

Efficient Tomography of Non-Interacting-Fermion States

Scott Aaronson  

The University of Texas at Austin, TX, USA

Sabee Grewal   

The University of Texas at Austin, TX, USA

Abstract

We give an efficient algorithm that learns a non-interacting-fermion state, given copies of the state. For a system of n non-interacting fermions and m modes, we show that $O(m^3 n^2 \log(1/\delta)/\epsilon^4)$ copies of the input state and $O(m^4 n^2 \log(1/\delta)/\epsilon^4)$ time are sufficient to learn the state to trace distance at most ϵ with probability at least $1 - \delta$. Our algorithm empirically estimates one-mode correlations in $O(m)$ different measurement bases and uses them to reconstruct a succinct description of the entire state efficiently.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum computation theory; Theory of computation \rightarrow Quantum information theory; Mathematics of computing \rightarrow Probabilistic inference problems; Theory of computation \rightarrow Quantum complexity theory

Keywords and phrases free-fermions, Gaussian fermions, non-interacting fermions, quantum state tomography, efficient tomography

Digital Object Identifier 10.4230/LIPIcs.TQC.2023.12

Related Version *Full Version*: <https://arxiv.org/abs/2102.10458>

Funding *Scott Aaronson*: Supported by a Vannevar Bush Fellowship from the US Department of Defense, the Berkeley NSF-QLCI CIQC Center, a Simons Investigator Award, and the Simons “It from Qubit” collaboration.

Sabee Grewal: Supported by Scott Aaronson’s Vannevar Bush Fellowship from the US Department of Defense, Berkeley NSF-QLCI CIQC Center, Simons Investigator Award, and Simons “It from Qubit” collaboration.

Acknowledgements We thank Andrew Zhao for notifying us that the previous version of this manuscript contained an error and providing other insightful comments. We also thank Yuxuan Zhang, Alex Kulesza, Ankur Moitra, William Kretschmer, Dax Enshan Koh, Andrea Rocchetto, and Patrick Rall for helpful discussions, and Alex Arkhipov, William Kretschmer, and Daniel Liang for helpful comments on a previous version of this manuscript.

1 Introduction

There are two types of particles in the universe: bosons and fermions. Bosons include force carriers, such as photons and gluons, and fermions include matter particles like quarks and electrons. Each particle can be in a certain mode (e.g., a position or state). For a system of n particles, a configuration of the system is described by specifying how many particles are in each of m modes. Bosons are particles where multiple occupancy of a mode is allowed, whereas fermions are particles where multiple occupancy is forbidden; that is, two or more fermions cannot occupy the same mode at once (this is the *Pauli exclusion principle*). It follows that a system of n fermions and m modes has $\binom{m}{n}$ possible configurations; we denote the set of possible configurations by $\Lambda_{m,n}$.



© Scott Aaronson and Sabee Grewal;
licensed under Creative Commons License CC-BY 4.0

18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023).

Editors: Omar Fawzi and Michael Walter; Article No. 12; pp. 12:1–12:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Our main result is an efficient algorithm (both in copy complexity and time complexity) for learning a non-interacting-fermion state (also called a free-fermion state or a Gaussian fermion state), which is a superposition over configurations in $\Lambda_{m,n}$. Even though a non-interacting-fermion state lives in an exponentially large Hilbert space, we show how to exploit its structure to output a succinct description of the state efficiently. A non-interacting-fermion state can be completely specified by an $m \times n$ column-orthonormal matrix A . Our algorithm measures copies of the input state in $O(m)$ different measurement bases, and uses the measurement data to reconstruct an $m \times n$ matrix \hat{A} in polynomial time. We prove that a polynomial number of copies of the input state is enough for the output state to be ϵ -close to the original state in trace distance.

► **Theorem 1 (Main result).** *Let $|\Psi\rangle$ be a state of n non-interacting fermions and m modes. There exists an algorithm that uses $O(m^3 n^2 \log(1/\delta)/\epsilon^4)$ copies of $|\Psi\rangle$, $O(m^4 n^2 \log(1/\delta)/\epsilon^4)$ classical time, and $O(m)$ measurement bases, and outputs a succinct description of a non-interacting-fermion state $|\hat{\Psi}\rangle$ that is ϵ -close in trace distance distance to $|\Psi\rangle$ with probability at least $1 - \delta$.*

Our algorithm can also be adapted to conventional quantum state tomography, which we explain in Section 5.

1.1 Main Ideas

Here and throughout, let $U \in \mathbb{C}^{m \times m}$ be the unitary that prepares the unknown non-interacting-fermion state $|\Psi\rangle$ from the standard initial state $|1_n\rangle$ (the state where the first n modes are occupied and the remaining are unoccupied), and let $A \in \mathbb{C}^{m \times n}$ be the column-orthonormal matrix corresponding to the first n columns of U . Define $K = (k_{ij}) := AA^\dagger \in \mathbb{C}^{m \times m}$. We refer to K as the *kernel matrix* due to the connection between determinantal point processes and non-interacting fermions (which we discuss further in Section 1.2.1). In the physics literature, the kernel matrix is also called the one-body reduced density matrix (1-RDM) or the correlation matrix.

The elements of $\Lambda_{m,n}$ are the possible configurations of a system of n non-interacting fermions and m modes. Formally, $\Lambda_{m,n}$ is the set of all lists $S = (s_1, \dots, s_m)$ such that $s_i \in \{0, 1\}$ and $\sum_{i \in [m]} s_i = n$. The set $\{|S\rangle_{S \in \Lambda_{m,n}}\}$ is a basis for n -fermion and m -mode systems, which we refer to as the standard basis. The $m \times n$ column-orthonormal matrix A describes the state

$$|\Psi\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S) |S\rangle,$$

where, for $S = (s_1, \dots, s_m) \in \Lambda_{m,n}$, A_S is the $n \times n$ submatrix obtained by removing row i of A if $s_i = 0$. Therefore, upon measurement, we observe the configuration $S \in \Lambda_{m,n}$ with probability

$$|\langle S|\Psi\rangle|^2 = |\det(A_S)|^2 = \det(K_S),$$

where, for $S = (s_1, \dots, s_m)$, K_S is the $n \times n$ submatrix obtained by removing row and column i of K if $s_i = 0$. In other words, upon measurement, we observe a configuration $S \in \Lambda_{m,n}$ with probability equal to the corresponding principal minor of the kernel matrix K . The probability that any subset of k modes is occupied corresponds to a principal minor of order k , obtained as above (remove the rows and columns of K corresponding to unoccupied modes and compute the determinant of the resulting submatrix). For example, the diagonal

entries k_{ii} correspond to the one-mode correlations (i.e., k_{ii} is the probability that mode i is occupied). Passing $|\Psi\rangle$ through a unitary transformation $V \in \mathbb{C}^{m \times m}$ maps K to VKV^\dagger . Given copies of the unknown state $|\Psi\rangle$, our goal is to output a column-orthonormal matrix \hat{A} such that $|\hat{\Psi}\rangle = \sum_S \det(\hat{A}_S) |S\rangle$ is ϵ -close to $|\Psi\rangle$ in trace distance.

