

# Learning quantum states prepared by shallow circuits in polynomial time

Zeph Landau\*

Yunchao Liu†

## Abstract

We give a polynomial time algorithm that, given copies of an unknown quantum state  $|\psi\rangle = U|0^n\rangle$  that is prepared by an unknown constant depth circuit  $U$  on a finite-dimensional lattice, learns a constant depth quantum circuit that prepares  $|\psi\rangle$ . The algorithm extends to the case when the depth of  $U$  is  $\text{polylog}(n)$ , with a quasi-polynomial run-time. The key new idea is a simple and general procedure that efficiently reconstructs the global state  $|\psi\rangle$  from its local reduced density matrices. As an application, we give an efficient algorithm to test whether an unknown quantum state on a lattice has low or high quantum circuit complexity.

## 1 Introduction

We consider the following fundamental question in quantum complexity theory: is there an efficient algorithm to learn quantum states of low circuit complexity? That is, given copies of an unknown quantum state  $|\psi\rangle$  that was generated from a constant depth quantum circuit acting on a finite-dimensional lattice, whether there exists an efficient algorithm that learns a circuit to prepare  $|\psi\rangle$ . In this paper we answer this question in the affirmative.

**Theorem 1** (Main result, simplified version of [Theorem 4](#)). *There is an algorithm that, given copies of an unknown state  $|\psi\rangle$ , with the promise that  $|\psi\rangle = U|0^n\rangle$  where  $U$  is an unknown depth- $d$  circuit acting on a  $k$ -dimensional lattice (using arbitrary 2-qubit gates), outputs a depth- $(2k+1)d$  circuit  $W$  that prepares  $|\psi\rangle$  up to 0.01 trace distance, with success probability 0.99. The algorithm uses  $M$  copies of  $|\psi\rangle$  and runs in time  $T$ , where*

$$M = \tilde{O}(n^4) \cdot 2^{O(c)}, \quad T = \tilde{O}(n^4) \cdot 2^{O(c)} + (nkd \cdot c)^{O(d \cdot c)}. \quad (1)$$

Here,  $c = O((3k)^{k+2}d)^k$ , and  $W$  uses  $r \cdot n$  ancilla qubits where  $r > 0$  can be chosen to be an arbitrarily small constant.

Note that the running time is polynomial when  $d = O(1)$ , and quasi-polynomial when  $d = \text{polylog}(n)$ . In addition, the dominating term in the running time (second term in  $T$ ) can be significantly improved when assuming a discrete gate set.

For quantum states prepared by shallow circuits, it is known that learning (sufficiently large) local reduced density matrices suffices to information-theoretically reconstruct the state [[RF21](#); [YW23](#)]. The question is whether the reconstruction can be computationally efficient. A naive approach finds small circuits for different local regions and stitch them together into a global circuit by checking local consistency, but this runs into a seemingly hard constraint satisfaction problem (in two and higher dimensions). Recent work [[HLB+24](#)] developed an efficient algorithm

---

\*UC Berkeley. [zeph.landau@gmail.com](mailto:zeph.landau@gmail.com)

†UC Berkeley. [yunchao.liu@berkeley.edu](mailto:yunchao.liu@berkeley.edu)

in 2D by showing that to learn a 2D state it suffices to solve a 1D constraint satisfaction problem which is efficient; however this approach runs into the same issue at three and higher dimensions. Here we develop new techniques that do not rely on solving any consistency problem; this enables our algorithm to work at arbitrary dimensions.

**Testing quantum circuit complexity.** As an application of our result, we also give an algorithm to test whether an unknown state on a  $k$ -dimensional lattice has low or high quantum circuit complexity.

**Corollary 1** (Simplified version of [Theorem 5](#)). *Fix some constant  $L > 0$ . Given copies of an unknown state  $|\psi\rangle$  on a  $k$ -dimensional lattice, with the promise that one of the following two cases hold:*

- **Case 1: Low complexity.**  $|\psi\rangle = U |0^n\rangle$  where the depth of  $U$  is at most  $L$ ;
- **Case 2: High complexity.** Any state prepared by a constant depth circuit using  $O(n)$  ancilla qubits is at least 0.01-far from  $|\psi\rangle$  in trace distance.

*There is an algorithm that decides which is the case, with success probability at least 0.99, with polynomial sample and time complexity.*

Next we give a brief review of the background and motivations of this work.

**Learning phases of matter.** Quantum systems defined on finite-dimensional lattices are a central subject in condensed matter physics, where quantum states are classified into different “phases of matter” [[CGW10](#)]. Quantum circuit complexity plays an important role in the definition of phases of matter: the “trivial” phase is typically defined as quantum states prepared by a constant (or polylog( $n$ )) depth circuit acting on a  $k$ -dimensional lattice (e.g. [[AT18](#); [PSC21](#)]), while quantum states in a topologically ordered phase have high circuit complexity [[BHV06](#)]. Our result therefore shows that:

- Quantum states in the “trivial” phase can be learned in polynomial (or quasi-polynomial) time.
- Given an arbitrary quantum state, there is an efficient algorithm to test whether it is in the “trivial” phase or some other high-complexity phase.

