POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Anytime-Valid Quantum State Tomography via Confidence Sequences

Laurea Magistrale in Mathematical Engineering - Ingegneria matematica

Author: **Aldo Cumitini**

Advisor: **Prof. Luca Barletta**

Co-advisor: **Prof. Osvaldo Simeone**

Academic year: **2025-2026**

---

1. Introduction

Quantum state tomography (QST) is a fundamental procedure for characterizing quantum systems, enabling the estimation of an unknown quantum state from measurement outcomes. As quantum technologies mature and scale toward practical applications in computing, communication, and sensing, the ability to accurately estimate and validate quantum states becomes increasingly critical. Traditional QST methods provide point estimates or credible regions based on accumulated data, but they lack rigorous guarantees in the presence of misspecified models or when inference is performed sequentially. As illustrated in Fig. 1, the goal of this work is to augment existing QST techniques by associating current state estimates with confidence sets that are guaranteed to contain the true quantum state with a user-defined probability *uniformly* over time.

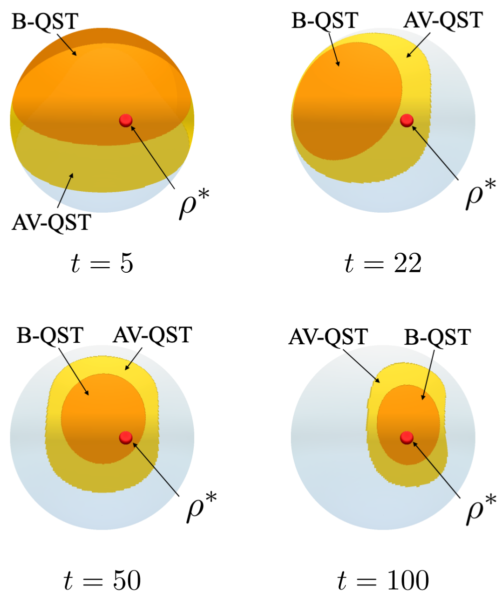

Figure 1: Confidence regions obtained as measurements are carried out on copies of a true state ρ^* over discrete time $t = 1, 2, \ldots$ using Bayesian credible regions (B-QST) [1] and the proposed anytime-valid QST (AV-QST). Bayesian QST (B-QST) does not provide time-uniform frequentist coverage guarantees (highlighted at $t = 22$); while the proposed AV-QST ensures that the confidence region includes the true state ρ^* (red dot) with probability no smaller than $1 - \alpha$ *uniformly* over *all* time steps t.

QST has been extensively studied from both frequentist and Bayesian perspectives. Maximum likelihood estimation (MLE) for quantum states provides consistent point estimates but no uncertainty quantification. The works [1] and [2] developed Bayesian methods using Markov chain Monte Carlo (MCMC) and sequential importance sampling (SIS) to construct posterior estimates and credible regions. The statistical validity of these techniques relies on the correct specification of the prior distribution. Recent advances have explored compressed sensing techniques, adaptive measurement strategies [2], and machine learning approaches. There is also an important line of work on estimating specific properties of quantum states, rather than recovering the full state.

The problem of providing rigorous error bounds for QST has been addressed by several recent works. Notably, the paper [3] introduced confidence regions based on the likelihood ratio that offer rigorous frequentist coverage guarantees. While this method provides reliable error bounds, it is designed for a *fixed sample size*. That is, the confidence regions are theoretically valid only for a predetermined number of measurements decided before data collection begins. This limitation poses challenges in sequential or adaptive scenarios where the experimenter may wish to monitor the state estimate in real-time, make decisions based on intermediate results, or stop the experiment when a desired accuracy is reached.

In contrast, this work aims at developing *anytime-valid* confidence sets that maintain their coverage guarantees uniformly over all time steps, including at data-dependent stopping times. This property is critical for sequential decision-making. The proposed method, *anytime-valid quantum state tomography* (AV-QST), augments existing QST methods with rigorous, time-uniform confidence guarantees. AV-QST builds on *confidence sequences*, a recent advance in sequential inference [4].