At a high level, our algorithm constructs \hat{K} (an approximation of the kernel matrix), computes a decomposition $\hat{K} = \hat{A}\hat{A}^\dagger$, and then outputs \hat{A} . Our algorithm begins by measuring $O(\log(1/\delta)/\gamma^2)$ copies of the input state in the standard basis to empirically estimate the one-mode correlations of the state to accuracy $\pm\gamma$. The estimates are obtained simply by computing the average number of times each mode was occupied and are the diagonal entries of \hat{K} . One can then estimate the (i, j) entry of K as follows. Apply the beamsplitter

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

on modes i and j , which maps the diagonal entry k_{ii} to $\frac{1}{2}(k_{ii} + k_{jj} + 2\text{Re}(k_{ij}))$, and measure the resulting state in the standard basis. Repeat this $O(\log(1/\delta)/\gamma^2)$ times. As we did before, average the number of times mode i is occupied to obtain an estimate for $\frac{1}{2}(k_{ii} + k_{jj} + 2\text{Re}(k_{ij}))$ to accuracy $\pm\gamma$. Finally, using the previously obtained estimates for k_{ii} and k_{jj} , solve for $\text{Re}(k_{ij})$ (up to accuracy $\pm\gamma$). Repeat this process with the beamsplitter

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}$$

to estimate the imaginary part of k_{ij} to accuracy $\pm\gamma$.

Our algorithm proceeds as follows. Simultaneously execute the process above on the pairs of modes $(1, 2), (3, 4), \dots, (m-1, m)$, then $(1, 3), (2, 4), \dots, (m-2, m)$, and so on, until all the off-diagonal entries are recovered. It is easy to check that $O(m)$ measurement bases are needed to recover all off-diagonal entries of K . Then, we compute $Q\Lambda Q^\dagger$, an eigendecomposition of \hat{K} . Finally, we set \hat{A} to be the $m \times n$ matrix corresponding to the first n columns of Q , and return \hat{A} . Overall, the algorithm requires $O(m/\gamma^2)$ copies of the input state and $O(m^2/\gamma^2)$ time.

The technical part is to understand how far $|\hat{\Psi}\rangle$ is from $|\Psi\rangle$ in trace distance, given that our algorithm begins by learning the entries of K to within γ in magnitude. To do this, we give a new proof that learning the kernel matrix is enough to learn the state, despite the kernel matrix only consisting of one- and two-mode correlations.

► **Theorem 2** (Informal version of Theorem 5). *Let $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ be n -fermion and m -mode non-interacting-fermion states with kernel matrices K and \hat{K} , respectively. Then*

$$d_{\text{tr}}(|\hat{\Psi}\rangle, |\Psi\rangle) \leq \sqrt{n \|\hat{K} - K\|_2},$$

where $d_{\text{tr}}(\cdot, \cdot)$ is the trace distance and $\|\cdot\|_2$ is the spectral norm.

While this may seem surprising, this has been known to physicists since the 1960s and is (in some sense) the content of the Hohenberg-Kohn theorems [13] and Kohn-Sham equations [17], which form the theoretical foundations of density functional theory. These results paved the way for computational methods in quantum chemistry, earning Walter Kohn and John Pople the 1998 Nobel Prize in Chemistry. Additionally, it is well-known that Wick's theorem [33] can be used to write higher-order correlations in terms of one- and two-mode correlations.

Although this topic has received intense study for decades, we claim that our error analysis offers two improvements. First, this is the first “physics-free” proof that kernel matrices suffice to learn the state: we make no mention of creation/annihilation operators, energy potentials, Hamiltonians, or the like. We believe our proof can be understood by any mathematician or theoretical computer scientist without a physics background and perhaps even undergraduates with a linear algebra background. Second, our theorem quantitatively relates the distance between the states and the distance between the kernel matrices, which (to our knowledge) has never been done.

In Section 4 we use this theorem to show that the trace distance between $|\hat{\Psi}\rangle$ and $|\Psi\rangle$ is at most $\sqrt{2nm\gamma}$. Therefore, the trace distance will be ϵ -close if we set γ to $\epsilon^2/2nm$.

1.2 Related Work

Previous work showed how to *simulate* non-interacting fermions efficiently. In 2002, Valiant [30] introduced a class of quantum circuits called matchgate circuits and showed that they can be simulated classically in polynomial time. Soon after, Terhal and DiVincenzo [28] (see also Knill [16]) showed that evolutions of non-interacting fermions give rise to unitary matchgate circuits. (See also [4][Appendix 13] for a simpler and faster simulation algorithm.) Thus, the contribution of this paper is to complement these classical simulation results with an efficient *learnability* result.

More broadly, *quantum state tomography* is the task of constructing a classical description of a d -dimensional quantum mixed state, given copies of the state. With entangled measurements, the optimal number of copies for quantum state tomography is known to be $\Theta(d^2)$ due to Haah et al. [12] and O’Donnell and Wright [22]. With unentangled measurements, the optimal number of copies is $\Theta(d^3)$ [18, 12, 11].

Quantum state tomography can be computationally efficient in restricted settings. Montanaro [21] showed that *stabilizer states* are efficiently learnable using measurements in the Bell basis. Cramer et al. [9] showed that states approximated by *matrix product states* are efficiently learnable. Arunachalam et al. [6] showed that some classes of *phase states* are efficiently learnable. With this work, non-interacting-fermion states is an additional class of quantum states for which we know computationally efficient learning algorithms.

Different models for learning properties of mixed states ρ have been studied. For example, Aaronson [1] showed that such states are learnable under the Probably Approximately Correct (PAC) model, using training sequences of length only logarithmic in the Hilbert space dimension. Since our goal is simply to reconstruct a distribution, we have no need for the PAC framework. Aaronson also introduced *shadow tomography* [2, 7], where, given a list of known two-outcome observables and copies of an unknown state, the goal is to estimate the expectation value of each observable with respect to the unknown state to additive accuracy. Although computationally inefficient, Aaronson showed that the number of copies of the input state scales polylogarithmically with both the number of observables and the Hilbert space dimension. Soon after, Huang, Kueng, and Preskill [15] introduced *classical shadows*, a shadow tomography algorithm that is computationally efficient for certain problem instances. For example, with random Clifford measurements, the classical time cost in classical shadows is dominated by computing quantities of the form $\langle s|O|s\rangle$, where O is an observable and $|s\rangle$ is some stabilizer state, which is computationally efficient for certain observables.

Recently, there have been several results that extend the classical shadows protocol to fermionic states and circuits [34, 31, 19, 23]. In particular, after the original version of this paper appeared but before the current version, O’Gorman [23] gave an algorithm for learning non-interacting-fermion states, which is based on classical shadows. His learning

algorithm uses $O(m^7 n^2 \log(m/\delta)/\epsilon^4)$ samples and $O(m^9 n^2 \log(m/\delta)/\epsilon^4)$ time to learn a non-interacting-fermion state to trace distance at most ϵ with probability at least $1 - \delta$. Our work has substantially better sample and time complexities, and makes no use of randomized measurements.

Finally, [25, Appendix C] proposes an algorithm for reconstructing a kernel matrix that involves iterating over $O(m)$ perfect matchings, just as our algorithm does, which we were unaware of until the final stages of our work. However, we note that their circuits/measurements differ from ours, and they do not provide an error analysis for their algorithm.

1.2.1 Determinantal Point Processes

For reasons having nothing to do with non-interacting fermions, problems extremely close to ours have already been studied in classical machine learning, in the field of *Determinantal Point Processes* (DPPs). A DPP is a model specified by an $m \times m$ matrix K (typically symmetric or Hermitian), such that the probabilities of various events are given by various principal minors of K , exactly as for non-interacting fermions. The connection between DPPs and fermions has been known for decades [20].

Two results in particular are directly relevant to us: Rising et al. [27] and Urschel et al. [29]. Rising et al. give an efficient algorithm for the *symmetric principal minor assignment problem*: given a list of all 2^m principal minors of an unknown $m \times m$ symmetric matrix K , reconstruct K . Their algorithm is based on constructing an m -vertex graph with an edge from i to j whenever $K_{ij} \neq 0$, and then analyzing the minimum spanning trees and chordless cycles in that graph. Rising et al., however, do not do an error analysis (they assume exact knowledge of the principal minors), and they solve the problem for real and complex *symmetric* matrices, whereas in our problem, the matrix is complex and Hermitian. This difference turns out to be surprisingly important, as the determinants of Hermitian matrices are always real, so much of the phase information vanishes – making the Hermitian case much harder.

