Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Qubit Ensembles from Directed Distinguishability with Uncalibrated Loss

Idrees Oreibi

idrees114af@gmail.com

Directorate of Education Babylon, Ministry of Education, Iraq, Babil

Rehab Shather Abdul Hamza

Directorate of Education Babylon, Ministry of Education, Iraq, Babil

Research Article

Keywords: Quantum information theory, Quantum tomography, Quantum measurements

Posted Date: September 15th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11029264/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Qubit Ensembles from Directed Distinguishability with Uncalibrated Loss

Idrees Oreibi[1,*] and Rehab Shather Abdul Hamza[1,†]

[1] *Directorate of Education Babylon, Ministry of Education, Iraq, Babil*

Can pairwise distinguishability determine an ensemble of quantum states? We show that the complete directed max-relative-entropy matrix determines four generic faithful states of a single qubit up to a common unitary or antiunitary transformation. For arbitrary normalized Lorentz-ball ensembles we classify every compatible realization by one scalar; its endpoint uniquely minimizes embedding dimension and, for qubits, maximizes every state's entropy. We establish local stability, a nonuniform-stability obstruction, sparse extension from four anchors, and acquisition rules that cancel specified loss and throughput nuisance parameters. Seven independent 100×100 superconducting-qubit frequency tables at each of 5 and 10 μs provide an end-to-end finite-setting test: experimental frequencies are converted directly to a directed divergence and then to reconstructed qubit geometry, with no tomography in the reconstruction path. Median four-state Gram errors are about 2.4×10^{-2} for pooled data and 3.5×10^{-2} under strict split-sample selection/evaluation, with monotonic improvement as the number of measurement directions increases. Independent full-Stokes data separately test acquisition-layer nuisance cancellation. The protocol provides calibration robustness rather than a universal sample advantage.

An ensemble is conventionally specified relative to calibrated measurements. A different description retains only how strongly every preparation can be distinguished from every other preparation. Whether such relational information fixes the ensemble depends on which notion of distinguishability is retained. We study the directed max-relative entropy

$$D_{ij} = D_{\max}(\rho_i \| \rho_j) = \log \inf\{\lambda : \rho_i \le \lambda \rho_j\}, \qquad (1)$$

for faithful trace-one states and natural logarithms. Max-relative entropy is a standard one-shot divergence [1]; its dual cone expression is a worst-case ratio of positive-effect probabilities [2]. The symmetric sum $D_{ij} + D_{ji}$ is Hilbert's projective metric, whose qubit geometry is hyperbolic [2–4]. Hyperbolic distance-matrix reconstruction and open-domain Funk rigidity are established [5, 6]; related self-calibrating tomography addresses different calibration models [7, 8]. Projective distances alone do not fix state purities: a common Lorentz boost changes the normalized section while preserving projective distances. The central question is whether the antisymmetric sector of the finite directed matrix supplies the missing normalization.

Write a qubit state as $\rho_i = (I + \boldsymbol{r}_i \cdot \boldsymbol{\sigma})/2$, with $|\boldsymbol{r}_i| < 1$, and define $\Delta_i = 1 - |\boldsymbol{r}_i|^2$. The boundary of $\rho_i \le \lambda \rho_j$ gives

$$D_{ij} = \eta_{ij} + s_i - s_j, \qquad s_i = \tfrac{1}{2} \log \Delta_i, \qquad (2)$$

$$\cosh \eta_{ij} = \frac{1 - \boldsymbol{r}_i \cdot \boldsymbol{r}_j}{\sqrt{\Delta_i \Delta_j}}. \qquad (3)$$

For a proposed directed matrix set

$$H = \tfrac{1}{2}(D + D^T), \quad A = \tfrac{1}{2}(D - D^T), \quad G_{ij} = \cosh H_{ij}. \qquad (4)$$

Cycle consistency is $A_{ij} + A_{jk} + A_{ki} = 0$. Choose the gauge $\tilde{s}_i = A_{i1}$, let $w_i = e^{-\tilde{s}_i}$ and $W = \mathrm{diag}(w)$. If G has inertia $(1, r-1, n-r)$, then a normalized Lorentz-ball realization exists exactly when

$$w \in \mathrm{Range}(G), \qquad q = w^T G^+ w > 0, \qquad (5)$$

and every compatible Euclidean Gram matrix is

$$B(\tau) = \mathbf{1}\mathbf{1}^T - \tau W^{-1} G W^{-1}, \qquad 0 < \tau \le q. \qquad (6)$$