A shorter version of the thesis has been submitted to *IEEE Signal Processing Letters*. The submission distills the core methodology and theoretical guarantees underlying AV-QST and reports a subset of the numerical results presented in the original manuscript. A preprint is available at `https://arxiv.org/abs/2601.20761`.

2. Problem Definition

QST aims to reconstruct an unknown quantum state $\rho^* \in \mathcal{D}(\mathcal{H})$, where $\mathcal{D}(\mathcal{H})$ denotes the set of density matrices acting on a Hilbert space $\mathcal{H}$. For a D-dimensional quantum system, where $D = 2^m$ for an m-qubit system, the state ρ is represented by a $D \times D$ complex-valued density matrix. By definition, this matrix is positive semidefinite and has unit trace, i.e., $\rho \succeq 0$ and $\mathrm{tr}(\rho) = 1$.

In a tomographic experiment, information about the true state ρ^* is inferred from a sequence of measurement outcomes obtained from identically prepared copies of the system. Each measurement is described by a positive operator-valued measure (POVM) $\mathcal{M} = \{\Pi_x\}_{x=1}^M$, where each element Π_x is a positive semidefinite matrix and the equality $\sum_{x=1}^M \Pi_x = I$ is satisfied, with I being the $D \times D$ identity matrix. Accordingly, the probability of observing outcome x when the system is in state ρ is given by Born's rule:

$$\mathbb{P}(x|\rho; \mathcal{M}) = \mathrm{tr}(\Pi_x \rho). \tag{1}$$

At each discrete time step $t = 1, 2, \ldots$, a POVM $\mathcal{M}_t = \{\Pi_{t,x}\}_{x=1}^M$ is chosen, and an observation $X_t \sim \mathbb{P}(x_t|\rho; \mathcal{M}_t)$ is collected. Given the current sequence of measurement outcomes $\{X_{t'}\}_{t'=1}^t$, the goal of QST is to maintain an estimate of the underlying true state ρ^*.

Standard frequentist methods provide point estimates $\hat{\rho}_t$. In contrast, Bayesian methods not only produce a point estimate $\hat{\rho}_t$, but they also aim to quantify the uncertainty associated with this estimate via a credible set $\mathcal{C}_t \subseteq \mathcal{D}(\mathcal{H})$ (see Fig. 1). Specifically, the credible sets should ideally cover the true state ρ^* with some user-defined probability. However, this statistical validity property hinges on the correct specification of prior distribution which can be practically difficult [1].

Thus motivated, in this work, we study the problem of complementing frequentist or Bayesian point estimates $\hat{\rho}_t$ by constructing a sequence of uncertainty sets $\{\mathcal{C}_t\}_{t\geq 1}$, with $\mathcal{C}_t \subseteq \mathcal{D}(\mathcal{H})$ that are *anytime-valid* [5]. That is, at any point in the measurement sequence, including at random stopping times, the reported uncertainty set $\mathcal{C}_t$ maintains a statistically rigorous confidence level $1 - \alpha$ for containing the true quantum state ρ^*, where $1 - \alpha$ is a user-defined probability. Mathematically, this requirement is expressed by the inequality

$$\mathbb{P}(\forall t \geq 1,\ \rho^* \in \mathcal{C}_t) \geq 1 - \alpha. \tag{2}$$

Note that this is a frequentist guarantee, in the sense that the probability is evaluated with respect to the distribution of the observations for any fixed, unknown, true state ρ^*.

3. Conventional Methods

In this section, we summarize baseline frequentist and Bayesian QST methods.

3.1. Frequentist Quantum State Tomography

The standard frequentist methodology is based on *maximum likelihood estimation* (MLE). Given the available measurements $\{X_{t'}\}_{t'=1}^{t}$, the MLE maximizes the probability of the observed measurement outcomes, producing a point estimate of the true state.

The probability of observing the data under a candidate state ρ is

$$\mathcal{L}_t(\rho) = \prod_{t'=1}^{t} \mathbb{P}(X_{t'}|\rho; \mathcal{M}_{t'}) = \prod_{t'=1}^{t} \mathrm{tr}(\Pi_{X_{t'}}\rho), \tag{3}$$

where we have denoted for simplicity $\Pi_{t,X_t} = \Pi_{X_t}$. The MLE estimate $\hat{\rho}^{\text{MLE}}$ is then evaluated as the maximizer

$$\hat{\rho}_t^{\text{MLE}} = \underset{\rho\succeq 0,\, \mathrm{tr}(\rho)=1}{\arg\max} \; \log \mathcal{L}_t(\rho), \tag{4}$$

where $\log \mathcal{L}_t(\rho)$ is the log-likelihood function.