Urschel et al. [29] further exploit connections between DPPs and graph theory to give an algorithm that recovers the entries of K , given samples from an unknown DPP. Their focus is on parameter learning (i.e., approximately recovering the entries of K), rather than learning the distribution induced by K in, say, total variation distance. They again assume that the DPP is described by a real symmetric matrix.

Our work provides insight on the Hermitian versions of these problems. Since there are many non-interacting-fermion states with the same distribution over the standard basis, there must be many kernel matrices that are consistent with the same list of principal minors. The goal for the Hermitian principal minor assignment problem is to output any such matrix.

It is also clear that the Hermitian version of Urschel et al.'s problem is *impossible*. Samples from an unknown DPP correspond to standard basis measurement outcomes, and learning the entries of a DPP corresponds to learning the kernel matrix. However, in Theorem 5, we show that learning the kernel matrix suffices to learn the entire state, which is impossible when restricted to standard basis measurements. (Even distinguishing $|+\rangle$ from $|-\rangle$ is information-theoretically impossible when given only standard basis measurements.) What one could hope for, and which we leave open, is to learn *some* kernel matrix that gives rise to a distribution close in variation distance to the observed one.

1.2.2 Errors in Previous Version

A previous version of this manuscript [5] – where we claimed to recover a non-interacting fermion distribution using only standard basis measurements – had serious errors, which we explain below.¹

In the previous manuscript, we sought to recover the rows $v_1, \dots, v_m \in \mathbb{C}^n$ of the $m \times n$ column-orthonormal matrix A up to isometry (see Section 1.1 for the definition of A). By estimating the two-mode correlations (i.e., the probability of finding a fermion in both mode i and mode j), one can deduce the approximate value of $|\langle v_i, v_j \rangle|$, i.e., the absolute value of the inner product, for any $i \neq j$. From that information, our goal was to recover v_1, \dots, v_m (or, more precisely, their relative configuration in n -dimensional space up to isometry).

The approach was as follows: if we knew $\langle v_i, v_j \rangle$ for all $i \neq j$, then we would get linear equations that iteratively constrained each v_i in terms of $\langle v_i, v_j \rangle$ for $j < i$, so all that would be required would be to solve those linear systems, and then show that the solution is robust with respect to small errors in our estimates of each $\langle v_i, v_j \rangle$. While it is true that the measurements only reveal $|\langle v_i, v_j \rangle|$ rather than $\langle v_i, v_j \rangle$ itself, the “phase information” in $\langle v_i, v_j \rangle$ seemed manifestly irrelevant, since it in any case depended on the irrelevant global phases of v_i and v_j themselves.

Alas, it turns out that the phase information *does* matter. As an example, suppose we only knew the following about three unit vectors $u, v, w \in \mathbb{R}^3$:

$$|\langle u, v \rangle| = |\langle u, w \rangle| = |\langle w, v \rangle| = \frac{1}{2}.$$

This is not enough to determine these vectors up to isometry! In one class of solution, all three vectors belong to the same plane, like so:

$$u = (1, 0, 0), \quad v = \left(\frac{1}{2}, \frac{\sqrt{3}}{2}, 0\right), \quad w = \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0\right).$$

In a completely different class of solution, the three vectors do not belong to the same plane, and instead look like three edges of a tetrahedron meeting at a vertex:

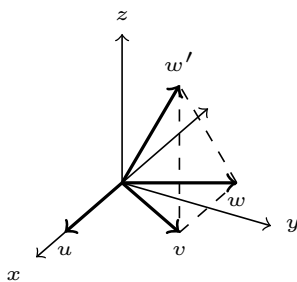
$$u = (1, 0, 0), \quad v = \left(\frac{1}{2}, \frac{\sqrt{3}}{2}, 0\right), \quad w = \left(\frac{1}{2}, \frac{\sqrt{3}}{6}, \sqrt{\frac{2}{3}}\right).$$

Both classes of solutions are shown in Figure 1. These solutions correspond to different sign choices for $|\langle u, v \rangle|$, $|\langle u, w \rangle|$, and $|\langle w, v \rangle|$ – choices that collectively matter, even though each one is individually irrelevant.

It follows that, even in the special case where the vectors are all real, the two-mode correlations are not enough to determine the vectors’ relative positions. And alas, the situation is even worse when, as for us, the vectors can be complex. Indeed, even for systems of 2 fermions and 4 modes, it is possible to exhibit distributions that require complex vectors. For example, let

$$A = \begin{pmatrix} \frac{1}{2} & 0 \\ \frac{1}{2\sqrt{2}} & \sqrt{\frac{2}{5}} \\ \frac{1}{2\sqrt{2}} & \sqrt{\frac{2}{5}}i \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{10}} - \frac{1}{\sqrt{10}}i \end{pmatrix}.$$

¹ These errors were previously explained in <https://scottaaronson.blog/?p=5706>.



■ **Figure 1** The vectors u, v , and w belong to the same plane, while v, w , and w' are edges of a tetrahedron that meet at a vertex. Both sets of vectors satisfy the same inner product constraints up to phase information, yet belong to two distinct classes of solutions.

Denote the one-mode correlations by p_i and the two-mode correlations by p_{ij} . Then by explicit calculation (e.g., compute K , then compute the appropriate principal minors of K), one can verify that

$$p_1 = \frac{1}{4}, \quad p_2 = p_3 = \frac{21}{40}, \quad p_{12} = p_{13} = p_{23} = \frac{1}{10}.$$
²

Hence, to represent the corresponding distribution with real vectors only, one must find three vectors in \mathbb{R}^2 with squared lengths $\frac{1}{4}$, $\frac{21}{40}$, and $\frac{21}{40}$, respectively, such that the squared area of the parallelogram created by any pair of vectors is $\frac{1}{10}$. One can verify that this is not possible.

We conclude that any possible algorithm for learning fermionic distributions from standard-basis measurements will have to solve a system of nonlinear equations (albeit, a massively overconstrained system that is guaranteed to have a solution); it will have to use three-mode correlations (i.e., statistics of triples of fermions), and indeed (one can show) in some exceptional cases four-mode correlations and above; it will sometimes have to output complex solutions even when all the input data is real (which rules out a purely linear-algebraic approach); and it will have to learn the phase information relevant to the *distribution* (rather than the entire state).

Finally, separate from the issues above, we are grateful to Andrew Zhao for identifying an error in the most recent version of this manuscript (v3 on the arXiv). In our error analysis, we assumed that the $m \times m$ matrix output by our algorithm is rank- n , when, in fact, it is full-rank. In short, we handle this issue by computing an eigendecomposition of said matrix and discarding the smallest $m - n$ eigenvalues and eigenvectors.

2 Preliminaries

Throughout this work, we use the following notation. $[n] := \{1, \dots, n\}$. Let $X \in \mathbb{C}^{n \times n}$ and $v \in \mathbb{C}^n$. Then $\|v\|_p := (\sum_{i \in [n]} |v_i|^p)^{1/p}$ is the ℓ_p -norm, and $\|X\|_2 := \sup_{\|v\|_2=1} \|Xv\|_2$ is the spectral norm. Let ρ and σ be two quantum mixed states. Then $d_{\text{tr}}(\rho, \sigma) := \frac{1}{2} \text{tr} \left(\sqrt{(\rho - \sigma)^2} \right) = \frac{1}{2} \sum_i |\lambda_i|$ is the trace distance, where the λ_i 's are eigenvalues of the error matrix $\rho - \sigma$.