**Quantum algorithms in NISQ.** Shallow quantum circuits provide a natural theoretical model for noisy intermediate-scale quantum (NISQ) computation, and optimism about quantum advantage from NISQ devices relies on the fact that shallow quantum circuits can generate nontrivial probability distributions that are hard to simulate classically [[TD04](#); [GWD17](#); [BHS+18](#); [HE23](#)]. This motivates many NISQ algorithms which heuristically use shallow quantum circuits as an ansatz; the goal is to optimize a parameterized shallow quantum circuit to solve interesting problems [[BCK+22](#)]. In practice, these algorithms run into issues such as local minima [[AK22](#)]. Our learning algorithm addresses a simple question in this direction: provably learning a shallow quantum circuit to prepare an unknown quantum state, which potentially provides a useful primitive for new NISQ algorithms.

**Other related works.** The classical question of learning low-depth Boolean circuits (e.g.  $\text{NC}^0$  and  $\text{AC}^0$ ) is well-studied [MOS03; LMN93; CIKK16]. Recent work [HLB+24] also gave an efficient algorithm for learning shallow quantum circuits from random input/output samples, which can be viewed as a quantum analog of learning shallow Boolean circuits. Our problem of learning quantum states prepared by shallow circuits can be viewed as another analog, but here there is no access of “input” to the shallow circuit and the learning algorithm can only make measurements to the quantum state.

Separately, developing efficient algorithms to learn interesting families of quantum states is a major direction in quantum learning theory, and we refer to [AA24] for a recent survey.

**Discussion.** An interesting question is whether our algorithm can be generalized to other geometries (or interaction graphs) beyond finite-dimensional lattices. In fact, we show that our algorithm works for any geometry which has a property called “covering scheme” (Definition 4), and we construct covering schemes for finite-dimensional lattices. It is interesting to study what other geometries can have this covering scheme.

**Note added.** We recently became aware of independent related work of Hyun-Soo Kim, Isaac Kim, and Daniel Ranard [KKR24], achieving similar results via a different approach.

## 2 Learning algorithm

**Notations.** Our goal is to learn a quantum state  $|\psi\rangle$ , with the promise that  $|\psi\rangle = U|0^n\rangle$  where  $U$  is an unknown depth- $d$  circuit acting on a  $k$ -dimensional lattice. We do not assume knowledge of the circuit architecture: each layer of the circuit consists of non-overlapping 2-qubit gates, where each qubit could interact with any of its neighbors. The following concepts are used throughout the argument.

- **Ball:**  $\mathcal{B}(A, r)$  denotes the radius- $r$  neighborhood on the lattice for a set of vertices  $A$  (including  $A$ ). For example, the dotted region in Fig. 1b shows a ball around region  $A$  on the 2D lattice.
- **Lightcone:**  $\mathcal{L}(A, d)$  denotes the volume of locations that can be reached by a depth  $d$  circuit starting from region  $A$ . In particular, the support of the lightcone at top layer equals  $\mathcal{B}(A, d)$ . For example, the green region in Fig. 1a denotes the lightcone of the leftmost qubit for a circuit defined on 1D lattice.

A lightcone can be viewed as propagating or spreading causal influence in a circuit from input to output. Later we will define a dual notion of *backward* lightcone, which spreads from output to input.

### 2.1 Overview of technical challenges and new ideas

We first give an overview of the technical challenges in developing a learning algorithm, and our new ideas to overcome these challenges. We start by defining a key concept of local inversion.

**Definition 1** (Local inversion). *Given a state  $|\psi\rangle$  and a subset  $A \subseteq [n]$  of qubits, a unitary operator  $V$  is called a local inversion of  $A$  if  $V$  acts on a ball of  $A$  and  $V|\psi\rangle = |0\rangle_A \otimes |\phi\rangle$  for some arbitrary pure state  $|\phi\rangle$  on  $n - |A|$  qubits.*

**Fact 1.** *Suppose  $|\psi\rangle = U|0^n\rangle$  where  $U$  is a depth- $d$  circuit acting on a  $k$ -dimensional lattice. Then for any  $A \subseteq [n]$ , there exists a local inversion  $V$  of  $A$  satisfying:*




Figure 1: Local inversions for quantum states prepared by shallow circuits. (a) 1D lattice; (b) 2D lattice, where a local inversion of  $A$  can be constructed by applying a depth- $d$  circuit on  $AB$ .