3.2. Bayesian Quantum State Tomography

Like frequentist QST, Bayesian QST produces a point estimate of the underlying quantum state, but it also quantifies the uncertainty of this estimate via a credible set [1]. To start, Bayesian QST requires the specification of a prior distribution $\pi_0(\rho)$ over quantum states. For pure states, a typical choice is given by the Haar measure, which is uniform over the unit sphere in Hilbert space, while for density matrices one can induce uniform distributions via partial tracing.

Given the available data $\{X_{t'}\}_{t'=1}^{t}$, at time t, the posterior distribution $\pi_t(\rho)$ can be evaluated recursively by updating the previous posterior distribution $\pi_{t-1}(\rho)$ with the new observation X_t and the corresponding measurement probability $\mathbb{P}(X_t|\rho; \mathcal{M}_t) = \mathrm{tr}(\Pi_{X_t}\rho)$ as

$$\pi_t(\rho) \propto \pi_{t-1}(\rho)\, \mathrm{tr}(\Pi_{X_t}\rho). \tag{5}$$

A point estimator of the state is typically defined in *Bayesian QST* (B-QST) via the *posterior mean*

$$\hat{\rho}_t^{\text{B-QST}} = \int \rho\, \pi_t(\rho)\, d\rho. \tag{6}$$

An efficient, approximate, numerical approach to evaluate (6) via a *Markov chain Monte Carlo* (MCMC) method is detailed in [1], while an adaptive numerical update of the posterior via *sequential importance sampling* (SIS) is described in [2]. These methods are based on maintaining a set of particles over time t as an approximation of the posterior distribution (6).

Moreover, B-QST augments the point estimate (6) with a *credible set* $\mathcal{C}_t$ targeting a coverage level $1-\alpha$ separately for each time t, i.e., $\mathbb{P}(\rho^* \in \mathcal{C}_t | \{X_{t'}\}_{t'=1}^{t}) \geq 1-\alpha$. Note that this is a Bayesian, not frequentist, guarantee, in the sense that the probability is evaluated for a fixed dataset $\{X_{t'}\}_{t'=1}^{t}$ under the given prior distribution for the random set $\rho^* \sim \pi_0(\rho)$. This set can be in principle obtained by finding the highest density region in the density matrix space $\mathcal{D}(\mathcal{H})$ under the posterior distribution $\pi_t(\rho)$.

However, in practice, evaluating the set is computationally prohibitive and one leverages heuristic methods, e.g., based on Gaussian approximations of the posterior distribution [1]. The threshold used to bound the Gaussian density can then be obtained by ensuring that a fraction $1-\alpha$ of the particles produced by MCMC [1] or SIS [2] are included in the set.

Overall, B-QST generally fails to satisfy an anytime-validity requirement akin to (2) for a number of reasons: (*i*) the prior distribution may be misspecified; (*ii*) the posterior distribution is only asymptotically Gaussian; and (*iii*) even in the ideal case of correct prior specification and exact posterior calculation, guarantees are only provided *separately* for each time t.

4. Anytime-Valid Quantum Tomography

We now introduce anytime-valid QST (AV-QST), a framework for online uncertainty quantification in QST that augments any point predictor of the quantum state with anytime-valid confidence sets satisfying the condition (2). The approach builds on the statistical methodology of confidence sequences [4].