² We do not need to look at the fourth mode to show that complex numbers are necessary.

2.1 Non-Interacting Fermions

We briefly review non-interacting fermions. For more detail, see e.g. [3, Chapters 7-10].

As mentioned above, $\Lambda_{m,n}$ is the set of all lists $S = (s_1, \dots, s_m)$ such that $s_i \in \{0, 1\}$ and $\sum_{i \in [m]} s_i = n$. One can easily verify that the total number of configurations $|\Lambda_{m,n}| = \binom{m}{n}$. Define $M := \binom{m}{n}$. The set $\{|S\rangle\}_{S \in \Lambda_{m,n}}$ is a basis for systems of n non-interacting fermions in m modes, which we refer to as the standard basis (also called the Fock basis or occupation number basis). Hence, a non-interacting-fermion state is a unit vector in an M -dimensional complex Hilbert space with the form

$$|\Psi\rangle = \sum_{S \in \Lambda_{m,n}} \alpha_S |S\rangle, \quad \text{where} \quad \sum_{S \in \Lambda_{m,n}} |\alpha_S|^2 = 1,$$

and, upon measurement, we observe an element $S \in \Lambda_{m,n}$ with probability

$$\Pr[S] = |\langle \Psi | S \rangle|^2 = |\alpha_S|^2.$$

The transformations on non-interacting-fermion states can be described by $m \times m$ unitary matrices, which can always be constructed with $O(m^2)$ elementary operations called *beam splitters* and *phaseshifters* (see [26] for a proof of this statement). A beamsplitter acts on two modes and has the form

$$\begin{pmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & \cos \theta & -\sin \theta & & & \\ & & \sin \theta & \cos \theta & & & \\ & & & & \ddots & & \\ & & & & & & 1 \end{pmatrix},$$

while a phaseshifter applies a complex phase to a single mode and has the form

$$\begin{pmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & e^{i\theta} & & & \\ & & & \ddots & & \\ & & & & & 1 \end{pmatrix}.$$

An $m \times m$ unitary U describes how a single fermion would be scattered from one mode to the others. This induces an $M \times M$ unitary transformation $\varphi(U)$ on the Hilbert space of n -fermion states, where φ is the homomorphism that lifts $m \times m$ unitary transformations to n -fermion unitary transformations. For an $m \times m$ unitary matrix U , one way to define φ is

$$\langle S | \varphi(U) | T \rangle = \det(U_{S,T}), \tag{1}$$

for all $S = (s_1, \dots, s_m), T = (t_1, \dots, t_m) \in \Lambda_{m,n}$, where $U_{S,T}$ is the $n \times n$ submatrix obtained by removing row i of U if $s_i = 0$ and removing column j of U if $t_j = 0$. One can also view $\varphi(U)$ as a unitary transformation on m -qubit states. In this case, $\varphi(U)$ is a Hamming-weight-preserving matchgate unitary [30]. This follows directly from Terhal and DiVincenzo's work [28], which shows the equivalence between evolutions of non-interacting fermions and matchgate circuits.

Intuitively, the reason the determinant arises in Equation (1) is that $\langle S | \varphi(U) | T \rangle$ is the sum over the $n!$ permutations that take n fermions in configuration S to configuration T , each permutation contributing to the overall amplitude. When a permutation from S to T is odd, its contribution to the overall amplitude has a phase factor of -1 , while when the permutation is even the phase factor is 1 . This is the antisymmetry property of fermions: swapping two fermions picks up a -1 phase in the amplitude.³

2.2 Problem Setup

We use the following notation: $|1_n\rangle := |1, \dots, 1, 0, \dots, 0\rangle$ is the standard initial state (where n fermions occupy the first n modes), $U \in \mathbb{C}^{m \times m}$ is the unitary that prepares the unknown non-interacting-fermion state (i.e., $|\Psi\rangle = \varphi(U) |1_n\rangle$), and $A \in \mathbb{C}^{m \times n}$ is the $m \times n$ column-orthonormal matrix corresponding to the first n columns of U . Define $K := AA^\dagger$.

For each element $S = (s_1, \dots, s_m) \in \Lambda_{m,n}$, let A_S be the $n \times n$ submatrix obtained by removing row i of A if $s_i = 0$, and let K_S be the $n \times n$ submatrix obtained by removing row and column i of K if $s_i = 0$. Then from Equation (1), it follows that

$$|\Psi\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S) |S\rangle,$$

and \mathcal{D}_K , the probability distribution over $S \in \Lambda_{m,n}$ obtained by measuring the state $|\Psi\rangle = \varphi(U) |1_n\rangle$ in the standard basis, is given by

$$|\langle 1_n | \varphi(U) | S \rangle|^2 = |\det(A_S)|^2 = \det(K_S),$$

where the last equality uses the fact that, for any square matrices X and Y , $\det(X)^* = \det(X^\dagger)$ and $\det(X) \det(Y) = \det(XY)$. Further, for any list $S = (s_1, \dots, s_m)$ where $\sum_i s_i = k < n$, the marginal probability that those k modes are occupied is $\det(K_S)$.

It is easy to verify that K is Hermitian, positive semi-definite, and a projector ($K^2 = K$); and that $\text{tr}(K) = n$. Additionally, observe that the (i, j) entry is the inner product between the i th and j th rows of A (i.e., K is a *Gram matrix*). Finally, note that A is a *highly non-unique* description of $|\Psi\rangle$ and K . Let R be any $n \times n$ unitary matrix. Then A and AR describe the same state:

$$|\Psi\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S R) |S\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S) \det(R) |S\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S) |S\rangle,$$

where, in the last equality, we use the fact that the determinant of a unitary matrix is a complex unit, which only adds an irrelevant global phase. Meanwhile, the kernel matrix of $|\Psi\rangle$ is unchanged: $K = (AR)(AR)^\dagger = AA^\dagger$.

Applying a unitary V maps $|\Psi\rangle = \varphi(U) |1_n\rangle$ to $\varphi(VU) |1_n\rangle$. It is easy to check that the probability that we observe S upon measuring $\varphi(VU) |1_n\rangle$ in the standard basis is

$$|\langle 1_n | \varphi(VU) | S \rangle|^2 = \det((VKV^\dagger)_S),$$

and in general, applying V has the following effect on the matrix K : $K \mapsto VKV^\dagger$.

Given copies of $|\Psi\rangle$ and the ability to apply beamsplitter networks before measurement, our goal is to output a matrix $\hat{A} \in \mathbb{C}^{m \times n}$ such that $|\hat{\Psi}\rangle = \sum_S \det(\hat{A}_S) |S\rangle$ is ϵ -close to $|\Psi\rangle$ in trace distance.

³ Meanwhile, bosons are symmetric under transpositions, so no minus signs show up. This is precisely why permanents arise when computing amplitudes for non-interacting bosons, while determinants show up for non-interacting fermions.

3 Learning Algorithm

In this section, we present our learning algorithm, which is given copies of an unknown non-interacting-fermion state, and outputs an $m \times n$ matrix \hat{A} such that the corresponding state $|\hat{\Psi}\rangle = \sum_S \det(\hat{A}_S) |S\rangle$ is close to the original in trace distance. The algorithm has three phases: first, we learn the diagonal entries of the kernel matrix K with standard basis measurements, then we learn the off-diagonal entries by measuring the unknown state in $O(m)$ different bases. Finally, we decompose the reconstructed kernel matrix into the $m \times n$ output matrix \hat{A} .