At the endpoint $\tau = q$, $\mathrm{rank}\, B = r - 1$; for $\tau < q$, $\mathrm{rank}\, B = r$. Thus four generic qubit preparations have $r = 4$, and only the endpoint fits in the three-dimensional Bloch ball. Three generic states retain the scalar ambiguity. The endpoint also maximizes every qubit entropy because

$$\mathrm{Tr}\, \rho_i^2 = 1 - \frac{\tau}{2w_i^2}, \qquad \mathrm{Tr}(\rho_i \rho_j) = 1 - \frac{\tau G_{ij}}{2 w_i w_j}. \qquad (7)$$

The complete proof, singular cases, and range-restricted pseudoinverse argument are given in the Supplemental Material [9].

Stability and sparse extension.— For an invertible four-state anchor, the inverse map is analytic locally. With $v = G^{-1} w$,

$$dq = 2 v^T dw - v^T (dG) v, \qquad dG_{ij} = \sinh(H_{ij})\, dH_{ij}. \qquad (8)$$

If $\max_{ij} |\widehat{D}_{ij} - D_{ij}| \le \epsilon$, then $\|\widehat{G} - G\|_2 \le 4\epsilon \sinh(h + \epsilon)$ with $h = \max H_{ij}$, and a resolvent bound gives $\|\widehat{B} - B\|_2 = O(\epsilon)$ whenever the inverse gap remains open. No uniform modulus exists over all generic faithful ensembles: two sequences can approach the same planar directed limit while their Gram matrices remain separated. Numerically, the controlled phase diagram reaches a 90th-percentile Gram-error amplification of 22.49 relative to directed-entry noise near affine degeneracy, reduced to 6.90 after cycle projection.

Once four anchors are known, a new preparation needs its four directed distances to the anchors and one reverse distance. Hence

$$12 + 5(n-4) \tag{9}$$

directed entries suffice for an n-state ensemble containing a full-span anchor set; no optimality of this count is claimed.

Loss-robust acquisition.— The dual cone identity gives

$$e^{D_{ij}} = \max_{|\boldsymbol{u}|=1} \frac{1 + \boldsymbol{r}_i \cdot \boldsymbol{u}}{1 + \boldsymbol{r}_j \cdot \boldsymbol{u}}. \tag{10}$$

For a one-port analyzer $E_k = \eta_k(I + \boldsymbol{u}_k \cdot \boldsymbol{\sigma})/2$, the click probability is $p_{ik} = \eta_k(1 + \boldsymbol{r}_i \cdot \boldsymbol{u}_k)/2$, so p_{ik}/p_{jk} cancels unknown setting efficiency when the same setting loss applies to every preparation. With K iid uniform directions and T trials per state-setting pair, a conservative sufficient finite-copy bound is

$$K \geq \frac{4}{\mu\epsilon} \log \frac{2n(n-1)}{\delta}, \qquad T \geq \frac{192}{\eta_0 \mu \epsilon^2} \log \frac{4nK}{\delta}, \tag{11}$$

where $\lambda_{\min}(\rho_i) \geq \mu$ and $\eta_k \geq \eta_0$.

Recording both complementary outputs removes the shared-across-preparations restriction. If

$$P(+|i,k) = \gamma_{ik} f_{ik}, \qquad P(-|i,k) = \gamma_{ik}(1 - f_{ik}), \tag{12}$$

then $C^+_{ik}/(C^+_{ik} + C^-_{ik})$ estimates f_{ik} independently of arbitrary common-mode throughput γ_{ik}. For paired measurements exchanging the two outcomes between channels, background-subtracted counts

$$C^{(a)}_h = \gamma_a g_h p_-, \qquad C^{(a)}_v = \gamma_a g_v p_+, \tag{13}$$

$$C^{(b)}_h = \gamma_b g_h p_+, \qquad C^{(b)}_v = \gamma_b g_v p_- \tag{14}$$

give the exact double ratio

$$Q = \frac{C^{(a)}_v C^{(b)}_h}{C^{(a)}_h C^{(b)}_v} = \left(\frac{p_+}{p_-}\right)^2, \qquad \boldsymbol{r}_i \cdot \boldsymbol{u}_k = \tanh\left(\tfrac{1}{4}\log Q\right). \tag{15}$$

Thus γ_a, γ_b and unequal stable channel gains g_h, g_v cancel. Additive backgrounds must be subtracted and relative gains must remain stable within a swap pair; no universal nuisance cancellation is claimed.

