

Uniform single-copy tomography of fermionic Gaussian states

Yuxuan Zhang

Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL),
CH-1015 Lausanne, Switzerland

Department of Physics, Princeton University,
Princeton, New Jersey 08544, USA

Version 1.0 — 12 October 2026

Abstract

How many preparations of a fermionic Gaussian state are needed to reconstruct it when each preparation must be measured separately? We give an adaptive protocol that learns every mixed Gaussian state on m modes to trace-norm error ε , with probability at least $2/3$, using $\tilde{O}(m^3/\varepsilon^2)$ copies. The guarantee is uniform over the covariance spectrum, including pure modes and arbitrary pairing. It improves the general quartic single-copy upper bound and, together with Rubin’s cubic lower bound, determines the worst-case copy complexity up to logarithmic factors. The main difficulty is that small unweighted covariance error need not imply small state error near pure modes. We construct a relative covariance estimate at successively finer scales, then measure in its canonical basis and project in a fixed quantum Fisher metric. A weak Gaussian attenuation channel handles the boundary at a controlled trace-norm cost. Each experiment uses one input copy and fresh ancillas; only classical information persists. The same method learns the Gaussian state associated with the covariance of any parity-even input. The bounds concern copy complexity; a finite-bit running-time guarantee is not established here.

1 Introduction

Fermionic Gaussian states describe noninteracting fermions and the states generated by fermionic linear optics. Their density matrices act on a space of dimension 2^m , but their two-point correlations form a matrix of size $2m$. This makes covariance estimation an attractive route to tomography. The statistical cost of recovering the whole state, however, depends on how covariance errors propagate to distinguishability.

That propagation is delicate near pure modes. A small change in each of many rarely occupied modes can add up to an appreciable change in the state. An estimate that treats every covariance entry on the same absolute scale can therefore spend too many copies on the wrong accuracy criterion. The problem is to learn the relevant scales without assuming a known mode basis or a lower bound on the distance from purity.

Bittel, Mele, Eisert and Leone gave a computationally efficient single-copy protocol for general mixed fermionic Gaussian states using $O(m^4\varepsilon^{-2}\log(m^2/\delta))$ copies [1, Theorem 5]. Chen and coauthors established the optimal quadratic copy complexity when measurements can be entangled across input copies [2]. Rubin proved an $\Omega(m^3/\varepsilon^2)$ lower bound for arbitrary adaptive single-copy measurements with fresh ancillas and no retained quantum memory. His cubic upper bound applies when the covariance stays a fixed distance from the boundary; Conjecture 10.1 asks for a spectrum-uniform cubic upper bound [3].

We prove that upper bound. The protocol first learns a physical covariance that is accurate in relative matrix order after adding a small regularizer. It does so by repeatedly measuring in the current estimate's mode basis. In that basis, nearly empty and nearly full sectors have correspondingly small measurement variance. This allows the regularization scale to decrease geometrically at total cost proportional to its inverse, rather than its inverse square.

The final estimate uses the same variance information in the symmetric logarithmic derivative (SLD) metric. The Gaussian SLD identity is established prior work [4]; we include a derivation to fix the normalization and handle degenerate occupations. Relative matrix-order control makes this metric comparable along the entire segment joining the true regularized covariance to its estimate. A projection in that fixed metric then converts an unbiased, possibly unphysical estimate into a physical Gaussian state without losing the error bound. A weak attenuation channel supplies the regularized input using fresh ancillas on each shot.

These ingredients remove the purity-gap dependence. Relative to Rubin's proposed occupation-sensitive strategy, the additional steps are a uniform adaptive construction of the mode basis and a global conversion from its variance-weighted covariance error to trace-norm error. The lower bound is Rubin's; the contribution here is the matching upper bound up to logarithmic factors. We allow arbitrary operations within a single copy and unrestricted classical post-processing. Neither a local-detector guarantee nor a finite-bit polynomial running time follows from the copy bound.