■ **Algorithm 1** Efficient tomography of non-interacting-fermion states.

Input: Black-box access to copies of $|\Psi\rangle$ (the input state), $\gamma \in (0, 1)$ (accuracy parameter), and $\delta \in (0, 1)$ (confidence parameter).

Output: an $m \times n$ matrix $\hat{A} \in \mathbb{C}^{m \times n}$.

- 1: Measure $O(\log(1/\delta)/\gamma^2)$ copies of $|\Psi\rangle$ in the standard basis and estimate the one-mode correlations for all m modes. Set \hat{k}_{ii} to the empirical estimate that mode i is occupied.
- 2: Choose $O(m)$ different perfect matchings among the m modes, which together cover all possible (i, j) pairs.
- 3: **for** each perfect matching **do**
- 4: Apply the beamsplitter

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

to each pair of modes in the perfect matching and measure in the standard basis. Repeat this $O(\log(1/\delta)/\gamma^2)$ times, and use the measurement data to estimate the one-mode correlations.

- 5: Repeat the previous step with the beamsplitter

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}.$$

- 6: For each pair (i, j) in the perfect matching, the beamsplitter network in step 4 maps the one-mode correlation k_{ii} to $k'_{ii} := \frac{1}{2}(k_{ii} + k_{jj} + 2\text{Re}(k_{ij}))$. Denote the i th estimate obtained in step 4 by \hat{k}'_{ii} . Let $\text{Re}(\hat{k}_{ij})$ be the estimate of $\text{Re}(k_{ij})$ obtained by solving the following equation:

$$\hat{k}'_{ii} = \frac{1}{2}(\hat{k}_{ii} + \hat{k}_{jj} + 2\text{Re}(\hat{k}_{ij})),$$

where \hat{k}_{ii} and \hat{k}_{jj} are the estimates from step 1.

- 7: Repeat the previous step with the estimates from step 5 to obtain estimate $\text{Im}(\hat{k}_{ij})$ for each pair (i, j) . (Note that the beamsplitter network in step 5 maps k_{ii} to $\frac{1}{2}(k_{ii} + k_{jj} + 2\text{Im}(k_{ij}))$, for each pair (i, j) in the perfect matching.)
- 8: For each pair (i, j) in the perfect matching, set $\hat{k}_{ij} = \text{Re}(\hat{k}_{ij}) + i\text{Im}(\hat{k}_{ij})$ and $\hat{k}_{ji} = \hat{k}_{ij}^*$.
- 9: Let $\hat{K} = (\hat{k}_{ij}) \in \mathbb{C}^{m \times m}$, and let $Q\Lambda Q^\dagger$ be an eigendecomposition of \hat{K} . Set \hat{A} to be the $m \times n$ matrix corresponding to the first n columns of Q .
- 10: **return** \hat{A} .

For each measurement basis, we estimate $O(m)$ entries of K to within γ in magnitude, which can be accomplished with $O(\log(1/\delta)/\gamma^2)$ copies [8, Theorem 9] and $O(m \log(1/\delta)/\gamma^2)$ time. To estimate the off-diagonal entries to the target accuracy, an additional constant factor appears in the sample complexity, which is absorbed into the $O(\log(1/\delta)/\gamma^2)$. We use $O(m)$ measurement bases in total, so the overall copy and time complexities are $O(m \log(1/\delta)/\gamma^2)$ and $O(m^2 \log(1/\delta)/\gamma^2)$ respectively. Our algorithm then computes an eigendecomposition of an $m \times m$ matrix, which requires $O(m^3)$ time, but computing this decomposition is not the bottleneck in our algorithm.

We note that \hat{K} is clearly Hermitian by construction, so the eigendecomposition of \hat{K} exists. Following convention, it is assumed that the eigenvalues are ordered from largest to smallest (i.e., the first column of Q corresponds to the largest eigenvalue of \hat{K} , and so on). The output of our algorithm \hat{A} is column-orthonormal because the columns of the unitary Q are orthonormal. Therefore, \hat{A} describes a non-interacting-fermion state $|\hat{\Psi}\rangle = \sum_S \det(\hat{A}_S) |S\rangle$. In the next section, we show that if $\gamma = \frac{\epsilon^2}{2nm}$, then the trace distance between $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ is at most ϵ . Hence, for the two states to be ϵ -close in trace distance, $O(m^3 n^2 \log(1/\delta)/\epsilon^4)$ copies and $O(m^4 n^2 \log(1/\delta)/\epsilon^4)$ time suffice.

4 Error Analysis

In this section, we show that the trace distance between $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ is at most $\sqrt{2nm}\gamma$. Therefore, if $\gamma = \frac{\epsilon^2}{2nm}$, then the trace distance between $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ is at most ϵ .

The error analysis is presented in two parts. First, we show that the trace distance between any two non-interacting-fermion states is bounded above the spectral distance between their kernel matrices. Then we show that the output of our algorithm is close to the original state in trace distance.

4.1 The Kernel Matrix Suffices

We prove that the trace distance between two non-interacting-fermion states is upper bounded by the spectral difference between their kernel matrices. To show this, we need the following two lemmas.

► **Lemma 3.** *Let $a_1, a_2, \dots, a_n \in [0, 1]$. Then*

$$1 - \prod_i a_i \leq n \max_i 1 - a_i.$$

Proof. For any $x, y \in [0, 1]$,

$$1 - xy = 1 - x + x - xy = (1 - x) + x(1 - y) \leq (1 - x) + (1 - y).$$

We can inductively apply this to get

$$1 - \prod_i a_i \leq \sum_i 1 - a_i \leq n \max_i 1 - a_i. \quad \blacktriangleleft$$

► **Lemma 4.** *Let $A, \hat{A} \in \mathbb{C}^{m \times n}$ ($m \geq n$) be $m \times n$ matrices, and let A be column-orthonormal. Define $K := AA^\dagger$ and $\hat{K} := \hat{A}\hat{A}^\dagger$. Let $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ where σ_i are the singular values of $\hat{A}^\dagger A$. Then*

$$\|I - \Sigma^2\|_2 \leq \|\hat{K} - K\|_2.$$

12:12 Efficient Tomography of Non-Interacting-Fermion States

Proof. Since A is column-orthonormal, $A^\dagger A = I$ (I is the identity matrix). Let $Q\Sigma V^\dagger$ be a singular value decomposition of $\hat{A}^\dagger A$. First, note that

$$V\Sigma^2 V^\dagger = (Q\Sigma V^\dagger)^\dagger Q\Sigma V^\dagger = (\hat{A}^\dagger A)^\dagger \hat{A}^\dagger A = A^\dagger \hat{A} \hat{A}^\dagger A.$$

Therefore,

$$\|I - \Sigma^2\|_2 = \|I - A^\dagger \hat{A} \hat{A}^\dagger A\|_2 = \|A^\dagger A - A^\dagger \hat{A} \hat{A}^\dagger A\|_2 = \|A^\dagger (A - \hat{A} \hat{A}^\dagger A)\|_2 \leq \|A - \hat{A} \hat{A}^\dagger A\|_2,$$

where the first step uses the fact that the spectral norm is unitarily invariant and the final step follows from the submultiplicativity of matrix norms and that $\|A^\dagger\|_2 = \|A\|_2 = 1$ since A is column-orthonormal. Finally,

$$\|A - \hat{A} \hat{A}^\dagger A\|_2 = \|AA^\dagger A - \hat{A} \hat{A}^\dagger A\|_2 \leq \|AA^\dagger - \hat{A} \hat{A}^\dagger\|_2 = \|K - \hat{K}\|_2. \quad \blacktriangleleft$$

We are now ready to show that if two kernel matrices are close, then the corresponding states will also be close.

