

Quantum tomography

The authors are grateful to Richard Kueng for reviewing this chapter.

Rough overview (in words)

In quantum tomography, often also termed quantum state estimation, we are given repeated copies of an unknown quantum state (or quantum channel) and the goal is to obtain a full classical description of the quantum state (or quantum channel) by extracting information by means of performing measurements. Here, we focus on quantum state tomography, with multiple independent and identical copies of an unknown quantum state ρ provided—that is of fixed and of known dimension—and the task is to find an estimate of the density matrix of the quantum state up to an approximation error in some distance measure (and up to some failure probability). We are then typically interested in the optimal sample complexity in terms of the number of copies n , the quantum state dimension d , the approximation error ε , and the overall failure probability δ . Additionally, algorithmic complexity aspects of the used schemes might be of importance as well.

Rough overview (in math)

Given (many copies of) an unknown quantum state ρ of known dimension d , the goal is to give a description of $\tilde{\rho}$ with the statistical estimate $\tilde{\rho} \approx_{\varepsilon} \rho$, up to some distance measure with corresponding approximation parameter $\varepsilon \geq 0$. This is achieved by extracting classical information by applying measurements $\mathcal{M}^n(\cdot)$ via $\rho^{\otimes n}$. To start with, one has to distinguish tomography schemes based on different types of measurements used. This includes in particular:

- (i) Independent and identical (IID) measurements, where the choice of measurement $\mathcal{M}^n = \mathcal{M}^{\otimes n}$ is fixed and the same for each copy.

- (ii) Adaptive measurements, where the choice of measurement $\mathcal{M}_2$ on the second copy can depend on the outcomes of measurement $\mathcal{M}_1$ on the first copy, and so on.
- (iii) Entangled measurements, where one measurement $\mathcal{M}^k$ with $1 < k \leq n$ is performed on k copies at once.

Further, if one has some information about the type of quantum state provided, then tomography schemes can become more efficient. This includes, for example, pure state tomography, low-rank- k state tomography, matrix product state tomography, or ground/thermal state tomography of Hamiltonians (some references on tight schemes are given later on). For some schemes, one a priori has certain information about the state in question and under this assumption the scheme is then promised to work (e.g., low-rank tomography [462]). Other schemes work generally, but are only *a posteriori* guaranteed to be more efficient if the unknown state happens to be approximately of the type sought after (e.g., matrix product state tomography [317]). Finally, for maximum likelihood estimates or Bayesian statistical estimates and alike, priors could be added as well.

Note that the best understood case of pure state tomography can also be used for general quantum states, if one has access to the relevant purification. Specifically for pure state tomography, one then also needs to specify in what form access is given to the quantum state. Possible access models for pure state tomography include:

- Via samples of computational basis measurements $p(x) = \langle x|\rho|x\rangle$ for estimating the probabilities in the computational basis (not yet the pure state amplitudes).
- Via copies of the state that can be processed before measurement.
- Via the state preparation unitary $U|0^n\rangle\langle 0^n|U^\dagger = \rho$ (with ρ pure).
- Via the controlled version of aforementioned state preparation unitary U .
- Via aforementioned state preparation unitary U and its inverse $U^\dagger$.

Typically studied distance metrics to measure closeness of the statistical estimate to the true quantum state are the trace distance $T(\rho, \sigma) = \frac{1}{2}\text{tr}[|\rho - \sigma|]$, the quantum fidelity $F(\rho, \sigma) = \left(\text{tr}\left[\sqrt{\sqrt{\rho}\sqrt{\sigma}}\right]\right)^2$, and for pure quantum states also the ℓ_2 -norm of the difference $\| |\psi\rangle - |\phi\rangle \|^2$ of the pure states $|\psi\rangle, |\phi\rangle$ corresponding (up to global phase) to ρ and σ , respectively.

Dominant resource cost (gates/qubits)

Besides some potential ancilla qubits (few for typical tomographic schemes), the number of qubits is fixed by the dimension of the quantum state (of course,

whenever entangled measurements are used, the corresponding number of copies is needed). As such, the sample complexity is typically the relevant figure of merit. In the following, the notation $\Theta(\cdot)$ stands for simultaneous upper $O(\cdot)$ and lower $\Omega(\cdot)$ bounds on the asymptotic sample complexity, and the variant $\widetilde{\Theta}(\cdot)$ denotes the same up to factors that scale polylogarithmically in the relevant parameters. Tight sample and query complexity characterizations, in terms of an approximation error $\varepsilon \in [0, 1]$, then include the following noteworthy results:

- $\widetilde{\Theta}(d\varepsilon^{-2})$ sample complexity for pure state tomography in ℓ_2 -norm up to global phase [49] ([610] gave an algorithm with similar complexity, but requiring a state preparation unitary). The main idea is to use computational basis measurements to recover the absolute values of the amplitudes and then create some interference pattern for learning the phases.
- $\widetilde{\Theta}(d\varepsilon^{-1})$ query complexity for pure state tomography in ℓ_2 -norm with access to controlled state preparation unitary and its inverse [49], featuring a quadratic speedup in $1/\varepsilon$ reminiscent of amplitude estimation. The achievability results are based on the subroutine of quantum gradient estimation via an unbiased version of quantum phase estimation. Note that [267] used an alternative simpler algorithm based on iterative refinement and amplitude amplification that achieves the same query complexity but comes with improved gate complexity.
- $\Theta(dk^2\varepsilon^{-2})$ sample complexity for rank- k state tomography in trace distance with IID measurements [656, 478, 471]. The achievability results are based on low-rank matrix recovery techniques, where semidefinite programs have to be solved for reconstructing the quantum state from the collected measurement statistics. Note that the special case $k = 1$ corresponds to pure state tomography as in the setting of the first bullet point.
- $\widetilde{\Theta}(dk\varepsilon^{-2})$ sample complexity for rank- k state tomography in trace distance with entangled measurements [807, 478, 1075]. The achievability results are based on representation-theoretic techniques around the Schur transform.
- $\widetilde{\Theta}(dk\varepsilon^{-1})$ query complexity for rank- k state tomography in trace distance with access to controlled state preparation unitary of a purification and its inverse [49], featuring a quadratic speedup in $1/\varepsilon$ reminiscent of amplitude estimation.

For variations of the above, additional results in terms of lower and upper bounds are known. The derivations of the sample complexity lower bounds are often based on information-theoretic methods, exploiting the monotonicity of quantum-entropy-based measures. For sample complexity upper bounds, it is in practice additionally important that the algorithmic complexities of

the underlying schemes become efficient (in particular for entangled measurements performed on all n copies at once). Relevant metrics for the algorithmic complexity include quantum gate depth and number of measurement outcomes needed, as well as runtime and memory requirements of the classical postprocessing stage. We refer to [723] for a recent discussion on these computational aspects.

Caveats

As shown by the presented information-theoretic lower bounds, the sample complexity for general quantum state tomography grows exponentially in the number of qubits. As such, whenever quantum tomography is invoked as a subroutine in quantum algorithms, one has to carefully analyze if this step does not eliminate any claimed speedups of the quantum algorithm compared to state-of-the-art classical methods. One also has the inverse polynomial scaling in terms of the approximation parameter from the finite statistics, which is often prohibitively expensive for certain applications.

Additionally, on top of sample complexity for tomography schemes, the accompanying gate complexity should be considered as well. We refer to [49] for a discussion.

An alternative is to resort to only revealing partial classical information about quantum states, which might still be informative for the (algorithmic) task at hand. One such example with favorable scaling is shadow tomography, where the task is to not estimate the density matrix itself, but (very) many observables thereof. Shadow tomography—also known as quantum data analysis—can achieve exponential sample complexity improvements in terms of Hilbert space dimension compared to full state tomography and is guaranteed to yield exponential improvements in the number of target observables (compared to directly measuring all of them sequentially). The strongest result of this kind [2, 5, 209] requires entangled measurements across many state copies, as well as prohibitively large gate counts. More hardware-friendly protocols have been derived, known as classical shadow tomography [540, 823, 373]. In more detail, there exist algorithmically efficient and universal schemes that can simultaneously ε -approximate M linear functions $\text{tr}[O_i \rho]$ of an unknown quantum state ρ by only using $O(\log(M) \cdot \max_i \|O_i\|_s^2 \varepsilon^{-2})$ IID measurements. Note the scaling with $\log(M)$ instead of the standard M scaling. The shadow norm term $\|O_i\|_s^2$ scales in general as d , leading to the worst-case query complexity $O(d \log(M) \varepsilon^{-2})$. However, for observables with bounded Hilbert–Schmidt norm or for local observables, the overall dimension-free query complexity $O(\log(M) \varepsilon^{-2})$ is achievable.

Example use cases

Quantum tomographic or related data collection schemes are omnipresent in quantum algorithms. Some applications include:

- Quantum linear system solvers that output the full classical solution vector, where such solvers are, for example, employed for quantum interior point methods or for solving differential equations.
- Classical data about quantum states for variational quantum algorithms.
- Characterizing the performance of physical devices.
- Probing entanglement dynamics throughout a quantum simulation.
- Characterizing quantum processes.

Further reading

- Short perspective article, entitled “Focus on quantum tomography” [87].
- Recent overview on query complexity aspects [49].
- Recent overview on computational complexity aspects [723].
- Shadow tomography of quantum states [2].
- Review article on classical shadows and randomized measurements [373].