2 Model and result

Let $a_1, \dots, a_m$ satisfy the canonical anticommutation relations. Write

$$x = (a_1, \dots, a_m, a_1^\dagger, \dots, a_m^\dagger), \quad R_{ij} = \text{Tr}(\rho x_i^\dagger x_j).$$

With $C_{ij} = \langle a_i^\dagger a_j \rangle$ and $F_{ij} = \langle a_i^\dagger a_j^\dagger \rangle$, this is

$$R = \begin{pmatrix} C & F \\ F^\dagger & I - C^T \end{pmatrix}, \quad 0 \leq R \leq I, \quad J \bar{R} J = I - R, \quad J = \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}. \quad (1)$$

The last relation expresses particle-hole symmetry. Every matrix satisfying (1) is the covariance of a unique even Gaussian state, denoted ρ_R . Indeed, its Majorana representation is a real antisymmetric contraction; its orthogonal block form gives canonical occupations $p_i \in [0, 1]$. In the associated modes,

$$R = \text{diag}(p_1, \dots, p_m, 1 - p_1, \dots, 1 - p_m),$$

and ρ_R is the product of the Bernoulli occupation states. This includes endpoints. Canonical changes of Nambu coordinates preserve matrix order and the particle-hole relation. No ordering or nondegeneracy of the p_i is required.

We use the full trace norm $\|\rho - \sigma\|_1$; the trace distance is half of this quantity. One-copy measurements may use fresh ancillary systems and depend on earlier classical outcomes. They cannot act jointly on different input copies or retain a quantum register between copies. All complexity bounds below count input copies.

Theorem 1. *For every $m \geq 1$, $0 < \varepsilon \leq 1$, and even m -mode fermionic Gaussian state ρ_R , there is an adaptive single-copy protocol that outputs a classical physical covariance $\tilde{R}$ such that*

$$\Pr\left(\left\|\rho_R - \rho_{\tilde{R}}\right\|_1 \leq \varepsilon\right) \geq \frac{2}{3}.$$

Its worst-case copy count is

$$O\left(\frac{m^3}{\varepsilon} \log\left(24m \left\lceil \log_2 \frac{8m}{\varepsilon} \right\rceil\right) + \frac{m^3}{\varepsilon^2}\right). \quad (2)$$

The constants are universal, with no dependence on an unknown covariance gap or mode basis. All quantum operations act on one input copy and fresh ancillas, and only classical information persists between copies.

Together with the lower bound in [3, Theorem 1.1], valid in its stated small-error range, this gives worst-case single-copy complexity $\tilde{\Theta}(m^3/\varepsilon^2)$. It establishes the upper bound requested in [3, Conjecture 10.1].

3 A uniform relative covariance estimate

The preliminary estimate will control both particles and holes. This section uses only the anticommutation relations and occupation probabilities; the input need not be Gaussian.

Lemma 2. *For every covariance R of a parity-even m -mode state, $0 < \tau \leq 1$ and $0 < \eta < 1/3$, an adaptive single-copy protocol returns a physical covariance $\hat{R}$ satisfying, with probability at least $1 - \eta$,*

$$\begin{aligned} \frac{1}{2}(R + \tau I) &\leq \hat{R} + \tau I \leq \frac{3}{2}(R + \tau I), \\ \frac{1}{2}(I - R + \tau I) &\leq I - \hat{R} + \tau I \leq \frac{3}{2}(I - R + \tau I). \end{aligned} \quad (3)$$

Its worst-case cost is

$$O\left(\frac{m^2}{\tau} \log \frac{4m \max\{1, \lceil \log_2(1/\tau) \rceil\}}{\eta}\right).$$

At $\tau = 1$ no copies are needed. For $\tau < 1$, the successful comparison factors can be taken as $1 \pm 3/32$.