► **Theorem 5.** *Let $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ be non-interacting-fermion states of n fermions and m modes described by the $m \times n$ column-orthonormal matrices $A, \hat{A} \in \mathbb{C}^{m \times n}$, respectively. Define $K := AA^\dagger$ and $\hat{K} := \hat{A} \hat{A}^\dagger$. Then*

$$d_{\text{tr}}(|\hat{\Psi}\rangle, |\Psi\rangle) \leq \sqrt{n \|\hat{K} - K\|_2}.$$

Proof. Recall that $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ can be written as

$$|\Psi\rangle = \sum_{S \in \Lambda_{m,n}} \det(A_S) |S\rangle \quad \text{and} \quad |\hat{\Psi}\rangle = \sum_{S \in \Lambda_{m,n}} \det(\hat{A}_S) |S\rangle$$

for the column-orthonormal matrices $A, \hat{A} \in \mathbb{C}^{m \times n}$. Then

$$\begin{aligned} d_{\text{tr}}(|\hat{\Psi}\rangle, |\Psi\rangle) &= \sqrt{1 - |\langle \hat{\Psi} | \Psi \rangle|^2} \\ &= \sqrt{1 - \left| \sum_{S \in \Lambda_{m,n}} \det(\hat{A}_S)^* \det(A_S) \right|^2} \\ &= \sqrt{1 - \left| \sum_{S \in \Lambda_{m,n}} \det(\hat{A}_S^\dagger) \det(A_S) \right|^2} \\ &= \sqrt{1 - |\det(\hat{A}^\dagger A)|^2}, \end{aligned}$$

where the second-to-last step follows because, for any square matrix X , $\det(X^T) = \det(X)$ and $\det(X)^* = \det(X^*)$; and the final step follows from the Cauchy-Binet formula.

Let $Q\Sigma V^\dagger$ be a singular value decomposition of $\hat{A}^\dagger A$, where $\sigma_1, \dots, \sigma_n$ are the singular values on the diagonal of Σ . Then

$$|\det(\hat{A}^\dagger A)|^2 = |\det(Q\Sigma V^\dagger)|^2 = |\det(Q) \det(\Sigma) \det(V^\dagger)|^2 = \det(\Sigma)^2 = \prod_i \sigma_i^2.$$

Note that, for all $i \in [n]$, $\sigma_i \leq \sigma_i(\hat{A}^\dagger)\|A\|_2 = 1$ [14, Chapter 3], where $\sigma_i(\hat{A}^\dagger)$ is the i th singular value of \hat{A}^\dagger . Plugging this into our bound on the trace distance, we get

$$\begin{aligned}
d_{\text{tr}}(|\hat{\Psi}\rangle, |\Psi\rangle) &= \sqrt{1 - |\det(\hat{A}^\dagger A)|^2} \\
&= \sqrt{1 - \prod_i \sigma_i^2} \\
&\leq \sqrt{n \left(\max_i 1 - \sigma_i^2\right)} && \text{(By Lemma 3).} \\
&= \sqrt{n\|I - \Sigma^2\|_2} \\
&\leq \sqrt{n\|\hat{K} - K\|_2} && \text{(By Lemma 4).} \quad \blacktriangleleft
\end{aligned}$$

4.2 Completing the Analysis

We prove that the fermionic state output by Section 3 is close to the input state in trace distance. To do so, we make use of Weyl's inequality, which implies that the spectrum of a Hermitian matrix is stable under small perturbations.

► **Theorem 6** (A Consequence of Weyl's Inequality [32]). *Let $M, N, R \in \mathbb{C}^{n \times n}$ be $n \times n$ Hermitian matrices such that $M = N + R$. Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of M , and let μ_1, \dots, μ_n be the eigenvalues of N . Then, for all $i \in [n]$,*

$$|\lambda_i - \mu_i| \leq \|R\|_2.$$

We are now ready to prove that Section 3 successfully learns a non-interacting-fermion state.

► **Theorem 7.** *Let \hat{A} be the output of Section 3 when given $|\Psi\rangle$ as input, and let $|\hat{\Psi}\rangle$ be the non-interacting-fermion state described by \hat{A} . Then*

$$d_{\text{tr}}(|\hat{\Psi}\rangle, |\Psi\rangle) \leq \sqrt{2nm\gamma}.$$

Proof. It is convenient to recall the last steps of Section 3:

In our algorithm, once the quantum measurements are complete, we have a matrix \hat{K} whose entries are within γ in magnitude of K . We then compute the eigendecomposition $Q\Lambda Q^\dagger$ of \hat{K} , where the first column of Q is the eigenvector corresponding to the largest eigenvalue of \hat{K} and so on. Finally, we set \hat{A} to be the $m \times n$ matrix corresponding to the first n columns of Q , and output \hat{A} . Therefore, the kernel matrix of $|\hat{\Psi}\rangle$ is $\hat{A}\hat{A}^\dagger$, and, by Theorem 5, the trace distance between $|\Psi\rangle$ and $|\hat{\Psi}\rangle$ is bounded above by $\sqrt{n\|\hat{A}\hat{A}^\dagger - K\|_2}$, where $K = AA^\dagger$ is the kernel matrix corresponding to the input state $|\Psi\rangle$. To complete the proof, we must bound $\|\hat{A}\hat{A}^\dagger - K\|_2$. To that end, by the triangle inequality,

$$\|\hat{A}\hat{A}^\dagger - K\|_2 \leq \|\hat{A}\hat{A}^\dagger - \hat{K}\|_2 + \|\hat{K} - K\|_2. \quad (2)$$

Observe that $\hat{K} = K + E$, where E is a perturbation of K whose entries have magnitude at most γ . Therefore, $\|\hat{K} - K\|_2 = \|K + E - K\|_2 = \|E\|_2$. The error matrix E is Hermitian because K and \hat{K} are Hermitian, and Hermitian matrices are closed under addition/subtraction. Therefore, since \hat{K} , K , and E are all Hermitian, we can use Theorem 6 to upper bound the absolute difference between the eigenvalues of K and \hat{K} . In particular, the absolute difference between the i th eigenvalues of K and \hat{K} is at most $\|E\|_2$, for all $i \in [m]$.

12:14 Efficient Tomography of Non-Interacting-Fermion States

Let $\mathbb{1}_n = \text{diag}(1, \dots, 1, 0, \dots, 0) \in \mathbb{R}^{m \times m}$ be the diagonal matrix whose first n diagonal entries are 1 and the rest 0. Observe that $\hat{A}\hat{A}^\dagger = Q\mathbb{1}_nQ^\dagger$, since \hat{A} is the first n columns of Q , and recall that $\hat{K} = Q\Lambda Q^\dagger$. Therefore,

$$\|\hat{A}\hat{A}^\dagger - \hat{K}\|_2 = \|Q\mathbb{1}_nQ^\dagger - Q\Lambda Q^\dagger\|_2 = \|\mathbb{1}_n - \Lambda\|_2,$$

where, in the last equality, we use the fact that the spectral norm is unitarily invariant. Note that Λ contains the eigenvalues of \hat{K} , and $\mathbb{1}_n$ contains the eigenvalues of K (since K is a trace- n , rank- n projector). Therefore, $\|\mathbb{1}_n - \Lambda\|_2$ is the maximum absolute difference between the eigenvalues of K and \hat{K} , which is at most $\|E\|_2$, as we argued in the previous paragraph.

The trivial bound on the spectral norm of E is the Frobenius norm of E , which is maximum when all entries of E have magnitude γ . Specifically, $\|E\|_2 \leq \|E\|_F \leq m\gamma$. Note $\|E\|_2 = \|E\|_F$ when the entries of E are all γ , so we cannot hope for a tighter bound on $\|E\|_2$.

