# Paranoid tomography: Confidence regions for quantum hardware

Robin Blume-Kohout
*Theoretical Division, Los Alamos National Laboratory*[*]

Many "paranoid" quantum information processing protocols, such as fault tolerance and cryptography, require rigorously validated quantum hardware. Such hardware (elementary state- and gate-producing components) is calibrated using tomography, by combining many measurement results into a single estimate of the state $\rho$ or process $\chi$. But a *point estimate* – a single matrix ($\hat{\rho}$ or $\hat{\chi}$) that is probably close to the true state $\rho$ or process $\chi$ – cannot provide the rigorous validation needed for paranoid protocols. *Region estimators*, on the other hand, provide just such a guarantee.

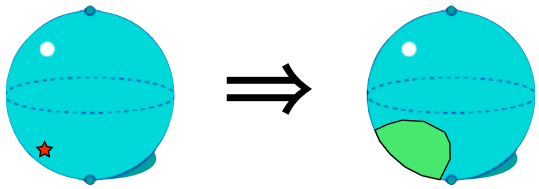

This paper presents *likelihood-ratio confidence regions*, a powerful, reliable, and convenient tool for rigorous state (and process[1]) tomography.

**A REGION ESTIMATOR** is a map from data $D$ to regions $\hat{\mathcal{R}}$ in state space. Its defining property is *correctness*: **"$\hat{\mathcal{R}}(D)$ contains the true state/process with probability $\alpha = 1 - \epsilon$."** Some popular ad hoc estimators, such as bootstrapped standard errors, generally fail to satisfy any such condition. A good region estimator should also assign the smallest achievable regions. Other ad hoc estimators, like those derived from large deviation bounds, assign much larger regions than necessary.

The two *good* candidates differ mainly in how they define correctness. *Credible regions* are inherently Bayesian, and define correctness as

$$Pr(\rho_{\text{true}} \in \hat{\mathcal{R}}(D)|D) \geq \alpha.$$

But this condition requires assuming a prior distribution $Pr(\rho_{\text{true}})$ – and its validity depends critically on the "truth" of the prior. As a result, credible regions can be explicitly broken by adversarial choice of $\rho_{\text{true}}$. This renders them unsuitable for paranoid applications. *Confidence region estimators* (CREs), define correctness as

$$\begin{aligned} Pr(\rho_{\text{true}} \in \hat{\mathcal{R}}(D)|\rho_{\text{true}}) &\geq \alpha \text{ for all } \rho_{\text{true}} \\ &\Rightarrow Pr(\rho_{\text{true}} \in \hat{\mathcal{R}}(D)) \geq \alpha, \end{aligned} \quad (1)$$

and are guaranteed to produce a region containing $\rho_{\text{true}}$ with probability $\alpha$, even if $\rho_{\text{true}}$ is chosen adversarially! Eq. 1) leaves us a great deal of freedom in defining a CRE that (i) assigns small, *powerful* regions, while (2) minimizing the complexity of computing $\hat{\mathcal{R}}(D)$.

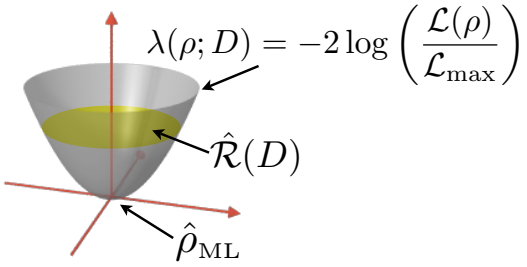


**Likelihood ratio (LR) confidence regions** achieve correctness, near-optimal power, and simplicity all at once. LR regions are defined and computed using the likelihood function, $\mathcal{L}(rho) \equiv Pr(D|\rho)$, where $D$ is the actual observed data. Given data $D$, we report the region

$$\hat{\mathcal{R}}(D) = \left\{\text{all } \rho \text{ such that } -2\log\left(\frac{\mathcal{L}(\rho)}{\mathcal{L}_{\max}}\right) \leq \lambda_c\right\}. \quad (2)$$

$\mathcal{L}_{\max}$ is the maximum value of $\mathcal{L}(\rho)$. The cutoff $\lambda_c$ (see "Correctness" below) is set by the system dimension $d$ and the desired confidence $\alpha$.