3.1 One measurement stage

Fix a physical pilot R_0 and work in its canonical basis. For $u > 0$, let

$$D = R_0 + uI = \text{diag}(d_1, \dots, d_m, d_{\bar{1}}, \dots, d_{\bar{m}}), \quad \bar{i} = m + i.$$

Suppose $R + uI \leq KD$. Write $q_i = \langle n_i \rangle$, where $n_i = a_i^\dagger a_i$. Measure all occupations jointly on n fresh copies. Edge-color the complete graph on m vertices into $M \leq m$ matchings, with $M = 0$ for $m = 1$. For each matching, use n fresh copies for each of the four joint spectral measurements

$$\begin{aligned} H_x^{ij} &= a_i^\dagger a_j + a_j^\dagger a_i, & H_y^{ij} &= -i(a_i^\dagger a_j - a_j^\dagger a_i), \\ P_x^{ij} &= a_i^\dagger a_j^\dagger + a_j a_i, & P_y^{ij} &= -i(a_i^\dagger a_j^\dagger - a_j a_i). \end{aligned} \quad (4)$$

Even observables on disjoint mode pairs commute. These measurements therefore use $(4M + 1)n$ copies. Their sample means give

$$\hat{C}_{ij} = \frac{1}{2}(\hat{H}_x^{ij} + i\hat{H}_y^{ij}), \quad \hat{F}_{ij} = \frac{1}{2}(\hat{P}_x^{ij} + i\hat{P}_y^{ij}).$$

Complete C by Hermitian symmetry and F by antisymmetry, use the occupation means on the top diagonal and their complements on the bottom diagonal, and denote the result by R_{raw} . It is unbiased, Hermitian and particle-hole symmetric.

Here is a matrix variance bound that does not require independence within one shot. A quadrature outcome z on a pair contributes two disjoint Hermitian edges to R_{raw} , with coefficients of magnitude $z/2$. After whitening by $D^{-1/2}$, the magnitudes are $z/(2\sqrt{d_a d_b})$. Distinct matching pairs have disjoint four-dimensional supports, so the square of the entire one-shot contribution Y is diagonal, with entries $z^2/(4d_a d_b)$.

The observables have eigenvalues in $\{-1, 0, 1\}$, and

$$\begin{aligned}(H_x^{ij})^2 &= (H_y^{ij})^2 = n_i(1 - n_j) + (1 - n_i)n_j, \\ (P_x^{ij})^2 &= (P_y^{ij})^2 = n_i n_j + (1 - n_i)(1 - n_j).\end{aligned}\tag{5}$$

Their expectations obey

$$s_H \leq \min(q_i + q_j, 2 - q_i - q_j), \quad s_P \leq \min(q_i + 1 - q_j, q_j + 1 - q_i).$$

The order assumption gives $q_i + u \leq K d_i$ and $1 - q_i + u \leq K d_i$. For each corresponding edge, the appropriate sum in these bounds is at most $K(d_a + d_b)$. Since $d_a, d_b \geq u$,

$$\mathbb{E}Y^2 \leq \frac{K}{2u}I, \quad \|Y - \mathbb{E}Y\|_{\text{op}} \leq \frac{1}{u}.$$

Also $\mathbb{E}(Y - \mathbb{E}Y)^2 = \mathbb{E}Y^2 - (\mathbb{E}Y)^2 \leq \mathbb{E}Y^2$. The centered whitened occupation contribution is diagonal, with entries $(z_i - q_i)/d_i$ and $-(z_i - q_i)/d_i$. Its norm is at most $1/u$ and its second moment at most $(K/u)I$, using $q_i(1 - q_i) \leq q_i$ or $1 - q_i$. Correlations between occupations do not enter a diagonal square.