Plugging this into Equation (2),

$$\begin{aligned} \|\hat{A}\hat{A}^\dagger - K\|_2 &\leq \|\hat{A}\hat{A}^\dagger - \hat{K}\|_2 + \|\hat{K} - K\|_2 \\ &\leq \|E\|_2 + \|E\|_2 \\ &\leq 2m\gamma. \end{aligned}$$

5 Connections to Quantum State Tomography

Although physically different, our problem is closely related to the quantum state tomography problem. In quantum state tomography, we want to recover an unknown Hermitian matrix, namely a d -dimensional mixed state $\rho \in \mathbb{C}^{d \times d}$, and applying a quantum circuit V to ρ maps ρ to $V\rho V^\dagger$. In our problem, applying a unitary V to $|\Psi\rangle$ maps K to VKV^\dagger , where K is an unknown Hermitian matrix, and our algorithm is able to recover the entries of K to within γ in magnitude. Therefore, our algorithm can also be viewed as a state tomography algorithm: measure copies of ρ in the $O(d)$ measurement bases obtained by choosing perfect matchings that cover all (i, j) pairs and output the resulting matrix $\hat{\rho}$ (skipping the last few steps of the algorithm that involve computing an eigendecomposition). As before, the algorithm requires $O(d \log(1/\delta)/\gamma^2)$ copies and $O(d^2 \log(1/\delta)/\gamma^2)$ time.

The error analysis is slightly different than for learning a fermionic state. For the state tomography problem, we want the output matrix $\hat{\rho}$ to be close to ρ in trace distance, which is proportional to $\|\hat{\rho} - \rho\|_1$, whereas, for fermionic tomography, our trace distance upper bound is proportional to $\|\hat{K} - K\|_2$ (see Theorem 5). Hence, to analyze the performance of our algorithm for state tomography, we must upper bound $\|\hat{\rho} - \rho\|_1$.

Recall that the error matrix $E = \hat{\rho} - \rho$ is Hermitian and has entries with magnitude at most γ . For $i \in [m]$, let λ_i denote the eigenvalues of E . Then

$$d_{\text{tr}}(\hat{\rho}, \rho) = \frac{1}{2} \|\hat{\rho} - \rho\|_1 = \frac{1}{2} \|E\|_1 = \frac{1}{2} \sum_i |\lambda_i| \leq \frac{\sqrt{d}}{2} \sqrt{\sum_i |\lambda_i|^2} \leq \frac{1}{2} d^{3/2} \gamma.$$

The first inequality follows from the fact that the arithmetic mean is bounded above by the quadratic mean, and the second inequality follows from the fact that $\sqrt{\sum_i |\lambda_i|^2} = \|E\|_F \leq d\gamma$. For $d_{\text{tr}}(\hat{\rho}, \rho) \leq \epsilon$, we must set $\gamma = 2d^{-3/2}\epsilon$. The resulting copy complexity is $O(d^4 \log(1/\delta)/\epsilon^2)$ and time complexity is $O(d^5 \log(1/\delta)/\epsilon^2)$. Note that the optimal copy complexity for quantum state tomography *with unentangled measurements* is $\Theta(d^3/\epsilon^2)$ [18, 12, 11], compared to $\Theta(d^2/\epsilon^2)$ with entangled measurements [22, 12].

Finally, we show that our upper bound on $\|\hat{\rho} - \rho\|_1$ is tight. Let F be the $d \times d$ Fourier transform whose (i, j) entry is $\exp(2\pi ij/d)/\sqrt{d}$. Then F scaled by a factor of $\sqrt{d}\gamma$ is a valid error matrix whose 1-norm is equal to $d^{3/2}\gamma$. Indeed, any $d \times d$ unitary matrix scaled by a factor of $\gamma\sqrt{d}$ will match our upper bound.

6 Open Problems

Perhaps the most interesting open problem is to give an algorithm to learn fermionic distributions using only standard basis measurements. Specifically, the following problems remain open:

1. **Learn real DPPs in variation distance.** Given sample access to a distribution induced by an $m \times m$ symmetric matrix $K \in \mathbb{R}^{m \times m}$, output an $m \times m$ matrix \hat{K} such that the induced distribution is close in variation distance. (See Section 1.2.1 for detail on DPPs.)
2. **Hermitian principal minor assignment problem.** Given a list of all 2^m principal minors of an unknown Hermitian matrix $K \in \mathbb{C}^{m \times m}$, reconstruct any Hermitian matrix that is consistent with that list.
3. **Learn non-interacting fermion distributions with standard basis measurements.** Given sample access to a non-interacting fermion distribution, efficiently learn the distribution in total variation distance.

The third problem is in some sense a combination of the first and second. To solve the first problem, we believe that the connections between DPPs and graph theory used in Rising et al. [27] and Urschel et al. [29] should be enough to develop an efficient algorithm. As discussed in Section 1.2.1, our work shows that there are many kernel matrices that have the same principal minors (indeed, any set non-interacting states that have the same distribution over the standard basis will give rise to a set of kernel matrices that are consistent with the same list of principal minors). The goal for the second and third problems is to output any one of the valid matrices.

For the second problem, however, the following example shows that some combinatorial information about the kernel matrix K , above and beyond the obvious complex conjugation ambiguities, is *not* determined even in principle by K 's principal minors. Consider the following 4×4 Hermitian matrix:

$$K = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & a & b^* \\ 1 & a^* & 1 & c \\ 1 & b & c^* & 1 \end{pmatrix}.$$

Suppose we have learned, by looking at the 2×2 and 3×3 principal minors, that

$$a = e^{ix}, \quad b = e^{iy}, \quad \text{and} \quad c = e^{iz},$$

where $|x| = |y| = |z| = w$ for some w that is known. That is, we have determined a, b , and c up to complex conjugation, and up to complex conjugation they are all equal. By looking at the bottom-most 3×3 principal minor, we can learn $\text{Re}(abc)$ and hence $|x + y + z|$. Suppose that this is also w . From the 4×4 minor, combined with the 2×2 and 3×3 minors, we get one additional piece of information, namely:

$$\text{Re}(ab) + \text{Re}(ac) + \text{Re}(bc).$$

12:16 Efficient Tomography of Non-Interacting-Fermion States

Suppose that, as expected, this is $2 + \cos(2w)$. Then even though we have extracted all information from the principal minors, there are *still* three essentially different solutions possible. Namely,

$$(1) x = y = w \text{ and } z = -w, \quad (2) x = z = w \text{ and } y = -w, \quad (3) y = z = w \text{ and } x = -w.$$

Of course, for fermionic distributions, the matrix K must be Hermitian, positive semi-definite, and a projector, and $\text{rank}(K) = n$. The example above is neither positive semi-definite nor a projector. However, we conjecture that this example can be embedded into a larger matrix that does satisfy these constraints.

Other directions for future work include improving the copy and time complexities of our algorithm, or giving conditional or unconditional lower bounds. Currently, the best lower bound we know is that $\Omega(m/\log m)$ measurements are needed, just from an information-theoretic argument (each measurement gives at most $n \log m$ bits of information and the state is characterized by $2nm$ real parameters). Similarly, $\Omega(mn)$ time is needed just to write down the output.

Since $\text{rank}(K) = n$, it should be possible to reduce the number of measurement bases from $O(m)$ to $O(n)$ (perhaps with the low-rank matrix recovery techniques used in [18]). Doing so would yield an immediate improvement in the copy and time complexities of our algorithm.