**POWER** can be quantified by volume.

$$V(\mathcal{R}) = \int_{\rho\in\mathcal{R}} \mathrm{d}\rho.$$

A powerful estimator assigns small regions. Remarkably, it does not matter what measure $\mathrm{d}\rho$ is chosen – the same estimator minimizes *all* notions of volume! But $V(\hat{\mathcal{R}}(D))$ cannot be simultaneously minimized for every dataset $D$. Nor can we simultaneously minimize *expected* volume,

$$\overline{V}(\rho) = \sum_D Pr(D|\rho)V(\hat{\mathcal{R}}(D)),$$

for every true state $\rho$. So instead of a unique "best" CRE, we find a whole class of *admissible* CREs (ones that are not strictly dominated by any other). Each admissible estimator minimizes *average* expected volume with respect to some averaging measure $Pr(\rho)\mathrm{d}\rho$,

$$\langle\overline{V}\rangle_{Pr(\rho)\mathrm{d}\rho} = \int \overline{V}(\rho)Pr(\rho)\mathrm{d}\rho,$$

and is (see technical paper for a simple proof) a *probability-ratio estimator*:

$$\hat{\mathcal{R}}(D) = \left\{\text{all } \rho \text{ such that } \frac{Pr(D|\rho)}{Pr(D)} \geq r_c(\rho)\right\}. \quad (3)$$

[1] The Choi-Jamiolkowski isomorphism makes process estimation formally equivalent to state estimation, on a larger system

The averaging measure determines $Pr(D) = \int Pr(D|\rho)Pr(\rho)\mathrm{d}\rho$, while the cutoff $r_c(\rho)$ is chosen explicitly to satisfy Eq. 1.

If $\rho$ were distributed according to a known measure, then obviously we should use the corresponding probability-ratio estimator. Lacking (in general) reliable prior knowledge of this sort, we want an estimator with good behavior for *all* states $\rho$. We get it (see technical paper) by choosing $Pr(\rho)\mathrm{d}\rho$ so that

$$Pr(D) \propto \max_{\rho} Pr(D|\rho).$$

Inserting this choice into Eq. 3 – and observing that $Pr(D|\rho) = \mathcal{L}(\rho)$ – yields the LR prescription given in Eq. 2.

**CORRECTNESS** relies on a wise choice of the cutoff $\lambda_c$. Too low, and Eq. 1 is violated. Too high, and overly large regions are assigned. The perfect choice of $\lambda_c$ varies with $\rho$, and barely satisfies

$$Pr\left[\lambda(D,\rho) \leq \lambda_c(\rho)\right] \geq \alpha, \tag{4}$$

where $\lambda(D,\rho) = -2\log\left(\mathcal{L}(\rho)/\mathcal{L}_{\max}\right)$ is the *loglikelihood ratio* statistic. But allowing $\lambda_c$ to vary with $\rho$ yields disconnected confidence regions. So, instead, we replace $\lambda_c(\rho)$ with a constant upper bound

$$\lambda_c \gtrsim \max_{\rho} \lambda_c(\rho),$$

which (at a small cost in power) ensures connected *and* convex regions.

For any given $\rho_{\text{true}}$, $\lambda$ is a statistic (i.e., a random variable depending on the data $D$). We calculate $\lambda_c$ by studying the worst-case distribution of $\lambda$ and then making sure that Eq. 4 is always satisfied. If the data were Gaussian (rather than multinomial), with $K = d^2 - 1$ degrees of freedom (corresponding to the $d^2 - 1$ independent parameters in $\rho$), then $\lambda$ would be a $\chi^2_K$ random variable with cumulative distribution

$$Pr(\lambda \leq x) = \frac{\gamma\left(\frac{K}{2}\right), \frac{x}{2}}{\Gamma\left(\frac{K}{2}\right)}.$$

