

Can machine learning improve quantum-state reconstruction?

Our transformer reconstructs a full two-qubit state from measurements that are noisy, incomplete or both.

THE TASK

Qubits store quantum information, but their state is not directly observable.

Physicists probe it through repeated measurements. Each type checks a chosen property, like looking from another angle. Tomography turns the answers into a density matrix that encodes the quantum state.

FROM IMPERFECT MEASUREMENTS TO A DENSITY MATRIX

1 QUANTUM MEASUREMENTS

For n qubits, full tomography uses $4n$ measurement types. For two qubits, that makes 16, shown as an image-like matrix.



A measurement gives one random result. One shot means preparing and measuring the qubits once. Repeating the same measurement type improves the estimate but costs more lab time.

2 ML RECONSTRUCTION

Noise and missing measurements can make several states fit the data. Our transformer-based model learns from simulated examples to reconstruct one complete density matrix.

ONE TRANSFORMER

The same trained network works with complete measurements, noisy results, missing measurements or both. Its output is constrained to be a possible quantum state.

THE IMAGE-RESTORATION IDEA

Like removing noise or filling a hole in an image, the model learns the structure of complete examples and uses it to repair imperfect measurement data.


3 FULL DENSITY MATRIX

25-63%

LOWER RECONSTRUCTION ERROR

Across four headline tests vs the best tested classical reference

CORRECTED BENCHMARK



We uncovered a metric error in a leading published benchmark. After correction, our model reaches the lowest errors over most of the reproduced learning-method sweep, including the previously best LSTM comparison.

1-16 measurements on the x-axis. Our model is magenta and teal. Lower is better.
Note: plot uses exact simulated probabilities consistent with the previously best LSTM benchmark

WHY THIS MATTERS

Fewer measurements

Using 12 instead of all 16 measurement types saves experiment time. In that test, error is 63% lower than the tested classical reference.

Improved statistics

Real experiments do not reveal clean probabilities. On synthetic tests, where the answer is known, the model is 37% closer than raw frequencies at 100 shots per measurement type.

One model

One trained network handles complete, noisy and missing measurements, including both at once.

More qubits

Current continuation work scales the transformer to larger synthetic systems and explores a reusable tomography foundation model.