

Efficient Learning of Quantum States Prepared With Few Fermionic Non-Gaussian Gates

Antonio Anna Mele^{1,*} and Yaroslav Herasymenko^{2,3,4,†}

¹*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin 14195, Germany*

²*QuSoft and CWI, Science Park 123, Amsterdam 1098 XG, Netherlands*

³*QuTech, TU Delft, P.O. Box 5046, Delft 2600 GA, Netherlands*

⁴*Delft Institute of Applied Mathematics, TU Delft, Delft 2628 CD, Netherlands*



(Received 21 August 2024; revised 24 November 2024; accepted 5 December 2024; published 28 January 2025)

The experimental realization of increasingly complex quantum states underscores the pressing need for new methods of state learning and verification. In one such framework, quantum state tomography, the aim is to learn the full quantum state from data obtained by measurements. Without prior assumptions on the state, this task is prohibitively hard. Here, we present an efficient algorithm for learning states on n fermion modes prepared by any number of Gaussian and at most t non-Gaussian gates. By Jordan-Wigner mapping, this also includes n -qubit states prepared by nearest-neighbor matchgate circuits with at most t SWAP gates. Our algorithm is based exclusively on single-copy measurements and produces a classical representation of a state, guaranteed to be close in trace distance to the target state. The sample and time complexity of our algorithm is $\text{poly}(n, 2^t)$; thus if $t = \mathcal{O}(\log(n))$, it is efficient. We also show that, if t scales *slightly* more than logarithmically, any learning algorithm to solve the same task must be inefficient, under common cryptographic assumptions. We also provide an efficient property-testing algorithm that, given access to copies of a state, determines whether such a state is far or close to the set of states for which our learning algorithm works. In addition to the outputs of quantum circuits, our tomography algorithm is efficient for some physical target states, such as those arising in time dynamics and low-energy physics of impurity models. Beyond tomography, our work sheds light on the structure of states prepared with few non-Gaussian gates and offers an improved upper bound on their circuit complexity, enabling an efficient circuit-compilation method.

DOI: [10.1103/PRXQuantum.6.010319](https://doi.org/10.1103/PRXQuantum.6.010319)

I. INTRODUCTION

Quantum state tomography is the task of reconstructing a classical description of a quantum state from experimental data [1,2]. Beyond its foundational significance in quantum information theory, it stands as the gold standard for verification and benchmarking of quantum devices [2]. However, in the absence of any prior assumptions on the state to be learned, one encounters necessarily the *curse of dimensionality* of the Hilbert space: learning the classical description of a generic quantum state demands resources that grow exponentially with the number of qubits [1,3]. Simply storing and outputting the density

matrix of a state already results in an exponential cost in time. This raises the crucial question of identifying interesting classes of quantum states that can be efficiently learned using a number of state copies and time scaling at most *polynomially* with the system size. Only a few classes of states are currently known to be efficiently learnable—in particular, matrix-product states [2,4], finitely correlated states [5], high-temperature Gibbs states [6], states prepared by shallow quantum circuits [7–9], stabilizer states [10], quantum phase states [11], and fermionic Gaussian states [12,13]. The latter class of states comprises those prepared by fermionic Gaussian circuits [14], also referred to as free fermionic (noninteracting) evolutions or fermionic linear-optics circuits [15,16]. Via Jordan-Wigner mapping, such states on n fermionic modes can also be viewed as n -qubit states, prepared by generalized matchgate circuits [16–18]. Fermionic Gaussian states play a key role in condensed-matter physics and quantum chemistry, via the Hartree-Fock method and in the context of Fermi liquid and Bardeen-Cooper-Schrieffer theories [19–22]. These states are also essential in understanding many

*Contact author: a.mele@fu-berlin.de

†Contact author: yaroslav@cw.nl

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

exactly solvable spin models [23–25]. In quantum computing, fermionic Gaussian states are primarily recognized for their efficient classical simulability [15,17,18]. As in the case of Clifford circuits, for which the introduction of magic gates, such as T gates, allows universal quantum computation to be reached [26], also for the case of Gaussian circuits the inclusion of certain magic gates [27–29], for example, SWAP gates [28], allows universality to be reached. If the number t of T gates in a Clifford circuit is low, the resulting states can still be efficiently simulated classically [30–32]; it has also been recently demonstrated that such states, termed as t -doped stabilizer states [33,34], are still efficiently learnable [35–37]. Similarly, in the past year, it has been shown that Gaussian circuits with a few magic gates are also classically simulable [38–40]. However, the learnability of such “ t -doped fermionic Gaussian states” remains unknown and this motivates the core question of our work:

Can we efficiently learn states prepared by Gaussian operations (e.g., matchgates) and a few magic gates?

We answer it by proposing a quantum algorithm of polynomial time and sample complexity that uses only single-copy measurements and learns a succinct classical description of a t -doped fermionic Gaussian state; the learned state is guaranteed to be close to the true state in trace distance. Our presentation is framed in the language of qubits, but the results seamlessly translate into the fermionic formalism. Our learning algorithm may also be feasible to implement in near-term fermionic analog quantum simulators [41,42], like cold atoms in optical lattices [43], since we utilize only time evolutions of simple few-body fermionic Hamiltonians [44]. The core of our algorithm relies on a result of independent interest, elucidating the structure of states in question. In particular, for any t -doped fermionic Gaussian state $|\psi\rangle$ we show that there exists a Gaussian operation G such that $G^\dagger |\psi\rangle = |\phi\rangle \otimes |0^{n-\kappa t}\rangle$, where $|\phi\rangle$ is supported on κt qubits and κ is a small constant. Informally, this says that all the magic (non-Gaussianity) of such states can be *compressed* to a few qubits via a Gaussian operation. The proof of our compression theorem is constructive, which has implications for the circuit complexity of $|\psi\rangle$ and for improved preparation of doped fermionic Gaussian states.

The high-level idea of the learning algorithm is to first learn a Gaussian unitary G , which *compresses* the magic, apply it to the state, and then perform full-state tomography on the first few qubits alone. Our learning algorithm has a time complexity $\mathcal{O}(\text{poly}(n, 2^t))$, i.e., it scales polynomially in the system size n and exponentially in the number of non-Gaussian gates t . Thus it is efficient as long as the number of non-Gaussian gates is $t = \mathcal{O}(\log(n))$. Furthermore, we establish that the task of learning such states is computationally intractable when the number of non-Gaussian gates scales *slightly*

more than logarithmically, under a common cryptography assumption [46–49]. We show the latter result using the theory of pseudorandom quantum states [50,51] and qubit-to-fermion mappings [25]. In doing that, we bring pseudorandom quantum states, so far explored only for qubit-based systems, to the fermionic realm. Our learning algorithm generalizes the one presented by Aaronson *et al.* [12], which is tailored to learn only those states prepared by *particle-number conserving* Gaussian gates and $t = 0$ (in our work we relax both of these assumptions).

Furthermore, our algorithm extends to all *Gaussian-compressible* states, i.e., those states, which can be written as $G(\sigma \otimes |0^{n-t}\rangle\langle 0^{n-t}|)G^\dagger$, where G is a Gaussian unitary and σ is a possibly mixed quantum state supported on the first t qubits. We also propose an efficient method to *test* if a given state is close or far from the set of compressible states, by showing an efficiently estimatable quantity that lower bounds the distance to this set. Moreover, we also demonstrate that our learning algorithm can learn states that are close to being compressible—this feature is particularly significant as it emphasizes the noise robustness of the algorithm. Additionally, we note that ground states of impurity models, a well-regarded class of quantum states in condensed-matter physics [45,52–54], are *approximately* compressible, making our algorithm suitable for efficient tomography of such physically relevant states. We further provide numerical evidence showing that states prepared by time evolutions governed by impurity model Hamiltonians [45] remain approximately compressible until constant evolution times, thus making them learnable by our tomography algorithm.

It should be noted that the concept of magic compression was first introduced in the context of Clifford+ T circuits by Leone *et al.* and Oliviero *et al.* [55,56] and later exploited for learning t -doped stabilizer states [35,36]. Our strategy of proving non-Gaussianity compression and applying it to quantum state tomography was inspired by these earlier works. It is an intriguing fact that a similar compression theorem holds in our context, even though the mathematical structures of stabilizer states and fermionic Gaussian states appear quite different.

See Fig. 1 for the visual summary of our results. In the next sections we summarize our findings, stating more precisely our results and the essential ideas that underlie them. In the Appendix, we provide the technical details.

A. Preliminaries

Our work can be applied to two distinct and naively separate settings: a system of n qubits with one-dimensional (1D) matchgates circuits and their magic gates (e.g., SWAP gates), or a native fermionic system of n modes with states prepared by fermionic Gaussian evolutions and local non-Gaussian evolutions. These two perspectives are mathematically related through the Jordan-Wigner mapping.

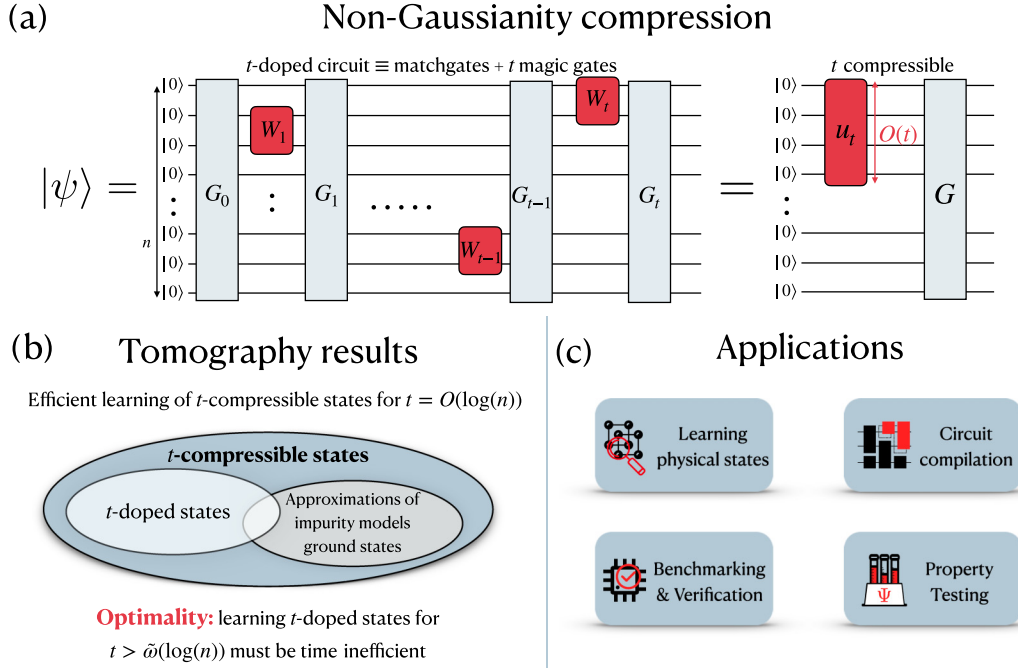


FIG. 1. A visual summary of our work. (a) We show that any t -doped state, i.e., a state prepared by an arbitrary number of Gaussian layers and at most t local fermionic non-Gaussian gates (i.e., a t -doped matchgate circuit), can be equivalently prepared by a single non-Gaussian unitary u_t acting solely on $O(t)$ fermionic modes (or qubits) and a global Gaussian unitary G . This shows that non-Gaussianity can be compressed into a localized part of the system. (b) We provide a tomography algorithm to learn such t -compressible states, i.e., states of the form $G(|\phi\rangle \otimes |0^{n-t}\rangle)$, where G is an arbitrary Gaussian unitary and $|\phi\rangle$ is an arbitrary quantum state supported solely on t modes. The algorithm is efficient as long as $t = O(\log(n))$. This also directly implies an algorithm to learn t -doped states. Furthermore, up to common cryptographic assumptions, we prove that any tomography algorithm capable to learn a state prepared by slightly more than $\log(n)$ gates must be computationally inefficient. Moreover, we show that physical states, such as ground states of impurity models [45], are approximately $O(\log(n))$ compressible and thus efficiently learnable by our algorithm. (c) Our results have various applications, including learning physical states; efficient circuit compilation (implied by the non-Gaussianity compression); benchmarking and verification of digital and analog quantum devices; and testing if the underlying unknown state is t compressible—of the form for which our algorithm is efficient.

We will use it now as a *definition* of Majorana operators, thus directly aligning our discussion with the qubit language. Majorana operators, denoted as γ_{2k-1} and γ_{2k} for $k \in [n] := \{1, \dots, n\}$, are defined in terms of standard Pauli operators as $\gamma_{2k-1} := (\prod_{j=1}^{k-1} Z_j) X_k$ and $\gamma_{2k} := (\prod_{j=1}^{k-1} Z_j) Y_k$. Alternatively, they can be defined in the fermionic language through their anticommutation relations [15,57]. A fermionic Gaussian unitary G is a unitary that satisfies $G^\dagger \gamma_\mu G = \sum_{\nu=1}^{2n} O_{\mu,\nu} \gamma_\nu$ for any $\mu \in [2n]$, where $O \in O(2n)$ is an orthogonal matrix. The product of two Gaussian unitaries is Gaussian. Notably, a one-to-one correspondence exists between Gaussian unitaries up to a global phase and $O(2n)$ orthogonal matrices. Given an orthogonal matrix, it is known how to exactly implement the associated Gaussian unitary using $O(n^2)$ 2-local qubits or 2-local fermionic Gaussian operations [38,58,59]. A pure fermionic Gaussian state can be defined as $|\psi\rangle = G|0^n\rangle$, where G is a Gaussian unitary and $|0^n\rangle$ denotes the zero computational basis state. Given a quantum state ρ , its correlation matrix $C(\rho)$ is defined

as the real antisymmetric $2n \times 2n$ matrix with elements $[C(\rho)]_{j,k} := -\frac{i}{2} \text{Tr}(\gamma_j \gamma_k \rho)$, for any $j < k \in [2n]$. We have that $C(G\rho G^\dagger) = OC(\rho)O^T$, for any Gaussian unitary G associated with $O \in O(2n)$. A well-known result in linear algebra [60] asserts that any real antisymmetric matrix C can be decomposed in the so-called “normal form”:

$$C = O \bigoplus_{j=1}^n \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix} O^T, \quad (1)$$

where O is an orthogonal matrix in $O(2n)$ and $\lambda_j \geq 0$, for any $j \in [n]$, are dubbed as “normal” eigenvalues, ordered in increasing order. We denote the trace distance between two quantum states $|\psi\rangle$ and $|\phi\rangle$ as $d_{\text{tr}}(|\psi\rangle, |\phi\rangle) := \frac{1}{2} \|\psi\langle\psi| - |\phi\rangle\langle\phi|\|_1$. Given a matrix A , its operator norm $\|A\|_\infty$ is defined as its largest singular value. We refer to the Appendix for more preliminaries.

B. Structure of t -doped Gaussian states

States prepared by Gaussian circuits applied to a computational basis state are efficiently simulable classically. However, by incorporating “non-Gaussian”, or “magic” operations, such as SWAP gates [17,28], one can render Gaussian circuits universal for quantum computation. The term “magic gate” comes from a loose parallel to Clifford circuits, which are efficiently simulable per se but become universal upon introduction of “magic” non-Clifford T gates.

Here we consider non-Gaussian operations generated by κ Majorana operators $\{\gamma_{\mu(r)}\}_{r=1}^{\kappa}$, where $\mu(1), \dots, \mu(\kappa) \in [2n]$. Examples of such non-Gaussian operations for $\kappa = 4$ are the SWAP gate or a unitary $\exp(i\theta\gamma_1\gamma_5\gamma_6\gamma_8)$ for $\theta \in \mathbb{R}$; for $\kappa = 3$, an example is $\exp(\theta\gamma_2\gamma_6\gamma_7)$. We refer to κ as the maximum Majorana locality of the employed non-Gaussian gates.

Definition 1 (t -doped fermionic Gaussian state). A state $|\psi\rangle$ is a (t, κ) -doped Gaussian state if it can be prepared by Gaussian unitaries $\{G_i\}_{i=0}^t$ and t non-Gaussian κ -local gates $\{W_i\}_{i=1}^t$, specifically

$$|\psi\rangle = G_t W_t \cdots G_1 W_1 G_0 |0^n\rangle, \quad (2)$$

where κ local means that each non-Gaussian gate involves at most κ Majorana operators. Informally, a state is t -doped Gaussian if it is (t, κ) -doped Gaussian for some fixed constant κ .

Similarly, we denote the unitary $U_t := G_t W_t \cdots G_1 W_1 G_0$ as a t -doped Gaussian unitary. We now present our main result concerning the structure of t -doped Gaussian states: it is possible to compress all the “non-Gaussianity” of the state into a localized region of the system via a Gaussian operation. This motivates the following definition.

Definition 2 (t -compressible Gaussian state). Let $t \in [n]$. A state $|\psi\rangle$ is (Gaussian) t compressible if and only if

$$|\psi\rangle = G(|\phi\rangle \otimes |0^{n-t}\rangle), \quad (3)$$

where G is a Gaussian operation, and $|\phi\rangle$ is a state supported solely on the first t qubits.

In the following, we assume $\kappa t \leq n$.

Theorem 1 (Magic compression in t -doped Gaussian states). Any (t, κ) -doped Gaussian state is κt compressible.

Proof. Let $U_t |0^n\rangle$ be the t -doped state, where $U_t = (\prod_{i=1}^t G_i W_i) G_0$ is the t -doped unitary. We rearrange U_t as $U_t = \tilde{G}_t G_{\text{aux}} \prod_{i=1}^t (G_{\text{aux}}^\dagger \tilde{W}_i G_{\text{aux}}) G_{\text{aux}}^\dagger$, introducing a

Gaussian operation G_{aux} to be fixed and defining $\tilde{W}_i := \tilde{G}_{i-1}^\dagger W_i \tilde{G}_{i-1}$ and $\tilde{G}_i := G_i \cdot G_0$. We require that G_{aux} satisfies $G_{\text{aux}}^\dagger |0^n\rangle = |0^n\rangle$, and that each $G_{\text{aux}}^\dagger \tilde{W}_i G_{\text{aux}}$ is supported nontrivially only on the first κt qubits. The latter is enforced by demanding that the Heisenberg evolution, via $\tilde{G}_{i-1} G_{\text{aux}}$, of each Majorana operator involved in the Hamiltonian generating W_i , has nontrivial support exclusively on the first κt qubits. The existence of G_{aux} is shown by demonstrating the existence of its associated orthogonal matrix O_{aux} . The requirements on G_{aux} translate into the demand that O_{aux} must be symplectic and such that it sends κt fixed vectors to the span of the first $2\kappa t$ canonical basis vectors. The existence of such O_{aux} can be proven via the isomorphism between real $2n \times 2n$ symplectic orthogonal matrices and $n \times n$ unitaries [61]. Additional details are provided in the Appendix (see Theorem 5). ■

This compressibility should not be confused with the compressed quantum computation result described in Ref. [62], which pertains to free-fermionic circuits ($t = 0$) and demonstrates how such computations can be equivalently expressed in logarithmic space.

Note that while a (t, κ) -doped Gaussian state is a κt -compressible Gaussian state, the reverse implication does not hold due to circuit complexity arguments. Similarly to Theorem 1, we show that any t -doped Gaussian unitary can be represented as

$$U_t = G_A(u_t \otimes I) G_B, \quad (4)$$

where G_A and G_B denote Gaussian operations, and u_t is a unitary operator supported on $\lceil \kappa t/2 \rceil$ qubits (with $\lceil \cdot \rceil$ denoting rounding to the next integer), as elaborated in the Appendix (Theorem 4). Notably, if U_t is a particle-number-conserving unitary [15], then G_A , u_t , and G_B can also be chosen as such.

Our proof of Theorem 1 is constructive, i.e., given a classical description of the circuit that prepares $|\psi\rangle$, it provides an efficient procedure for finding the compressing Gaussian circuit G and the state $|\phi\rangle$. The decomposition of t -doped Gaussian states (unitaries) reveals also that they have a circuit complexity, i.e., number of local gates needed for implementing the state (unitary), upper bounded by $\mathcal{O}(n^2 + t^3)$ (Proposition 3 in the Appendix). This provides a better circuit complexity upper bound compared to the naive $\mathcal{O}(n^2 t)$ implied by Definition 1 for $\kappa = \mathcal{O}(1)$. Hence, our construction reveals also a method to compress the circuit depth (and not only the magic), which might be used in practice for efficient circuit compilation of doped matchgate circuits. Remarkably, analogous results hold for the Clifford+ T gate circuits [56].

In the context of Clifford operations, the notions of *stabilizer dimension* [35] and *stabilizer nullity* [63,64] were introduced to quantify the degree of stabilizerness of a

quantum state, and their connection to magic monotones was investigated. Analogously, we define the *Gaussian dimension* of a state as the number of normal eigenvalues of its correlation matrix that are equal to one, and the *Gaussian nullity* as the number of normal eigenvalues of its correlation matrix that are strictly less than one.

Using Eq. (1), we can show that a state has a Gaussian nullity of at most t (a Gaussian dimension of at least $n - t$) if and only if it is t compressible. Consequently, (t, κ) -doped Gaussian states have a Gaussian nullity of at most κt (a Gaussian dimension of at least $n - \kappa t$).

C. Learning algorithm

We present an algorithm for learning t -compressible Gaussian states, or, equivalently, quantum states with at least $n - t$ Gaussian dimension. Note that this is a broader class than t -doped Gaussian states; as an example unrelated to t -doped states, ground states of quantum impurity models are approximately of this form (as shown in Ref. [45] and elaborated further in the subsequent subsections). By definition, any t -compressible Gaussian state $|\psi\rangle$ can be factorized as $G^\dagger |\psi\rangle = |\phi\rangle \otimes |0^{n-t}\rangle$, where G^\dagger is Gaussian and $|\phi\rangle$ is a state on t qubits. At a high level, our strategy is to learn the Gaussian unitary G^\dagger , apply it to $|\psi\rangle$, and then perform full-state tomography solely on the first t qubits to learn $|\phi\rangle$. Since full-state tomography algorithms scale exponentially with the number of qubits [65], for $t = \mathcal{O}(\log(n))$ our algorithm will be efficient.

To delve deeper, the initial phase of our learning algorithm entails estimating the correlation matrix entries through single-copy measurements. This can be achieved using different methods outlined in the Appendix, such as measurements in the Pauli basis, global Clifford Gaussian measurements [66], or fermionic classical shadows

[58,67,68]. The estimated correlation matrix \hat{C} is subsequently transformed into its normal form in Eq. (1) to yield the corresponding orthogonal matrix \hat{O} associated with the Gaussian operation \hat{G} . (We use the hat symbol to denote the objects estimated from the measurements.) Applying the inverse operation \hat{G}^\dagger to $|\psi\rangle$ results in a state that exhibits high fidelity with a state, which is tensor product of an arbitrary state on the first t qubits and the zero computational basis state on the remaining $n - t$ qubits. Consequently, the learning algorithm queries multiple copies of $|\psi\rangle$ (one at a time), applies \hat{G}^\dagger to them and measures the last $n - t$ qubits. If the outcome of such measurements correspond to $|0^{n-t}\rangle$, then the algorithm proceeds with a step of pure-state tomography [65,69] on the t -qubits state. The state tomography routine performed in the compressed space yields the state $|\hat{\phi}\rangle$. The final output of the learning algorithm is $|\hat{\psi}\rangle := \hat{G}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle)$, and an efficient classical representation can be provided if $t = \mathcal{O}(\log(n))$. Namely, to specify $|\hat{\psi}\rangle$, it is sufficient to provide the complete description of the t -qubit state $|\hat{\phi}\rangle$ and the orthogonal matrix $\hat{O} \in \text{O}(2n)$ associated with \hat{G} .