Conditional on the pilot, different shots and batches are independent. Thus $D^{-1/2}(R_{\text{raw}} - R)D^{-1/2}$ is a sum of centered Hermitian matrices, each of norm at most $L = 1/(nu)$, with total variance parameter

$$V \leq \frac{K(2M + 1)}{nu} \leq \frac{K(2m + 1)}{nu}.$$

The Hermitian matrix Bernstein inequality [5] gives

$$\Pr\left(\left\|D^{-1/2}(R_{\text{raw}} - R)D^{-1/2}\right\|_{\text{op}} > \delta\right) \leq 4m \exp\left(-\frac{\delta^2}{2(V + L\delta/3)}\right).\tag{6}$$

In particular, failure probability at most η_s follows from

$$n \geq \frac{2K(2m + 1) + 2\delta/3}{u\delta^2} \log \frac{4m}{\eta_s}.\tag{7}$$

3.2 Physical updates and decreasing scales

Given R_{raw} , consider the classical convex feasibility problem

$$0 \leq X \leq I, \quad X = X^\dagger, \quad J\bar{X}J = I - X, \quad -\delta D \leq X - R_{\text{raw}} \leq \delta D.\tag{8}$$

If the set is empty, flag failure and output a fixed physical covariance. Otherwise choose its unique minimum-Frobenius-norm element R_{new} . The set is compact and convex, and the squared Frobenius norm is strictly convex. On the concentration event in (6), R is feasible, and

$$-2\delta D \leq R_{\text{new}} - R \leq 2\delta D.\tag{9}$$

Particle-hole conjugation gives the corresponding bound with $I - R_0 + uI$.

To prove Lemma 2, set $\delta = 1/64$. At $\tau = 1$, return $R_0 = I/2$, which satisfies (3). For $\tau < 1$, set

$$S = \lceil \log_2(1/\tau) \rceil, \quad u_s = \max\{\tau, 2^{-s}\} \quad (0 \leq s \leq S).$$

Initialize $R_0 = I/2$ at scale $u_0 = 1$. Suppose the previous pilot satisfies (3) at $u = u_{s-1}$. Its lower comparison gives $R + uI \leq 2(R_{s-1} + uI)$, so the preceding stage applies with $K = 2$, failure allowance η/S , and

$$n_s = \left\lceil \frac{4(2m+1) + 2\delta/3}{u\delta^2} \log \frac{4mS}{\eta} \right\rceil.$$

On success, (9) and the previous upper comparison bound the update error between $\pm 3\delta(R + uI)$. Since $v = u_s \geq u/2$ and $R \geq 0$, it is bounded between $\pm 6\delta(R + vI)$. Adding $R + vI$ proves the next invariant with factors $1 \pm 6\delta = 1 \pm 3/32$. Particle-hole conjugation supplies its complementary invariant. Return to the fixed original coordinates after each update.

A first-failing-stage union bound gives total failure probability at most η . The sample counts are fixed independently of outcomes, and all measurements have finitely many outcomes. Thus the history tree is finite. Fix a canonical mode basis at every feasible node and the unique optimizer in (8); this defines the protocol on all histories without an oracle or a measurability assumption. Only classical pilots connect stages. Finally,

$$\sum_{s=1}^S \frac{1}{u_{s-1}} = 2^S - 1 < \frac{2}{\tau},$$

so the total cost is at most

$$(4m+1) \left\{ \frac{2[4(2m+1) + 2\delta/3]}{\tau\delta^2} \log \frac{4mS}{\eta} + S \right\}.$$

Since $S = O(1/\tau)$, this proves the lemma.

4 From relative covariance error to state error

For an interior physical covariance R , define a positive superoperator on matrices by

$$\mathcal{S}_R(Y) = RY(I - R) + (I - R)YR.$$

For a Hermitian particle-hole tangent X , set

$$\mathcal{F}_R(X) = \text{Tr}[X\mathcal{S}_R^{-1}(X)]. \quad (10)$$

The Gaussian SLD identity identifies this with the SLD Fisher quadratic form of the density tangent $D = (d/dt)\rho_{R+tX}|_{t=0}$ [4]. We give the details used here.