But this Gaussian approximation is *not* a rigorous lower bound on the cumulative distribution of $\lambda$. Deriving such a bound is necessary, but tedious (see technical paper). One series of approximations yields

$$Pr(\lambda \leq x) \geq \frac{\gamma\left(\frac{K}{2}\right), \frac{x}{2}}{\Gamma\left(\frac{K}{2}\right)} - e^{-x/2}\left[\left(1 + \frac{\sqrt{3ex}}{\pi}\right)^K - \frac{\sqrt{3ex}}{\pi}^K\right] \tag{5}$$

Although numerics suggest that this bound is loose by a factor of $\sqrt{x}$ (Fig. 1), the $e^{-x/2}$ term dominates. More importantly, it *is* a strict bound. So regions based on Eq. 5 are slightly loose, but have guaranteed coverage probability of at least $\alpha$.

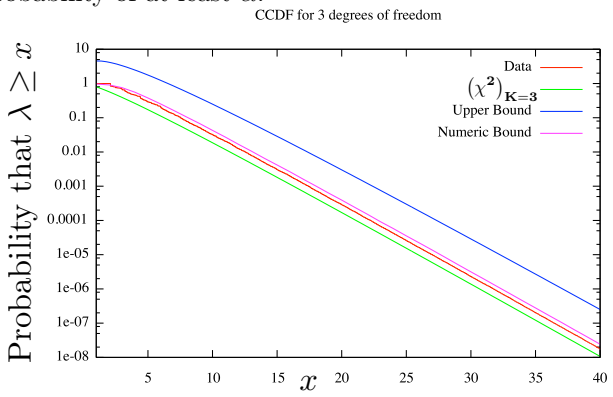


FIG. 1: The complementary cumulative distribution function of $\lambda$ – $Pr(\lambda \geq x)$ – for $K = d^2 - 1 = 3$ degrees of freedom, corresponding to a qubit state. **Red**: the empirical worst-case distribution, computed numerically. **Green:** The $\chi^2$ prediction, clearly a fatal underestimate. **Blue:** Eq. 5, a loose but reliable upper bound. **Purple:** a more sophisticated version of Eq. 5 requiring some number-crunching.

**SIMPLICITY**: The simple prescription given in Eq. 2 yields convex regions defined by level sets of an efficiently computable likelihood function. If an explicit region is needed, $\hat{\mathcal{R}}(D)$ can be efficiently described by (i) sampling from its boundary, and (ii) computing its minimum-volume bounding ellipsoid as $O(d^4)$ numbers.

Usually, however, $\hat{\mathcal{R}}(D)$ can be used implicitly via convex programming. *Example:* A convex program can find the point $\hat{\rho}$ that maximizes $\min_{\rho\in\hat{\mathcal{R}}} F(\hat{\rho},\rho)$, which yields a guaranteed upper bound on the infidelity. But there are easier ways to validate existing protocols. Region estimates can be used to design *better* protocols, tailored to known errors and uncertainties. *Example:* each fault tolerance protocol defines a "witness" hypersurface separating "good" states from "fail" states. Under certain assumptions, a convex program can search for tailored protocols that work for every $\rho \in \hat{\mathcal{R}}$.

LR regions also provide an elegant *theoretical* framework for analyzing errors, in terms of the derivatives of $\mathcal{L}(\rho)$ at its maximum ($\hat{\rho}_{\text{MLE}}$). When $\hat{\rho}_{\text{MLE}}$ is full rank, and $\vec{\nabla}\mathcal{L}$ vanishes, confidence regions are ellipsoidal and we recover the standard Fisher information. When $\hat{\rho}_{\text{MLE}}$ is rank-deficient, $\vec{\nabla}\mathcal{L} \neq 0$ and Fisher information goes haywire. But the LR-region framework remains robust, and implies truncated-ellipsoid confidence regions described by the first *and* second derivatives of $\mathcal{L}(\rho)$.

* Electronic address: robin@blumekohout.com