We now present a theorem, which formalizes and proves the efficiency of the discussed procedure, outlined in Algorithm 1, to learn doped Gaussian states or, more generally, t -compressible Gaussian states.

Theorem 2 (Learning algorithm guarantees). Let $|\psi\rangle$ be a t -compressible Gaussian state, and $\varepsilon, \delta \in (0, 1]$. Utilizing $\mathcal{O}(\text{poly}(n, 2^t))$ single-copy measurements and computational time, Algorithm 1 outputs a classical representation of a state $|\hat{\psi}\rangle$, such that $d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq \varepsilon$ with probability $\geq 1 - \delta$.

ALGORITHM 1. Learning algorithm for t -compressible fermionic Gaussian states.

Input: Accuracy ε , failure probability δ , $N := \lceil \frac{256n^5}{\varepsilon^4} \log\left(\frac{12n^2}{\delta}\right) + 2N_{\text{tom}}(t, \frac{\varepsilon}{2}, \frac{\delta}{3}) + 24 \log(\frac{3}{\delta}) \rceil$ copies of the t -compressible state $|\psi\rangle$, where N_{tom} is the number of copies needed for t -qubit pure state tomography with accuracy $\frac{\varepsilon}{2}$ and failure probability $\frac{\delta}{3}$.

Output: A classical description of $|\hat{\psi}\rangle$, ensuring $d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq \varepsilon$ with probability at least $1 - \delta$.

- 1 Estimate the correlation matrix of $|\psi\rangle$ using $\lceil \frac{256n^5}{\varepsilon^4} \log\left(\frac{12n^2}{\delta}\right) \rceil$ single-copy measurements (see Lemma 38), obtaining \hat{C} ;
 - 2 Express \hat{C} in its normal form $\hat{C} = \hat{O} \hat{\Lambda} \hat{O}^T$ (Eq.(1)) and find the Gaussian unitary $G_{\hat{O}}$ associated with $\hat{O} \in \text{O}(2n)$;
 - 3 **for** $i \leftarrow 1$ **to** $\lceil 2N_{\text{tom}}(t, \frac{\varepsilon}{2}, \frac{\delta}{3}) + 24 \log(\frac{3}{\delta}) \rceil$ **do**
 - 4 Apply $G_{\hat{O}}^\dagger$ to $|\psi\rangle$;
 - 5 Measure the last $n - t$ qubits in the computational basis;
 - 6 **if** the outcome corresponds to $|0^{n-t}\rangle$ **then**
 - 7 Proceed;
 - 8 **else**
 - 9 Discard and move to the next iteration;
 - 10 Perform a step of pure state tomography [65] on the remaining t qubits;
 - 11 Obtain the t -qubit state $|\hat{\phi}\rangle$ from tomography;
 - 12 **return** \hat{O} and $|\hat{\phi}\rangle$, which identify $|\hat{\psi}\rangle := G_{\hat{O}}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle)$;
-

Proof sketch. Using $\mathcal{O}(\text{poly}(n))$ copies of $|\psi\rangle$, we estimate its correlation matrix C , yielding \hat{C} such that $\|\hat{C} - C\|_\infty \leq \varepsilon_c$ with a failure probability $\leq \delta/3$, where $\varepsilon_c := \varepsilon^2/(4(n-t))$. Expressing \hat{C} in its normal form [Eq. (1)], we find the Gaussian unitary \hat{G} associated to $\hat{O} \in \text{O}(2n)$. Let $|\psi'\rangle := \hat{G}^\dagger |\psi\rangle$. As detailed in the Appendix, we derive $\langle \psi' | Z_k | \psi' \rangle \geq 1 - 2\varepsilon_c$ for each $k \in \{t+1, \dots, n\}$ and, by quantum union bound [70] we get $d_{\text{tr}}(|\phi\rangle \otimes |0^{n-t}\rangle, |\psi'\rangle) \leq \varepsilon/2$, where $|\phi\rangle \otimes |0^{n-t}\rangle$ corresponds to the state obtained by measuring the last $n-t$ qubits of $\hat{G}^\dagger |\psi\rangle$ in the computational basis and obtaining the outcome corresponding to $|0^{n-t}\rangle$, an event that occurs with probability $\geq 1 - \varepsilon^2/4$. By querying $\lceil 2N_{\text{tom}}(t, \varepsilon/2, \delta/3) + 24 \log(3/\delta) \rceil$ copies of $|\psi\rangle$, and, for each copy, applying \hat{G}^\dagger and measuring the last $n-t$ qubits, we get the outcome $|0^{n-t}\rangle$ at least $N_{\text{tom}}(t, \varepsilon/2, \delta/3)$ times, with failure probability $\leq \delta/3$ due to Chernoff bound. Here, $N_{\text{tom}}(t, \varepsilon/2, \delta/3)$ is the number of copies sufficient for full-state tomography [65] of a t -qubit state with an $\varepsilon/2$ accuracy and a failure probability $\leq \delta/3$. Performing the t -qubit tomography on all the copies where the outcome $|0^{n-t}\rangle$ occurred yields $|\hat{\phi}\rangle$ such that $d_{\text{tr}}(|\hat{\phi}\rangle, |\phi\rangle) \leq \varepsilon/2$, with a failure probability $\leq \delta/3$. Defining $|\hat{\psi}\rangle := \hat{G}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle)$, we have $d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq d_{\text{tr}}(|\hat{\phi}\rangle, |\phi\rangle) + d_{\text{tr}}(|\phi\rangle \otimes |0^{n-t}\rangle, \hat{G}^\dagger |\psi\rangle)$. This is $\leq \varepsilon$ if the algorithm does not fail, an event occurring with probability $\geq 1 - \delta$ due to the union bound. ■

Theorem 2 is restated and rigorously proven in the Appendix as Theorem 6. The sample, time and memory complexity of our algorithm for learning t -compressible states exhibits a polynomial dependence on n and an exponential dependence on t : specifically, the $\text{poly}(n)$ contribution [specifically an $\mathcal{O}(n^5)$ scaling] arises solely from estimating and postprocessing the correlation matrix, while the $\exp(t)$ contribution arises from full-state tomography on t qubits. It is easy to see that the dependence on t is optimal, because learning t -compressible states is at least as hard as learning an arbitrary pure state on t qubits and thus requires at least $\exp(\Omega(t))$ copies of the state [3].

However, if we focus on the subclass of t -doped Gaussian states, a classical shadow-tomography-based algorithm presented in Refs. [71,72] achieves $\mathcal{O}(\text{poly}(n, t))$ sample complexity. Specifically, this algorithm requires a number of copies that scales polynomially with the circuit complexity of the state, and t -doped states have a circuit complexity $\mathcal{O}(\text{poly}(n, t))$. However, the time complexity of the algorithm in Refs. [71,72] scales exponentially with the number of qubits n , while our algorithm's time complexity scales only polynomially (although always exponentially in t). This observation also applies to t -doped stabilizer states learning analyzed in recent works [35,36].

In our Appendix, we extend our learning algorithm to handle mixed states. Specifically, we provide an algorithm

to learn, in trace distance, possibly mixed quantum states that have at least $n-t$ normal eigenvalues of their correlation matrix equal to one. In the more general mixed-state scenario, the algorithm in Table 1 becomes significantly simpler. Notably, measuring the last $n-t$ qubits and post-selecting on the outcome $|0^{n-t}\rangle$ is not necessary, as we do not require the output state to be pure. Thus, it is sufficient to perform full state tomography [1] on the first t qubits right after the Gaussian operation \hat{G}^\dagger is applied to the state $|\psi\rangle$. More details are given in Sec. F of the Appendix.

Additionally, in Sec. F, we analyze the noise robustness of our algorithm, which is crucial for practical experimental scenarios where the unknown state may not be exactly t -compressible, but approximately so. Specifically, our analysis reveals that our algorithm allows for efficient tomography of states that are possibly mixed and (sufficiently) approximately compressible, which are the types of states one would expect to get when running a t -doped matchgate circuit on a noisy quantum device.

D. Time-complexity lower bound

It is natural to wonder whether there exist algorithms for learning t -doped Gaussian states with time-complexity scaling in t as $\mathcal{O}(\text{poly}(t))$. We establish that the answer is no (see Proposition 8 in the Appendix), relying on a widely believed cryptography assumption. Specifically, we show that certain families of pseudorandom quantum states [50,51] can be generated using a polynomial number of local non-Gaussian gates. This implies that if there were an algorithm with polynomial time complexity in t for learning t -doped Gaussian states, quantum computers could solve [73] in polynomial time, which is considered unlikely [46–48,73–75]. While this rules out the existence of efficient algorithms if t scales polynomially with the number of qubits n , it does not yet preclude the existence of efficient algorithms if t grows *slightly* more than logarithmically, for example, $t = \mathcal{O}((\log n)^2)$. However, we can rule out this possibility by making the stronger assumption that quantum computers cannot solve in subexponential time [46–49]. This implies that the time complexity of any algorithm to learn $\tilde{\mathcal{O}}(t)$ -doped Gaussian states [where $\tilde{\mathcal{O}}(\cdot)$ hides polylogarithmic factors] would necessarily be $\exp(\Omega(t))$. In other words, the following holds.

Theorem 3 (Time-complexity lower bound, informal). Assuming that quantum computers cannot solve in subexponential time, then there is no time-efficient algorithm to learn the $\tilde{\omega}(\log(n))$ -doped Gaussian state, which outputs a description of an efficiently preparable quantum state. Here, $\tilde{\omega}(\log(n)) := \omega(\log(n)) \text{polylog}(\log(n))$.

This would prove that the time complexity in t of our algorithm is essentially optimal, because our algorithm is efficient as long as $t = \mathcal{O}(\log(n))$. We show Theorem 3

by efficiently encoding the pseudorandom quantum state constructions [72] via a specific qubits-to-fermions mapping [25] into other states produced by the same number of gates, all of which are now local non-Gaussian. Crucially for our construction, this mapping sends local qubit operations to local fermionic operations with only a constant overhead in the number of qubits. We refer the reader to the Appendix (Proposition 9) for more details.

E. Testing Gaussian dimension

We have introduced an algorithm for efficiently learning states with a high Gaussian dimension (or, equivalently, small Gaussian nullity), specifically those promised to be t compressible with a small t . A natural question arises: how can we test the Gaussian dimension of a state? In other words, how can we determine if the underlying state is close or far from the set of t -compressible states? In our Appendix, using ideas developed in Ref. [76] for the case of Gaussian states ($t = 0$), we establish that the minimum-trace distance between a quantum state $|\psi\rangle$ and the set of t -compressible Gaussian states, denoted by \mathcal{G}_t , satisfies

$$\frac{1 - \lambda_{t+1}}{2} \leq \min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \leq \sqrt{\sum_{k=t+1}^n \frac{1 - \lambda_k}{2}}, \quad (5)$$

where $\{\lambda_k\}_{k=1}^n$ represents the normal eigenvalues of the correlation matrix of $|\psi\rangle$ ordered in increasing order. These inequalities imply that $|\psi\rangle$ is close in trace distance to the set \mathcal{G}_t if and only if λ_{t+1} is close to one. In particular, assuming that $|\psi\rangle$ is either a state in \mathcal{G}_t or $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \geq \varepsilon$, we can determine with at least $1 - \delta$ probability, which of the two cases is true by accurately estimating λ_{t+1} . Specifically, $\mathcal{O}((n^5/\varepsilon^4) \log(n^2/\delta))$ copies of the state suffice for this purpose. Notably, this complexity scaling for testing is independent from t , in contrast to learning. More details are provided in Sec. D of the Appendix, along with a generalization to the mixed-state scenario.

F. Numerical simulations: t -compressible states in many-body physics

In the previous section, we introduced an algorithm to learn quantum states that are t compressible. In Sec. F of the Appendix, we extended the algorithm to learn states that are not exactly t compressible but *approximately*. This connects to the work of Bravyi and Gosset [45], who demonstrated that ground states of so-called quantum impurity models can be well approximated by compressible states (see Ref. [45], Corollary 1). In this section we discuss the learnability of such states and of the states produced by the time dynamics of impurity models.

Quantum impurity models represent a bath of free fermions coupled with a small localized interacting subsystem referred to as an impurity. These models are of paramount importance in condensed-matter physics and have been extensively studied over the past half-century. The exploration of such Hamiltonians gained significant traction in the 1960s and 1970s, largely due to the pioneering work of Anderson, Kondo, and Wilson. Their focus was on the behavior of magnetic impurities within metals [52–54,77], providing a theoretical basis for the Kondo effect [53], which had been experimentally observed much earlier. Quantum impurity models are also crucial computational tools for analyzing the electronic structure of strongly correlated materials, including transition metal compounds and high-temperature superconductors. They are integral to the dynamical mean field theory (DMFT) [78], a method used to study these materials. In DMFT, the impurity typically represents a cluster of atoms within a unit cell, while the bath models the bulk of the material.

Formally, a quantum impurity model Hamiltonian is a Hamiltonian H_m that can be expressed as $H_m = H_{\text{free}} + H_{\text{imp},m}$, where H_{free} is a free-fermionic Hamiltonian (i.e., a Hamiltonian quadratic in the Majorana operators $H_{\text{free}} = i \sum_{p,q=1}^{2n} h_{p,q} \gamma_p \gamma_q$, with h being a real antisymmetric matrix) and $H_{\text{imp},m}$ is an arbitrary Hamiltonian, potentially nonfree fermionic, acting nontrivially on at most m distinct Majorana operators, where m is chosen to be $\mathcal{O}(1)$ relative to the number of fermionic modes n .

Restating Corollary 1 from Ref. [45], Bravyi and Gosset demonstrated that ground states $|\psi_{\text{GS}}\rangle$ of impurity models can be approximated, for any $\varepsilon > 0$, by a t -compressible Gaussian state with $t := \mathcal{O}(\log(\varepsilon^{-1}))$. Specifically, there exists a Gaussian unitary G and a state $|\phi\rangle$ supported on t modes such that $\|\psi_{\text{GS}} - \psi_t\|_1 \leq \varepsilon$, where $|\psi_t\rangle := G(|\phi\rangle \otimes |0^{n-t}\rangle)$ [79]. Thus, our algorithm can effectively be applied to learn ground states of impurity models with arbitrary ε precision, with a sample complexity scaling as $\text{poly}(n, \varepsilon^{-1})$.

An intriguing question arises: can *evolutions* governed by such Hamiltonians maintain approximate t compressibility up to certain evolution times, thus making our algorithm applicable to such nonequilibrium quantum states? To explore this, we conducted numerical simulations with an interacting impurity embedded in the background of two distinct free models: the 1D *transverse-field Ising* model (TFIM) and an *expander graph* model. The respective Hamiltonians are $H_{\text{TFIM}} = wZ_1Z_2 + \sum_{j=1}^n (gX_jX_{\text{mod}(j+1,n)} + Z_j)$, where w and g are real numbers, and $H_{\text{graph}} = w\gamma_1\gamma_2\gamma_3\gamma_4 + \sum_{e \in \mathcal{G}} i v_e \gamma_{e(1)} \gamma_{e(2)}$, where \mathcal{G} represents a random four-regular graph, with coupling constants v_e uniformly set to 1 for simplicity. Without the impurity term (i.e., for $w = 0$), both Hamiltonians become quadratic in the Majorana operators (up to a long Z -string term, corresponding to a fixed-parity contribution), implying that the continuous evolution they realize is

Gaussian. We focus on the cases $g = w = 1$, so that the energy scales of the impurity term and the quadratic part are comparable.

We consider the continuous evolution $|\psi(T)\rangle = e^{-iHT} |0^n\rangle$ with $H = H_{\text{graph}}$ and H_{TFIM} and investigate whether the resulting state $|\psi(T)\rangle$ is approximately t compressible. To determine this, we check if there exists a t such that

$$\sqrt{\sum_{m=t+1}^n \frac{1}{2}(1 - \lambda_m)} \leq \varepsilon_{\text{trunc}}, \quad (6)$$

where $\{\lambda_j\}_{j=1}^n$ are the normal eigenvalues of the correlation matrix of $|\psi(T)\rangle$ ordered in increasing order, and $\varepsilon_{\text{trunc}}$ is a fixed truncation error, which captures the approximate t compressibility of the underlying quantum state. This implies that the state $|\psi(T)\rangle$ can be learned with our algorithm up to an error $\varepsilon + \varepsilon_{\text{trunc}}$ in trace distance, where ε is the accuracy chosen in the tomography algorithm (see Theorem 9 of the Appendix for more details). This learning error can be brought closer to $\varepsilon_{\text{trunc}}$ by decreasing ε at the respective cost. In Fig. 2, we present the dynamics of the minimal Gaussian nullity t that allows a compression error below $\varepsilon_{\text{trunc}} = 0.05$, as a function of the evolution time T .

For small evolution times T , as shown in Fig. 2 (left panel), $|\psi(T)\rangle$ exhibits small approximate Gaussian nullity. One might wonder if this approximate t compressibility of impurity dynamics scales well with the system size, or whether, on the contrary, maintaining a fixed error in trace distance becomes increasingly difficult for larger n .

Our numerical analysis suggests the former. For instance, consider the maximal time T_4 that allows maintaining approximate t compressibility for $t = 4$ and a fixed $\varepsilon_{\text{trunc}}$. The question is whether T_4 decreases to zero with increasing n or stabilizes at a constant value. We find that T_4 indeed saturates at a certain constant for both the TFIM and the expander models (Fig. 2, right panel).

We also examine whether the compressibility shown in Fig. 2 can be derived from our Theorem 1 for t -doped circuits using the Trotter approximation. Our numerical analysis rules out this gate-counting hypothesis, because this hypothesis predicts high approximate nullity already for minuscule times (vertical dotted lines in Fig. 2, left panel). A detailed explanation of this analysis is given in Sec. G of the Appendix. We do not have an analytical derivation of approximate t compressibility of impurity dynamics as observed in our numerical investigation, and leave this question open for future work. The failure of explanation via Trotter expansion highlights the difficulty of this problem. For the TFIM model, the approximate t compressibility may be derivable using 1D Lieb-Robinson bounds. However, such a derivation should not work for the expander graph model—and at present we do not know how this case could be analytically handled. This open problem also connects to the open questions about the classical simulation of impurity dynamics, raised earlier by Bravyi and Gosset [45].

To run the simulations of $|\psi(T)\rangle = e^{-iHT} |0^n\rangle$ presented in Fig. 2, we employed sparse matrix multiplication routines. More details about our numerical simulations are given in Sec. G of the Appendix.




FIG. 2. (Left) Dynamics of the Gaussian nullity in an impurity model with expander graph and 1D TFIM backgrounds as defined in the text. The admissible deviation from t compressibility, $\varepsilon_{\text{trunc}}$, is set to 0.05; system size is $n = 14$. The deviation from t compressibility is upper bounded using the normal eigenvalues of the correlation matrix in the state $|\psi(T)\rangle$, as in Eq. (6). Vertical dotted lines show times T , which are too large to derive $t < 14$ nullity from Trotter approximation and Theorem 1. The time axis is resolved up to $\Delta T = 0.05$. (Right) Maximal evolution time T_4 allowing approximate $t = 4$ compression (up to $\varepsilon_{\text{trunc}} = 0.05$). One observes that T_4 saturates at a constant value as a function of n for both models.

G. Conclusions

In this work, we have presented an algorithm for efficiently learning t -doped fermionic Gaussian states, with sample and time complexity scaling as $\mathcal{O}(\text{poly}(n, 2^t))$. Additionally, we have established, under standard cryptography assumptions, that there is no learning algorithm for such class of states with a polynomial dependence on t in the time complexity. Crucially, our algorithm utilizes solely experimentally feasible single-copy measurements. Its working idea is based on a theorem that we prove, which says that all the non-Gaussianity in a t -doped fermionic Gaussian state can be efficiently compressed onto $\mathcal{O}(t)$ qubits through a Gaussian operation. This observation carries potential significance beyond the scope of learning, particularly within the context of quantum many-body theory or within efficient circuit compilation. Thus, the results presented in this work, besides being directly relevant to device verification and benchmarking, among other tasks, hold fundamental significance for quantum information theory, as they reveal more about the structure of Gaussian states with fermionic magic gates. Additionally, we introduce a variety of useful analytical techniques, such as new ways to leverage pseudorandom quantum states [50,51] in the context of fermionic systems, which are likely to find applications in future research. We also observe that the results presented in our study for fermionic systems have already found application in continuous-variable systems and bosonic Gaussian states [80], where the concept of t -doped bosonic Gaussian states is explored, highlighting the broader applicability of our ideas and techniques to a different context.

Our work offers new directions for further research. An open question arising from this work is whether t -doped Gaussian unitaries can be efficiently learned in a scenario where both the input states to the unitary and the measurements at the end can be chosen. The specific case of $t = 0$ has already been addressed in Ref. [81], and it would be interesting to generalize this to $t > 0$. Another promising direction is the study of the resource theory of fermionic non-Gaussianity, where the notion of Gaussian nullity we introduced may play a significant role, similar to stabilizer nullity in the Clifford context [63,64]. Finally, an analytical justification of the approximate t compressibility of impurity-model dynamics remains an intriguing open problem, which we intend to explore in future research.

ACKNOWLEDGMENTS

We thank Mauro Morales, Lorenzo Leone, Lennart Bittel, Janek Denzler, Tommaso Guaita, Ansgar Burchards, Julio Magdalena de la Fuente, Daniel Liang, Matthias Caro, Alexander Nietner, Marten Folkertsma, Francisco Escudero Gutiérrez, Simona Etinski, Stefano Polla, Jonas Helsen, Barbara Terhal, and Jens Eisert for insightful discussions. We acknowledge the PCMI 2023 program for

providing a collaborative space where the authors brainstormed this project. A.A.M. acknowledges fundings from BMBF (FermiQP).

APPENDIX

In this Appendix, we provide a more comprehensive level of detail and explanation for certain statements covered in the main text.

APPENDIX A: PRELIMINARIES

1. Notation and basics

In this work, we employ the following notation. $\mathcal{L}(\mathbb{C}^d)$ denotes the set of linear operators acting on the d -dimensional complex vector space \mathbb{C}^d . Additionally, we use $[d]$ to represent the set of integers from 1 to d , i.e., $[d] := \{1, \dots, d\}$. We denote with $\text{Mat}(d, \mathbb{F})$ the set of $d \times d$ matrices over the field \mathbb{F} . Let $v \in \mathbb{C}^d$ be a vector, and let $p \in [1, \infty]$. The p -norm of v is denoted by $\|v\|_p$, defined as $\|v\|_p := (\sum_{i=1}^d |v_i|^p)^{1/p}$. The Schatten p -norm of a matrix $A \in \mathbb{C}^d$, with $p \in [1, \infty]$, is given by $\|A\|_p := \text{Tr}((\sqrt{A^\dagger A})^p)^{1/p}$, corresponding to the p -norm of the vector of singular values of A . The trace norm and the Hilbert-Schmidt norm are important instances of Schatten p -norms, denoted as $\|\cdot\|_1$ and $\|\cdot\|_2$, respectively. The Hilbert-Schmidt norm is induced by the Hilbert-Schmidt scalar product $\langle A, B \rangle_{HS} := \text{Tr}(A^\dagger B)$ for $A, B \in \mathcal{L}(\mathbb{C}^d)$. The infinity norm, $\|\cdot\|_\infty$, of a matrix is defined as its largest singular value. This norm can be interpreted as the limit of the Schatten p -norm of the matrix as p approaches infinity. For any unitaries U and V , and a matrix A , we have the unitary invariance property $\|UAV\|_p = \|A\|_p$. Also, $\|A \otimes B\|_p = \|A\|_p \|B\|_p$ for $A, B \in \mathcal{L}(\mathbb{C}^d)$. We denote with $U(n)$ the group of $n \times n$ unitary matrices. We denote $O(2n)$ as the group of real orthogonal $2n \times 2n$ matrices. $\text{Sp}(2n, \mathbb{R})$ denotes the group of symplectic matrices over the real field, defined as

$$\text{Sp}(2n, \mathbb{R}) := \{S \in \text{Mat}(2n, \mathbb{R}) : S\Omega S^T = \Omega\}, \quad (\text{A1})$$

where $\Omega := \bigoplus_{i=1}^n \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. The n -qubit Pauli operators are represented as elements of the set $\{I, X, Y, Z\}^{\otimes n}$, where I, X, Y, Z represent the standard single-qubit Pauli. Pauli operators are traceless, Hermitian, they square to the identity, and form an orthogonal basis with respect to the Hilbert-Schmidt scalar product for the space of linear operators. We define the set of quantum states as $\mathcal{S}(\mathbb{C}^d) := \{\rho \in \mathcal{L}(\mathbb{C}^d) : \rho \geq 0, \text{Tr}(\rho) = 1\}$. The trace distance between two quantum states ρ, σ is defined as $d_{\text{tr}}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1$. For a function $f(n)$, if there exists a constant c and a specific input size n_0 such that $f(n) \leq c \cdot g(n)$ for all $n \geq n_0$, where $g(n)$ is a well-defined function, then we express it as $f(n) = \mathcal{O}(g(n))$. This notation signifies the upper limit of how fast a function grows in relation to $g(n)$.