Also, just as our algorithm can be adapted to quantum state tomography, it is possible that the converse holds. Can quantum state tomography algorithms (in the entangled or unentangled measurement setting) be adapted to non-interacting-fermion state tomography? Also, do quantum state tomography lower bounds imply lower bounds for learning non-interacting-fermion state? In analogy with quantum state tomography, perhaps $\Theta(mn)$ copies are optimal to learn non-interacting-fermion states with entangled measurements and $\Theta(mn^2)$ copies are optimal with unentangled measurements.

It would be interesting to generalize our algorithm – for example, to superpositions over different numbers of fermions, or fermionic circuits that take inputs – and to find other classes of quantum states that admit efficient learning algorithms (for example, perhaps low-entanglement states or the outputs of small-depth circuits or low-stabilizer-complexity states [10]). We remark that [24] gives evidence that generalizing our algorithm to superpositions over different numbers of fermions may not be possible unless one limits the number of terms in the superposition. On the other hand, [24, Theorem 8] might be useful in developing learning algorithms for matchgate circuits.

Finally, what can be said about learning non-interacting *boson* states? The goal would be to reconstruct an $m \times n$ column-orthonormal matrix A given copies of a non-interacting boson state. However, the boson case is even trickier than the fermion case. In particular, boson statistics no longer depend only on the inner products between the rows of A , the way fermion statistics do. Indeed, even if we collected enough information to reconstruct a bosonic state, it seems that any algorithm would have to solve a quite complicated set of nonlinear equations.

References

- 1 Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, 2007. doi:10.1098/rspa.2007.0113.
- 2 Scott Aaronson. Shadow Tomography of Quantum States. *SIAM Journal on Computing*, 49(5):STOC18–368–STOC18–394, 2020. doi:10.1145/3188745.3188802.
- 3 Scott Aaronson. Introduction to Quantum Information Science II Lecture Notes, 2022. scottaaronson.com/qisii.pdf.

- 4 Scott Aaronson and Alex Arkhipov. BosonSampling is Far From Uniform. *Quantum Information and Computation*, 14(15-16):1383–1423, 2014.
- 5 Scott Aaronson and Sabeel Grewal. Efficient Learning of Non-Interacting Fermion Distributions. *arXiv preprint arXiv:2102.10458v2*, 2021.
- 6 Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. Optimal algorithms for learning quantum phase states, 2022. doi:10.48550/arxiv.2208.07851.
- 7 Costin Bădescu and Ryan O’Donnell. Improved Quantum Data Analysis. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1398–1411, 2021. doi:10.1145/3406325.3451109.
- 8 Clément L. Canonne. A short note on learning discrete distributions. *arXiv preprint arXiv:2002.11457*, 2020. doi:10.48550/arXiv.2002.11457.
- 9 Marcus Cramer, Martin B. Plenio, Steven T. Flammia, Rolando Somma, David Gross, Stephen D. Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature communications*, 1(1):1–7, 2010. doi:10.1038/ncomms1147.
- 10 Sabeel Grewal, Vishnu Iyer, William Kretschmer, and Daniel Liang. Low-Stabilizer-Complexity Quantum States Are Not Pseudorandom. *arXiv preprint arXiv:2209.14530*, 2022. doi:10.48550/arXiv.2209.14530.
- 11 Madalin Guță, Jonas Kahn, Richard Kueng, and Joel A. Tropp. Fast state tomography with optimal error bounds. *Journal of Physics A: Mathematical and Theoretical*, 53(20):204001, 2020. doi:10.1088/1751-8121/ab8111.
- 12 Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-Optimal Tomography of Quantum States. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017. doi:10.1109/TIT.2017.2719044.
- 13 Pierre Hohenberg and Walter Kohn. Inhomogeneous Electron Gas. *Physical Review*, 136(3B):B864, 1964. doi:10.1103/PhysRev.136.B864.
- 14 Roger A. Horn and Charles R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1994. doi:10.1017/CB09780511840371.
- 15 Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020. doi:10.1038/s41567-020-0932-7.
- 16 Emanuel Knill. Fermionic Linear Optics and Matchgates. *arXiv preprint quant-ph/0108033*, 2001. doi:10.48550/arXiv.quant-ph/0108033.
- 17 Walter Kohn and Lu Jeu Sham. Self-Consistent Equations Including Exchange and Correlation Effects. *Physical Review*, 140(4A):A1133, 1965. doi:10.1103/PhysRev.140.A1133.
- 18 Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low Rank Matrix Recovery From Rank One Measurements. *Applied and Computational Harmonic Analysis*, 42(1):88–116, 2017. doi:10.1016/j.acha.2015.07.007.
- 19 Guang Hao Low. Classical shadows of fermions with particle number symmetry, 2022. doi:10.48550/arxiv.2208.08964.
- 20 Odile Macchi. The Coincidence Approach to Stochastic Point Processes. *Advances in Applied Probability*, 7(1):83–122, 1975. doi:10.2307/1425855.
- 21 Ashley Montanaro. Learning stabilizer states by Bell sampling. *arXiv:1707.04012*, 2017. doi:10.48550/arXiv.1707.04012.
- 22 Ryan O’Donnell and John Wright. Efficient Quantum Tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pages 899–912, 2016. doi:10.1145/2897518.2897544.
- 23 Bryan O’Gorman. Fermionic tomography and learning, 2022. doi:10.48550/arxiv.2207.14787.
- 24 Michał Oszmaniec, Ninnat Dangniam, Mauro E.S. Morales, and Zoltán Zimborás. Fermion Sampling: A Robust Quantum Computational Advantage Scheme Using Fermionic Linear Optics and Magic Input States. *PRX Quantum*, 3:020328, 2022. doi:10.1103/PRXQuantum.3.020328.

- 25 Google AI Quantum and Collaborators. Hartree-Fock on a superconducting qubit quantum computer. *Science*, 369(6507):1084–1089, 2020. doi:10.1126/science.abb9811.
- 26 Michael Reck, Anton Zeilinger, Herbert J. Bernstein, and Philip Bertani. Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.*, 73:58–61, 1994. doi:10.1103/PhysRevLett.73.58.
- 27 Justin Rising, Alex Kulesza, and Ben Taskar. An Efficient Algorithm for the Symmetric Principal Minor Assignment Problem. *Linear Algebra and its Applications*, 473:126–144, 2015. doi:10.1016/j.laa.2014.04.019.
- 28 Barbara M. Terhal and David P. DiVincenzo. Classical simulation of noninteracting-fermion quantum circuits. *Physical Review A*, 65(3):032325, 2002. doi:10.1103/PhysRevA.65.032325.
- 29 John Urschel, Victor-Emmanuel Brunel, Ankur Moitra, and Philippe Rigollet. Learning Determinantal Point Processes with Moments and Cycles. In *International Conference on Machine Learning*, pages 3511–3520. PMLR, 2017.
- 30 Leslie G. Valiant. Quantum Circuits That Can Be Simulated Classically in Polynomial Time. *SIAM Journal on Computing*, 31(4):1229–1254, 2002. doi:10.1137/S0097539700377025.
- 31 Kianna Wan, William J. Huggins, Joonho Lee, and Ryan Babbush. Matchgate Shadows for Fermionic Quantum Simulation, 2022. doi:10.48550/arxiv.2207.13723.
- 32 Hermann Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912. doi:10.1007/BF01456804.
- 33 Gian-Carlo Wick. The Evaluation of the Collision Matrix. *Physical Review*, 80(2):268, 1950. doi:10.1103/PhysRev.80.268.
- 34 Andrew Zhao, Nicholas C. Rubin, and Akimasa Miyake. Fermionic Partial Tomography via Classical Shadows. *Phys. Rev. Lett.*, 127:110504, 2021. doi:10.1103/PhysRevLett.127.110504.