1.  $V$  is supported on  $\mathcal{B}(A, d)$ ;
2.  $V$  is a depth  $d$  circuit, whose inverse shape is contained within the lightcone  $\mathcal{L}(A, d)$ .

This fact follows from *undoing* the gates in the lightcone of  $A$ . For example, Fig. 1a shows a state prepared by a circuit acting on a 1D lattice, and the green region denotes the lightcone of the leftmost qubit. There exists a local inversion  $V_1$  (for example, by inverting the gates in the lightcone) which disentangles the leftmost qubit. Note that the shape of  $V_1$  is the inverse of the shape of the lightcone. Fig. 1b shows a 2D lattice, where the dotted region  $AB$  equals  $\mathcal{B}(A, d)$ . There exists a local inversion of  $A$  which is a depth- $d$  circuit acting on  $AB$ .

**Learning local inversions.** Local inversions provide the basic tool for a learning algorithm because they are easy to learn. For example, consider a constant-size region  $A$  in Fig. 1b. We first perform quantum state tomography to learn the reduced density matrix on  $AB$ , and then brute force search over all depth- $d$  circuits acting on  $AB$ . For each circuit, we can efficiently check whether it correctly inverts the region  $A$  to  $|0\rangle_A$ . In this way we can find possibly a set of local inversion operators of  $A$ .

For the remainder of Section 2, we will assume that we have access to (exact) local inversion operators. In reality, we can only learn *approximate* local inversion operators. This issue is addressed in Section 3.

Our algorithm has a simple framework: (1) Quantum learning: learn local reduced density matrices; (2) Classical processing: find local inversion operators for local regions, and combine them into a circuit. Below we review the challenges in realizing this framework.

**Challenges.** The first issue is that there are multiple local inversion operators for a given local region. A naive approach to find a circuit is the following: pick a local inversion operator for each local region, such that local inversions for neighboring regions are *consistent*. Consistency demands that two quantum circuits share the same gates where they overlap, so that they can be merged together as a bigger circuit. This is precisely a constraint satisfaction problem that is hard in general in 2D and higher dimensions.

Recent work [HLB+24] addresses this problem in 2D, by showing that one can efficiently solve such a constraint satisfaction problem in 1D, to find 1D circuits that disentangle the 2D state into many 1D stripes [HLB+24, Eq. (9)], and argue that the remaining 1D stripes are easy to learn.

This approach fails at three and higher dimensions as the disentangling step still requires solving a hard constraint satisfaction problem.

The key challenge in solving the learning problem on finite dimensional lattices is to find a way to avoid any constraint satisfaction problem. In this paper we give such a way.

Here is a basic idea: by definition, applying a local inversion  $V$  to a state  $|\psi\rangle$  gives  $V|\psi\rangle = |0\rangle_A \otimes |\phi\rangle$ , where  $|\phi\rangle$  is a *smaller* state on  $n - |A|$  qubits. Can we repeat the process by further removing qubits from  $|\phi\rangle$ ? The issue here is that now the state has been *disturbed*. In particular, the local inversion  $V$  we applied may not be the “ground truth” which undoes the gates in the lightcone of  $A$ , and  $V$  may just be an arbitrary depth- $d$  circuit that happens to invert  $A$ . Therefore the state  $|\phi\rangle$  suffers from a potential quantum circuit complexity blow-up: we no longer have the guarantee that  $|\phi\rangle$  is prepared by a depth- $d$  circuit. Repeating this process will further increase the blow-up.

**Key ideas.** Overcoming these obstacles requires two insights. The first is to observe that we can *undo* the local inversion after applying it, so the state  $|\psi\rangle$  is not disturbed. The utility of this observation is that in rewriting the state in this way, we have replaced part of the unknown state with a partial piece of known operations. The second insight is to realize that with a careful choice of local inversions based on the geometry of the lattice, these partial pieces can be layered together in a way that the *backward* lightcone of the final state is covered completely by these partial pieces and does not depend on the unknown initial state. The result is the learned constant depth circuit consisting of parts of local inversions and their inverses.


## 2.2 Replacement process

As discussed above, our first key idea is the following, which at first glance seems useless: apply a local inversion, and then undo it. We formally define this as a replacement process.

**Definition 2** (Replacement process). *Given a state  $|\psi\rangle$  and a region  $A$ , define the  $A$ -replacement process as follows: take a local inversion  $V$  of region  $A$ , then perform the following operations on  $|\psi\rangle$ :*

1. *apply  $V$ ,*
2. *trace out the qubits in  $A$ , replace each qubit with  $|0\rangle$ ,*
3. *apply  $V^\dagger$ .*

Note that step 2 is in fact not doing anything (since step 1 already inverted the region  $A$  to  $|0\rangle_A$ ) and is included here to help the illustration of the argument. Next, since step 2 is effectively the identity operation, step 1 and step 3 cancel with each other, and therefore the state  $|\psi\rangle$  remains unchanged. This is illustrated as follows (for simplicity, here we draw local inversions as boxes without the wedges).