For a function $f(n)$, if there exists a constant c and a specific input size n_0 such that $f(n) \geq c \cdot g(n)$ for all $n \geq n_0$, where $g(n)$ is a well-defined function, then we express it as $f(n) = \Omega(g(n))$. This notation signifies the lower limit of how fast a function grows in relation to $g(n)$.

For a function $f(n)$, if for any constant c , there exists an input size n_0 such that $f(n) > c \cdot g(n)$ for all $n \geq n_0$, where $g(n)$ is a well-defined function, then we express it as $f(n) = \omega(g(n))$. This notation implies that the function grows strictly faster than the provided lower bound.

2. Basics of probability theory

In this section, we present fundamental results from probability theory useful in our work.

Lemma 1 (Union bound). Let A_1, A_2, \dots, A_M be events in a probability space. The probability of the union of these events is bounded by the sum of their individual probabilities:

$$\Pr\left(\bigcup_{i=1}^M A_i\right) \leq \sum_{i=1}^M \Pr(A_i).$$

Lemma 2 (Chernoff bound). Consider a set of independent and identically distributed random variables $\{X_i\}_{i=1}^N$ with binary outcomes, taking values in $\{0, 1\}$. Define $X := \sum_{i=1}^N X_i$ and $\mu := \mathbb{E}[X]$. For any $\alpha \in (0, 1)$, the probability of X being less than $(1 - \alpha)$ times its expected value is exponentially bounded as follows:

$$\Pr[X \leq (1 - \alpha)\mu] \leq \exp\left(-\frac{\alpha^2 \mu}{2}\right).$$

Lemma 3 (Hoeffding's inequality). Let $\{X_i\}_{i=1}^N$ be independent and identically distributed (IID) random variables with values in $[a, b] \subseteq \mathbb{R}$. For any $\varepsilon > 0$, the probability of the deviation of $\hat{X} := (\sum_{i=1}^N X_i)/N$ from its expected value is exponentially bounded as follows:

$$\Pr\left(\left|\frac{1}{N} \sum_{i=1}^N X_i - \mathbb{E}[\hat{X}]\right| \geq \varepsilon\right) \leq 2 \exp\left(-\frac{2N\varepsilon^2}{(b-a)^2}\right).$$

Corollary 1. For any $\varepsilon > 0$ and $\delta > 0$, let $\{X_i\}_{i=1}^N$ be IID random variables with values in $[a, b] \subseteq \mathbb{R}$ and $\hat{X} := (\sum_{i=1}^N X_i)/N$. According to Hoeffding's inequality, a sample size N satisfying

$$N \geq \frac{(b-a)^2}{2\varepsilon^2} \log\left(\frac{2}{\delta}\right)$$

suffices to guarantee that $|1/N \sum_{i=1}^N X_i - \mathbb{E}[\hat{X}]| < \varepsilon$ with a probability of at least $1 - \delta$.

3. Fermionic Gaussian states

In this section, we explore the definitions and essential properties of fermionic Gaussian states. We focus on a system consisting of n qubits or n fermionic modes, resulting in a Hilbert space dimension of 2^n . More precisely, our work is approached from two perspectives: examining a system of n qubits with 1D matchgate circuits and their magic gates (e.g., SWAP gates), or an equivalent native fermionic system of n modes with states prepared by fermionic Gaussian evolutions and local non-Gaussian evolutions. These perspectives are mathematically connected through the Jordan-Wigner mapping, which we use now for defining Majorana operators in terms of Pauli operators.

Definition 3 (Majorana operators). For each $k \in [n]$, Majorana operators are defined as

$$\gamma_{2k-1} := \left(\prod_{j=1}^{k-1} Z_j\right) X_k, \quad \gamma_{2k} := \left(\prod_{j=1}^{k-1} Z_j\right) Y_k. \quad (\text{A2})$$

Majorana operators can also be defined directly in the fermionic language through their anticommutation relations [15,57]. Majorana operators are Hermitian, traceless, and their squares yield the identity, as deducible from their definitions. Moreover, distinct Majorana operators exhibit anticommutativity and orthogonality with respect to the Hilbert-Schmidt inner product.

Definition 4 (Majorana ordered products). Given a set $S := \{\mu_1, \dots, \mu_{|S|}\} \subseteq [2n]$ with $1 \leq \mu_1 < \dots < \mu_{|S|} \leq 2n$, we define the Majorana product operator as $\gamma_S = \gamma_{\mu_1} \cdots \gamma_{\mu_{|S|}}$ if $S \neq \emptyset$, and $\gamma_\emptyset = I$ otherwise.

The 4^n distinct ordered Majorana products are orthogonal to each other with respect the Hilbert-Schmidt inner product, therefore they form a basis for the linear operators $\mathcal{L}(\mathbb{C}^{2^n})$.

Definition 5 (Fermionic Gaussian unitary (FGU)). A fermionic Gaussian unitary G_O is a unitary operator satisfying

$$G_O^\dagger \gamma_\mu G_O = \sum_{v=1}^{2n} O_{\mu,v} \gamma_v \quad (\text{A3})$$

for any $\mu \in [2n]$, where $O \in O(2n)$ is an orthogonal matrix.

Since the ordered products of Majorana operators γ_μ with $\mu \in [2n]$ form a basis for the linear operators, it is sufficient to specify how a unitary acts under conjugation on the $2n$ Majorana operators γ_μ , where $\mu \in [2n]$, to uniquely

determine the unitary up to a phase. Thus, there is a one-to-one mapping between n -qubit fermionic Gaussian unitaries (up to a global phase) and orthogonal matrices $O(2n)$. In particular, given $O \in O(2n)$, it is possible to build the associated unitary using at most $\mathcal{O}(n(n-1)/2)$ two-qubit FGU operations. For a more detailed explanation on how to map an $O(2n)$ matrix to a fermionic Gaussian unitary, refer to Refs. [38,58,82]. From the previous definition, it readily follows that $G_O^\dagger = G_{O^T}$. Moreover we have that the product of two Gaussian unitaries is Gaussian, namely $(G_{O_1} G_{O_2})^\dagger \gamma_\mu G_{O_1} G_{O_2} = \sum_{\nu=1}^{2n} (O_1 O_2)_{\mu,\nu} \gamma_\nu$. To streamline notation, we will frequently refer to fermionic Gaussian unitaries G_O simply as G when there is no need to specify the associated orthogonal matrix. Now, we define a fermionic Gaussian state.

Definition 6 (Fermionic Gaussian state). An n -qubit state $|\psi\rangle$ is a (pure) fermionic Gaussian state if it can be expressed as $|\psi\rangle = G|0^n\rangle$, where G is a fermionic Gaussian unitary.

It is noteworthy that any computational basis state $|x\rangle$ is a Gaussian state. This stems from the observation that each Pauli- X_i gate acting on the i th qubit, where $i \in [n]$, is a fermionic Gaussian unitary. An additional useful identity is $Z_j = -i\gamma_{2j-1}\gamma_{2j}$. Thus, the density matrix of a pure fermionic Gaussian state associated to an orthogonal matrix $O \in O(2n)$ can be written as

$$G_O |0^n\rangle\langle 0^n| G_O^\dagger = G_O \left(\prod_{j=1}^n \frac{I - i\gamma_{2j-1}\gamma_{2j}}{2} \right) G_O^\dagger = \prod_{j=1}^n \left(\frac{I - i\tilde{\gamma}_{2j-1}\tilde{\gamma}_{2j}}{2} \right), \quad (\text{A4})$$

where $\tilde{\gamma}_\mu := G_O \gamma_\mu G_O^\dagger = \sum_{\nu=1}^{2n} O_{\mu,\nu}^T \gamma_\nu$ for each $\mu \in [2n]$.

It can be shown that free-fermionic Hamiltonians, which are quadratic Hamiltonians in the Majorana operators, i.e., $H_{\text{free}} := i \sum_{p,q=1}^{2n} h_{p,q} \gamma_p \gamma_q$, where h is a real antisymmetric matrix, have fermionic Gaussian states as their ground states. Moreover, it can be shown that their time evolutions are fermionic Gaussian unitaries.

We now proceed to define the correlation matrix for any (possibly non-Gaussian) state.

Definition 7 (Correlation matrix). For any n -qubit quantum state ρ , its correlation matrix $C(\rho)$ is defined as

$$[C(\rho)]_{j,k} := -\frac{i}{2} \text{Tr}([\gamma_j, \gamma_k] \rho), \quad (\text{A5})$$

where $j, k \in [2n]$.

The correlation matrix of any state is real and antisymmetric, possessing eigenvalues in pairs $\pm i\lambda_j$ for $j \in [2n]$,

where λ_j are real numbers such that $|\lambda_j| \leq 1$. The correlation matrix of a quantum state, when evolved using fermionic Gaussian unitaries, undergoes a transformation through conjugation with the corresponding orthogonal matrix, as articulated in the following lemma.

Lemma 4 (Transformation of the correlation matrix under FGU). For a given n -qubit state ρ , we have

$$C(G_O \rho G_O^\dagger) = O C(\rho) O^T, \quad (\text{A6})$$

for any orthogonal matrix $O \in O(2n)$ and associated fermionic Gaussian unitary G_O .

This result is readily verified through the definitions of the correlation matrix and fermionic Gaussian unitary. The state $|x\rangle$ is characterized by a correlation matrix of the following form:

$$C(|x\rangle\langle x|) = \bigoplus_{j=1}^n \begin{pmatrix} 0 & (-1)^{x_i} \\ -(-1)^{x_i} & 0 \end{pmatrix}. \quad (\text{A7})$$

Hence, for a fermionic Gaussian state $|\psi\rangle := G_O |0^n\rangle$, the correlation matrix takes the form:

$$C(|\psi\rangle\langle\psi|) = O \bigoplus_{j=1}^n \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} O^T. \quad (\text{A8})$$

In the subsequent discussion, we will use $C(|\psi\rangle)$ to denote the correlation matrix of a pure state $|\psi\rangle$. If the state $|\psi\rangle$ is a pure Gaussian state, then each of the eigenvalues of $C(|\psi\rangle)$ is one in absolute value. Moreover, it is worth noting that every real antisymmetric matrix can be decomposed in the following form.

Lemma 5 (Normal form of real antisymmetric matrices [60]). Any real antisymmetric matrix C can be decomposed in the so-called “normal form”:

$$C = O \bigoplus_{j=1}^n \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix} O^T, \quad (\text{A9})$$

where O is an orthogonal matrix in $O(2n)$ and $\lambda_j \geq 0 \in \mathbb{R}$, for any $j \in [n]$, are ordered in increasing order. The eigenvalues of C are $\pm i\lambda_j$ where $\lambda_j \in \mathbb{R}$ for any $j \in [n]$.

Definition 8 (Normal eigenvalues). Given a real-antisymmetric matrix decomposed as in the previous Lemma 5, $\{\lambda_j\}_{j=1}^n$ are dubbed as the “normal eigenvalues” of the matrix.

4. Particle-number-preserving unitaries

In this section, we introduce the concept of particle-number-preserving fermionic unitaries and establish definitions and facts useful for subsequent discussions. We begin by defining creation and annihilation operators.

Definition 9 (Creation and annihilation operators). The annihilation operators are defined as

$$a_j := \frac{\gamma_{2j-1} + i\gamma_{2j}}{2}, \quad (\text{A10})$$

for any $j \in [n]$. The creation operators $\{a_j^\dagger\}_{j=1}^n$ are defined as the adjoints of the annihilation operators.

Definition 10 (Particle-number operator). The operator $\hat{N} := \sum_{i=1}^n a_i^\dagger a_i$ is denoted as the particle number operator.

The computational basis forms a set of eigenstates for the particle-number operator:

$$\hat{N} |x_1, \dots, x_n\rangle = (x_1 + \dots + x_n) |x_1, \dots, x_n\rangle, \quad (\text{A11})$$

where $x_1, \dots, x_n \in \{0, 1\}$.

Definition 11 (Particle-number-preserving unitaries). A unitary U is said to be particle-number preserving if and only if

$$U^\dagger \hat{N} U = \hat{N}, \quad (\text{A12})$$

where $\hat{N} := \sum_{i=1}^n a_i^\dagger a_i$ is the particle-number operator.

Definition 12 (Symplectic group). The group of real symplectic matrices, denoted as $\text{Sp}(2n, \mathbb{R})$, is defined as

$$\text{Sp}(2n, \mathbb{R}) := \{S \in \text{Mat}(2n, \mathbb{R}) : S\Omega S^T = \Omega\}, \quad (\text{A13})$$

where $\Omega := \bigoplus_{i=1}^n \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = \bigoplus_{i=1}^n iY$.

It is often convenient to express Ω as $\Omega = I_n \otimes iY$, where I_n denotes the $n \times n$ identity matrix. Note that Ω in the literature is sometimes defined (see, e.g., Ref. [61]) as $iY \otimes I_n = \begin{pmatrix} 0_n & I_n \\ -I_n & 0_n \end{pmatrix}$, but the two definitions are equivalent up to orthogonal transformation.

Now, we state an important proposition that will be useful in the subsequent section.

Proposition 1 ($U(n)$ is isomorphic to $O(2n) \cap \text{Sp}(2n, \mathbb{R})$). The set of unitaries $U(n)$ is isomorphic to the set of real symplectic orthogonal matrices $O(2n) \cap \text{Sp}(2n, \mathbb{R})$.

In particular, any orthogonal symplectic matrix $O \in O(2n) \cap \text{Sp}(2n, \mathbb{R})$ can be written as follows:

$$O = \text{Re}(u) \otimes I + \text{Im}(u) \otimes iY, \quad (\text{A14})$$

where $u \in U(n)$ is an $n \times n$ unitary.

Proof. We refer the reader to Appendix B.1 of the book [61] for a detailed proof. However, by inspection, it can be verified that the matrix O defined in this way is both orthogonal and symplectic. This follows from the unitarity of u , which implies the relations:

$$\begin{aligned} \text{Re}(u)\text{Re}(u)^t + \text{Im}(u)\text{Im}(u)^t &= I, \\ \text{Re}(u)\text{Im}(u)^t - \text{Im}(u)\text{Re}(u)^t &= 0. \end{aligned} \quad (\text{A15})$$

■

We now present a lemma, which shows (some) equivalent definitions of a particle-number-preserving fermionic Gaussian unitary.

Lemma 6 (Particle-number-preserving Gaussian unitary). Let G be a fermionic Gaussian unitary associated with the orthogonal matrix $O \in O(2n)$. The following points are equivalents:

- (1) G is particle-number-preserving,
- (2) $G|0^n\rangle\langle 0^n|G^\dagger = |0^n\rangle\langle 0^n|$,
- (3) O is symplectic orthogonal, i.e., $O \in O(2n) \cap \text{Sp}(2n, \mathbb{R})$.

Proof. If G is particle-number-preserving, then

$$G^\dagger \hat{N} G |0^n\rangle = \hat{N} |0^n\rangle = 0, \quad (\text{A16})$$

where the first equality uses Definition 11. This implies $\hat{N} G |0^n\rangle = 0$. Since the ground space corresponding to the zero eigenvalue of the particle-number operator \hat{N} is one-dimensional and spanned by $|0^n\rangle$, it follows that $G |0^n\rangle$ is equal to $|0^n\rangle$, up to a phase. Thus, (1), implies (2).

Noting that $|0^n\rangle\langle 0^n|$ is a Gaussian state with a correlation matrix $\Omega = \bigoplus_{j=1}^n iY$, we deduce that the correlation matrix of $G|0^n\rangle\langle 0^n|G^\dagger$ is $O\Omega O^T$. Therefore, the condition $G|0^n\rangle\langle 0^n|G^\dagger = |0^n\rangle\langle 0^n|$ is equivalent to $O\Omega O^T = \Omega$, i.e., O is a real symplectic orthogonal matrix. This proves that (2) is equivalent to (3).

Now, let us assume that $O \in O(2n) \cap \text{Sp}(2n, \mathbb{R})$. Then, we have for $l \in [n]$:

$$G^\dagger a_l G = G^\dagger \left(\frac{\gamma_{2l-1} + i\gamma_{2l}}{2} \right) G = \sum_{j=1}^{2n} (O_{2l-1,j} + iO_{2l,j}) \frac{\gamma_j}{2} \quad (\text{A17})$$

$$\begin{aligned} &= \sum_{j=1}^n (O_{2l-1,2j-1} + iO_{2l,2j-1}) \frac{\gamma_{2j-1}}{2} \\ &\quad + \sum_{j=1}^n (O_{2l-1,2j} + iO_{2l,2j}) \frac{\gamma_{2j}}{2}, \end{aligned} \quad (\text{A18})$$

$$= \sum_{j=1}^n (\text{Re}(u)_{lj} + i\text{Im}(u)_{lj}) \frac{\gamma_{2j-1}}{2} + \sum_{j=1}^n (-\text{Im}(u)_{lj} + i\text{Re}(u)_{lj}) \frac{\gamma_{2j}}{2}, \quad (\text{A19})$$

$$= \sum_{j=1}^n u_{lj} \frac{(\gamma_{2j-1} + i\gamma_{2j})}{2} = \sum_{j=1}^n u_{lj} a_j, \quad (\text{A20})$$

where in the fourth equality we used that $O \in \text{O}(2n) \cap \text{Sp}(2n, \mathbb{R})$, and so, because of Proposition 1, it can be written as $O = \text{Re}(u) \otimes I + \text{Im}(u) \otimes iY$ where $u \in \text{U}(n)$ is a $n \times n$ unitary. Similarly, we have $G^\dagger a_l^\dagger G = \sum_{j=1}^n u_{lj}^* a_j^\dagger$. This implies

$$G^\dagger \hat{N} G = \sum_{l=1}^n G^\dagger a_l^\dagger a_l G = \sum_{l,j,k=1}^n u_{lj}^* u_{lk} a_j^\dagger a_k = \sum_{j=1}^n a_j^\dagger a_j = \hat{N}, \quad (\text{A21})$$

where we used the unitarity of u in the last step. This proves that (3) and implies (1). ■

APPENDIX B: STRUCTURE OF t -DOPED GAUSSIAN UNITARIES AND STATES

In this section, we analyze the concept of t -doped fermionic Gaussian unitaries and states.

Definition 13 (t -doped fermionic Gaussian unitary). A unitary U_t is a (t, κ) -doped fermionic Gaussian unitary if it can be decomposed in terms of Gaussian unitaries $\{G_i\}_{i=0}^t$ and at most t non-Gaussian κ -local gates $\{W_i\}_{i=1}^t$, specifically

$$U_t = G_t W_t \cdots G_1 W_1 G_0. \quad (\text{B1})$$

Here κ local refers to the number of distinct Majorana operators that generate each non-Gaussian gate. Informally, a unitary is t -doped Gaussian if it is (t, κ) -doped Gaussian for some fixed constant κ .

In our work, we consider non-Gaussian gates $W_{t'}$ for $t' \in [t]$, each generated by $\kappa \leq 2n$ in κ fixed Majorana operators $\{\gamma_{\mu(t',j)}\}_{j=1}^\kappa$, where $\mu(t', 1), \dots, \mu(t', \kappa) \in [2n]$. For $\kappa = 3$, an example of a non-Gaussian gate is $\exp(\theta \gamma_1 \gamma_2 \gamma_3) = \exp(-iY_1 \theta)$, where $\theta \in \mathbb{R}$, and for $\kappa = 4$, an example is the SWAP gate.

Definition 14 (t -doped fermionic Gaussian state). An n -qubit state $|\psi\rangle$ is a (t, κ) -doped (informally, t -doped) fermionic Gaussian state if it can be expressed as $|\psi\rangle = U_t |0^n\rangle$, where U_t is a (t, κ) -doped (t -doped) fermionic Gaussian unitary.

1. Compression of t -doped Gaussian unitaries and states

We start by presenting a theorem, which shows how all non-Gaussianity in a t -doped unitary can be “compressed” or “moved” to the first few qubits.

Theorem 4 (Compression of non-Gaussianity in t -doped unitaries). Any (t, κ) -doped fermionic Gaussian unitary U_t can be expressed as

$$U_t = G_A(u_t \otimes I) G_B, \quad (\text{B2})$$

where G_A, G_B are Gaussian unitaries, and u_t is a unitary supported exclusively on $\lceil \kappa t/2 \rceil$ qubits.

Proof. Let us denote $M := \kappa t$. We express the t -doped unitary as $U_t = (\prod_{t'=1}^t G_{t'} W_{t'}) G_0$. Rearranging it, we have

$$U_t = \tilde{G}_t \prod_{t'=1}^t \tilde{W}_{t'}, \quad (\text{B3})$$

where $\tilde{W}_{t'} := \tilde{G}_{t'-1}^\dagger W_{t'} \tilde{G}_{t'-1}$, and $\tilde{G}_{t'} := G_{t'} \cdots G_0$. Informally, the idea behind this rewriting is that Gaussian operations act nicely under conjugation. Next, we rewrite

$$U_t = \tilde{G}_t G_{\text{aux}} \prod_{t'=1}^t (G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}) G_{\text{aux}}^\dagger, \quad (\text{B4})$$

by introducing a Gaussian operation G_{aux} that we fix later and that will be responsible for moving all the non-Gaussian gates to the first qubits. Now, we set

$$G_A := \tilde{G}_t G_{\text{aux}} \quad (\text{B5})$$

$$u_t := \prod_{t'=1}^t (G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}) \quad (\text{B6})$$

$$G_B := G_{\text{aux}}^\dagger. \quad (\text{B7})$$

Note that G_A so defined is Gaussian because the product of Gaussian unitaries is Gaussian, and G_B is clearly Gaussian because the adjoint of a Gaussian unitary is Gaussian. We need to show that it is possible to choose G_{aux} such that u_t is supported only on the first $\lceil M/2 \rceil$ qubits.