4.1. Confidence Sequences

At each time step t, AV-QST assumes access to *any* point predictor $\hat{\rho}_{t-1}$, e.g., the MLE in (4) or the posterior mean in (6). We initialize the predictor with a randomly chosen quantum state $\hat{\rho}_0$. Based on this estimate, given the available data $\{X_{t'}\}_{t'=1}^{t}$, AV-QST constructs the *likelihood ratio*

$$R_t(\rho) = \prod_{t'=1}^{t} \frac{\mathrm{tr}(\Pi_{X_{t'}} \hat{\rho}_{t'-1})}{\mathrm{tr}(\Pi_{X_{t'}} \rho)}. \quad (7)$$

The likelihood ratio $R_t(\rho)$ in (7) quantifies the available evidence against the candidate state ρ relative to the estimates $\hat{\rho}_{t'-1}$ evaluated at each time $t' = 1, \dots, t$. Specifically, a lower value of the ratio $R_t(\rho)$ indicates that state ρ is more likely to represent the true state ρ^* as compared to these estimates.

Having evaluated this ratio, AV-QST obtains the uncertainty set

$$\mathcal{C}_t^{\text{AV-QST}}(\alpha) = \{\rho \in \mathcal{D}(\mathcal{H}) : R_t(\rho) \leq 1/\alpha\}. \quad (8)$$

The set $\mathcal{C}_t^{\text{AV-QST}}(\alpha)$ is the collection of all quantum states whose likelihood ratio is sufficiently low, with the threshold $1/\alpha$ depending on the desired target coverage probability in (2). Based on the definition (7), the set $\mathcal{C}_t^{\text{AV-QST}}(\alpha)$ thus contains states for which there is sufficient evidence in their favor as compared to the current best estimate.

The likelihood ratio (7), instantiated with the MLE as the point estimate, was also used in [3] to construct a confidence set $\mathcal{C}_T(\alpha)$ for a *batch* setting with a fixed sample size T. The scheme, referred to as *likelihood ratio-based QST* (LR-QST), returns the single set

$$\mathcal{C}_T^{\text{LR-QST}}(\alpha) = \{\rho \in \mathcal{D}(\mathcal{H}) : R_T^{\text{LR-QST}}(\rho) \leq \lambda_\alpha\} \quad (9)$$

at the last time step, where

$$R_T^{\text{LR-QST}}(\rho) = -2 \log \left(\frac{\mathcal{L}_T(\rho)}{\mathcal{L}_T(\hat{\rho}_T^{\text{MLE}})} \right) \quad (10)$$

is a likelihood ratio, and the threshold λ_α is obtained by solving numerically an implicit equation (see [3, Equation (14)]). LR-QST was derived with the aim of offering only coverage guarantees for the last time step T, and it does not allow for the use of more powerful point estimates such as the posterior mean (6).

4.2. Theoretical Properties

As stated in the following proposition, which follows directly from the properties of confidence sequences [4], the sequence of sets $\{\mathcal{C}_t^{\text{AV-QST}}(\alpha)\}_t$ produced by AV-QST satisfies the desired anytime validity property (2).

Proposition 4.1. *For any sequence of point estimates $\{\hat{\rho}_t\}_t$, the sequence of AV-QST uncertainty sets $\{\mathcal{C}_t^{\text{AV-QST}}(\alpha)\}_t$ in (8) satisfies the anytime-valid coverage property (2).*

Proof. Given the definition of the AV-QST set (8), the target coverage probability (2) can be written as

$$\begin{aligned} &\mathbb{P}\Big(\forall t \geq 1,\ \rho^* \in \mathcal{C}_t^{\text{AV-QST}}(\alpha)\Big) = \\ &\quad = \mathbb{P}\Big(\forall t \geq 1,\ R_t(\rho^*) \leq 1/\alpha\Big). \end{aligned} \quad (11)$$

where the probability is evaluated with respect to the distribution of the measurements, i.e., $\prod_{t'=1}^{t} \mathbb{P}(x_{t'}|\rho^*; \mathcal{M}_{t'})$. Furthermore, the likelihood ratio $R_t(\rho^*)$ can be readily verified to be a test martingale under the measurement distribution (see Appendix). By Ville's inequality [5], this directly implies the inequality

$$\mathbb{P}(\forall t \geq 1, R_t(\rho^*) \leq 1/\alpha) \geq 1 - \alpha. \quad (12)$$