Lemma 3. *For every interior R and physical tangent X , the SLD Fisher information of D is (10), and $\|D\|_1 \leq \sqrt{\mathcal{F}_R(X)}$.*

Proof. For any full-rank density matrix with eigenvalues λ_u , the equation $D = (\rho L + L\rho)/2$ gives

$$\text{Tr}(\rho L^2) = \sum_{u,v} \frac{2|D_{uv}|^2}{\lambda_u + \lambda_v}.$$

In canonical Gaussian modes the density matrix is a product of Bernoulli occupations. Its diagonal occupation score is $\sum_i dp_i(n_i - p_i)/[p_i(1 - p_i)]$, and its Fisher information is $\sum_i dp_i^2/[p_i(1 - p_i)]$.

First take generic occupations, with $p_i \neq p_j$ and $p_i + p_j \neq 1$ for all pairs. A hopping quadratic generator couples configurations 10 and 01 on a pair, with probabilities $A = p_i(1 - p_j)$ and $B = (1 - p_i)p_j$. For a spectator configuration of probability r_s , the covariance derivative is the generator coefficient times $A - B$, while the density derivative has the same factor times r_s and a fermionic sign of modulus one. Hence its entry has magnitude $|dC_{ij}|r_s$. The two directed entries in the SLD sum contribute $4|dC_{ij}|^2 r_s / (A + B)$. Summing over spectators gives $4|dC_{ij}|^2 / w_H^{ij}$, where

$$w_H^{ij} = p_i + p_j - 2p_i p_j.$$

A pairing generator couples 00 and 11 and similarly gives $4|dF_{ij}|^2 / w_P^{ij}$, where $w_P^{ij} = 1 - w_H^{ij}$. Distinct unordered pairs change distinct sets of occupation bits; hopping and pairing on the same pair connect different configurations. These contributions have disjoint matrix-element supports. At generic occupations, they and the occupation derivatives span the physical tangent space.

In the canonical Nambu basis, $\mathcal{S}_R(E_{uv}) = [r_u(1 - r_v) + (1 - r_u)r_v]E_{uv}$. Each occupation derivative occurs twice, and each hopping or pairing derivative four times in the doubled covariance. Consequently (10) is exactly the Fisher form just computed, or, with quadrature derivatives as in (4),

$$\mathcal{F}_R(X) = \sum_i \frac{dp_i^2}{p_i(1 - p_i)} + \sum_{i < j} \left\{ \frac{dH_x^2 + dH_y^2}{w_H^{ij}} + \frac{dP_x^2 + dP_y^2}{w_P^{ij}} \right\}. \quad (11)$$

The identity extends to degenerate interior occupations by continuity. Indeed, Majorana monomials form an operator basis, and their Gaussian expectations are Pfaffians of covariance submatrices. Thus $R \mapsto \rho_R$ is polynomial. Wick's rule follows for a canonical product state and then under linear Majorana transformations. Its differential is continuous; both the density SLD inverse and $\mathcal{S}_R^{-1}$ are continuous in the interior. Approximating the occupations by generic ones with the same fixed tangent proves the identity without dividing by a degeneracy. Canonical coordinate changes preserve it.

Finally, for Hermitian B with $\|B\|_{\text{op}} \leq 1$,

$$|\text{Tr}(BD)| = |\text{Re Tr}(\rho BL)| \leq \sqrt{\text{Tr}(\rho B^2) \text{Tr}(\rho L^2)} \leq \sqrt{\mathcal{F}_R(X)}.$$

Trace-norm duality proves the final assertion. $\square$

Fix an interior pilot R_0 and constants $0 < a \leq 1 \leq b$. Define

$$\mathcal{K} = \{R : R = R^\dagger, J\bar{R}J = I - R, \quad aR_0 \leq R \leq bR_0, \quad a(I - R_0) \leq I - R \leq b(I - R_0)\}. \quad (12)$$

This is a nonempty compact convex set of interior physical covariances, containing R_0 .