More precisely, we will require that G_{aux} ensures that each $G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}$ is generated by the first M Majorana operators alone. We will achieve it by ensuring that the Heisenberg evolution under $\tilde{G}_{t'-1} G_{\text{aux}}$, of each Majorana that generates $W_{t'}$, has nontrivial support exclusively on the first M Majorana operators. To find G_{aux} with the desired property, we will find the associated orthogonal matrix O_{aux} . Let $\{\mu(t', r)\}_{t' \in [t], r \in [\kappa]}$, with $\mu(t', r) \in [2n]$, be the set of indices of Majorana operators generating the

t' th non-Gaussian gate (listed in increasing order). For example, consider $\kappa = 4$ and the t' th non-Gaussian gate $W_{t'} := \exp(i\gamma_2\gamma_4\gamma_6\gamma_8 + i\gamma_4)$. In this case, $\mu(t', 3) = 6$. For each such Majorana operator $\gamma_{\mu(t', r)}$, where $r \in [2n]$, its Heisenberg evolution yields

$$G_{\text{aux}}^\dagger \tilde{G}_{t'-1}^\dagger \gamma_{\mu(t', r)} \tilde{G}_{t'-1} G_{\text{aux}} = \sum_{m=1}^{2n} (O_{t'-1} O_{\text{aux}})_{\mu(t', r), m} \gamma_m, \quad (\text{B8})$$

where $O_{t'-1}$ is the orthogonal matrix associated with $\tilde{G}_{t'-1}$. Our demand on the support of Heisenberg-evolved $\gamma_{\mu(t', r)}$ implies

$$(O_{t'-1} O_{\text{aux}})_{\mu(t', r), m} = (O_{\text{aux}}^T O_{t'-1}^T)_{m, \mu(t', r)} = 0 \quad (\text{B9})$$

for any $m \in \{M+1, \dots, 2n\}$.

Let us denote as $\{\mathbf{e}_i\}_{i=1}^{2n}$ the canonical basis vectors of \mathbb{R}^{2n} . For easy notation, we now denote the unit norm vectors $\{O_{t'-1}^T \mathbf{e}_{\mu(t', r)}\}_{t' \in [t], r \in [\kappa]}$ with the set of vectors $\{\mathbf{v}_j\}_{j=1}^M$ (remember that $M = \kappa t$). We can prove the existence of such O_{aux} by proving the existence of its transpose $O := O_{\text{aux}}^T$. In such notation, the condition in Eq. (B9) reads as

$$\mathbf{e}_m^T O \mathbf{v}_j = 0, \quad (\text{B10})$$

for any $j \in [M]$ and $m \in \{M+1, \dots, 2n\}$. In other words, we need to prove the existence of an orthogonal matrix O that maps any given real vectors $\mathbf{v}_1, \dots, \mathbf{v}_M \in \mathbb{R}^{2n}$, where $M \leq 2n$, to the span of the first M canonical basis vectors of \mathbb{R}^{2n} . The existence of such a matrix is readily established by selecting an orthonormal basis for $W := \text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_M)$ and defining the orthogonal matrix that maps this orthonormal basis to the first $\dim(W) \leq M$ canonical basis vectors. This concludes the proof. ■

The subsequent Theorem 5 demonstrates the compression of t -doped Gaussian states.

Theorem 5 (Compression of non-Gaussianity in t -doped Gaussian states). Any (t, κ) -doped fermionic Gaussian state $|\psi\rangle$ can be represented as

$$|\psi\rangle = G(|\phi\rangle \otimes |0^{n-\kappa t}\rangle), \quad (\text{B11})$$

where G is a Gaussian unitary, and $|\phi\rangle$ is a state supported exclusively on κt qubits.

Proof. Let $|\psi\rangle := U_t |0^n\rangle$, where $U_t = (\prod_{t'=1}^t G_{t'} W_{t'})$ G_0 is a t -doped fermionic Gaussian unitary. The proof begins analogously to the one of the previous Theorem 4

and it uses the same notation. In particular, we have

$$U_t = \tilde{G}_t G_{\text{aux}} \prod_{t'=1}^t (G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}) G_{\text{aux}}^\dagger, \quad (\text{B12})$$

where $\tilde{W}_{t'} := \tilde{G}_{t'-1}^\dagger W_{t'} \tilde{G}_{t'-1}$ and $\tilde{G}_{t'} := G_{t'} \dots G_0$. We set, as before

$$G_A := \tilde{G}_t G_{\text{aux}} \quad (\text{B13})$$

$$u_t := \prod_{t'=1}^t (G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}) \quad (\text{B14})$$

$$G_B := G_{\text{aux}}^\dagger. \quad (\text{B15})$$

However, now, we require that u_t has support on the first M qubits, where $M := \kappa t$ (while in the previous proof of Theorem 4 we requested $\lceil M/2 \rceil$), or, equivalently, we request that the generators of $G_{\text{aux}}^\dagger \tilde{W}_{t'} G_{\text{aux}}$ for any $t' \in [t]$ involve only the first $2M$ Majorana operators.

This time we also impose that $G_{\text{aux}}^\dagger |0^n\rangle = |0^n\rangle$. This implies that $O_{\text{aux}}^T \in \text{Sp}(2n, \mathbb{R})$ (where O_{aux} is the orthogonal matrix associated to G_{aux}), i.e., O_{aux}^T must be a symplectic orthogonal matrix, because of Lemma 6. We now define $O := O_{\text{aux}}^T$.

Similarly to the previous theorem and using the same notation, we can ensure that u_t is supported only on the first M qubits by demonstrating the existence of an orthogonal, but this time also symplectic, matrix O that satisfies

$$\mathbf{e}_m^T O \mathbf{v}_j = 0, \quad (\text{B16})$$

for any $j \in [M]$ with arbitrary $\mathbf{v}_1, \dots, \mathbf{v}_M$ real vectors, and $m \in \{2M+1, \dots, 2n\}$. The existence of such O follows from the subsequent Lemma 7, which crucially uses the isomorphism between $2n \times 2n$ symplectic orthogonal real matrices and $n \times n$ unitaries [61]. ■

Lemma 7 (Compression via symplectic orthogonal transformations). Let $\{\mathbf{e}_i\}_{i=1}^{2n}$ be the canonical basis of \mathbb{R}^{2n} . Let $\mathbf{v}_1, \dots, \mathbf{v}_M \in \mathbb{R}^{2n}$ be a set of unit-norm real vectors, where $M \leq n$. There exists an orthogonal symplectic matrix $O \in \text{O}(2n) \cap \text{Sp}(2n, \mathbb{R})$ such that

$$\mathbf{e}_i^T O \mathbf{v}_j = 0, \quad (\text{B17})$$

for all $i \in \{2M+1, \dots, 2n\}$ and $j \in [M]$, meaning that all $\{O \mathbf{v}_j\}_{j=1}^M$ are exclusively supported on the span of the first $2M$ canonical basis vectors.

Proof. Orthogonal symplectic matrices $O \in \text{O}(2n) \cap \text{Sp}(2n, \mathbb{R})$ have a bijective correspondence with unitary matrices $U \in \text{U}(n)$ through a well-defined vector-space mapping [61] (see Proposition 1). Specifically, for a $2n$ -dimensional real vector $\mathbf{w} := (w_1, \dots, w_{2n})$, there exists

a bijective mapping to an n -dimensional complex vector $\mathbf{f}(\mathbf{w}) := (w_1 - iw_2, \dots, w_{2n-1} - iw_{2n})$. A unitary transformation U in this n -dimensional complex space corresponds to an orthogonal symplectic transformation O in the corresponding $2n$ -dimensional real space, and vice versa. Thus, finding a unitary U that maps the span of $\mathbf{f}(\mathbf{v}_1), \dots, \mathbf{f}(\mathbf{v}_M)$ to the span of the first M canonical basis vectors of this n -dimensional complex space implies the existence of a symplectic orthogonal matrix O that maps $\mathbf{v}_1, \dots, \mathbf{v}_M$ to the span of the first $2M$ canonical basis vectors of the $2n$ -dimensional real space. To establish the existence of such a unitary, consider the complex vector subspace $W := \text{Span}(\mathbf{f}(\mathbf{v}_1), \dots, \mathbf{f}(\mathbf{v}_M))$ of dimension at most M . By selecting an orthonormal basis for this subspace, we can construct a unitary matrix U that maps this basis to the first $\dim(W) \leq M$ canonical basis vectors. Hence, the existence of the required unitary matrix is confirmed, implying the existence of a symplectic orthogonal matrix O that maps $\mathbf{v}_1, \dots, \mathbf{v}_M$ to the span of the first $2M$ canonical basis vectors. This concludes our proof. ■

It is noteworthy that the “compressed size” obtained for t -doped unitaries, which is $\lceil \kappa t/2 \rceil$, is less than κt , which we proved for t -doped Gaussian states.

2. Compression of t -doped Gaussian particle-number-preserving unitaries

The subsequent Proposition 2 demonstrates that the compression of t -doped Gaussian unitaries can also be achieved in the particle-number-preserving case.

Proposition 2 (Particle-number-preserving t -doped unitaries). Let U_t be a t -doped fermionic Gaussian unitary, as per Definition 13, where all the unitaries that compose U_t are particle-number-preserving. Then U_t can be decomposed as

$$U_t := G_A(u_t \otimes I)G_B, \quad (\text{B18})$$

where G_A and G_B are Gaussian unitaries, which preserve the number of particles (see Definition 11) and u_t is a particle-number-preserving possibly non-Gaussian unitary supported on κt qubits.

Proof. The proof of such proposition follows the same lines as the one of Theorem 5. In fact, by inspecting the proof, it readily follows that the so-defined G_A is particle-number-preserving. The fact that G_B and u_t are particle-number-preserving follows from the condition $G_{\text{aux}}^\dagger |0^n\rangle = |0^n\rangle$, where $G_B := G_{\text{aux}}^\dagger$ and by Lemma 6. ■

3. Circuit complexity of t -doped Gaussian unitaries and states

The circuit complexity of a unitary (state) is defined as the minimum number of $\mathcal{O}(1)$ -local gates needed for

implementing the unitary (state). We will consider locality both in the qubit and in the fermionic sense; in the latter case it refers to the number of distinct Majorana operators that generate each non-Gaussian gate. Our subject of interest is the scaling of the complexity of a t -doped unitary. Throughout this section, we assume $\kappa = \mathcal{O}(1)$ and let $t = t(n)$ change in some way with n . By Definition 13, a t -doped Gaussian unitary U_t can be written as $U_t = G_t W_t \cdots G_1 W_1 G_0$, where $\{G_i\}_{i=0}^t$ are Gaussian unitaries and $\{W_i\}_{i=1}^t$ are, possibly non-Gaussian, κ -local fermionic gates. From this definition and using the fact that any Gaussian unitary can be decomposed as the product of $\leq 2n(2n-1)/2$ (fermionic) 2-local gates [38, 58, 59], the fermionic circuit complexity of t -doped Gaussian unitaries is upper bounded by $\mathcal{O}(n^2 t)$. The same can be shown for qubit circuit complexity (see below). But more importantly, we have proven earlier that a t -doped Gaussian unitary can be decomposed as $U_t = G_A(u_t \otimes I)G_B$, where G_A, G_B are Gaussians and u_t is a unitary on $\lceil \kappa t/2 \rceil$ qubits. In the following, we show that such decomposition reveals an improved upper bound on the circuit complexity of t -doped Gaussian unitaries.

Proposition 3 (Circuit complexity of t -doped Gaussian unitaries). The circuit complexity $\mathcal{C}(U_t)$ of a t -doped Gaussian unitary U_t is (both in the qubit and fermionic sense):

$$\mathcal{C}(U_t) = \begin{cases} \mathcal{O}(n^2 + t^3), & \text{if } \kappa t \leq n \\ \mathcal{O}(n^2 t), & \text{otherwise.} \end{cases} \quad (\text{B19})$$

Proof. Let us assume that $\kappa t \leq n$. Then, U_t can be written as $U_t = G_A(u_t \otimes I)G_B$, where G_A, G_B are Gaussians and u_t is a unitary on $\lceil \kappa t/2 \rceil$ qubits. In fact, u_t is itself a (t, κ) -doped Gaussian unitary on $\lceil \kappa t/2 \rceil$ qubits. It is not directly obvious, but will be shown momentarily; this will imply the desired circuit complexity $\mathcal{O}(t^3)$. We recall our definitions used in the proof of Theorem 4. We have $\tilde{G}_{t'} := G_{t'} \cdots G_0$ for $t' \in [t]$ and a Gaussian unitary G_{aux} , and set

$$u_t := \prod_{t'=1}^t (G_{\text{aux}}^\dagger \tilde{G}_{t'-1}^\dagger W_{t'} \tilde{G}_{t'-1} G_{\text{aux}}) = \prod_{t'=1}^t w_{t'}, \quad (\text{B20})$$

where we defined the unitaries $w_{t'} := G_{\text{aux}}^\dagger \tilde{G}_{t'-1}^\dagger W_{t'} \tilde{G}_{t'-1} G_{\text{aux}}$, which act only on the first $\lceil \kappa t/2 \rceil$ qubits (equivalently, fermionic modes). We note that $w_{t'}$ is generated by κ Majorana operator superpositions of form $\sum_{i=1}^{\kappa t} [\tilde{\mathbf{v}}_j]_i \gamma_i$, $j \in [\kappa(t'-1) + 1, \kappa t']$; here $\tilde{\mathbf{v}}_j := O \mathbf{v}_j$ (cf. notation O and \mathbf{v}_j from the proof of Theorem 4). Hence, for each of these (nonlocal) non-Gaussian unitaries $w_{t'}$, we can find a Gaussian operation $g_{t'}$ on the first $\lceil \kappa t/2 \rceil$ qubits whose associated orthogonal matrix rotates vectors $\tilde{\mathbf{v}}_j$ into the span of the first κ -basis vectors. As a result, we have $w_{t'} = g_{t'}^\dagger \tilde{w}_{t'} g_{t'}$,

where \tilde{w}_t is now a κ -local non-Gaussian unitary generated by the first κ Majorana operators. By implication, it is also a local qubit gate acting on the first $\lceil \kappa/2 \rceil$ qubits. From this it follows that the circuit complexity of each w_t scales as that of a Gaussian g_t . As g_t acts on the first $O(t)$ qubits and fermionic modes, its circuit complexity is $O(t^2)$ both in the qubit and the fermionic sense. Therefore, the circuit complexity of u_t is $tO(t^2) = O(t^3)$. Moreover, the circuit complexity (both qubit and fermionic) to implement G_A and G_B is $O(n^2)$. Putting the above observations together, it follows that the circuit complexity of U_t is $O(n^2 + t^3)$. As long as $\kappa t \leq n$, this upper bound is tighter than the one whose proof follows from the t -doped definition, namely $O(n^2 t)$.

The qubit (and not only fermionic) circuit complexity of $O(n^2 t)$ for $\kappa t > n$ can be found in a similar way as the complexity of $O(t^3)$ we showed for u_t above. In particular, consider any κ -local non-Gaussian fermionic unitary W_t which participates in U_t . Using auxiliary Gaussian rotations, its generating Majorana operators can be mapped to $\{\gamma_1, \dots, \gamma_\kappa\}$, resulting in a unitary supported by the first $\lceil \kappa/2 \rceil$ qubits alone. The asymptotic qubit complexity of U_t is thus determined by that of remaining t Gaussian layers, yielding $O(n^2 t)$ as promised. ■

This proposition reveals that t -doped fermion Gaussian unitaries allow not only a “spatial compression for the magic,” but also a compression of the circuit depth. Since our proof of Theorem 4 is constructive, this provides a useful method for compiling magic matchgate circuits, which might be used in practice to reduce the circuit depth.

4. t -compressible Gaussian states

We now introduce the notion of t -compressible fermionic Gaussian state, a class of states that includes the one of t -doped Gaussian states. We now reiterate Definition 2 for convenience. Throughout this section, we assume that $t \in [n]$.

Definition 15 (t -compressible Gaussian state). A state $|\psi\rangle$ is a t -compressible (Gaussian) state if and only if it can be represented as $|\psi\rangle = G(|\phi\rangle \otimes |0^{n-t}\rangle)$, where G is a Gaussian operation, and $|\phi\rangle$ is a pure state supported solely on the first t qubits.

A t -doped Gaussian state is also a κt -compressible Gaussian state because of Theorem 5. However, the reverse is not true because of circuit complexity arguments: t -doped Gaussian states exhibit a circuit complexity of at most $O(n^2 t)$. In contrast, a t -compressible state features a circuit complexity of $O(n^2 + \exp(t))$, representing the complexity needed for implementing a single Gaussian operation and preparing a generic state supported on t qubits.

In the subsequent proposition, we elucidate the structure of the correlation matrix of any t -compressible state, such as t -doped states.

Proposition 4 (Correlation matrix of a t -compressible Gaussian state). The correlation matrix $C(|\psi\rangle)$ of a t -compressible Gaussian state $|\psi\rangle$ can be expressed as

$$C(|\psi\rangle) = O \bigoplus_{j=1}^n \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix} O^T, \quad (\text{B21})$$

where $\lambda_j \leq 1$ for $j \in [t]$ and $\lambda_j = 1$ for $j \in \{t+1, \dots, n\}$, and $O \in O(2n)$ is an orthogonal matrix.

Proof. As per Definition 15, we represent $|\psi\rangle$ as $|\psi\rangle = G(|\phi\rangle \otimes |0^{n-t}\rangle)$, where G is a Gaussian operation and $|\phi\rangle$ is a pure state supported solely on the first t qubits. Utilizing Lemma 5, we can express the correlation matrix $C(|\psi\rangle)$ as follows:

$$C(|\psi\rangle) = Q C(|\phi\rangle \otimes |0^{n-t}\rangle) Q^T \quad (\text{B22})$$

$$= Q (C(|\phi\rangle) \oplus C(|0^{n-t}\rangle)) Q^T, \quad (\text{B23})$$

where $Q \in O(2n)$ is the orthogonal matrix associated to the Gaussian unitary G . By Eq. (A7), we have

$$C(|0^{n-t}\rangle) = \bigoplus_{j=1}^{n-t} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (\text{B24})$$

Since $C(|\phi\rangle)$ is an antisymmetric real matrix, we can decompose it into its normal form (Lemma 5):

$$C(|\phi\rangle) = O_t \bigoplus_{j=1}^t \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix} O_t^T. \quad (\text{B25})$$

The proof concludes by defining $O := Q(O_t \oplus I_{2n-2t})$. ■

The previous proposition reveals that t -compressible states exhibit at least $n - t$ normal eigenvalues, which are exactly one. This motivates the following definition, in analogy to the stabilizer dimension and nullity defined in the stabilizer case [35,63,64], which found applications in resource theory of magic.

Definition 16 (Gaussian dimension and Gaussian nullity of a state). The Gaussian dimension of a state is defined as the number of the normal eigenvalues of its correlation matrix, which are equal to one, while the Gaussian nullity is defined as the number of the normal eigenvalues of its correlation matrix, which are strictly less than one.

In the following, we show that this is also a sufficient condition for a state to be Gaussian t compressible.

Lemma 8 (Sufficient condition for t compressibility). Let $|\psi\rangle$ be an n -qubit quantum state. If $|\psi\rangle$ has Gaussian dimension $\geq n - t$ (or, equivalently, Gaussian nullity $\leq t$), then $|\psi\rangle$ is a t -compressible Gaussian state.

Proof. The correlation matrix of $|\psi\rangle$ can be written in its normal form as $C(|\psi\rangle) = O\Lambda O^T$, where $\Lambda := \bigoplus_{j=1}^n \begin{pmatrix} 0 & \lambda_j \\ -\lambda_j & 0 \end{pmatrix}$, and $O \in O(2n)$. The $\{\lambda_j\}_{j=1}^n$ are the normal eigenvalues such that the last $n - t$ are equal to one. Consider the state $|\psi'\rangle := G_O^\dagger |\psi\rangle$ where G_O is the Gaussian unitary associated with O . Then, $C(\psi') = O^T C(|\psi\rangle) O = \Lambda$. In particular,

$$\text{Tr}(|\psi'\rangle\langle\psi'| Z_k) = C(\psi')_{2k-1,2k} = \Lambda_{2k-1,2k} = 1, \quad (\text{B26})$$

for each $k \in \{t+1, \dots, n\}$. Therefore, $|\psi'\rangle$ must be of the form $|\psi'\rangle = |\phi\rangle \otimes |0^{n-t}\rangle$, where $|\phi\rangle$ is an arbitrary state on the first t qubits. Therefore, we have $|\psi\rangle = G_O(|\phi\rangle \otimes |0^{n-t}\rangle)$, which is a t -compressible state. ■

Hence, Proposition 4 and Lemma 8 prove the following.

Proposition 5 (Equivalence between t compressibility and $n - t$ Gaussian dimension). A n -qubit state is t compressible if and only if its Gaussian dimension is at least $n - t$ (or, equivalently, its Gaussian nullity is at most t).

Note that Proposition 5 also proves that a quantum state is a pure Gaussian state if and only if its Gaussian dimension is n . Furthermore, as a direct consequence of the proof of Lemma 8, we establish that the Gaussian unitary associated with a t -compressible Gaussian state can be selected as the Gaussian unitary corresponding to any orthogonal matrix placing its correlation matrix in the normal form (Lemma 5). This is summarized as follows.

Lemma 9. Every t -compressible Gaussian state $|\psi\rangle$ can be written as $|\psi\rangle := G_O(|\phi\rangle \otimes |0^{n-t}\rangle)$, where G_O is chosen as the Gaussian unitary associated with an orthogonal matrix $O \in O(2n)$ that arranges its correlation matrix in the normal form described in Lemma 5, and $|\phi\rangle$ is a state supported on t qubits.

APPENDIX C: TOMOGRAPHY ALGORITHM

In this section, we present a detailed and rigorous analysis of the tomography algorithm for t -compressible states outlined in the main text (Algorithm 1). Throughout this section, we assume that $t \in [n]$.

1. Useful lemmas and subroutines

Let us start with a lemma, which gives a sample complexity upper bound to estimate the correlation matrix of a state using single-qubit Pauli-basis measurements. We

recall that the correlation matrix of a state ρ is a real antisymmetric matrix, defined as

$$[C(\rho)]_{j,k} = \text{Tr}(O^{(j,k)} \rho), \quad (\text{C1})$$

where $O^{(j,k)} := -i\gamma_j \gamma_k$, for $j < k \in [2n]$ (and the other elements are given by the antisymmetry of the matrix). Note that $O^{(j,k)}$ are Pauli observables. Thus, we have a total of $M := n(2n - 1)$ Pauli expectation values to estimate.