**Fact 2.** *The state  $|\psi\rangle$  is invariant under any  $A$ -replacement process for any region  $A$ .*

Next we give an intuitive explanation of why the replacement process may be helpful for learning. First, we informally introduce the new concept of backward lightcone (green region in Eq. (2)). The formal definition is given in Definition 6.

**Definition 3** (Backward lightcone, informal). *The backward lightcone of a subset of qubits  $S$  at the output of a quantum circuit is the minimal part of the circuit diagram that determines the reduced density matrix of  $S$ .*

For example, in Eq. (2) the green region starts from  $A_0$  at the output of the circuit and keeps growing backwards (which looks like the inverse of a (forward) lightcone), until it hits a region of  $|0\rangle$ . There is no need to grow further because the input is completely determined. Now, note that the reduced density matrix of  $A_0$  at the output of the circuit (which equals the reduced density matrix of  $A_0$  in  $|\psi\rangle$ ) is determined by its backward lightcone: start from a region of  $|0\rangle$  which is larger than  $A_0$ , apply the green circuit, and trace out the qubits not in  $A_0$ . All quantum gates not in the green region are irrelevant since removing them does not affect the reduced density matrix of  $A_0$ .

Here is an interesting observation about the  $A$ -replacement process: suppose we choose  $A$  to be a (sufficiently large) ball of some smaller region  $A_0$  as in Eq. (2), then the backward lightcone of  $A_0$  ends at the freshly initialized qubits in step 2 of Definition 2. In particular, the backward lightcone does not reach the unknown state  $|\psi\rangle$ , which allows us to reconstruct the reduced density matrix of  $|\psi\rangle$  on  $A_0$  by a *known* circuit. And yet, due to the invariance of  $|\psi\rangle$  (Fact 2) we can pretend that nothing has happened to  $|\psi\rangle$  and repeat this process. In other words,

*Key observation: We have replaced part of the state with known operations, without disturbing the state.*

This observation suggests an approach to learning a circuit for  $|\psi\rangle$ : repeatedly apply  $A$ -replacement processes for different small regions  $A$ , and hope to have the backward lightcone of more and more output qubits be contained entirely within the replacement process, and hope that eventually this holds true for all of the output qubits. If we can manage this, it means we must have generated  $|\psi\rangle$  solely from the collection of replacement processes (which are constructed by quantum circuits that are known to us) and thus we can extract from them a circuit that can generate  $|\psi\rangle$ , simply via the backward lightcone of all output qubits.

It is not obvious that this can work, because we need to apply replacement processes not only in parallel, but also on top of each other, since a single layer cannot cover all output qubits (e.g. Fig. 2a). The issue is that applying replacement processes on top of each other changes the lightcone structure: for example, the backward lightcone of  $A_0$  may be much larger than in Eq. (2) if there are additional layers on top, because the backward lightcone starts from the output which is at the very top of the circuit.

In the next section we pin down the exact conditions for this approach to work.

### 2.3 Covering schemes and reconstruction circuits

It turns out that we can make this approach work if we can find a collection of small regions that satisfy the following conditions which we call a covering scheme. Our plan is:

1. In this section we show that a covering scheme implies a learning algorithm.
2. In Section 2.4 we show how to construct good covering schemes for  $k$ -dimensional lattices.

**Definition 4** (Covering scheme). A  $(\ell, c, d)$  covering scheme is a collection of subsets of qubits  $S_j^i \subseteq [n]$

$$S_1^1, S_2^1, \dots, S_{m_1}^1, S_1^2, S_2^2, \dots, S_{m_2}^2, \dots, S_1^\ell, S_2^\ell, \dots, S_{m_\ell}^\ell$$


which satisfy the following conditions.

1. The size of each  $\mathcal{B}(S_j^i, d)$  is upper bounded by  $c$ .
2. For every fixed  $i$ , the sets  $\mathcal{B}(S_j^i, d)$  are pairwise disjoint for  $1 \leq j \leq m_i$ .
3. For each qubit  $v \in [n]$ , there exists a  $S_j^i$  such that  $\mathcal{B}(\{v\}, (2\ell - 1)d) \subseteq S_j^i$ .

We can think of these subsets as being divided into  $\ell$  different layers: in each layer  $1 \leq i \leq \ell$ , there are subsets  $S_1^i, S_2^i, \dots, S_{m_i}^i$ . Condition 1 says that each of them is small (even after being enlarged by a radius of  $d$ ). Condition 2 says that the subsets in the same layer are disjoint, even after each of them is enlarged by a radius of  $d$ . Condition 3 says that for each qubit, a ball around that qubit (of radius  $(2\ell - 1)d$ ) is entirely contained within some subset.