Lemma 4. For $U, V \in \mathcal{K}$,

$$\|\rho_U - \rho_V\|_1 \leq a^{-1} \sqrt{\mathcal{F}_{R_0}(U - V)}.$$

Proof. Column vectorization represents $\mathcal{S}_R$ by $(I - R)^T \otimes R + R^T \otimes (I - R)$. If $A \geq aA_0 \geq 0$ and $B \geq aB_0 \geq 0$, then

$$A \otimes B - a^2 A_0 \otimes B_0 = (A - aA_0) \otimes B + aA_0 \otimes (B - aB_0) \geq 0.$$

Transposition preserves Hermitian positivity, so $\mathcal{S}_R \geq a^2 \mathcal{S}_{R_0}$ throughout $\mathcal{K}$. Taking inverses yields $\mathcal{F}_R(X) \leq a^{-2} \mathcal{F}_{R_0}(X)$. The segment from U to V remains in $\mathcal{K}$. Integrate the derivative bound of Lemma 3 along this segment. Polynomial reconstruction of the density matrix ensures differentiability. $\square$

5 Variance-weighted estimation in the pilot basis

Suppose the state being measured has covariance $Q \in \mathcal{K}$. Use the measurement batches in (4), now in the canonical basis of R_0 , with n shots per batch. Let its occupations be q_i and the pilot occupations p_i . The comparison inequalities imply

$$q_i \leq bp_i, \quad 1 - q_i \leq b(1 - p_i), \quad \text{Var}(n_i) \leq b^2 p_i(1 - p_i).$$

The hopping-sector second moment in (5) obeys

$$s_H \leq b \min(p_i + p_j, 2 - p_i - p_j) \leq 2bw_H^{ij}.$$

For the last inequality, when $s = p_i + p_j \leq 1$, use $p_i p_j \leq s^2/4$ to obtain $w_H^{ij} \geq s/2$; for $s \geq 1$ use complementary occupations. Applying the same inequality to p_i and $1 - p_j$ gives $s_P \leq 2bw_P^{ij}$.

The raw covariance estimate is unbiased. Equation (11) has no cross-coordinate terms, so correlations within a shot do not affect its expected quadratic error. The four quadratures per pair and the occupation coordinates give

$$\mathbb{E}_{\mathcal{F}_{R_0}}(R_{\text{raw}} - Q) \leq \frac{b^2 m + 8b \binom{m}{2}}{n}. \quad (13)$$

These bounds hold for arbitrary input states with covariance Q .

Project R_{raw} onto $\mathcal{K}$ in the fixed Hilbert norm $\|X\|_0^2 = \mathcal{F}_{R_0}(X)$ on the real particle-hole tangent space, and call the result $\tilde{R}$. This is a unique classical strictly convex quadratic minimization over a known compact convex set. The projection is nonexpansive: adding its two variational inequalities gives

$$\|P(y) - P(z)\|_0^2 \leq \langle P(y) - P(z), y - z \rangle_0.$$

As $P(Q) = Q$, $\|\tilde{R} - Q\|_0 \leq \|R_{\text{raw}} - Q\|_0$. Combining Markov's inequality with Lemma 4 proves

$$\Pr\left(\|\rho_{\tilde{R}} - \rho_Q\|_1 > t\right) \leq \frac{b^2 m + 8b \binom{m}{2}}{na^2 t^2}. \quad (14)$$

The copy cost is at most $(4m + 1)n$. The projection uses the fixed SLD norm, not an unweighted spectral clipping rule.

6 Regularization and the complete protocol

Proof of Theorem 1. Set $\tau = \varepsilon/(8m)$ and use Lemma 2 with $\eta = 1/6$. Define

$$R_0 = \frac{\hat{R} + \tau I}{1 + 2\tau}, \quad Q = \frac{R + \tau I}{1 + 2\tau}.$$

The pilot is always interior. On a successful preliminary estimate,

$$\frac{1}{2}Q \leq R_0 \leq \frac{3}{2}Q, \quad \frac{1}{2}(I - Q) \leq I - R_0 \leq \frac{3}{2}(I - Q).$$

In particular, $Q \in \mathcal{K}$ for $a = 1/2$ and $b = 2$.