Lemma 10 (Sample complexity for estimating the correlation matrix by Pauli measurements). Let $\varepsilon_c, \delta > 0$. Assume to have access to $N \geq N_c(n, \varepsilon_c, \delta)$, with

$$N_c(n, \varepsilon_c, \delta) := \left\lceil \frac{8n^3(2n-1)}{\varepsilon_c^2} \log\left(\frac{2n(2n-1)}{\delta}\right) \right\rceil, \quad (\text{C2})$$

copies of an n -qubit state ρ . Utilizing only N single-copies measurements in the Pauli basis, with probability $\geq 1 - \delta$, we can construct an antisymmetric real matrix \hat{C} such that it satisfies

$$\|\hat{C} - C(\rho)\|_\infty \leq \varepsilon_c. \quad (\text{C3})$$

Proof. Let $\varepsilon > 0$ an accuracy parameter to be fixed. For each $j < k \in [2n]$, we measure N' copies of ρ in the Pauli basis corresponding to $O^{(j,k)}$, obtaining outcomes $\{X_m^{(j,k)}\}_{m=1}^{N'}$, where $X_m^{(j,k)} \in \{-1, +1\}$. Let $\hat{C}_{j,k} := 1/N' \sum_{m=1}^{N'} X_m^{(j,k)}$. Hoeffding's inequality (specifically Corollary 1) implies that $N' \geq (4/(2\varepsilon^2)) \log(2M/\delta)$ suffices to guarantee that, with probability at least $1 - \delta/M$, we have $|\hat{C}_{j,k} - \text{Tr}(O^{(j,k)} \rho)| < \varepsilon$. By using the union bound, we conclude that the probability that this holds for any $j < k \in [2n]$ is at least $1 - \delta$. More specifically,

$$\begin{aligned} & \Pr\left(\forall j < k \in [2n] : |\hat{C}_{j,k} - \text{Tr}(O^{(j,k)} \rho)| < \varepsilon\right) \\ &= 1 - \Pr\left(\exists j < k \in [2n] : |\hat{C}_{j,k} - \text{Tr}(O^{(j,k)} \rho)| \geq \varepsilon\right) \end{aligned} \quad (\text{C4})$$

$$\geq 1 - \sum_{j < k \in [2n]} \Pr\left(|\hat{C}_{j,k} - \text{Tr}(O^{(j,k)} \rho)| \geq \varepsilon\right) \quad (\text{C5})$$

$$\geq 1 - \delta. \quad (\text{C6})$$

Therefore, the total number of measurements needed is $N = N'M$. Now, we can conclude by transferring the error to the operator norm. Let $A := \hat{C} - C(\rho)$. For the definition of the operator norm, we have $\|A\|_\infty :=$

$\sup_{|\psi\rangle} \sqrt{|\langle\psi|A^\dagger A|\psi\rangle|}$. Thus, we have

$$\begin{aligned} |\langle\psi|A^\dagger A|\psi\rangle| &\leq \sum_{i,j,k=1}^{2n} |\langle\psi|i\rangle\langle i|A^\dagger|j\rangle\langle j|A|k\rangle\langle k|\psi\rangle| \\ &= 2n\varepsilon^2 \sum_{i,k=1}^{2n} |\langle\psi|i\rangle| |\langle k|\psi\rangle| \leq 4n^2\varepsilon^2 \quad (C7) \end{aligned}$$

where, in the first step, we inserted the resolution of the identity and applied the triangle inequality, in the second step, we applied the upper bound on each matrix element, and, in the last step, we used the Cauchy-Schwartz inequality. Hence, we have $\|\hat{C} - C(\rho)\|_\infty \leq 2n\varepsilon$. We conclude by choosing $\varepsilon_c := \varepsilon/2n$. ■

While the sequential estimation of individual correlation matrix entries by measurements in the Pauli basis, as described above, may not be the most sample-efficient approach, it might be convenient to adopt in an experiment because of its easy implementation scheme. However, instead of independently estimating each correlation matrix entry, one could choose to simultaneously measure mutually commuting observables [66] or utilize the fermionic classical shadow protocol introduced in Refs. [58,67,68]. This refinement would lead to a reduction in sample complexity by a factor of n , at the cost of implementing a slightly more complicated measurement scheme.

For completeness, we present now a lemma, which gives a sample complexity upper bound for estimating the correlation matrix, using a commuting observables measurement scheme. The idea is to partition the observables $O^{(j,k)} := -i\gamma_j\gamma_k$, for $j < k \in [2n]$ into disjoint sets of commuting observables. Subsequently, one employs the fact that commuting Pauli observables can be measured simultaneously via a Clifford measurement [66,83]. A crucial observation is that two different Pauli observables of the form $-i\gamma_j\gamma_k$ commute if and only if they are associated with different Majorana operators. Using this observation, we can partition these $M = (2n-1)n$ observables into $2n-1$ disjoint sets, each containing n commuting Pauli observables. We refer to Ref. [66] Appendix C for details of such a partition, and we omit repeating the construction here. However, we point out that the required Clifford transformations can be chosen to be Gaussian as well.

Lemma 11 (Sample complexity for estimating the correlation matrix by grouping commuting observables). Let $\varepsilon_c, \delta > 0$. Assume to have access to $N \geq N_c(n, \varepsilon_c, \delta)$, with

$$N_c(n, \varepsilon_c, \delta) := \left\lceil \frac{8n^2(2n-1)}{\varepsilon_c^2} \log\left(\frac{2n(2n-1)}{\delta}\right) \right\rceil, \quad (C8)$$

copies of an n -qubit state ρ . Utilizing N single-copy (Gaussian) measurements, with probability $\geq 1 - \delta$, we can construct an antisymmetric real matrix \hat{C} such that it satisfies

$$\|\hat{C} - C(\rho)\|_\infty \leq \varepsilon_c. \quad (C9)$$

Proof. For each of the $2n-1$ sets of commuting Pauli, we find the Clifford U that allows us to simultaneously measure such commuting Pauli in the given set, i.e., we map each of the n Pauli to $\{Z_k\}_{k=1}^n$. Now this Clifford can also be chosen to be Gaussian. Indeed, the key constraint on U is that each of the different Paulis of the form $-i\gamma_j\gamma_k$ with $j < k \in [2n]$ [where pairs (j, k) are nonoverlapping since these Paulis commute] is mapped to $\{Z_k\}_{k=1}^n$ with $Z_k := -i\gamma_{2k-1}\gamma_{2k}$. This constraint can be satisfied by using a Gaussian operation associated to the orthogonal matrix, which is a permutation of the Majorana indices from the commuting Paulis into the Majorana indices from the Z Paulis. Consequently, we measure N' copies of $U\rho U^\dagger$ in the computational basis. Thus, for each $O^{(j,k)}$, we obtain outcomes $\{X_m^{(j,k)}\}_{m=1}^{N'}$, where $X_m^{(j,k)} \in \{-1, +1\}$. The unbiased estimators are $\hat{C}_{j,k} := 1/N' \sum_{m=1}^{N'} X_m^{(j,k)}$. As before, Hoeffding's inequality and union bound imply that $N' \geq (2/\varepsilon^2) \log(2M/\delta)$ suffices to guarantee that the probability of $|\hat{C}_{j,k} - \text{Tr}(O^{(j,k)}\rho)| < \varepsilon$ holding for each $j < k \in [2n]$ is at least $1 - \delta$. Therefore, the total number of measurements needed is $N = N'(2n-1)$. We can conclude as in the previous lemma. ■

Lemma 12 (Perturbation bounds on the normal eigenvalues of correlation matrices). Let A and B be two $2n \times 2n$ antisymmetric real matrices with normal eigenvalues $\{\lambda_k(A)\}_{k=1}^n$ and $\{\lambda_k(B)\}_{k=1}^n$, respectively, ordered in increasing order. Then, we have

$$|\lambda_k(A) - \lambda_k(B)| \leq \|A - B\|_\infty, \quad (C10)$$

for any $k \in [n]$.

Proof. This follows from the fact that $C := iA$ and $D := iB$ are Hermitian matrices. Applying Weyl's perturbation theorem (see Ref. [84], Sec. VI), which states that two $2n \times 2n$ Hermitian matrices C and D , with eigenvalues $c_1 \leq \dots \leq c_{2n}$ and $d_1 \leq \dots \leq d_{2n}$, satisfy

$$\|C - D\|_\infty \geq \max_{j \in [n]} |c_j - d_j|. \quad (C11)$$

Since A and B are antisymmetric, their eigenvalues are $\{\pm i\lambda_k(A)\}_{k=1}^n$ and $\{\pm i\lambda_k(B)\}_{k=1}^n$, respectively. Hence, the eigenvalues of C and D are $\{\pm \lambda_k(A)\}_{k=1}^n$ and $\{\pm \lambda_k(B)\}_{k=1}^n$,

respectively. This implies that

$$\begin{aligned} \|A - B\|_\infty &= \|C - D\|_\infty \geq \max_{j \in [2n]} |c_j - d_j| \\ &= \max_{k \in [n]} |\lambda_k(A) - \lambda_k(B)|. \end{aligned} \quad (\text{C12})$$

■

To formalize our learning algorithm, it is useful to invoke the following well-known lemma.

Lemma 13 (Quantum union bound [70,85,86]). Let $\varepsilon_1, \dots, \varepsilon_M > 0$, where $M \in \mathbb{N}$. Let $\{P_i\}_{i=1}^M$ be projectors, and ρ be a quantum state. If $\text{Tr}(P_i \rho) \geq 1 - \varepsilon_i$ for all $i \in [M]$, then

$$\left\| \rho - \frac{P_M \dots P_1 \rho P_1 \dots P_M}{\text{Tr}(P_M \dots P_1 \rho P_1 \dots P_M)} \right\|_1 \leq 2 \sqrt{\sum_{i \in [M]} \varepsilon_i}. \quad (\text{C13})$$

We now leverage this known lemma to prove the following.

Lemma 14. Let $|\psi\rangle$ be a t -compressible Gaussian state. Given an estimate \hat{C} for the correlation matrix $C(|\psi\rangle)$, there exists a Gaussian operation \hat{G} such that

$$d_{\text{tr}}(|\phi\rangle \otimes |0^{n-t}\rangle, \hat{G}^\dagger |\psi\rangle) \leq \sqrt{(n-t)\|\hat{C} - C(|\psi\rangle)\|_\infty} \quad (\text{C14})$$

where $|\phi\rangle \otimes |0^{n-t}\rangle$ corresponds to the postmeasurement state obtained by measuring the last $n-t$ qubits of the state $\hat{G}^\dagger |\psi\rangle$ in the computational basis and obtaining the outcome corresponding to $|0^{n-t}\rangle$. This event occurs with a probability of at least $1 - (n-t)\|\hat{C} - C\|_\infty$.

Proof. According to Proposition 4, the correlation matrix $C := C(|\psi\rangle)$ can be put in the form $C = O\Lambda O^T$, where $O \in \text{O}(2n)$ and $\Lambda = i \bigoplus_{j=1}^n \lambda_j(C) Y$. Here, $\lambda_j \leq 1$ for $j \in [t]$ and $\lambda_j = 1$ for $j \in \{t+1, \dots, n\}$, and Y represents the Y -Pauli matrix. Let $\varepsilon_c := \|\hat{C} - C\|_\infty$, then we have (because of Lemma 12) that $|\lambda_j(\hat{C}) - \lambda_j(C)| \leq \varepsilon_c$, where $\{\lambda_j(\hat{C})\}_{j=1}^n$ and $\{\lambda_j(C)\}_{j=1}^n$ are the normal eigenvalues of the matrices \hat{C} and C , respectively. Thus, we have

$$\lambda_m(\hat{C}) \geq 1 - \varepsilon_c, \quad (\text{C15})$$

for $m \in \{t+1, \dots, n\}$. We can now express the real anti-symmetric matrix \hat{C} in its normal form $\hat{C} = \hat{O} \hat{\Lambda} \hat{O}^T$, where $\hat{O} \in \text{O}(2n)$ is an orthogonal matrix and $\hat{\Lambda}$ is a matrix of the form $\hat{\Lambda} = i \bigoplus_{j=1}^n \lambda_j(\hat{C}) Y$, with $\lambda_j(\hat{C}) \in \mathbb{R}$ for any $j \in [n]$. Next, consider the state $|\psi'\rangle := \hat{G}^\dagger |\psi\rangle$, where \hat{G} is the

Gaussian unitary associated to \hat{O}^T . It holds that $|C(\psi')_{j,k} - (\hat{\Lambda})_{j,k}| \leq \varepsilon_c$, where we used that $C(\psi') = \hat{O}^T C(|\psi\rangle) \hat{O}$, $\hat{\Lambda} = \hat{O}^T \hat{C} \hat{O}$, Cauchy-Schwarz and the definition of infinity norm. Therefore, we have $C(\psi')_{j,k} \geq (\hat{\Lambda})_{j,k} - \varepsilon_c$. In particular, for $m \in \{t+1, \dots, n\}$, we get

$$\text{Tr}(Z_m \psi') = C(\psi')_{2m-1,2m} \quad (\text{C16})$$

$$\geq (\hat{\Lambda})_{2m-1,2m} - \varepsilon_c \quad (\text{C17})$$

$$= \lambda_m(\hat{C}) - \varepsilon_c \quad (\text{C18})$$

$$\geq 1 - 2\varepsilon_c, \quad (\text{C19})$$

where $Z_m = -i\gamma_{2m-1}\gamma_{2m}$ is the Z -Pauli operator acting on the m th qubit and in the last step we used Eq. (C15). Therefore, we also have $\text{Tr}(|0\rangle\langle 0|_m \psi') \geq 1 - \varepsilon_c$. By using the quantum union bound (Lemma 13), we have

$$d_{\text{tr}}(|\psi'\rangle, |\phi\rangle \otimes |0^{n-t}\rangle) \leq \sqrt{(n-t)\varepsilon_c}, \quad (\text{C20})$$

where $|\phi\rangle \otimes |0^{n-t}\rangle$ is the postmeasurement state after having measured the outcomes corresponding to $|0^{n-t}\rangle$ in the last $n-t$ qubits. By Lemma 13, this scenario occurs with probability at least $1 - (n-t)\varepsilon_c$. ■

In the following, we mention the guarantees of a full pure-state tomography algorithm, which demonstrates optimal dependence on the number of qubits and uses only single-copy measurements, albeit with a trade-off in accuracy compared to other algorithms [3,87]. This is an example of a procedure that we can utilize in our t -qubit full-state tomography step of our learning algorithm.

Lemma 15 (Fast-state tomography [65]). For any unknown n -qubit pure state $|\psi\rangle$, there exists a quantum algorithm that, utilizing $N_{\text{tom}}(n, \varepsilon, \delta) := \mathcal{O}(2^n n \log(1/\delta) \varepsilon^{-4})$ copies of $|\psi\rangle$ and $T_{\text{tom}}(n, \varepsilon, \delta) := \mathcal{O}(4^n n^3 \log(1/\delta) \varepsilon^{-5})$ time, generates a classical representation of a state $|\tilde{\psi}\rangle$ that is ε close to $|\psi\rangle$ in trace distance with probability at least $1 - \delta$. Furthermore, the algorithm requires only single-copy Clifford measurements and classical postprocessing.

Next, we provide a lemma that is useful in the proof of the subsequent Theorem 6.

Lemma 16 (Boosting the probability of success). Let $\delta > 0$ and $N' \in \mathbb{N}$. Consider an algorithm \mathcal{A} that succeeds with a probability of $p_{\text{succ}} \geq \frac{3}{4}$. If we execute \mathcal{A} a total of $m \geq \lceil 2N' + 24 \log(1/\delta) \rceil$ times, then \mathcal{A} will succeed at least N' times with a probability of at least $1 - \delta$.

Proof. We will employ a Chernoff bound 2 to establish this result. Define binary random variables $\{X_i\}_{i=1}^m$ as

follows:

$$X_i = \begin{cases} 1 & \text{if } \mathcal{A} \text{ succeeds,} \\ 0 & \text{if } \mathcal{A} \text{ fails.} \end{cases} \quad (\text{C21})$$

Define $\hat{X} := \sum_{i=1}^m X_i$. We have $\mathbb{E}[\hat{X}] = mp_{\text{succ}}$. Moreover, we aim to upper bound by δ the probability that \mathcal{A} succeeds fewer than N' times, which is $\Pr(\hat{X} \leq N')$. We first write it as $\Pr(\hat{X} \leq N') = \Pr(\hat{X} \leq (1 - \alpha) \mathbb{E}[\hat{X}])$, where we defined $\alpha := 1 - N'/mp_{\text{succ}}$. Note that α satisfies $\alpha \geq \frac{1}{3}$, if

$$m \geq 2N', \quad (\text{C22})$$

exploiting the fact that $p_{\text{succ}} \geq \frac{3}{4}$. Applying the Chernoff bound, we obtain

$$\begin{aligned} \Pr(\hat{X} \leq (1 - \alpha) \mathbb{E}[\hat{X}]) &\leq \exp\left(-\frac{\alpha^2 \mathbb{E}[\hat{X}]}{2}\right) \\ &= \exp\left(-\frac{\alpha^2}{2} p_{\text{succ}} m\right). \end{aligned} \quad (\text{C23})$$

This is upper bounded by δ if

$$m \geq \frac{2}{p_{\text{succ}} \alpha^2} \log\left(\frac{1}{\delta}\right). \quad (\text{C24})$$

Therefore, choosing m as follows satisfies Eqs. (C22) and (C24):

$$m \geq 2N' + \frac{2}{\left(\frac{3}{4}\right) \left(\frac{1}{3}\right)^2} \log\left(\frac{1}{\delta}\right) = 2N' + 24 \log\left(\frac{1}{\delta}\right), \quad (\text{C25})$$

where we used the fact that $p_{\text{succ}} \geq \frac{3}{4}$ and $\alpha \geq \frac{1}{3}$. \blacksquare

2. Joining the pieces: proof of correctness

We now present the main theorem, which puts together the lemmas we have discussed. It demonstrates that to learn t -doped fermionic Gaussian states, or more generally t -compressible Gaussian states, with $t = \mathcal{O}(\log(n))$, only resources scaling polynomially in the number of qubits are required.

Theorem 6 (Efficient learning of t -compressible Gaussian states). Let $|\psi\rangle$ be a t -compressible Gaussian state,

and consider $\varepsilon, \delta \in (0, 1]$. By utilizing

$$N \geq \frac{256n^5}{\varepsilon^4} \log\left(\frac{12n^2}{\delta}\right) + 2N_{\text{tom}}\left(t, \frac{\varepsilon}{2}, \frac{\delta}{3}\right) + 24 \log\left(\frac{3}{\delta}\right) \quad (\text{C26})$$

single-copy measurements and

$$T \geq \mathcal{O}(n^3) + T_{\text{tom}}\left(t, \frac{\varepsilon}{2}, \frac{\delta}{3}\right) \quad (\text{C27})$$

computational time, Algorithm 1 yields a classical representation of a state $|\hat{\psi}\rangle$, satisfying $d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq \varepsilon$ with probability $\geq 1 - \delta$.

Here, $N_{\text{tom}}(t, \varepsilon/2, \delta/3)$ and $T_{\text{tom}}(t, \varepsilon/2, \delta/3)$, respectively, denote the number of copies and computational time sufficient for full-state tomography of a t -qubit state with an $\varepsilon/2$ accuracy and a failure probability of at most $\delta/3$ (using the notation of Lemma 15).

Proof. The learning procedure is outlined in Algorithm 1 in the main text. We now establish its efficiency and correctness. According to Lemma 11, $N_c(n, \varepsilon_c, \delta/3)$ single copies of $|\psi\rangle$ are sufficient to construct an antisymmetric real matrix \hat{C} such that $\|\hat{C} - C\|_{\infty} \leq \varepsilon_c$ with a probability of at least $1 - \delta/3$, where C is the correlation matrix of ρ . Here, we set $\varepsilon_c := \varepsilon^2/(4(n-t))$. Then, we can find the orthogonal matrix $\hat{O} \in \text{O}(2n)$ such that it puts \hat{C} in its normal form Eq. (1), which can be performed in $\mathcal{O}(n^3)$ time. Employing this, we construct the associated Gaussian unitary \hat{G} [a task achievable in time $\mathcal{O}(n^3)$, see [38,58]]. Subsequently, we consider the state $\hat{G}^\dagger |\psi\rangle$. As per Lemma 14, we have

$$d_{\text{tr}}(|\phi\rangle \otimes |0^{n-t}\rangle, \hat{G}^\dagger |\psi\rangle) \leq \frac{\varepsilon}{2}, \quad (\text{C28})$$

where $|\phi\rangle \otimes |0^{n-t}\rangle$ corresponds to the postmeasurement state obtained by measuring the last $n-t$ qubits of the state $\hat{G}^\dagger |\psi\rangle$ in the computational basis and obtaining the outcome corresponding to $|0^{n-t}\rangle$. The probability of such an occurrence, as per Lemma 14, is at least $1 - \varepsilon^2/4 \geq 3/4$. Thus, the algorithm proceeds iteratively by querying a total of m copies of $|\psi\rangle$. In each iteration, it applies the unitary \hat{G}^\dagger to $|\psi\rangle$ and computational basis measurements on the last $n-t$ qubits. By choosing $m := \lceil 2N_{\text{tom}}(t, \varepsilon/2, \delta/3) + 24 \log(3/\delta) \rceil$, it is guaranteed that the measurement outcome corresponding to $|0^{n-t}\rangle$ occurred at least $N_{\text{tom}}(t, \varepsilon/2, \delta/3)$ with probability at least $1 - \delta/3$ (this follows by Lemma 16). Applying the tomography algorithm of Lemma 15 to the first t qubits of all the copies where we obtained the outcome corresponding to $|0^{n-t}\rangle$,

we obtain an output state $|\hat{\phi}\rangle$ such that it is guaranteed that

$$d_{\text{tr}}(|\hat{\phi}\rangle, |\phi\rangle) \leq \frac{\varepsilon}{2}, \quad (\text{C29})$$

with a probability of at least $1 - \delta/3$. Our output state is $|\hat{\psi}\rangle := \hat{G}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle)$, and the information about such a state is provided in the output by providing the orthogonal matrix $\hat{O} \in \text{O}(2n)$, which identifies \hat{G} , and the t -qubit state $|\hat{\phi}\rangle$.

Considering the trace distance between $|\hat{\psi}\rangle$ and $|\psi\rangle$ and applying the triangle inequality with $\hat{G}(|\phi\rangle \otimes |0^{n-t}\rangle)$ as the reference state, we have

$$d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq d_{\text{tr}}(|\hat{\phi}\rangle, |\phi\rangle) + d_{\text{tr}}(|\phi\rangle \otimes |0^{n-t}\rangle, \hat{G}^\dagger |\psi\rangle), \quad (\text{C30})$$

where in the last step we use the unitary invariance of the trace distance and $d_{\text{tr}}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle, |\phi\rangle \otimes |0^{n-t}\rangle) = d_{\text{tr}}(|\hat{\phi}\rangle, |\phi\rangle)$. The algorithm's overall failure probability is contingent on the potential failure of any of the three subroutines—specifically, correlation matrix estimation, measurement of the last $n - t$ qubits, and the tomography protocol. Each subroutine is associated with a failure probability of at most $\delta/3$. Consequently, by the union bound, the algorithm's total failure probability is at most δ . Utilizing Eqs. (C28), (C29), (C30) and assuming the case in which the algorithm does not fail, we deduce $d_{\text{tr}}(|\hat{\psi}\rangle, |\psi\rangle) \leq \varepsilon$. The overall sample complexity is determined by the number of copies needed to estimate the correlation matrix $N_c(n, \varepsilon_c, \delta/3)$ plus the copies m for tomography, i.e., a total number of copies

$$N = \left\lceil \frac{256n^5}{\varepsilon^4} \log\left(\frac{12n^2}{\delta}\right) + 2N_{\text{tom}}\left(t, \frac{\varepsilon}{2}, \frac{\delta}{3}\right) + 24 \log\left(\frac{3}{\delta}\right) \right\rceil, \quad (\text{C31})$$

suffices, which is $\mathcal{O}(\text{poly}(n) + \exp(t))$. On the other hand, the time complexity involves postprocessing of the estimated correlation matrix, requiring $\mathcal{O}(n^3)$ time, and the time complexity for full-state tomography, which is $\mathcal{O}(\exp(t))$. ■

Remark 1. The output state of Algorithm 1 is $|\hat{\psi}\rangle := \hat{G}(|\hat{\phi}\rangle \otimes |0^{n-t}\rangle)$. Specifically, to provide a classical representation of $|\hat{\psi}\rangle$ in the output, it suffices to give the orthogonal matrix $\hat{O} \in \text{O}(2n)$ associated with \hat{G} and the classical description of the t -qubit state $|\hat{\phi}\rangle$. Therefore, the memory necessary to store the classical description of the state outputted by Algorithm 1 is $\mathcal{O}(\text{poly}(n, 2^t))$, similarly to its time and sample complexity.