### 2.3.1 Examples in 1D

The example below shows a covering scheme in 1D with  $\ell = 2$  layers.




Note that in each layer, the sets have some distance between each other, because we require that they remain disjoint even after being enlarged (Condition 2). In addition, the first layer and the second layer have a lot of overlap. This ensures that any qubit must lie within the interior of some set, which implies Condition 3.

Why is this useful? Recall that our idea is to repeatedly apply replacement processes. The reason to introduce covering schemes is the following:


*Key idea: Applying replacement processes for a covering scheme allows us to reconstruct the entire state.*

We first illustrate this idea in 1D in Fig. 2, and then give a formal argument in Section 2.3.2.


- **Reconstruction process.** Suppose we apply  $S_1^1$ -replacement process, which is supported on  $\mathcal{B}(S_1^1, d)$  and looks exactly the same as Eq. (2). Note that all replacement processes corresponding to the first layer in the covering scheme can be implemented in parallel, due to Condition 2 in Definition 4. Next, we apply all replacement processes corresponding to the second layer in the covering scheme. We call the resulting diagram a reconstruction process, shown in Fig. 2a.
- **Backward lightcone.** Now we construct the backward lightcone for all output wires of the reconstruction process (Fig. 2b). The way to do this is to color all gates at the top layer in green, and then “spread” the green color backwards, until hitting a region of  $|0\rangle$ . Note that some of the spreading stops at the second layer of  $|0\rangle$ , but inevitably some of the spreading goes beyond the second layer and enters the first layer. However, all of the spreading stops at the first layer of  $|0\rangle$  and never touches the bottom unknown state  $|\psi\rangle$ .



(a) Reconstruction process



(b) Backward lightcone



(c) Reconstructed circuit

Figure 2: Illustration of the learning algorithm in 1D.

- **Reconstructed circuit.** By [Fact 2](#), the state  $|\psi\rangle$  remains invariant under the reconstruction process, and therefore the output state at the top layer equals  $|\psi\rangle$ . The backward lightcone consists entirely of known quantum circuits, and therefore it gives a reconstructed circuit that prepares  $|\psi\rangle$  ([Fig. 2c](#)). This circuit has the following features: it has depth  $3d$  and acts on the all- $|0\rangle$  input state, where we view the red qubits as ancilla qubits. After the circuit is applied, the entire state equals  $|\psi\rangle \otimes |\text{junk}\rangle$ , where the wires at the top correspond to  $|\psi\rangle$ , and the red wires (ancilla qubits) correspond to  $|\text{junk}\rangle$ .

### 2.3.2 Reconstruction theorem

In this section we give a rigorous argument for how to reconstruct the state in general. We start with a formal definition of the reconstruction process.

**Definition 5** (Reconstruction process). *The reconstruction process for  $|\psi\rangle$  is defined as follows. Let  $V_{i,j}$  be a local inversion for  $S_j^i$ . By [Fact 1](#),  $V_{i,j}$  is a depth- $d$  circuit acting on  $\mathcal{B}(S_j^i, d)$ . The*



reconstruction process is defined as follows:

For each  $1 \leq i \leq \ell$ , apply the  $S_j^i$ -replacement process using  $V_{i,j}$  for all  $j$  in parallel.

More specifically, for each fixed  $i$  we do the following:

1. Apply  $V_{i,j}$  for all  $1 \leq j \leq m_i$  in parallel;
2. Trace out all qubits in  $S_j^i$  for each  $1 \leq j \leq m_i$  and replace with  $|0\rangle$ ;
3. Apply  $V_{i,j}^\dagger$  for all  $1 \leq j \leq m_i$  in parallel.

As shown in the 1D example in Fig. 2, a two-step argument is used to show that a circuit for preparing the unknown state  $|\psi\rangle$  can be extracted from the reconstruction process:

1. The state  $|\psi\rangle$  is invariant under each  $S_j^i$ -replacement process (Fact 2), and therefore is invariant under the entire reconstruction process.
2. We can reconstruct the output state of the reconstruction process (which equals  $|\psi\rangle$ ) via its backward lightcone, because the backward lightcone consists of known quantum circuits and in particular does not touch the unknown input state  $|\psi\rangle$  at the bottom.

Below we elaborate on the second point. At this point a more precise definition of the backward lightcone is needed. Note that this definition needs to be applicable to slightly more general quantum circuits with reset gates.

**Definition 6** (Backward lightcone). *Let  $|\phi\rangle = W|0^n\rangle$  where  $W$  is a quantum circuit consists of 2-qubit unitary gates and 1-qubit reset gates (a reset gate traces out the input and initializes a  $|0\rangle$  state). Let  $A \subseteq [n]$  be a subset of qubits. The backward lightcone of  $A$  is a circuit diagram which is part of the circuit diagram of  $W$ , defined as the collection of green gates acting on all-0 inputs, constructed as follows:*

1. Color the output wires of  $W$  corresponding to  $A$  in blue.
2. Repeat the following process until no changes happen:  
If there exists a 2-qubit gate  $G$  with a blue wire on top, color this gate in green. Moreover, consider each of the two wires at the bottom of  $G$ . If it is not connected to a reset gate, color the wire in blue.