Copies of ρ_Q can be produced one at a time. Append a fresh maximally mixed m -mode Gaussian environment with covariance $I/2$. A Gaussian beam splitter maps each system annihilator to

$\sqrt{T} a_i + \sqrt{1-T} b_i$, where $T = 1/(1+2\tau)$. The environment has zero first moments and no cross covariance with the input. After discarding it,

$$C' = TC + (1-T)I/2, \quad F' = TF, \quad R' = TR + (1-T)I/2 = Q.$$

Gaussianity is preserved. This implements the required regularization even with pairing, using only one input copy and fresh ancillary modes. Equivalently, pull the final POVMs back through this fixed channel.

Because Q is affine in R , the associated Gaussian states have the same canonical modes and occupation probabilities p_i and $p'_i = (p_i + \tau)/(1+2\tau)$. Couple each Bernoulli pair using a common uniform random variable, independently across modes. A union bound gives total variation at most $\sum_i |p_i - p'_i| \leq m\tau/(1+2\tau)$. Hence, including $p_i = 0, 1$,

$$\|\rho_R - \rho_Q\|_1 \leq \frac{2m\tau}{1+2\tau} \leq \varepsilon/4. \quad (15)$$

Condition on each successful preliminary history. Apply (14) with

$$n = \lceil 1024m^2/\varepsilon^2 \rceil, \quad t = 3\varepsilon/4, \quad a = 1/2, \quad b = 2.$$

Its numerator is $4m + 16\binom{m}{2} \leq 8m^2$, so the conditional failure probability is at most $512m^2/(9n\varepsilon^2) \leq 1/18$. Together with preliminary failure at most $1/6$, total failure is at most $2/9 < 1/3$. On success, (15) and the triangle inequality prove the claimed error.

On exceptional histories the output is still well defined: flagged preliminary failure gives a fixed physical covariance, and otherwise $R_0 \in \mathcal{K}$ ensures a unique final projection. No history exceeds the prescribed budget. The preliminary cost is the first term of (2), and the final cost is at most $(4m+1)\lceil 1024m^2/\varepsilon^2 \rceil$. The argument also covers $m = 1$. All basis choices are functions of classical outcomes, and all quantum registers are discarded after their shot. $\square$

7 Learning the Gaussian covariance model of an arbitrary input

For an even state ρ , let $\mathcal{G}(\rho)$ be the unique even Gaussian state with the same covariance. The variance bounds above allow the Gaussian input promise to be removed if the target is $\mathcal{G}(\rho)$.

Corollary 5. *For every parity-even input ρ , the protocol of Theorem 1 returns a Gaussian state σ satisfying*

$$\|\sigma - \mathcal{G}(\rho)\|_1 \leq \varepsilon$$

with probability at least $2/3$ and the same copy bound. If, in addition, $D(\rho\|\mathcal{G}(\rho)) \leq \Delta$, with natural logarithms, then on that event

$$\|\rho - \sigma\|_1 \leq \varepsilon + \sqrt{2\Delta}.$$

Proof. Lemma 2 and the variance estimates (5)–(13) depend only on anticommutation relations and occupation inequalities. Attenuation of an arbitrary even input still gives covariance Q , though its output need not be Gaussian. The final samples therefore satisfy the same error bound relative to Q . Gaussian SLD geometry is applied only to the reconstructed Gaussian states. Equation (15) compares $\mathcal{G}(\rho) = \rho_R$ with ρ_Q and completes the first claim. The second follows from quantum Pinsker's inequality and the triangle inequality. $\square$

The relative-entropy promise is the same notion of near-Gaussianity used by Bittel and coauthors [1, Proposition 6]; the improvement here is the cubic copy bound. For completeness, the identity

$$D(\rho\|\mathcal{G}(\rho)) = S(\mathcal{G}(\rho)) - S(\rho)$$

also holds with pure covariance modes. In canonical modes, $\langle n_i \rangle = 0$ or 1 forces ρ to have that fixed occupation, by positivity. Thus ρ is supported on the support of $\mathcal{G}(\rho)$. On the remaining modes, $\log \mathcal{G}(\rho)$ is a constant plus a quadratic occupation operator, so it has the same expectation in the two states. This proves the identity.