Combining this inequality with (11) completes the proof. □

5. Numerical Results

In this section, we present numerical results for simulated tomography experiments. All simulations were implemented in Python and executed on a M3 Max MacBook Pro. Throughout, we perform local measurements using the minimal informationally complete POVM (MIC-POVM) separately on each qubit. For each simulated Monte Carlo run, we generate a random pure quantum state ρ^* following the Haar uniform distribution. LR-QST is run by evaluating the set (9) separately for each time $t = T$. The point estimate $\hat{\rho}_t$ used in AV-QST is the MLE. Fig. 1 displays the confidence regions produced by the proposed AV-QST, along with B-QST (Sec. 3.2) for a given run. The figure shows that the true state ρ^* is not included in the B-QST region at time $t = 22$, reflecting the theoretical property that B-QST does not satisfy the anytime-validity requirement, contrary to the AV-QST region.

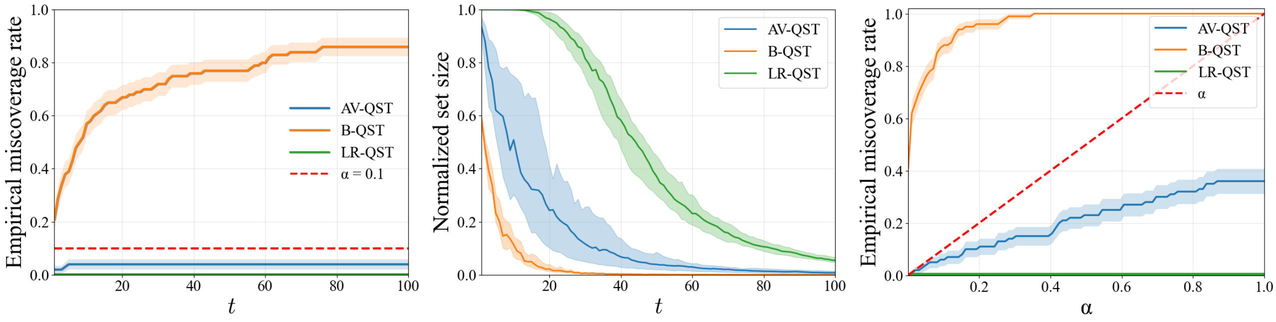

(a) (b) (c)

Figure 2: Empirical miscoverage and normalized set size for B-QST [1], LR-QST [3], and AV-QST (this work) for the two-qubit setting. The lines represent median values and the shaded areas correspond to the interquartile range (25th-75th percentiles).

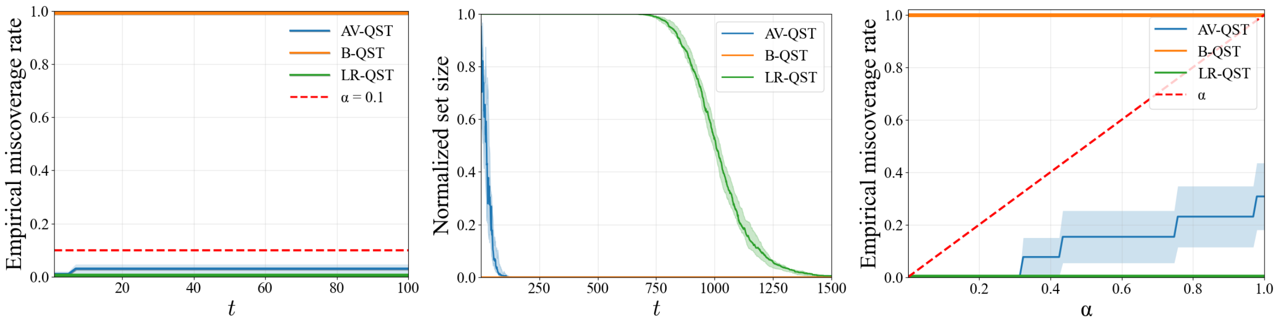

(a) (b) (c)

Figure 3: Empirical miscoverage and normalized set size for B-QST [1], LR-QST [3], and AV-QST (this paper) for the four-qubit setting. The lines represent median values and the shaded areas correspond to the interquartile range (25th-75th percentiles).