*In other words, the backward lightcone consists of all quantum gates that could influence the reduced density matrix of  $|\phi\rangle\langle\phi|$  on  $A$ . The running time to construct the backward lightcone is linear in the size of  $W$ .*

The process to construct the backward lightcone can be viewed as spreading the green color from top to bottom, as shown in Fig. 2b. The desired property for the backward lightcone, when applied to a reconstruction process, is that it stops entirely at intermediate reset gates and does not reach the bottom.

To help further illustrate this concept, we introduce a dual thought experiment. A different way of formulating our desired property is that we do not want the influence of the bottom input state to reach any top wire. The propagation of the influence can be viewed as forward lightcone spreading of the input state, that is, spreading the white color in Fig. 2b from bottom to top. Note that indeed, the white color stops entirely at intermediate reset gates and does not reach the top.

Now we are ready to prove that the desired property is guaranteed by the covering scheme.

**Theorem 2** (Covering scheme implies learning algorithm). *Suppose  $|\psi\rangle = U|0^n\rangle$  where  $U$  is a depth- $d$  circuit acting on a  $k$ -dimensional lattice. Let  $(\ell, c, d)$  be a covering scheme for the  $k$ -dimensional lattice, and let  $W$  be the reconstruction process for this covering scheme, which satisfies  $W|\psi\rangle = |\psi\rangle$ . Then the backward lightcone of all output wires in  $W|\psi\rangle$  is a depth- $(2\ell - 1)d$  circuit  $C$  which is part of  $W$ .*

*The circuit  $C$  can be used to prepare  $|\psi\rangle$  in the following sense: it acts on  $n$  qubits as well as  $m = O(n)$  ancilla qubits, and*

$$C|0^n\rangle \otimes |0^m\rangle = |\psi\rangle \otimes |\text{junk}\rangle. \quad (4)$$

*Proof.* To prove that the backward lightcone of all output wires in  $W|\psi\rangle$  is part of  $W$  and does not reach the input state  $|\psi\rangle$ , it suffices to show that the backward lightcone of each individual output wire in  $W|\psi\rangle$  must stop at some regions of reset gates in some replacement processes.

Keeping track of the backward lightcone can be tricky: the backward spreading process can stop at various different layers such as in Fig. 2b. However, here we give a simple and pessimistic argument which suffices for our purpose.

Focus on a single output wire  $v \in [n]$  of  $W|\psi\rangle$  and imagine its backward spreading process. Part of the spreading may stop early at some reset gates, while other parts of the spreading may continue further. Instead of working with  $W$ , suppose we consider a new quantum circuit  $W'$  where all reset gates in  $W$  is replaced by the identity gate, with one exception: according to Condition 3 of Definition 4, there exists a  $S_j^i$  such that  $\mathcal{B}(\{v\}, (2\ell - 1)d) \subseteq S_j^i$ ; we keep the reset gates in  $W'$  that correspond to the  $S_j^i$ -replacement process.

Clearly, the backward spreading process in  $W$  is entirely contained within the backward spreading process in  $W'$ , since  $W'$  removed some reset gates relative to  $W$  which can only help the spreading. Now, observe the following: *the backward spreading process in  $W'$  must stop at the reset gates in the  $S_j^i$ -replacement process.* This is because  $S_j^i$  contains a ball around  $v$  with radius  $(2\ell - 1)d$ . By the time the spreading process reaches those reset gates, the process cannot spread further than a distance of  $(2\ell - 1)d$  and therefore is entirely covered by those reset gates. The number  $(2\ell - 1)d$  comes from a worst case estimate: suppose  $S_j^i$  is within the very bottom layer of the covering scheme, then the spreading process has gone through the top  $\ell - 1$  layers of depth  $2d$  as well as  $V_{i,j}^\dagger$  of depth  $d$ , with a total depth of  $(2\ell - 1)d$ . This implies that the backward spreading process of  $v$  in  $W$  must be contained within  $W$ : if it cannot reach the bottom even in  $W'$ , then it also cannot reach the bottom in  $W$ .

This concludes the proof that  $C$  is a (at most) depth- $(2\ell - 1)d$  circuit and is part of  $W$ . Finally, note that  $C$  is a unitary quantum circuit which outputs a pure state. In addition, the reduced density matrix on the  $n$  output wires equals  $|\psi\rangle\langle\psi|$ . This implies that the system and ancilla must be tensor product pure states as in Eq. (4).  $\square$

To conclude this section we give a recap about the role of different parameters of a  $(\ell, c, d)$  covering scheme in the learning algorithm.