End-to-end experimental-data check.— We applied the inverse protocol directly to the public superconducting-qubit prepare-and-measure frequency tables of Aloy *et al.* [10]. The released time-evolution data contain seven independent 100×100 frequency tables at each tested delay, with 100 preparations against 100 binary measurements. From the pooled frequencies we form the finite-setting estimator

$$\widehat{D}^{(K)}_{ij} = \log \max\left\{1, \max_{k \leq K,\, a \in \{0,1\}} \frac{\widehat{p}(a|i,k)}{\widehat{p}(a|j,k)}\right\}. \tag{16}$$

Only this directed matrix enters cycle projection and the four-state inverse map. No density matrix or GPT coordinate enters the reconstruction. The authors' transformed rank-4 GPT geometry from the same experiment is used only afterward as a benchmark; it is therefore not an independent ground truth.

At 5 and 10 μs, respectively, 93.3% and 94.4% of 3000 fixed random quartets satisfy the Lorentz/physical-endpoint checks, with median benchmark Gram RMS 0.02393 and 0.02239. The sparse four-anchor plus one reverse-distance extension reconstructs all 100 preparations with pooled Gram RMS 0.02494 and 0.03359, respectively. To control maximization bias, a strict split-sample test selects the maximizing setting and outcome on runs 1–3 and evaluates that same effect on runs 4–7, then reverses the roles and averages the two held-out log-ratio estimates. Median quartet errors become 0.03513 and 0.03478, with physical-valid fractions 78.2% and 72.3%. Run-level bootstrap resampling gives 95% intervals $[0.02445, 0.03244]$ and $[0.02176, 0.02800]$ for the median pooled quartet error. Finally, restricting the number of measurement directions from $K = 100$ to 20 increases the median error from 0.02378 to 0.07043 at 5 μs and from 0.02358 to 0.05331 at 10 μs, with monotonic improvement as K grows. This separates the finite-direction approximation from the exact full-sphere theorem.

The superconducting data test the directed reconstruction itself; independent raw full-Stokes dual-track spectra test the complementary acquisition nuisance model [11, 12]. Track swapping suppresses false polarization by 36–50 relative to a single setting, while the paired-swap ratio remains invariant to injected common-mode throughput within numerical precision. Further optical, processor, tomography-regression, and dense-state consistency checks are retained in the Supplemental Material [9].

The directed matrix therefore carries a normalization datum absent from its symmetric projective geometry. Exact directed information fixes generic qubit ensembles at the four-state threshold, while the scalar family identifies both subthreshold ambiguity and the source of ill-conditioning. The experimental-data analysis demonstrates a finite-setting realization of the reconstruction pipeline, not a loophole-free measurement of the continuum maximum in Eq. (10). Acquisition variants remove explicit nuisance classes: setting loss, arbitrary common-mode state-setting throughput, and, with paired swapping, unequal static channel gains. The gain is calibration robustness rather than universal sample efficiency.

The authors thank the creators of the public optical, quantum-processor, and superconducting-qubit data sets used for validation. No claim is made to ownership of those external data.