APPENDIX D: TESTING t -COMPRESSIBLE STATES

In this section, we address the property testing problem, i.e., the problem of determining whether a state is close or far from the set of t -compressible states (or equivalently we test the Gaussian dimension of a state). We begin by establishing an upper bound on the trace distance between a state and the set of t -compressible Gaussian states. Subsequently, we present a lower bound on the same quantity. Finally, we leverage these two bounds to develop a testing algorithm for t -compressible Gaussian states.

We note that the testing problem in the context of fermionic Gaussian states ($t = 0$) was also unsolved. However, a forthcoming paper [76] addresses the testing problem for general, possibly mixed, fermionic Gaussian states. Here, we generalize the results presented in Ref. [76] regarding Gaussian testing to the scenario of t -compressible Gaussian states.

1. Upper bound on the trace distance from the set of t -compressible states

We observed in Lemma 8 that when $n - t$ normal eigenvalues of a state's correlation matrix are precisely one, the state is a t -compressible state. However, when these eigenvalues are close to one, we may inquire about the existence of a t -compressible state in close proximity. This inquiry is formalized in the subsequent Proposition 6.

Proposition 6 (Check the closeness to a t -compressible Gaussian state). Let $|\psi\rangle$ be a quantum state. Let $\{\lambda_i\}_{i=1}^n$ be the normal eigenvalues of its correlation matrix ordered in increasing order. Then, there exists a t -compressible Gaussian state $|\psi_t\rangle$ such that

$$d_{\text{tr}}(|\psi_t\rangle, |\psi\rangle) \leq \sqrt{\sum_{k=t+1}^n \frac{1}{2}(1 - \lambda_k)}. \quad (\text{D1})$$

In particular, $|\psi_t\rangle$ can be chosen as $|\psi_t\rangle := G_O(|\phi\rangle \otimes |0^{n-t}\rangle)$, where G_O is the Gaussian unitary associated with the orthogonal matrix $O \in \text{O}(2n)$ that puts the correlation matrix of $|\psi\rangle$ in its normal form (Lemma 5), and $|\phi\rangle \otimes |0^{n-t}\rangle$ is the state obtained by projecting the state $G_O^\dagger |\psi\rangle$ onto the zero state on the last $n - t$ qubits.

Proof. We can define the state $|\psi'\rangle := G_O^\dagger |\psi\rangle$. We have

$$\begin{aligned} \text{Tr}(|\psi'\rangle\langle\psi'| Z_k) &= C(\psi')_{2k-1, 2k} \\ &= \Lambda_{2k-1, 2k} = \lambda_k = 1 - (1 - \lambda_k), \end{aligned} \quad (\text{D2})$$

for each $k \in \{t+1, \dots, n\}$. By using that $Z_k = 2|0\rangle\langle 0|_k - I$, we have:

$$\text{Tr}(|\psi'\rangle\langle\psi'| 0|0\rangle_k) = 1 - \frac{(1 - \lambda_k)}{2} \quad (\text{D3})$$

By quantum union bound (Lemma 13), we have

$$d_{\text{tr}}(|\psi'\rangle, |\phi\rangle \otimes |0^{n-t}\rangle) \leq \sqrt{\sum_{k=t+1}^n \frac{(1-\lambda_k)}{2}}. \quad (\text{D4})$$

Therefore, by using the unitarity invariance of the trace-norm, we can conclude. ■

From this, it readily follows that trace distance between the state and the set of t -compressible pure Gaussian states \mathcal{G}_t is upper bounded by

$$\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \leq \sqrt{(n-t) \frac{(1-\lambda_{t+1})}{2}}. \quad (\text{D5})$$

2. Lower bound on the trace distance from the set of t -compressible states

In this section, we establish a lower bound on the trace distance of a state from the set of t -compressible Gaussian states. We denote the set of pure t -compressible Gaussian states as \mathcal{G}_t . In the following proof, we follow the derivation presented in Ref. [76] for the case of pure Gaussian states ($t = 0$), extending it to t -compressible states.

Proposition 7. Let $|\psi\rangle$ be a quantum state, and let $\{\lambda_i\}_{i=1}^n$ be the normal eigenvalues of its correlation matrix, ordered in increasing order. The lower bound on the trace distance between the state and the set of t -compressible pure Gaussian states \mathcal{G}_t is given by

$$\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \geq \frac{1}{2}(1 - \lambda_{t+1}). \quad (\text{D6})$$

Proof. Consider an arbitrary operator O with $\|O\|_{\infty} \leq 1$, to be fixed later. Let $|\phi_t\rangle \in \mathcal{G}_t$. Then, we have

$$d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) = \frac{1}{2} \| |\psi\rangle\langle\psi| - |\phi_t\rangle\langle\phi_t| \|_1 \quad (\text{D7})$$

$$\geq \frac{1}{2} |\text{Tr}(O(|\psi\rangle\langle\psi| - |\phi_t\rangle\langle\phi_t|))|, \quad (\text{D8})$$

where in the last step, we used Holder's inequality.

Let $C(|\psi\rangle)$ and $C(|\phi_t\rangle)$ be the correlation matrices of $|\psi\rangle$ and $|\phi_t\rangle$, respectively. Since $C(|\psi\rangle) - C(|\phi_t\rangle)$ is real and antisymmetric, it can be brought into its normal form $C(|\psi\rangle) - C(|\phi_t\rangle) = Q^T \Lambda' Q$, where $\Lambda' = \bigoplus_{j=1}^n i\sigma_j' Y$ and $\{\sigma_j'\}_{j=1}^n$ are the normal eigenvalues of $C(|\psi\rangle) - C(|\phi_t\rangle)$ and $Q \in \text{O}(2n)$ is an orthogonal matrix (Lemma 5).

Now, choose the operator O in the form $O = U_Q^\dagger i\gamma_{2k-1}\gamma_{2k} U_Q$, where $k \in [n]$ and U_Q is a Gaussian unitary associated with the orthogonal matrix $Q \in \text{O}(2n)$.

Note that $\|O\|_{\infty} = 1$. Fix k as a value of j that maximizes $|\sigma_j'|$. Thus, we have

$$|\text{Tr}(O(|\psi\rangle\langle\psi| - |\phi_t\rangle\langle\phi_t|))| = |[C(U_Q |\psi\rangle) - C(U_Q |\phi_t\rangle)]_{2k-1,2k}| \quad (\text{D9})$$

$$= |[Q(C(|\psi\rangle) - C(|\phi_t\rangle))Q^T]_{2k-1,2k}| \quad (\text{D10})$$

$$= |\sigma_k'| \quad (\text{D11})$$

$$= \max_{j \in [n]} |\sigma_j'| \quad (\text{D12})$$

$$= \|C(|\psi\rangle) - C(|\phi_t\rangle)\|_{\infty}, \quad (\text{D13})$$

where in the first step we used the definition of correlation matrix, in the second step we used Lemma 4 and in the last step we used the fact that the largest normal eigenvalues of an antisymmetric matrix corresponds to the infinity norm of the matrix. Therefore, combining with Eq. (D8), we have

$$d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \geq \frac{1}{2} \|C(|\psi\rangle) - C(|\phi_t\rangle)\|_{\infty}. \quad (\text{D14})$$

Now, by applying Lemma 12, we have

$$\|C(|\psi\rangle) - C(|\phi_t\rangle)\|_{\infty} \geq |\lambda(C(|\psi\rangle))_{t+1} - \lambda(C(|\phi_t\rangle))_{t+1}|, \quad (\text{D15})$$

where $\lambda(C)_{t+1}$ denotes the $t+1$ th smallest normal eigenvalue of a correlation matrix C . Since $|\phi_t\rangle$ is a t -compressible Gaussian state, its Gaussian dimension is $n-t$ (because of Proposition 5), hence its $t+1$ th smallest normal eigenvalue must be one. Therefore, the desired lower bound is obtained by taking the minimum over $|\phi_t\rangle \in \mathcal{G}_t$. ■

3. Testing the Gaussian dimension of a state

We present an efficient algorithm (Algorithm 2) for property testing of t -compressible states, where \mathcal{G}_t represents the set of n -qubits t -compressible Gaussian states, or equivalently the set of states with $n-t$ Gaussian dimension. The algorithm takes copies of a state $|\psi\rangle$ as input and determines whether $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \leq \varepsilon_A$ or $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \geq \varepsilon_B$, with the promise that one of these cases is true, where $\varepsilon_B > \varepsilon_A \geq 0$.

We provide the details of the testing algorithm in Algorithm 2. The correctness of this algorithm is established by the following theorem, the proof of which follows the same steps as proofs presented in Ref. [76].

Theorem 7 (Efficient t -compressible Gaussian testing). Let $|\psi\rangle$ be an n -qubit pure state. Assume $\varepsilon_B, \varepsilon_A \in [0, 1]$ such that $\varepsilon_B > \sqrt{(n-t)\varepsilon_A}$, $\delta \in (0, 1]$, and $\varepsilon_{\text{corr}} = (\varepsilon_B^2/(n-t) - \varepsilon_A)$. Assume that $|\psi\rangle$ is such that $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \leq \varepsilon_A$ or $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) > \varepsilon_B$. Then,

ALGORITHM 2. Property testing algorithm for t -compressible Gaussian states.

Input: Error thresholds $\varepsilon_A, \varepsilon_B$, failure probability δ . $N := \lceil 16(n^3/\varepsilon_{\text{corr}}^2) \log(4n^2/\delta) \rceil$ copies of the state $|\psi\rangle$, where

$$\varepsilon_{\text{corr}} = \left(\frac{\varepsilon_B^2}{n-t} - \varepsilon_A \right). \text{ Let } \varepsilon_{\text{test}} := \left(\frac{\varepsilon_B^2}{n-t} + \varepsilon_A \right).$$

Output: Output either “ ε_A -close to the set of t -compressible states” or “ ε_B -far from t -compressible states set”.

- 1 **Step 1:** Estimate the entries of the correlation matrix using N single-copy measurements, resulting in the estimated $2n \times 2n$ matrix $\hat{\Gamma}$;
- 2 **Step 2:** Find $\{\hat{\lambda}_k\}_{k=1}^n$, which corresponds to the ordered normal eigenvalues of $\hat{\Gamma}$;
- 3 **Step 3:** if $\hat{\lambda}_{t+1} \geq 1 - \varepsilon_{\text{test}}$ **then**
- 4 **Output:** “ ε_A -close to the set of t -compressible states”.
- 5 **else**
- 6 **Output:** “ ε_B -far from t -compressible states set.”

Algorithm 2 can discriminate between these two scenarios using $N = \lceil 16(n^3/\varepsilon_{\text{corr}}^2) \log(4n^2/\delta) \rceil$ single-copy measurements of the state $|\psi\rangle$ with a probability of success at least $1 - \delta$.

Proof. Let $\varepsilon_{\text{corr}} > 0$ be an accuracy parameter to be fixed later. By Lemma 11, with $N \geq 8(n^3/\varepsilon_{\text{corr}}^2) \log(4n^2/\delta)$ single-copy measurements, we can find a matrix $\hat{\Gamma}$ such that, with probability at least $1 - \delta$, it holds that $\|\hat{\Gamma} - \Gamma(|\psi\rangle)\|_\infty \leq \varepsilon_{\text{corr}}$. This implies that for all $k \in [n]$, we have $|\hat{\lambda}_k - \lambda_k| \leq \varepsilon_{\text{corr}}$, where $\{\hat{\lambda}_k\}_{k=1}^n, \{\lambda_k\}_{k=1}^n$ are the normal eigenvalues of $\hat{\Gamma}$ and $\Gamma(|\psi\rangle)$, respectively. Let $\varepsilon_{\text{test}}$ be a parameter to fix later. If $\hat{\lambda}_{t+1} \geq 1 - \varepsilon_{\text{test}}$, we aim to show that $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) \leq \varepsilon_B$, otherwise, we aim to show that $\min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) > \varepsilon_A$. Thus, we first assume that $\hat{\lambda}_{t+1} \geq 1 - \varepsilon_{\text{test}}$. From Lemma 6, we have that

$$\begin{aligned} \min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) &\leq \sqrt{\frac{(n-t)}{2}(1 - \lambda_{t+1})} \\ &\leq \sqrt{\frac{(n-t)}{2}(1 - \hat{\lambda}_{t+1} + \varepsilon_{\text{corr}})} \\ &\leq \sqrt{\frac{(n-t)}{2}(\varepsilon_{\text{test}} + \varepsilon_{\text{corr}})}. \end{aligned} \quad (\text{D16})$$

Therefore, we need to ensure that $\sqrt{(n-t)/2}(\varepsilon_{\text{test}} + \varepsilon_{\text{corr}}) \leq \varepsilon_B$. Let us analyze now the case in which $\hat{\lambda}_{t+1} < 1 - \varepsilon_{\text{test}}$. From Lemma 7, we have

$$\begin{aligned} \min_{|\phi_t\rangle \in \mathcal{G}_t} d_{\text{tr}}(|\psi\rangle, |\phi_t\rangle) &\geq \frac{1}{2}(1 - \lambda_{t+1}) \\ &\geq \frac{1}{2}(1 - \hat{\lambda}_{t+1} - \varepsilon_{\text{corr}}) > \frac{1}{2}(\varepsilon_{\text{test}} - \varepsilon_{\text{corr}}). \end{aligned} \quad (\text{D17})$$

Therefore, we impose that $(\varepsilon_{\text{test}} - \varepsilon_{\text{corr}})/2 \geq \varepsilon_A$. The two mentioned inequalities are satisfied by choosing $\varepsilon_{\text{test}} = (\varepsilon_B^2/(n-t) + \varepsilon_A)$ and $\varepsilon_{\text{corr}} = (\varepsilon_B^2/(n-t) - \varepsilon_A)$, by assuming that $\varepsilon_B > \sqrt{(n-t)\varepsilon_A}$. ■

APPENDIX E: PSEUDORANDOMNESS FROM T -DOPED GAUSSIAN STATES AND TIME-COMPLEXITY LOWER BOUND

We have presented an algorithm for learning t -doped fermionic Gaussian states with a time-complexity scaling as $\mathcal{O}(\text{poly}(n, 2^t))$ (assuming that the Majorana locality κ of each non-Gaussian gate is constant in the number of qubits). Thus, as long as $t = \mathcal{O}(\log(n))$, the algorithm is efficient, i.e., the total runtime is polynomial in the number of qubits. In this section, we delve into establishing lower bounds on the time complexity for learning t -doped fermionic Gaussian states. We begin by demonstrating that, under a well-believed cryptography assumption [46–48, 73–75], no algorithm can learn t -doped fermionic Gaussian states with time complexity scaling in t as $\mathcal{O}(\text{poly}(t))$. This rules out efficient algorithms if t scales polynomially with the number of qubits n . However, under a stronger cryptography assumption [46–49], we can establish that when t scales *slightly* faster than logarithmically in the number of qubits n , i.e., $t = \tilde{\omega}(\log(n))$, there exists no efficient algorithm to learn t -doped fermionic Gaussian states, where we defined $\tilde{\omega}(\log(n)) := \omega(\log(n))\text{polyloglog}(n)$.

1. t -doped Gaussian states cannot be learned in polynomial time in t

Our cryptography assumption relies on the conjecture that a specific problem, namely “learning with errors over rings” [73], is hard to solve by quantum computers [46–48, 73–75]. Detailed definitions and discussions about can be found in Ref. [73]. Informally, is a variant of the more general “learning with errors” (LWE) problem specialized to polynomial rings over finite fields, where the LWE problem is to distinguish random linear equations, perturbed by a small amount of noise, from truly uniform ones.

Crucial to our proof is a lemma, adapted from Ref. [72] (which strongly relies on previous work [46, 50, 51]), presented below.

Lemma 17 (Adapted from Theorem 14 in Ref. [72]). Assume that cannot be solved by quantum computers in polynomial time. Then, there exists a set \mathcal{S}_{PRS} of k -qubit

pure quantum states, known as pseudorandom quantum states, with the following properties:

- (1) Any state in \mathcal{S}_{PRS} can be prepared using $\tilde{O}(k)$ Toffoli and Hadamard gates [here, $\tilde{O}(\cdot)$ hides $\text{polylog}(\cdot)$ factors].
- (2) States in the set \mathcal{S}_{PRS} cannot be learned in time complexity $\mathcal{O}(\text{poly}(t))$ by quantum computers. More specifically, let $\rho \in \mathcal{S}_{\text{PRS}}$ be an unknown quantum state. Then there is no quantum algorithm that, using $\mathcal{O}(\text{poly}(k))$ copies of ρ and computational time, with probability at least $2/3$ outputs a classical description of a state $\hat{\rho}$, which can be prepared in polynomial time on a quantum computer such that $d_{\text{tr}}(\rho, \hat{\rho}) \leq 1/8$.

Now we use this lemma to show that there is no algorithm for learning t -doped states with a $\mathcal{O}(\text{poly}(t))$ computational time scaling in t .

Proposition 8 (No $\text{poly}(t)$ algorithm to learn t -doped states). Assume that cannot be solved by quantum computers in polynomial time. Then there is no quantum algorithm that, given access to copies of a t -doped fermionic Gaussian n -qubit state ρ and with a time complexity scaling in t as $\mathcal{O}(\text{poly}(t))$, outputs a classical description of a quantum state $\hat{\rho}$ such that it can be prepared in polynomial time on a quantum computer and, with probability at least $2/3$, it holds that $d_{\text{tr}}(\rho, \hat{\rho}) \leq 1/8$.

Proof. Consider an n -qubit state of the form $|\phi\rangle \otimes |0^{n-k}\rangle$, where $|\phi\rangle$ is a k -qubit state in the set \mathcal{S}_{PRS} defined in Lemma 17. Let $|\phi\rangle = U_{\text{PRS}} |0^k\rangle$, prepared by a unitary U_{PRS} that can be implemented using $\tilde{O}(k)$ Hadamard and Toffoli gates, as per Lemma 17. The Toffoli gates can be implemented in turn using Hadamard, CNOT, T gates, and T gates inverse. For each gate in the system, utilizing a standard SWAP-exchange trick, we can make any gate acting on some qubits of the system to act only on the first two qubits through a cascade of SWAP gates. This incurs a total overhead factor of $\mathcal{O}(k)$ in the total number of gates. The SWAP gates between nearest-neighbor qubits are local non-Gaussian gates, since they can be expressed as $\text{SWAP}_{i,i+1} = e^{-i\frac{\pi}{4}} \exp(i\frac{\pi}{4}(X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1}))$, which have Majorana locality equal to 4 (due to the Jordan-Wigner mapping). Also the other gates, now acting on the first two qubits, are (possibly) non-Gaussian gates with Majorana locality at most 4. This is because the Pauli operators in the generator of each gate can be expressed in terms of Majorana operators via the Jordan-Wigner transformation. This results in a total of $t = \tilde{O}(k^2)$ local non-Gaussian gates required to prepare the state. Thus, the $|\phi\rangle$ is a t -doped Gaussian state with $\kappa = 4$ Majorana local non-Gaussian gates. Now, due to Lemma 17, there exists no learning algorithm to learn such a state in time $\mathcal{O}(\text{poly}(k)) =$

$\mathcal{O}(\text{poly}(t))$, with error less than $1/8$ and probability of success at least $2/3$. ■

2. Learning $\tilde{\omega}(\log(n))$ -doped Gaussian states is hard

The previous proposition rules out efficient algorithms when $t = \Omega(\text{poly}(n))$. However, our proposed algorithm is not anymore efficient when $t = \tilde{\omega}(\log(n))$, while it is efficient for $t = \mathcal{O}(\log(n))$. If we were to make the stronger assumption, namely that cannot be solved by quantum computers in subexponential time [46–49], then we can rule out that the efficient algorithm for $t = \tilde{\omega}(\log(n))$ exists (which means faster than $\omega(\log(n))$ but hiding polyloglog factors). For showing this, we need first the following lemma adapted by Ref. [72].

Lemma 18 (Adapted from Theorem 14 and 15 in Ref. [72]). Assume that cannot be solved by quantum computers in subexponential time. Then, there exists a set \mathcal{S}_{PRS} of k -qubit pure quantum states, known as pseudorandom quantum states, with the following properties:

- (1) Any state in \mathcal{S}_{PRS} can be prepared using $\tilde{O}(k)$ Toffoli and Hadamard gates [here, $\tilde{O}(\cdot)$ hides $\text{polylog}(\cdot)$ factors].
- (2) Any algorithm to learn states from the set \mathcal{S}_{PRS} must have $\exp(\Omega(k))$ time complexity. More specifically, let $\rho \in \mathcal{S}_{\text{PRS}}$ be an unknown quantum state. Then any quantum algorithm that, by querying copies of ρ , with probability at least $2/3$ outputs a classical description of a state $\hat{\rho}$, which can be prepared in subexponential time on a quantum computer such that $d_{\text{tr}}(\rho, \hat{\rho}) \leq 1/8$, must have $\exp(\Omega(k))$ time complexity.

Now we use this to prove the following proposition, also stated informally in the main text as Theorem 3. It reaches a stronger conclusion than Proposition 8, but at the cost of stronger cryptography assumption. The idea of the following proof is to use a more compact qubits-to-fermion mapping than Jordan-Wigner, namely a modification of the one introduced by Kitaev [25], which would allow creation of pseudorandom quantum states with no overhead in the number of non-Gaussian gates compared to the number of Toffoli and Hadamard gates. In the following, we recall that $\tilde{\omega}(\log(n))$ is defined as $\tilde{\omega}(\log(n)) := \omega(\log(n)\text{polyloglog}(n))$.

Proposition 9 (Learning $\tilde{\omega}(\log(n))$ -doped Gaussian states is hard). Assume that cannot be solved by quantum computers in sub-exponential time. Then, there is no efficient (i.e., $\mathcal{O}(\text{poly}(n))$ time) quantum algorithm that, by querying copies of a $\tilde{\omega}(\log(n))$ -doped Gaussian state ρ , with probability at least $2/3$ outputs a description of a state $\hat{\rho}$, which can be prepared in polynomial time on a quantum computer such that $d_{\text{tr}}(\rho, \hat{\rho}) \leq 1/8$.