- $d$  is the promised depth of the unknown circuit that prepares  $|\psi\rangle$ .
- $\ell$  determines the depth of the learned circuit, which equals  $(2\ell - 1)d$ .
- $c$  determines the running time of the learning algorithm: the algorithm needs to find local inversions acting on a region of size  $c$ , which takes time exponential in  $c$  and  $d$ .

Clearly, these parameters play an important role and are in tension with each other.

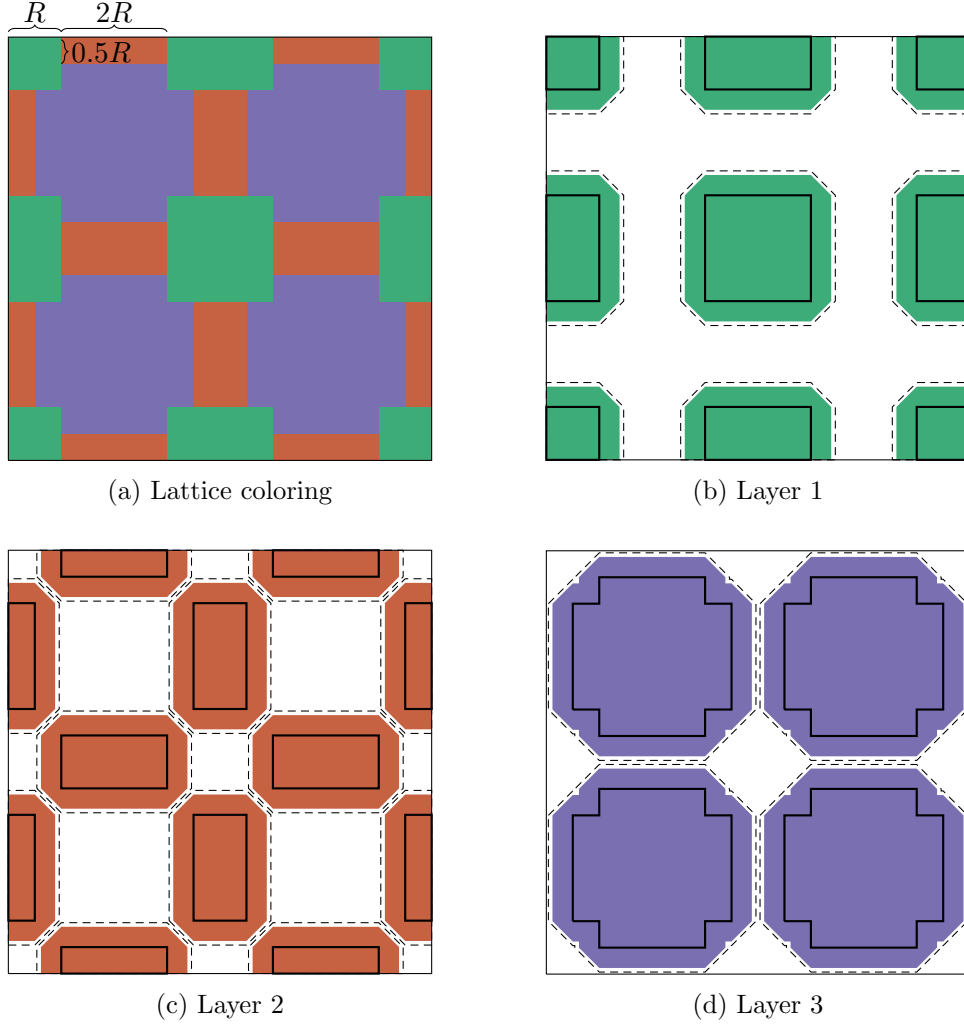


Figure 3: A covering scheme for the 2D lattice, which is derived from a lattice coloring. Here  $R = 13d$ .

## 2.4 Good covering schemes

In this section we construct good covering schemes for  $k$ -dimensional lattices.

**Theorem 3** (Lattice covering scheme). *For  $k$ -dimensional lattice there exists a  $(k + 1, c, d)$  covering scheme for any integer  $d > 0$ , where*

$$c \leq ((8k^2 + 14k + 2)d)^k. \quad (5)$$

Before giving the proof, we first discuss an example in 2D (Fig. 3). Our starting point is a new concept called lattice coloring.

**Definition 7** (Lattice coloring). *A  $(w, c, R)$  lattice coloring for the  $k$ -dimensional lattice is defined as follows. Suppose the lattice is divided into disjoint subsets where each subset is connected. Each subset is assigned a color. We demand the following properties:*

1. There are  $w$  different colors.

2. Each subset has size at most  $c$ .
3. Two different subsets with the same color must have distance at least  $R$ .

Fig. 3a shows a  $(3, 16R^2, R)$  coloring for the 2D lattice ( $16R^2$  is an overestimate). Here  $R$  is a free parameter. Similar coloring schemes have been widely used in the physics literature (e.g. [BK19]). Now, we use this coloring to construct a covering scheme for the 2D lattice.