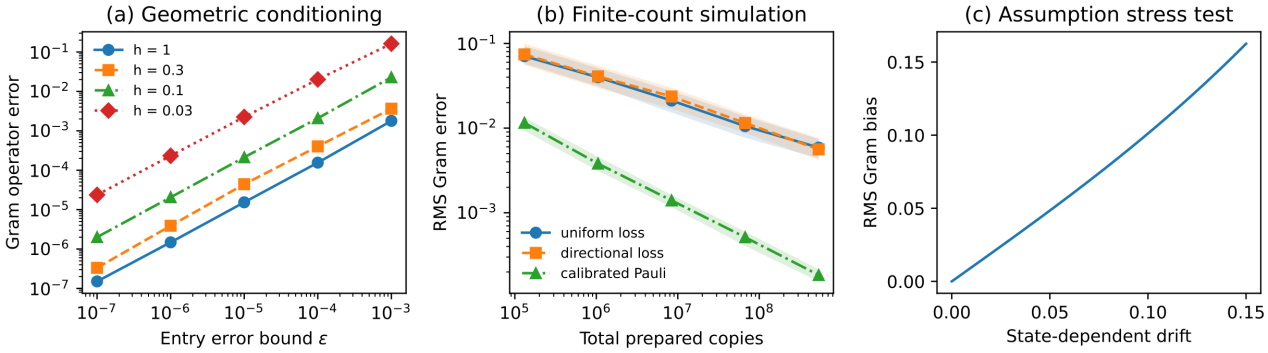


FIG. 1. **Baseline synthetic validation.** (a) Noise amplification as a faithful four-state ensemble approaches affine coplanarity. (b) Finite-count ratio reconstruction under uniform or direction-dependent shared loss, compared at equal prepared-copy budget with calibrated Pauli tomography. (c) Preparation-dependent one-port loss produces a bias floor. These tests delimit the original one-port assumption; dual-port and paired-swap estimators enlarge the admissible nuisance model.

DATA AND SOFTWARE AVAILABILITY

All author-generated numerical data, code, derived tables, and figure source data needed for the new end-to-end analysis are included in the accompanying reproducibility archive. External raw files are not redistributed; the archive records exact public sources, checksums, and analysis boundaries. The superconducting-qubit input is the public code-and-data record associated with Ref. [10] and its official repository snapshot. Other external sources are listed in the Supplemental Material.

COMPUTATIONAL ASSISTANCE DISCLOSURE

OpenAI ChatGPT (GPT-5.6 Sol) and the Codex Work interface were used substantively for proof exploration, literature checking, Python implementation, numerical stress tests, figure generation, and manuscript preparation during September 2026. The authors specified the research questions and assumptions and verified outputs using explicit derivations, generalized-eigenvalue checks, deterministic regression tests, bounded-perturbation tests, independent-data cross-checks, and reproducible simulations. The authors take full responsibility for the content.

* idrees114af@gmail.com
† alamerirehab300@gmail.com

[1] N. Datta, Min- and max-relative entropies and a new entanglement monotone, *IEEE Trans. Inf. Theory* **55**, 2816 (2009), doi:10.1109/TIT.2009.2018325.
[2] D. Reeb, M. J. Kastoryano, and M. M. Wolf, Hilbert's projective metric in quantum information theory, *J. Math. Phys.* **52**, 082201 (2011), doi:10.1063/1.3615729.
[3] Y. Lim, Hilbert's projective metric on Lorenz cones and Birkhoff formula for Lorentzian compressions, *Linear Algebra Appl.* **423**, 246 (2007), doi:10.1016/j.laa.2006.12.015.
[4] J. Park and S. Kim, Hilbert projective metric on a gyrogroup of qubit density matrices, *Rep. Math. Phys.* **76**, 389 (2015), doi:10.1016/S0034-4877(15)30039-2.
[5] P. Tabaghi and I. Dokmanić, Hyperbolic distance matrices, in *Proc. 26th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining*, 1728 (2020), doi:10.1145/3394486.3403224.
[6] A. Papadopoulos and M. Troyanov, From Funk to Hilbert geometry, in *Handbook of Hilbert Geometry*, 33 (2014), doi:10.4171/147-1/2.
[7] A. M. Brańczyk *et al.*, Self-calibrating quantum state tomography, *New J. Phys.* **14**, 085003 (2012), doi:10.1088/1367-2630/14/8/085003.
[8] J. Y. Sim, J. Shang, H. K. Ng, and B.-G. Englert, Proper error bars for self-calibrating quantum tomography, *Phys. Rev. A* **100**, 022333 (2019), doi:10.1103/PhysRevA.100.022333.
[9] See Supplemental Material for complete proofs, singular cases, finite-perturbation bounds, legacy validation protocols, end-to-end cross-fit/bootstrap details, data provenance, and the full reference audit.
[10] A. Aloy, M. Fadel, T. D. Galley, C. L. Jones, and M. P. Müller, Theory-independent monitoring of the decoherence of a superconducting qubit with generalized contextuality, *Nat. Commun.* **17**, 2474 (2026), doi:10.1038/s41467-026-69030-x; code and data, doi:10.5281/zenodo.17979407.
[11] A.-P. M. Reponen *et al.*, Broadband transient full-Stokes luminescence spectroscopy, *Nature* **643**, 675 (2025), doi:10.1038/s41586-025-09197-3.
[12] A.-P. M. Reponen *et al.*, Dataset for Broadband Transient Full-Stokes Luminescence Spectroscopy, Zenodo (2025), doi:10.5281/zenodo.15360970.