Proof. Following a similar approach as in the proof of Proposition 8, we begin by defining a state $|\phi\rangle \otimes |0^{n-k}\rangle$, where $|\phi\rangle = U_{\text{PRS}} |0^k\rangle$ represents a k -qubit state in the set \mathcal{S}_{PRS} as defined in Lemma 18. This state can be efficiently prepared using a unitary U_{PRS} with $\tilde{O}(k)$ Hadamard and Toffoli gates, which, in turn, can be implemented using Hadamard gates, CNOTs, T gates, and their inverses. In contrast to the previous proof in Proposition 8, an application of the same argument would not yield the desired conclusion due to the unfavorable quadratic overhead introduced by the SWAP-exchange trick in the number of non-Gaussian gates. That trick was necessary due to the use of Jordan-Wigner transformation. However, we can employ a more efficient qubit-to-fermion mapping, specifically a modified version of the one introduced by Kitaev [25]. Hereafter, we will refer to it as the “Kitaev encoding.” Our objective is to construct a fermionic state encoding $|\phi\rangle$ using a circuit of size $\tilde{O}(k)$ composed of Gaussian and $\kappa = 4$ local non-Gaussian gates. We employ a mapping of k qubits into $2k$ fermionic modes using Majorana operators $\{\gamma_{\alpha,j} | j \in [k], \alpha \in \{0, x, y, z\}\}$. These $4k$ Majorana operators are defined in terms of $4k$ Pauli strings via Jordan-Wigner, with an arbitrarily fixed operator ordering. We are now going to leverage the formalism and basics of stabilizer codes, for more in-depth information, refer to Ref. [88]. The Kitaev encoding involves defining a stabilizer code of $2k$ physical qubits encoded in k logical qubits, characterized by the following k stabilizer generators $\{s_j\}_{j=1}^k$ and logical Pauli operators $\{X_j^{\text{KE}}, Z_j^{\text{KE}}\}_{j=1}^k$:

$$s_j := \gamma_{0,j} \gamma_{x,j} \gamma_{y,j} \gamma_{z,j}, \quad (\text{E1})$$

$$X_j^{\text{KE}} := i\gamma_{y,j} \gamma_{z,j}, \quad (\text{E2})$$

$$Z_j^{\text{KE}} := i\gamma_{x,j} \gamma_{y,j}, \quad (\text{E3})$$

for each $j \in [k]$. Note that these operators explicitly satisfy the algebraic conditions on stabilizer generators and logical Pauli operators. The Kitaev encoding is associated to a Clifford transformation V_{KE} such that

$$V_{\text{KE}} X_j V_{\text{KE}}^\dagger = X_j^{\text{KE}}, \quad (\text{E4})$$

$$V_{\text{KE}} Z_j V_{\text{KE}}^\dagger = Z_j^{\text{KE}}, \quad (\text{E5})$$

for each $j \in [k]$, and

$$V_{\text{KE}} Z_j V_{\text{KE}}^\dagger = s_j, \quad (\text{E6})$$

for each $j \in \{k+1, \dots, 2k\}$. The last equation ensures that $V_{\text{KE}} |0^{2k}\rangle$ is an eigenstate with $+1$ eigenvalue for each of the stabilizer generators $\{s_j\}_{j=1}^k$, while Eq. (E5) implies that $V_{\text{KE}} |0^{2k}\rangle$ is an eigenstate with $+1$ eigenvalue for each $\{Z_j^{\text{KE}}\}_{j=1}^k$. Thus, $V_{\text{KE}} |0^{2k}\rangle$ is a valid “logical zero” stabilizer state. Exploiting the fact that $V_{\text{KE}} |0^{2k}\rangle$ is an eigenstate with

$+1$ eigenvalues of $\{Z_j^{\text{KE}}\}_{j=1}^k$ and $\{s_j\}_{j=1}^k$, its density matrix can be written as

$$\begin{aligned} V_{\text{KE}} |0^{2k}\rangle \langle 0^{2k}| V_{\text{KE}}^\dagger &= \prod_{j=1}^k \left(\frac{I + s_j Z_j^{\text{KE}}}{2} \right) \prod_{j=k+1}^{2k} \left(\frac{I + Z_j^{\text{KE}}}{2} \right) \\ &= \prod_{j=1}^k \left(\frac{I - i\gamma_{0,j} \gamma_{z,j}}{2} \right) \prod_{j=k+1}^{2k} \left(\frac{I + i\gamma_{x,j} \gamma_{y,j}}{2} \right). \end{aligned} \quad (\text{E7})$$

From this, we observe that $V_{\text{KE}} |0^{2k}\rangle$ is a fermionic Gaussian state because it can be written in the form of Eq. (A4), noting that signed permutation matrices are orthogonal matrices. Moreover, we denote $Y_j^{\text{KE}} := -iZ_j^{\text{KE}} X_j^{\text{KE}} = i\gamma_{x,j} \gamma_{z,j}$. The Kitaev encoding, as defined, ensures that the local qubit gates are mapped onto local fermionic ones. In particular, the gates of the circuit U_{PRS} that prepares the pseudorandom state $|\phi\rangle = U_{\text{PRS}} |0^k\rangle$ —the Hadamard H , CNOT, and T gate—are, up to an overall phase, mapped onto

$$\begin{aligned} H_j^{\text{KE}} &:= V_{\text{KE}} H_j V_{\text{KE}}^\dagger = V_{\text{KE}} \left(Z_j \frac{I + iY_j}{\sqrt{2}} \right) V_{\text{KE}}^\dagger \\ &= Z_j^{\text{KE}} \frac{I + iY_j^{\text{KE}}}{\sqrt{2}} = e^{-\frac{\pi}{2} \gamma_{x,j} \gamma_{y,j}} e^{-\frac{\pi}{4} \gamma_{z,j} \gamma_{x,j}}, \end{aligned} \quad (\text{E8})$$

$$\begin{aligned} \text{CNOT}_{j,l}^{\text{KE}} &:= V_{\text{KE}} \text{CNOT}_{j,l} V_{\text{KE}}^\dagger = e^{i\frac{\pi}{4} (1 - Z_j^{\text{KE}})(I - X_l^{\text{KE}})} \\ &= e^{i\frac{\pi}{4} (I - i\gamma_{x,j} \gamma_{y,j})(I - i\gamma_{y,l} \gamma_{z,l})}, \end{aligned} \quad (\text{E9})$$

$$T_j^{\text{KE}} := V_{\text{KE}} T_j V_{\text{KE}}^\dagger = e^{i\frac{\pi}{8} Z_j^{\text{KE}}} = e^{-\frac{\pi}{8} \gamma_{x,j} \gamma_{y,j}}, \quad (\text{E10})$$

for each $j \neq l \in [k]$. The only non-Gaussian among these is the encoding of the CNOT gate, which has Majorana locality $\kappa = 4$. This implies that the encoding of the circuit U_{PRS} , i.e., $U_{\text{PRS}}^{\text{KE}} := V_{\text{KE}} U_{\text{PRS}} V_{\text{KE}}^\dagger$, is a t -doped fermionic Gaussian unitary with local non-Gaussian gates and $t = \tilde{O}(k)$. Thus, we have that the Kitaev encoding of the pseudorandom state $|\phi\rangle = U_{\text{PRS}} |0^k\rangle$ is

$$\begin{aligned} |\phi\rangle_{\text{KE}} &:= V_{\text{KE}} |\phi\rangle \otimes |0^k\rangle = V_{\text{KE}} U_{\text{PRS}} |0^{2k}\rangle \\ &= U_{\text{PRS}}^{\text{KE}} V_{\text{KE}} |0^{2k}\rangle. \end{aligned} \quad (\text{E11})$$

Since $V_{\text{KE}} |0^{2k}\rangle$ is a Gaussian state and $U_{\text{PRS}}^{\text{KE}}$ is a t -doped Gaussian unitary with $t = \tilde{O}(k)$, then $V_{\text{KE}} |\phi\rangle \otimes |0^k\rangle$ is a t -doped Gaussian state with $t = \tilde{O}(k)$. The Majorana locality of each non-Gaussian gate is at most 4.

Using an arbitrary algorithm \mathcal{A} for learning a t -doped fermionic Gaussian state, we will now define a protocol for learning the pseudorandom state $|\phi\rangle$. Given a copy of a state $|\phi\rangle$ on k qubits, we use k auxiliary qubits in state $|0\rangle$ and apply the Clifford transformation V_{KE} . This means

it can be produced by a circuit with $O(k^2)$ two-qubit gates [89]. The resulting $|\phi\rangle_{\text{KE}}$ can be input to \mathcal{A} as a copy of a t -doped fermionic Gaussian state for $t = \tilde{O}(k)$. Using the number of copies of $|\phi\rangle$ given by sample complexity of \mathcal{A} , we learn a description of a state $\hat{\rho}_{\text{KE}}$, which, with probability at least $2/3$, satisfies

$$d_{\text{tr}}(\hat{\rho}_{\text{KE}}, \rho_{\text{KE}}) \leq \frac{1}{8}, \quad (\text{E12})$$

where we defined ρ_{KE} to be the density matrix associated with $|\phi\rangle_{\text{KE}}$. By defining $\hat{\rho} := \text{Tr}_{\{k+1, \dots, 2k\}}(V_{\text{KE}}^\dagger \hat{\rho}_{\text{KE}} V_{\text{KE}})$, where $\text{Tr}_{\{k+1, \dots, 2k\}}(\cdot)$ indicates the partial trace with respect to the qubits $\{k+1, \dots, 2k\}$, we also have

$$\begin{aligned} d_{\text{tr}}(\hat{\rho}, |\phi\rangle\langle\phi|) &\leq d_{\text{tr}}(V_{\text{KE}}^\dagger \hat{\rho}_{\text{KE}} V_{\text{KE}}, |\phi\rangle\langle\phi| \otimes |0^k\rangle\langle 0^k|) \\ &= d_{\text{tr}}(V_{\text{KE}}^\dagger \hat{\rho}_{\text{KE}} V_{\text{KE}}, V_{\text{KE}}^\dagger \rho_{\text{KE}} V_{\text{KE}}) \\ &= d_{\text{tr}}(\hat{\rho}_{\text{KE}}, \rho_{\text{KE}}) \leq \frac{1}{8}, \end{aligned} \quad (\text{E13})$$

where in the first step we used that the partial trace does not increase the trace distance between two states [90]. Hence, we found a state $\hat{\rho}$, which is in trace distance close to the target state $|\phi\rangle\langle\phi|$. To recap, we produced the learning algorithm for a pseudorandom state $|\phi\rangle$ from a learning algorithm for t -doped fermionic Gaussian states. The pseudorandom state learning algorithm has the same sample complexity as the fermionic one, and the time complexity $T_{\text{PRS}} = S_f \cdot O(k^2) + T_f$ where T_f and S_f are time and sample complexity of the fermionic learner. If there is a fermionic learner whose time and sample complexity scale subexponentially in k , the same property carries over to the pseudorandom states' learner. By Lemma 18, this would contradict the cryptographic assumption that cannot be solved by quantum computers in subexponential time. Hence, the time complexity of the fermionic learner needs to be $\exp(\Omega(k))$. If $k = \omega(\log(n))$, then this implies that any algorithm to learn t -doped fermionic Gaussian states with $t = \tilde{O}(k) = \omega(\log(n)\text{polyloglog}(n))$ must be inefficient, i.e., its time complexity must be $\omega(\text{poly}(n))$. ■

APPENDIX F: GENERALIZATION TO THE MIXED-STATE SCENARIO AND NOISE-ROBUSTNESS OF THE ALGORITHM

Many of the concepts introduced in our work can be extended to the more realistic mixed-state scenario. We start by providing the definition for t -compressible mixed states.

Definition 17 (Mixed t -compressible state). A possibly mixed state ρ is a t -compressible (Gaussian) state if and

only if it can be represented as

$$\rho = G(\sigma \otimes |0^{n-t}\rangle\langle 0^{n-t}|)G^\dagger, \quad (\text{F1})$$

where G is a Gaussian operation, and σ is a state supported solely on the first t qubits.

As in the pure-state scenario (Proposition 5), saying that a quantum state is t compressible is equivalent to the fact that at least $n - t$ normal eigenvalues of the state's correlation matrix are equal to one (i.e., its Gaussian dimension is $\geq n - t$ and its Gaussian nullity is $\leq t$).

Proposition 10 (Equivalent definition of t -compressible state). A possibly mixed n -qubit state ρ is t compressible if and only if at least $n - t$ normal eigenvalues of the correlation matrix of ρ are equal to one.

The proof is analogous to that of Proposition 5.

For the tomography algorithm, analogous guarantees extend from the t -compressible pure case to the t -compressible mixed-state scenario. In this case, the unknown state is promised to be a t -compressible Gaussian (possibly mixed) state. However, in this mixed-state scenario, the algorithm is simpler, since measuring the last $n - t$ qubits in the computational basis is not necessary because we do not require the output state to be pure, i.e., step 5 in Algorithm 1 is not necessary. Instead, the full-state tomography of the first t qubits can be applied right after having applied the Gaussian operation to the state. We explicitly show this in Algorithm 3 of the following subsection, whose proof of correctness is shown in the (more general) Theorem 9.

Furthermore, the following inequalities can be helpful to test how much a possibly mixed state ρ deviates (in trace distance) from the set of mixed t -compressible Gaussian states.

Proposition 11 (Minimum distance between a possibly mixed state and the set of t -compressible states). Let ρ be a possibly mixed quantum state, and let $\{\lambda_i\}_{i=1}^n$ be the normal eigenvalues of its correlation matrix, ordered in increasing order. The following trace distance upper and lower bounds between the state and the set of t -compressible mixed Gaussian states \mathcal{G}_t hold

$$\sqrt{\sum_{k=t+1}^n \frac{(1 - \lambda_k)}{2}} \geq \min_{\sigma_t \in \mathcal{G}_t} d_{\text{tr}}(\rho, \sigma_t) \geq \frac{1}{2}(1 - \lambda_{t+1}). \quad (\text{F2})$$

The proof of this theorem is analogous to those of Proposition 6 and Proposition 7, which held for the pure-state scenario but readily apply to the more general mixed-state scenario.

Thus, with the previous inequalities in hand, we can also provide a property testing algorithm that applies to a possibly mixed quantum state.

Theorem 8 (Efficient t -compressible Gaussian testing). Let ρ be an n -qubit possibly mixed quantum state. Assume $\varepsilon_B, \varepsilon_A \in [0, 1]$ such that $\varepsilon_B > \sqrt{(n-t)\varepsilon_A}$, $\delta \in (0, 1]$, and $\varepsilon_{\text{corr}} = (\varepsilon_B^2/(n-t) - \varepsilon_A)$. Assume that ρ is such that $\min_{\sigma_t \in \mathcal{G}_t} d_{\text{tr}}(\rho, \sigma_t) \leq \varepsilon_A$ or $\min_{\sigma_t \in \mathcal{G}_t} d_{\text{tr}}(\rho, \sigma_t) > \varepsilon_B$. Then, Algorithm 2 can discriminate between these two scenarios using $N = \lceil 16(n^3/\varepsilon_{\text{corr}}^2) \log(4n^2/\delta) \rceil$ single-copy measurements of the state ρ with a probability of success at least $1 - \delta$.

The proof of this theorem is similar to that of Theorem 7.

1. Tomography of approximately t -compressible states

In this subsection, we address the problem of tomography for “approximate t -compressible states,” i.e., states with $n - t$ normal eigenvalues of their correlation matrix that are not exactly one, but *almost*.

This approach is particularly useful for the tomography of states where most of the normal eigenvalues are very close to one, such as states obtained when considering impurity models [45], as we will explore in more detail in the next section. Moreover, when preparing a target t -compressible state in a quantum device (e.g., a target pure state prepared by a 1D matchgate circuit doped with a few SWAP gates), noise can cause an effective state ρ to be prepared instead of the target t -compressible state. Therefore, an algorithm that allows for perturbations from the set of t -compressible states is experimentally motivated.

We provide such an algorithm in Table 3. Note that this differs from Algorithm 1 because in this mixed-state scenario we do not need to measure the last $n - t$ qubits in the computational basis before performing full-state tomography on the first t qubits.

Theorem 9 (Efficient learning of approximate t -compressible Gaussian states). Let $\varepsilon, \delta \in (0, 1]$. Let ρ be a quantum state. By utilizing $N = \mathcal{O}(n^5/(\varepsilon^4) \log(n^2/\delta)) +$

$N_{\text{tom}}(t, \varepsilon/2, \delta/2)$ single-copy measurements and $T = \mathcal{O}(n^3) + T_{\text{tom}}(t, \varepsilon/2, \delta/2)$ computational time, Algorithm 3 yields a classical representation of a state $\hat{\rho}$, satisfying

$$d_{\text{tr}}(\hat{\rho}, \rho) \leq \varepsilon + \varepsilon_t, \quad (\text{F3})$$

with probability $\geq 1 - \delta$. Here, $\varepsilon_t := \sqrt{\sum_{j=t+1}^n (1 - \lambda_j)/2}$, where $\{\lambda_j\}_{j=1}^n$ are the normal eigenvalues of the correlation matrix of ρ ordered in increasing order. Moreover, $N_{\text{tom}}(t, \varepsilon/2, \delta/2)$ and $T_{\text{tom}}(t, \varepsilon/2, \delta/2)$, respectively, denote the number of copies and computational time sufficient for full-state tomography of a mixed t -qubit state with an $\varepsilon/2$ accuracy and a failure probability of at most $\delta/2$.

Proof. The learning procedure is outlined in Algorithm 3 and we now establish its correctness. According to Lemma 11, $N_c(n, \varepsilon_c, \delta/2)$ single copies of ρ are sufficient to construct an antisymmetric real matrix \hat{C} such that $\|\hat{C} - C\|_\infty \leq \varepsilon_c$ with a probability of at least $1 - \delta/2$, where C is the correlation matrix of ρ . Then, we can put \hat{C} in its normal form $\hat{C} = \hat{O}\hat{\Lambda}\hat{O}^T$, where $\hat{O} \in \text{O}(2n)$ is an orthogonal matrix and $\hat{\Lambda}$ is a matrix of the form $\hat{\Lambda} = i \bigoplus_{j=1}^n \lambda_j(\hat{C}) Y$, where $\{\lambda_j(\hat{C})\}_{j=1}^n$ are the normal eigenvalues of \hat{C} . We can then construct [38,58,59] the Gaussian unitary \hat{G} associated to \hat{O} . Subsequently, we consider the state $\rho' := \hat{G}^\dagger \rho \hat{G}$. For $m \in \{t+1, \dots, n\}$, we have

$$\text{Tr}(Z_m \rho') = C(\rho')_{2m-1,2m} \quad (\text{F4})$$

$$\geq (\hat{\Lambda})_{2m-1,2m} - \varepsilon_c \quad (\text{F5})$$

$$= \lambda_m(\hat{C}) - \varepsilon_c \quad (\text{F6})$$

$$\geq \lambda_m(C) - 2\varepsilon_c, \quad (\text{F7})$$

where $Z_m = -i\gamma_{2m-1}\gamma_{2m}$ is the Z -Pauli operator acting on the m th qubit. In the second step, we used that $|C(\psi')_{j,k} - (\hat{\Lambda})_{j,k}| \leq \varepsilon_c$, as follows by $C(\rho') = \hat{O}^T C(\rho) \hat{O}$, $\hat{\Lambda} = \hat{O}^T \hat{C} \hat{O}$, Cauchy-Schwarz and the definition of infinity norm. In the last step we used that, because of Lemma 12, we have $|\lambda_j(\hat{C}) - \lambda_j(C)| \leq \varepsilon_c$ for each $j \in [n]$, where

ALGORITHM 3. Learning algorithm for approximately t -compressible fermionic Gaussian states (possibly mixed).

Input: Accuracy ε , failure probability δ , $N = \mathcal{O}\left(\frac{n^5}{\varepsilon^4} \log\left(\frac{n^2}{\delta}\right)\right) + N_{\text{tom}}(t, \frac{\varepsilon}{2}, \frac{\delta}{2})$ copies of the state ρ , where N_{tom} is the number of copies needed for t -qubit state tomography with accuracy $\frac{\varepsilon}{2}$ and failure probability $\frac{\delta}{2}$.

Output: A classical description of $\hat{\rho}$, ensuring $d_{\text{tr}}(\hat{\rho}, \rho) \leq \varepsilon + \varepsilon_t$ with probability at least $1 - \delta$. Here, $\varepsilon_t := \sqrt{\sum_{j=t+1}^n \frac{(1-\lambda_j)}{2}}$, where $\{\lambda_j\}_{j=1}^n$ are the normal eigenvalues of the correlation matrix of ρ ordered in increasing order.

- 1 Estimate the correlation matrix of ρ using $\mathcal{O}\left(\frac{n^5}{\varepsilon^4} \log\left(\frac{n^2}{\delta}\right)\right)$ single-copy measurements (see Lemma 38), obtaining \hat{C} ;
 - 2 Express \hat{C} in its normal form $\hat{C} = \hat{O}\hat{\Lambda}\hat{O}^T$ (Eq.(1)) and find the Gaussian unitary $G_{\hat{O}}$ associated with $\hat{O} \in \text{O}(2n)$;
 - 3 Using $N_{\text{tom}}(t, \frac{\varepsilon}{2}, \frac{\delta}{2})$ copies, perform full-state tomography on the first t qubits of the state $G_{\hat{O}}^\dagger \rho G_{\hat{O}}$, and let $\hat{\sigma}$ be the t -qubit state output by tomography;
 - 4 **return** \hat{O} and $\hat{\sigma}$, which identify $\hat{\rho} := G_{\hat{O}}(\hat{\sigma} \otimes |0^{n-t}\rangle\langle 0^{n-t}|)G_{\hat{O}}^\dagger$;
-

$\{\lambda_j(C)\}_{j=1}^n$ are the normal eigenvalues of C . Therefore, we also have

$$\begin{aligned} \text{Tr}(|0\rangle\langle 0| 0_m \rho') &= \frac{1}{2} + \frac{1}{2} \text{Tr}(Z_m \rho') \\ &\geq 1 - \left(\varepsilon_c + \frac{1 - \lambda_m(C)}{2} \right). \end{aligned} \quad (\text{F8})$$

By using the quantum union bound (Lemma 13), we have

$$d_{\text{tr}}(\rho', \rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|) \leq \sqrt{(n-t)\varepsilon_c + \varepsilon_t^2}, \quad (\text{F9})$$

where $\varepsilon_t := \sqrt{\sum_{j=t+1}^n (1 - \lambda_j(C))/2}$, and $\rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|$ is the postmeasurement state after having measured the outcomes corresponding to $|0^{n-t}\rangle$ in the last $n-t$ qubits. Let ρ'_t be the reduced state of ρ' over the first t qubits. We have

$$\begin{aligned} d_{\text{tr}}(\rho', \rho'_t \otimes |0^{n-t}\rangle\langle 0^{n-t}|) &\leq d_{\text{tr}}(\rho', \rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|) \\ &\quad + d_{\text{tr}}(\rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho'_t \otimes |0^{n-t}\rangle\langle 0^{n-t}|) \end{aligned} \quad (\text{F10})$$

$$= d_{\text{tr}}(\rho', \rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|) + d_{\text{tr}}(\rho'_{\text{post}}, \rho'_t) \quad (\text{F11})$$

$$\leq 2d_{\text{tr}}(\rho', \rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|) \quad (\text{F12})$$

$$\leq 2\sqrt{(n-t)\varepsilon_c + \varepsilon_t^2}, \quad (\text{F13})$$

where in the first step we used the triangle inequality, in the second step we used the unitary invariance of the trace distance, in the third step we used the data-processing inequality [90] (with respect to the operation of tracing out a subsystem), i.e., $d_{\text{tr}}(\rho'_{\text{post}}, \rho'_t) \leq d_{\text{tr}}(\rho'_{\text{post}} \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho')$, and in the last step we used Eq. (F9). Using now $N_{\text{tom}}(t, \varepsilon/2, \delta/2)$ copies of the state $\rho' = \hat{G}^\dagger \rho \hat{G}$, we perform full-state tomography on the first t qubits, obtaining $\hat{\sigma}$ such that

$$d_{\text{tr}}(\hat{\sigma}, \rho'_t) \leq \varepsilon/2, \quad (\text{F14})$$

with $\geq 1 - \delta/2$ probability. Our output state is $\hat{\rho} := \hat{G}(\hat{\sigma} \otimes |0^{n-t}\rangle\langle 0^{n-t}|)\hat{G}^\dagger$, and the information about such a state is provided in the output by providing the orthogonal matrix $\hat{O} \in \text{O}(2n)$, which identifies \hat{G} , and the t -qubit state $\hat{\sigma}$. Considering the trace distance between $\hat{\rho}$ and ρ ,

we have

$$d_{\text{tr}}(\hat{\rho}, \rho) = d_{\text{tr}}(\hat{\sigma} \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \hat{G}^\dagger \rho \hat{G}) \quad (\text{F15})$$

$$= d_{\text{tr}}(\hat{\sigma} \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho') \quad (\text{F16})$$

$$\leq d_{\text{tr}}(\hat{\sigma} \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho'_t \otimes |0^{n-t}\rangle\langle 0^{n-t}|) + d_{\text{tr}}(\rho'_t \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho') \quad (\text{F17})$$

$$\leq d_{\text{tr}}(\hat{\sigma}, \rho'_t) + d_{\text{tr}}(\rho'_t \otimes |0^{n-t}\rangle\langle 0^{n-t}|, \rho') \quad (\text{F18})$$

$$\leq \frac{\varepsilon}{2} + 2\sqrt{(n-t)\varepsilon_c + \varepsilon_t}, \quad (\text{F19})$$

where in the last step we used Eqs. (F13) and (F14). By setting $\varepsilon_c := \varepsilon^2/(16(n-t))$, we get

$$d_{\text{tr}}(\hat{\rho}, \rho) \leq \frac{\varepsilon}{2} + 2\sqrt{\frac{\varepsilon^2}{16} + \varepsilon_t^2} \leq \varepsilon + \varepsilon_t, \quad (\text{F20})$$

where we used that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for each $a, b \geq 0$. The total failure probability of the algorithm, by union bound, is $\leq \delta$. ■

As a consequence of the previous theorem, we present the following proposition, which states that if the unknown quantum state is sufficiently close in trace distance to the set of t -compressible Gaussian states, then our learning algorithm (Algorithm 3) can still be effectively applied.