Small trace-norm contamination alone behaves differently. Let

$$\rho_d = (1 - d)|\text{vac}\rangle\langle\text{vac}| + d|\text{full}\rangle\langle\text{full}|, \quad 0 \leq d \leq 1.$$

Its covariance has all occupations equal to d , so $\mathcal{G}(\rho_d)$ is a product of Bernoulli occupations. Comparing the diagonal probabilities gives exactly

$$\|\rho_d - \mathcal{G}(\rho_d)\|_1 = 2[1 - (1 - d)^m - d^m], \quad \|\rho_d - |\text{vac}\rangle\langle\text{vac}|\|_1 = 2d.$$

For fixed $c > 0$ and $d = c/m$, the first tends to $2(1 - e^{-c})$ while the second tends to zero. Thus Gaussianification does not have dimension-independent trace-noise robustness around Gaussian inputs. This example concerns that target map, not every possible learner.

8 Scope and implementation

The result specifies finite measurements, a finite adaptive history tree and physical estimates on every history. The convex updates are well defined, but the theorem treats their classical solution and the choice of canonical mode bases exactly. An unconditional finite-bit running-time analysis, including certified optimization and measurement synthesis, remains to be supplied. The sample bound is consequently not a claim of an immediately efficient laboratory implementation.

The protocol allows general within-copy measurements. It does not settle the optimal rate for restricted local detectors, or whether a nonadaptive scheme attains the same spectrum-uniform bound. These are distinct measurement-resource questions. Its main conclusion is that approaching pure modes does not force the general single-copy worst-case rate above cubic, up to logarithmic factors.

Author and AI contributions

The author set the research directions, designed the harness that coordinated model tasks, tools, memory, and checking, supplied physical and mathematical intuitions and proof skeletons, and guided the development and revision of the research ideas. AI models developed detailed candidate arguments and derivations, searched the literature, wrote and ran computational checks, and drafted and revised the manuscript under this supervision.

The author takes responsibility for the mathematical claims, source attribution, and final manuscript. Internal proof reviews and finite numerical diagnostics accompany the source release; they are not external peer review or formal proof certification.

References

- [1] L. Bittel, A. A. Mele, J. Eisert, and L. Leone, *Optimal trace-distance bounds for free-fermionic states: Testing and improved tomography*, PRX Quantum **6**, 030341 (2025). [doi:10.1103/prx6-nkfb](https://doi.org/10.1103/prx6-nkfb); [arXiv:2409.17953v3](https://arxiv.org/abs/2409.17953).
- [2] S. Chen, M. Fanizza, F. Girardi, L. Lami, A. A. Mele, M. Walter, and F. Witteveen, *Optimal tomography of bosonic and fermionic Gaussian states* (2026). [arXiv:2607.11847v1](https://arxiv.org/abs/2607.11847).
- [3] R. Rubin, *Entangled measurements are necessary for optimal tomography of mixed fermionic Gaussian states and of bosonic Gaussian states near the vacuum* (2026). [arXiv:2609.23189v1](https://arxiv.org/abs/2609.23189).
- [4] A. Carollo, B. Spagnolo, and D. Valenti, *Symmetric logarithmic derivative of fermionic Gaussian states*, Entropy **20**, 485 (2018). [doi:10.3390/e20070485](https://doi.org/10.3390/e20070485).
- [5] J. A. Tropp, *User-friendly tail bounds for sums of random matrices*, Foundations of Computational Mathematics **12**, 389–434 (2012). [doi:10.1007/s10208-011-9099-z](https://doi.org/10.1007/s10208-011-9099-z).