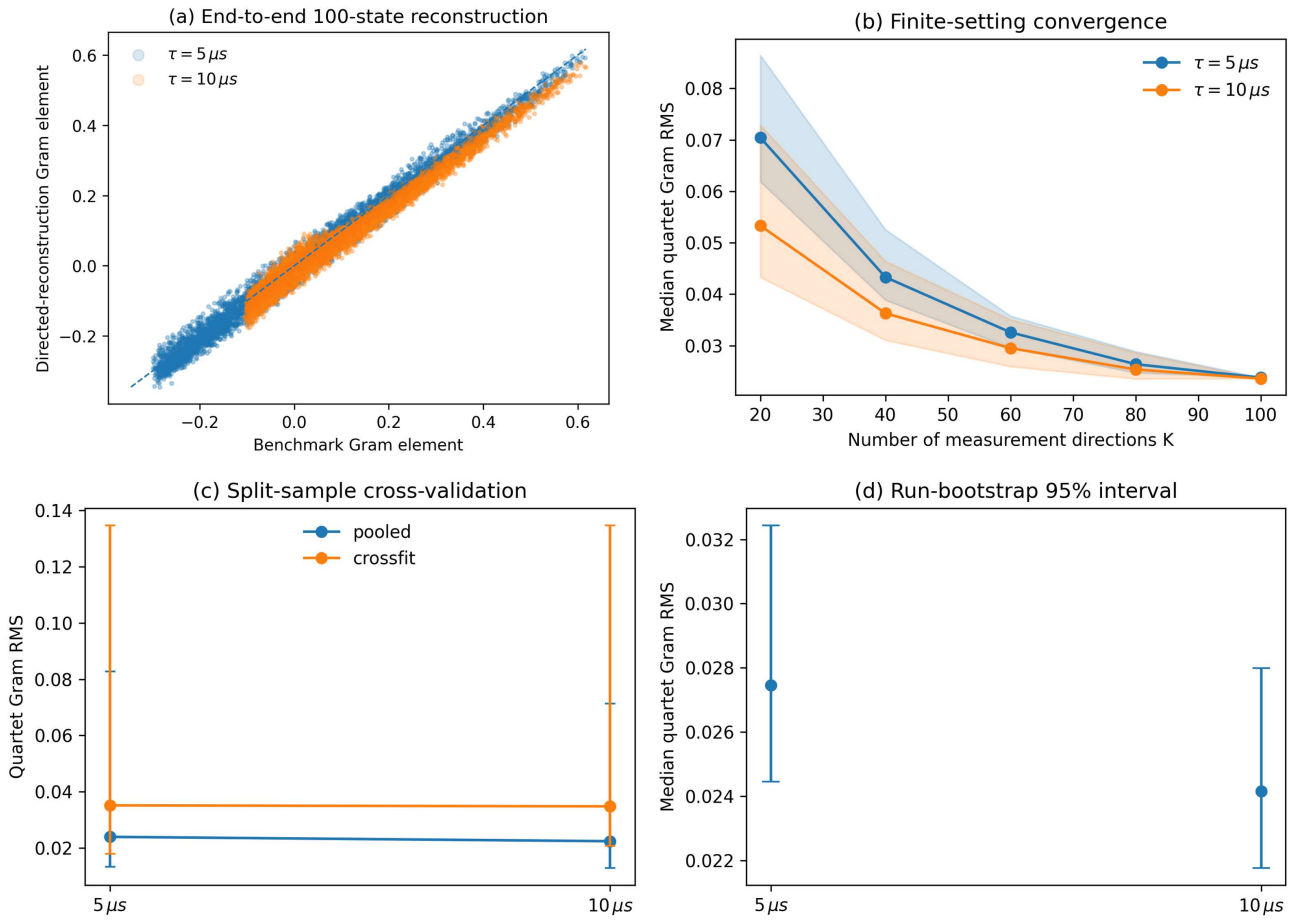


FIG. 2. **End-to-end experimental-data validation.** (a) Off-diagonal Gram elements reconstructed from pooled experimental frequency tables using the sparse directed protocol, compared only afterward with the authors' rank-4 GPT benchmark geometry at 5 and 10 μs. (b) Finite-setting convergence: median quartet Gram RMS versus the number K of randomly retained measurement directions; bands show the 10th–90th percentile over random direction subsets. (c) Pooled and strict split-sample cross-fit quartet errors; whiskers show the 10th–90th percentile over physically valid quartets. (d) Run-bootstrap 95% intervals for the median quartet error. No tomography or density-matrix estimate enters the directed reconstruction path.