Theorem 10 (Noise robustness of learning t -compressible Gaussian states). Let $\varepsilon, \delta \in (0, 1]$. Let ρ be a quantum state such that

$$\min_{\sigma_t \in \mathcal{G}_t} d_{\text{tr}}(\rho, \sigma_t) \leq \frac{\varepsilon^2}{4(n-t)}, \quad (\text{F21})$$

i.e., it is sufficiently close to the set of possibly mixed t -compressible states \mathcal{G}_t . Then, there exists a learning algorithm, which, utilizing

$$N = \mathcal{O}\left(\frac{n^5}{\varepsilon^4} \log\left(\frac{n^2}{\delta}\right)\right) + N_{\text{tom}}\left(t, \frac{\varepsilon}{4}, \frac{\delta}{2}\right) \quad (\text{F22})$$

single-copy measurements, yields a classical representation of a state $\hat{\rho}$, satisfying $d_{\text{tr}}(\hat{\rho}, \rho) \leq \varepsilon$ with probability $\geq 1 - \delta$.

Here, $N_{\text{tom}}(t, \varepsilon/4, \delta/2)$ denotes the number of copies sufficient for full-state tomography of a mixed t -qubit state with an $\varepsilon/4$ accuracy and a failure probability of at most $\delta/2$.

Proof. By applying the previous Theorem 9 with accuracy $\varepsilon/2$ (instead of ε), we get

$$d_{\text{tr}}(\hat{\rho}, \rho) \leq \frac{\varepsilon}{2} + \varepsilon_t, \quad (\text{F23})$$

where $\varepsilon_t := \sum_{j=t+1}^n (1 - \lambda_j)/2$, and $\{\lambda_j\}_{j=1}^n$ are the normal eigenvalues of the correlation matrix of ρ ordered in

increasing order. Let us upper bound ε_t . First of all, by assumption, it follows that there exists a t -compressible Gaussian state σ_t such that

$$d_{\text{tr}}(\sigma_t, \rho) \leq \frac{\varepsilon^2}{4(n-t)}. \quad (\text{F24})$$

Then, we have

$$\varepsilon_t := \sqrt{\sum_{j=t+1}^n \frac{(1-\lambda_j)}{2}} \quad (\text{F25})$$

$$\leq \sqrt{(n-t) \frac{(1-\lambda_{t+1})}{2}} \quad (\text{F26})$$

$$\leq \sqrt{\frac{(n-t)}{2} \|C(\sigma_t) - C(\rho)\|_{\infty}} \quad (\text{F27})$$

$$\leq \sqrt{(n-t) d_{\text{tr}}(\sigma_t, \rho)} \quad (\text{F28})$$

$$\leq \frac{\varepsilon}{2}, \quad (\text{F29})$$

where in the third step we used Lemma 12, in the fourth step we used that the one-norm difference between two (possibly non-Gaussian) quantum states upper bounds the infinity norm difference of their correlation matrix, i.e.,

$$\|C(\sigma_t) - C(\rho)\|_{\infty} \leq 2d_{\text{tr}}(\sigma_t, \rho), \quad (\text{F30})$$

as it can be seen by the variational definition of trace distance [see the derivation of Eq. (D14) for an analogous and explicit proof of the latter inequality]. In the last step, we used Eq. (F24). Thus, substituting in Eq. (F23), we conclude. ■

APPENDIX G: DETAILS OF NUMERICAL SIMULATIONS

This section concerns a few additional details about the numerical results on impurity model dynamics, presented in Sec. 1F and Fig. 2 of the main text.

In the main text we touch upon the question: can the compressibility shown in Fig. 2 be derived from our Theorem 1 for t -doped circuits simply using the Trotter approximation? If true, this would render our numerical results less surprising; the reasoning for such a derivation could go as follows. The evolution unitary e^{-iHT} can be approximated with a product formula using some number t' of Trotter steps, which for any impurity model yields a t' -doped circuit. Such a Trotter-approximated state $|\psi(T)\rangle_{\text{Trot};t'}$ would then be t compressible for $t = \kappa t'$ (where κ is the locality of the impurity term; for H_{graph} and H_{TFIM} , $\kappa = 4$) due to Theorem 1. To test whether this explains the observed compressibility, we numerically check the number of Trotter steps needed to approximate $|\psi(T)\rangle$ such that $d_{\text{tr}}(|\psi(T)\rangle, |\psi(T)\rangle_{\text{Trot};t'}) \leq$

$\varepsilon_{\text{trunc}}^2/(4(n-\kappa t'))$. Because of Eq. (F29), this would imply that $\sqrt{\sum_{m=t+1}^n \frac{1}{2}(1-\lambda_m)} \leq \varepsilon_{\text{trunc}}$ [i.e., Eq. (6)] is satisfied. We find that the prediction from the Trotter approximation and compression theorem, as explained above, does not capture the dynamics of the Gaussian nullity we have observed. For example, at minimal resolved time $T = 0.05$ for both impurity models, the Trotter approximation already predicts $t = 14$ as the minimal Gaussian nullity (see vertical dotted lines in Fig. 2, left panel). In reality, however, the expander model dynamics still allows approximate t -compression with $t = 4$ all the way up to $T = 0.4$ (and the TFIM model even up to $T = 1.1$).

One may ask, why exactly is this analysis, by Trotterization and Theorem 1, coming so short of explaining the observed compressibility. In our opinion, this should be related to the fact that the non-Gaussian gates in the Trotterized evolution circuits are close to identity. One intuitively expects that states produced in such circuits would be much more compressible (approximately), than if the involved non-Gaussian gates were generic, far from identity. However, Theorem 1 does not account for this difference, and gives the same compressibility guarantee if the number of non-Gaussian gates is the same; also, importantly, Theorem 1 deals only with exact and not approximate compressibility. It is at present unclear how to account for the nature of non-Gaussian evolution beyond gate counting, so that our numerical results on compressibility of impurity model dynamics could be explained analytically.

To run the simulations of $|\psi(T)\rangle = e^{-iHT} |0^n\rangle$ presented in Fig. 2, we employed sparse matrix multiplication routines. Specifically, we approximated the exponential operator as $e^{-iHT} \simeq (1 - iH\delta T)^{T/\delta T}$, followed by state




FIG. 3. Extraction of truncation error $\varepsilon_{\text{trunc}}$ from sparse simulation via $\delta T \rightarrow 0$ extrapolation, used to produce Fig. 2. The displayed extrapolation is for the most challenging case of the expander model on $n = 16$ qubits. Here the evolution time is $T = 0.4$ and the truncation nullity is $t = 4$. The relevant scale for the error is the truncation value $\varepsilon_{\text{trunc}} = 0.05$, used in Fig. 2.

renormalization. To ensure that our results are as accurate as possible, the value of $\varepsilon_{\text{trunc}}$ for each $|\psi(T)\rangle$ was obtained from a $\delta T \rightarrow 0$ extrapolation. In particular, it was computed $\varepsilon_{\text{trunc}}$ for $\delta T \in \{0.0005, 0.001, 0.0015, 0.002, 0.0025\}$ and then extended to $\delta T = 0$ via a parabolic fit (see Fig. 3). This approach enabled simulations for systems up to $n = 16$ qubits in Fig. 2.

-
- [1] A. Anshu and S. Arunachalam, A survey on the complexity of learning quantum states, *Nat. Rev. Phys.* **6**, 59 (2024).
 - [2] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
 - [3] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-optimal tomography of quantum states, *IEEE Trans. Inf. Theory* **63**, 1 (2017).
 - [4] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos, Efficient tomography of a quantum many-body system, *Nat. Phys.* **13**, 1158 (2017).
 - [5] M. Fanizza, N. Galke, J. Lumbrellas, C. Rouzé, and A. Winter, Learning finitely correlated states: stability of the spectral reconstruction, [arXiv:2312.07516](https://arxiv.org/abs/2312.07516) [quant-ph].
 - [6] C. Rouzé and D. S. França, Learning quantum many-body systems from a few copies, *Quantum* **8**, 1319 (2024).
 - [7] H.-Y. Huang, Y. Liu, M. Broughton, I. Kim, A. Anshu, Z. Landau, and J. R. McClean, Learning shallow quantum circuits, [arXiv:2401.10095](https://arxiv.org/abs/2401.10095) [quant-ph].
 - [8] Z. Landau and Y. Liu, Learning quantum states prepared by shallow circuits in polynomial time, [arXiv:2410.23618](https://arxiv.org/abs/2410.23618) [quant-ph].
 - [9] H.-S. Kim, I. H. Kim, and D. Ranard, Learning state preparation circuits for quantum phases of matter, [arXiv:2410.23544](https://arxiv.org/abs/2410.23544) [quant-ph].
 - [10] A. Montanaro, Learning stabilizer states by bell sampling, [arXiv:1707.04012](https://arxiv.org/abs/1707.04012) [quant-ph].
 - [11] S. Arunachalam, S. Bravyi, A. Dutt, and T. J. Yoder, Optimal algorithms for learning quantum phase states, [arXiv:2208.07851](https://arxiv.org/abs/2208.07851) [quant-ph].
 - [12] S. Aaronson and S. Grewal, Efficient tomography of non-interacting fermion states, [arXiv:2102.10458](https://arxiv.org/abs/2102.10458) [quant-ph].
 - [13] M. Gluza, M. Kliesch, J. Eisert, and L. Aolita, Fidelity witnesses for fermionic quantum simulations, *Phys. Rev. Lett.* **120**, 190501 (2018).
 - [14] J. Surace and L. Tagliacozzo, Fermionic Gaussian states: An introduction to numerical approaches, *SciPost Physics Lecture Notes*, 54 (2022).
 - [15] B. M. Terhal and D. P. DiVincenzo, Classical simulation of noninteracting-fermion quantum circuits, *Phys. Rev. A* **65**, 032325 (2002).
 - [16] E. Knill, Fermionic linear optics and matchgates, [arXiv:quant-ph/0108033](https://arxiv.org/abs/quant-ph/0108033) [quant-ph].
 - [17] R. Jozsa and A. Miyake, Matchgates and classical simulation of quantum circuits, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **464**, 3089 (2008).
 - [18] L. G. Valiant, Quantum circuits that can be simulated classically in polynomial time, *SIAM J. Comput.* **31**, 1229 (2002).
 - [19] P. Echenique and J. L. Alonso, A mathematical and computational review of Hartree–Fock SCF methods in quantum chemistry, *Mol. Phys.* **105**, 3057 (2007).
 - [20] R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, Cambridge, 2004).
 - [21] G. Giuliani and G. Vignale, *Quantum Theory of the Electron Liquid* (Cambridge University Press, Cambridge, 2008).
 - [22] J. R. Schrieffer, *Theory of Superconductivity* (CRC Press, Boca Raton, 2018).
 - [23] R. J. Baxter, *Exactly solved models in statistical mechanics* (Dover Pubns, Mineola, New York, 1982).
 - [24] L. Fidkowski and A. Kitaev, Topological phases of fermions in one dimension, *Phys. Rev. B* **83**, 075103 (2011).
 - [25] A. Kitaev, Anyons in an exactly solved model and beyond, *Ann. Phys. (N. Y.)* **321**, 2 (2006).
 - [26] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).
 - [27] M. Hebenstreit, R. Jozsa, B. Kraus, S. Strelchuk, and M. Yoganathan, All pure fermionic non-Gaussian states are magic states for matchgate computations, *Phys. Rev. Lett.* **123**, 080503 (2019).
 - [28] D. J. Brod and E. F. Galvão, Extending matchgates into universal quantum computation, *Phys. Rev. A* **84**, 022310 (2011).
 - [29] D. J. Brod, The computational power of non-interacting particles, [arXiv:1412.7637](https://arxiv.org/abs/1412.7637) [quant-ph].
 - [30] D. Gottesman, The Heisenberg representation of quantum computers, [arXiv:quant-ph/9807006](https://arxiv.org/abs/quant-ph/9807006) [quant-ph].
 - [31] S. Bravyi and D. Gosset, Improved classical simulation of quantum circuits dominated by Clifford gates, *Phys. Rev. Lett.* **116**, 250501 (2016).
 - [32] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
 - [33] S. F. Oliviero, L. Leone, and A. Hamma, Transitions in entanglement complexity in random quantum circuits by measurements, *Phys. Lett. A* **418**, 127721 (2021).
 - [34] L. Leone, S. F. E. Oliviero, Y. Zhou, and A. Hamma, Quantum chaos is quantum, *Quantum* **5**, 453 (2021).
 - [35] S. Grewal, V. Iyer, W. Kretschmer, and D. Liang, Efficient learning of quantum states prepared with few non-Clifford gates, [arXiv:2305.13409](https://arxiv.org/abs/2305.13409) [quant-ph].
 - [36] L. Leone, S. F. E. Oliviero, and A. Hamma, Learning t-doped stabilizer states, *Quantum* **8**, 1361 (2024).
 - [37] D. Hangleiter and M. J. Gullans, Bell sampling from quantum circuits, *Phys. Rev. Lett.* **133**, 020601 (2024).
 - [38] B. Dias and R. Koenig, Classical simulation of non-Gaussian fermionic circuits, *Quantum* **8**, 1350 (2024).
 - [39] O. Reardon-Smith, M. Oszmaniec, and K. Korzekwa, Improved simulation of quantum circuits dominated by free fermionic operations, *Quantum* **8**, 1549 (2024).
 - [40] J. Cudby and S. Strelchuk, Gaussian decomposition of magic states for matchgate computations, [arXiv:2307.12654](https://arxiv.org/abs/2307.12654) [quant-ph].

- [41] J. Vijayan, P. Sompet, G. Salomon, J. Koepsell, S. Hirth, A. Bohrdt, F. Grusdt, I. Bloch, and C. Gross, Time-resolved observation of spin-charge deconfinement in fermionic Hubbard chains, *Science* **367**, 186 (2020).
- [42] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, A cold-atom Fermi–Hubbard antiferromagnet, *Nature* **545**, 462 (2017).
- [43] T. Esslinger, Fermi-Hubbard physics with atoms in an optical lattice, *Annu. Rev. Condens. Matter Phys.* **1**, 129 (2010).
- [44] P. Naldesi, A. Elben, A. Minguzzi, D. Clément, P. Zoller, and B. Vermersch, Fermionic correlation functions from randomized measurements in programmable atomic quantum devices, *Phys. Rev. Lett.* **131**, 060601 (2023).
- [45] S. Bravyi and D. Gosset, Complexity of quantum impurity problems, *Commun. Math. Phys.* **356**, 451 (2017).
- [46] S. Arunachalam, A. B. Grilo, and A. Sundaram, Quantum hardness of learning shallow classical circuits, *Electron. Colloquium Comput. Complex.* **TR19-041** (2019).
- [47] I. Diakonikolas, D. M. Kane, P. Manurangsi, and L. Ren, Cryptographic hardness of learning halfspaces with massart noise, [arXiv:2207.14266](https://arxiv.org/abs/2207.14266) [cs.LG].
- [48] P. Ananth, A. Poremba, and V. Vaikuntanathan, Revocable cryptography from learning with errors, [arXiv:2302.14860](https://arxiv.org/abs/2302.14860) [quant-ph].
- [49] A. Gupte, N. Vafa, and V. Vaikuntanathan, Continuous LWE is as hard as LWE and applications to learning Gaussian mixtures, [arXiv:2204.02550](https://arxiv.org/abs/2204.02550) [cs.CR].
- [50] Z. Ji, Y.-K. Liu, and F. Song, in *Advances in Cryptology – CRYPTO 2018* (Springer International Publishing, Cham, 2018), p. 126.
- [51] Z. Brakerski and O. Shmueli, in *Theory of Cryptography*, edited by D. Hofheinz and A. Rosen (Springer International Publishing, Cham, 2019), pp. 229–250.
- [52] J. Kondo, Resistance minimum in dilute magnetic alloys, *Progress of Theoretical Physics* **32**, 37 (1964).
- [53] J. R. Schrieffer and P. A. Wolff, Relation between the Anderson and Kondo Hamiltonians, *Phys. Rev.* **149**, 491 (1966).
- [54] P. W. Anderson, Localized magnetic states in metals, *Phys. Rev.* **124**, 41 (1961).
- [55] S. F. E. Oliviero, L. Leone, S. Lloyd, and A. Hamma, Unscrambling quantum information with Clifford decoders, *Phys. Rev. Lett.* **132**, 080402 (2024).
- [56] L. Leone, S. F. E. Oliviero, S. Lloyd, and A. Hamma, Learning efficient decoders for quasichaotic quantum scramblers, *Phys. Rev. A* **109**, 022429 (2024).
- [57] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Ann. Phys. (N. Y.)* **298**, 210 (2002).
- [58] A. Zhao, Learning, Optimizing, and simulating fermions with quantum computers, [arXiv:2312.10399](https://arxiv.org/abs/2312.10399) [quant-ph].
- [59] A. Zhao and A. Miyake, Group-theoretic error mitigation enabled by classical shadows and symmetries, *Npj Quantum Inf.* **10**, 57 (2024).
- [60] B. Zumino, Normal forms of complex matrices, *J. Math. Phys.* **3**, 1055 (2004).
- [61] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Boca Raton, 2017).
- [62] R. Jozsa, B. Kraus, A. Miyake, and J. Watrous, *Proceedings of the Royal Society A: Mathematical, Phys. Eng. Sci.* **466**, 809 (2010).
- [63] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Lower bounds on the non-Clifford resources for quantum computations, *Quantum Sci. Technol.* **5**, 035009 (2020).
- [64] J. Jiang and X. Wang, Lower bound for the T count via unitary stabilizer nullity, *Phys. Rev. Appl.* **19**, 034052 (2023).
- [65] D. S. França, F. G. L. Brandão, and R. Kueng, in *16th Conference on the Theory of Quantum Computation, Communication and Cryptography* (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2021).
- [66] X. Bonet-Monroig, R. Babbush, and T. E. O’Brien, Nearly optimal measurement scheduling for partial tomography of quantum states, *Phys. Rev. X* **10**, 031064 (2020).
- [67] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [68] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Matchgate shadows for fermionic quantum simulation, *Commun. Math. Phys.* **404**, 629 (2023).
- [69] M. Gutá, J. Kahn, R. Kueng, and J. A. Tropp, Fast state tomography with optimal error bounds, *J. Phys. A Math. Theor.* **53**, 204001 (2020).
- [70] J. Gao, Quantum union bounds for sequential projective measurements, *Phys. Rev. A* **92**, 052331 (2015).
- [71] A. Abbas, R. King, H.-Y. Huang, W. J. Huggins, R. Movassagh, D. Gilboa, and J. R. McClean, in *Thirty-seventh Conference on Neural Information Processing Systems* (2023), <https://openreview.net/forum?id=HF6bnhfSqH>.
- [72] H. Zhao, L. Lewis, I. Kannan, Y. Quek, H.-Y. Huang, and M. C. Caro, Learning quantum states and unitaries of bounded gate complexity, *PRX Quantum* **5**, 040306 (2024).
- [73] V. Lyubashevsky, C. Peikert, and O. Regev, in *Advances in Cryptology – EUROCRYPT 2010*, edited by H. Gilbert (Springer Berlin Heidelberg, Berlin, Heidelberg, 2010), p. 1.
- [74] O. Regev, On lattices, learning with errors, random linear codes, and cryptography, *J. ACM* **56**, 34 (2009).
- [75] D. Aggarwal, H. Bennett, Z. Brakerski, A. Golovnev, R. Kumar, Z. Li, S. Peters, N. Stephens-Davidowitz, and V. Vaikuntanathan, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023* (Association for Computing Machinery, New York, NY, USA, 2023), pp. 1516–1526.
- [76] L. Bittel, A. A. Mele, J. Eisert, and L. Leone, Optimal trace-distance bounds for free-fermionic states: Testing and improved tomography, [arXiv:2409.17953](https://arxiv.org/abs/2409.17953) [quant-ph].
- [77] K. G. Wilson, The renormalization group: Critical phenomena and the Kondo problem, *Rev. Mod. Phys.* **47**, 773 (1975).
- [78] D. Vollhardt, K. Byczuk, and M. Kollar, in *Strongly Correlated Systems* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011), p. 203.

- [79] The asymptotic notation $t = O(\log(\varepsilon^{-1}))$ hides a logarithmic dependence on the inverse of the spectral gap of the free-fermionic Hamiltonian H_{free} . For more details, see Ref. [45].
- [80] F. A. Mele, A. A. Mele, L. Bittel, J. Eisert, V. Giovannetti, L. Lami, L. Leone, and S. F. E. Oliviero, Learning quantum states of continuous variable systems, [arXiv:2405.01431](#) [quant-ph].
- [81] M. Oszmaniec, N. Dangniam, M. E. Morales, and Z. Zimborás, Fermion sampling: A robust quantum computational advantage scheme using fermionic linear optics and magic input states, [PRX Quantum 3, 020328 \(2022\)](#).
- [82] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, Quantum algorithms to simulate many-body physics of correlated fermions, [Phys. Rev. Appl. 9, 044036 \(2018\)](#).
- [83] D. Miller, L. E. Fischer, K. Levi, E. J. Kuehnke, I. O. Sokolov, P. K. Barkoutsos, J. Eisert, and I. Tavernelli, Hardware-tailored diagonalization circuits, [Npj Quantum Inf. 10, \(2024\)](#).
- [84] R. Bhatia, *Matrix Analysis*, Graduate Texts in Mathematics (Springer New York, 1996).
- [85] R. O'Donnell and R. Venkateswaran, The quantum union bound made easy, [arXiv:2103.07827](#) [quant-ph].
- [86] S. Aaronson, QMA/qpoly is contained in PSPACE/poly: De-merlinizing quantum protocols, [arXiv:quant-ph/0510230](#) [quant-ph].
- [87] R. O'Donnell and J. Wright, Efficient quantum tomography, [arXiv:1508.01907](#) [quant-ph].
- [88] J. Roffe, Quantum error correction: An introductory guide, [Contemp. Phys. 60, 226 \(2019\)](#).
- [89] J. Dehaene and B. De Moor, Clifford group, stabilizer states, and linear and quadratic operations over GF(2), [Phys. Rev. A 68, 042318 \(2003\)](#).
- [90] M. M. Wilde, in *Quantum Information Theory* (Cambridge University Press, Cambridge, 2017), p. xi.