Fig. 2 reports the results averaged over multiple runs for the two-qubit case ($D = 4$), and Fig. 3 for the four-qubit case ($D = 16$). For the target miscoverage rate $\alpha = 0.1$, we plot the empirical miscoverage rate and the normalized set size as a function of time t in panels (a) and (b). The empirical miscoverage rate at any time t is estimated by counting the fraction of runs at which the true state ρ^* was not covered by the confidence set up to time t. The normalized set size is estimated by evaluating the fraction of the space of density matrices covered by the confidence set via sampling.

In both two-qubit and four-qubit settings, the set sizes for all schemes shrink as t increases and more measurement data are collected. Moreover, the empirical miscoverage rate for the proposed AV-QST remains below the required upper bound α specified in (12) at all times, whereas the miscoverage rate of the B-QST credible region grows well above this bound. In contrast, LR-QST almost never exhibits miscoverage because it produces an overly large confidence region, which makes it substantially less informative than AV-QST. This result highlights a general conclusion from our experiments: While LR-QST cannot theoretically guarantee anytime coverage, in practice it is seen to be extremely conservative, yielding regions whose size is much larger than for AV-QST, especially for a larger number of qubits.

Panels (c) in Fig. 2 and Fig. 3 show the empirical miscoverage rate as a function of the target miscoverage rate α at $t = 100$. B-QST sets are consistently smaller than the corresponding confidence sets produced by AV-QST, failing to meet the desired miscoverage rate, while the LR-QST sets are always larger than the sets produced by AV-QST, confirming the over-conservativeness of LR-QST sets.

6. Conclusion

This thesis has introduced anytime-valid confidence sets for quantum state tomography. The theoretical properties of the proposed approach, termed AV-QST, were seen via experiments to yield practical benefits in terms of coverage and set size. This work contributes to a recent line of research aiming at applying recent advances in statistics to quantum information processing. Among directions for future work, we mention the development of statistically valid change-detection mechanisms and the design of adaptive measurement strategies, which are supported by the proposed framework.

Appendix: Proof of Proposition

By defining $R_0(\rho^*) = 1$, we can write (7) via the recursive relation

$$R_t(\rho^*) = \frac{\operatorname{tr}(\Pi_{X_t}\hat{\rho}_{t-1})}{\operatorname{tr}(\Pi_{X_t}\rho^*)} R_{t-1}(\rho^*). \tag{13}$$

Let $\mathscr{F}_t = \sigma(X_1, \dots, X_t)$ be the natural filtration. By the properties of conditional expectation, we then have

$$\begin{aligned}
&\mathbb{E}[R_t(\rho^*) \mid \mathscr{F}_{t-1}] = \\
&\quad = R_{t-1}(\rho^*)\, \mathbb{E}\left[\frac{\operatorname{tr}(\Pi_{X_t}\hat{\rho}_{t-1})}{\operatorname{tr}(\Pi_{X_t}\rho^*)} \,\middle|\, \mathscr{F}_{t-1}\right] \\
&\quad = R_{t-1}(\rho^*) \sum_{x=1}^{M} \operatorname{tr}(\Pi_x\rho^*) \frac{\operatorname{tr}(\Pi_x\hat{\rho}_{t-1})}{\operatorname{tr}(\Pi_x\rho^*)} \\
&\quad = R_{t-1}(\rho^*).
\end{aligned} \tag{14}$$

The derivation (14) shows that the likelihood ratio sequence satisfies the martingale property. Furthermore, in a similar way, one can see that the expected value equals 1, i.e., $\mathbb{E}[R_t(\rho^*)] = \mathbb{E}[R_0(\rho^*)] = 1$ for all times t. Thus the sequence $R_t(\rho^*)$ is a test martingale for the null hypothesis that the state is ρ^*.

References

[1] Robin Blume-Kohout. Optimal, reliable estimation of quantum states. *New Journal of Physics*, 12(4):043034, April 2010.

[2] Ferenc Huszár and Neil MT Houlsby. Adaptive Bayesian quantum tomography. *Physical Review A*, 85(5):052120, 2012.

[3] Robin Blume-Kohout. Robust error bars for quantum tomography. *arXiv preprint arXiv:1202.5270*, 2012.

[4] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1085, 2021.

[5] Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values. *Foundations and Trends® in Statistics*, 1(1–2):1–390, 2025.