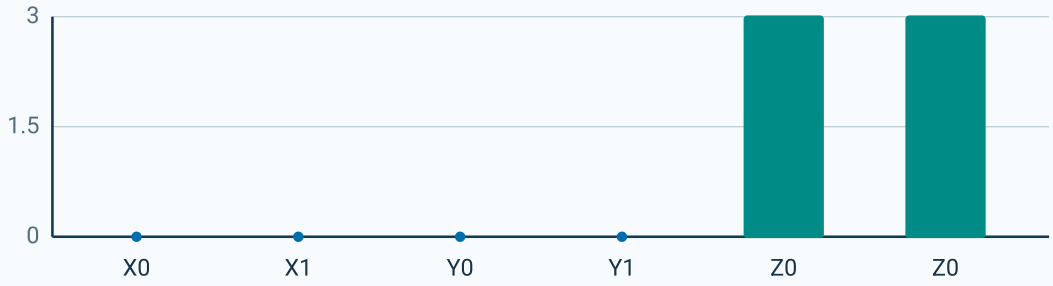
Choose a measurement basis

Classical shadows: save the basis as well as the outcome



Select X, Y or Z uniformly for each fresh preparation.

Stored shot = (basis, outcome). Z score = 3s in Z; zero in X or Y.



Z estimate = (0 + 0 + 0 + 0 + 3 + 3)/6 = 1

Actual X, Y and Z measurement circuits, followed by the six-shot example's Z estimator values.

To measure Z, measure directly. To measure X, apply H first. To measure Y, apply $S^\dagger$ and then H. All three finish with the device's Z measurement. A classical random selector chooses the basis for each shot, independently on each qubit. The example plot shows the six individual Z estimator scores; their average is the final estimate.

Exercises

Generate fresh records for $|0\rangle$, $|+\rangle$ and $(|0\rangle + i|1\rangle)/\sqrt{2}$. Choose target observables after collecting the record. Keep the dataset fixed while changing the target; this demonstrates reuse. Repeat the whole collection with a new random seed to see sampling variation. Add a known readout error and distinguish its bias from finite-shot fluctuations.

1. A Z-basis shot returns bit 1. What is its Z score?
2. An XY shot is used to estimate ZZ. What is its score?
3. Why save the settings as well as the bits?

Worked answers

1. -3.
2. zero; the required bases did not both match.
3. a bit's interpretation depends on the measured basis.

Check your understanding

PAUSE & THINK

What is the reusable object?

- [ ] A classical record of measurement settings and outcomes
- [ ] Unlimited physical copies of the state
- [ ] Exact values of all observables after a few shots

Answer: A classical record of measurement settings and outcomes. Different observables can be estimated by reprocessing this record within the method's guarantees.

Classical shadows trade full reconstruction for targeted prediction. Their value is clearest when many accessible observables share one carefully recorded experiment.

Sources

Huang, Kueng and Preskill, original classical-shadows paper and local-Pauli analysis. The six-shot table, single-qubit variance calculation, activities and figure above are original teaching examples.