**Lemma 1.** *For 2D lattice there exists a  $(3, 3844d^2, d)$  covering scheme for any integer  $d > 0$ .*

This covering scheme is shown in Figs. 3b to 3d. Below we explain it in detail.

The idea is the following: each color in the lattice coloring corresponds to a layer in the covering scheme (so  $\ell = 3$ ). We choose  $R$  to be large enough, and then

- Fix a color in the lattice coloring (say red regions in Fig. 3a), then for each small colored subset  $A$  in Fig. 3a, we assign a subset  $S = \mathcal{B}(A, 5d)$  to the covering scheme (red regions in Fig. 3c).

Now, we choose  $R$  to make sure that Condition 2 in Definition 4 is satisfied. Say we consider two red regions  $A$  and  $B$  in Fig. 3a that are separated by distance  $R$ . The corresponding subsets in the covering scheme are  $S_1 = \mathcal{B}(A, 5d)$  and  $S_2 = \mathcal{B}(B, 5d)$ . Condition 2 in Definition 4 demands that  $S_1$  and  $S_2$  are disjoint even after being enlarged further by distance  $d$  (dashed regions in Fig. 3c). That is,  $\mathcal{B}(A, 6d)$  and  $\mathcal{B}(B, 6d)$  must be disjoint. We therefore choose  $R = 13d$  to ensure this property.

Next, we show that Condition 3 in Definition 4 is satisfied. Take any qubit  $v \in [n]$ , then it must lie within *some* colored region in Fig. 3a. Suppose the region is red and call it  $A$ . Then  $\mathcal{B}(\{v\}, 5d) \subseteq \mathcal{B}(A, 5d)$  which is one of the subsets in the covering scheme in Fig. 3c.

Finally, note that each small colored region in Fig. 3a is contained within a box of  $4R \times 4R$ . The length scale of each subset in the covering scheme is at most  $4R + 2 \times 5d = 62d$ . Therefore, all subsets in the 2D covering scheme has size at most  $62d \times 62d = 3844d^2$ .

*Proof of Theorem 3.* The proof essentially repeats the above example. We first quote a coloring scheme for the  $k$ -dimensional lattice (see e.g. [HLB+24, below Definition 17]).

**Lemma 2** ( $k$ -dimensional lattice coloring). *For  $k$ -dimensional lattice there exists a  $(k + 1, (2kR)^k, R)$  lattice coloring. Each small colored region is contained in a  $k$ -dimensional box of length scale  $2kR$ .*

Each color in the lattice coloring corresponds to a layer in the covering scheme (so  $\ell = k + 1$ ). We choose  $R$  to be large enough, and then for each small colored subset  $A$  in the lattice coloring, we assign a subset  $S = \mathcal{B}(A, (2k + 1)d)$  to the corresponding layer of the covering scheme.

- The depth of learned circuit is  $(2\ell - 1)d = (2k + 1)d$ .
- To ensure Condition 2 in Definition 4 is satisfied, we choose  $R = 2 \times (2k + 2)d + d = (4k + 5)d$ .
- The length scale of each subset in the covering scheme is at most  $2kR + 2 \times (2k + 1)d = (8k^2 + 14k + 2)d$ . Therefore,  $c$  is at most  $((8k^2 + 14k + 2)d)^k$ .

□

**Remark 1.** *In the proof of Theorem 3 we have chosen  $R = (4k + 5)d$  which leads to the stated scaling of  $c$ . Note that choosing  $R$  to be any number larger than that also gives a valid covering scheme, with a larger  $c$ . This is useful for the discussions in Section 3.3.*

### 3 Detailed analysis

In this section we give a detailed analysis which leads to the following main result.

**Theorem 4** (Main result). *There is an algorithm that, given copies of an unknown state  $|\psi\rangle$ , with the promise that  $|\psi\rangle = U|0^n\rangle$  where  $U$  is an unknown depth- $d$  circuit acting on a  $k$ -dimensional lattice (using arbitrary 2-qubit gates), outputs a depth- $(2k+1)d$  circuit  $W$  that prepares  $|\psi\rangle$  up to  $\varepsilon$  trace distance, with success probability  $1 - \delta$ . The algorithm uses  $M$  copies of  $|\psi\rangle$  and runs in time  $T$ , where*

$$M = \frac{n^4 \cdot 2^{O(c)}}{\varepsilon^4} \log \frac{n}{\delta}, \quad T = \frac{n^4 \cdot 2^{O(c)}}{\varepsilon^4} \log \frac{n}{\delta} + \left( \frac{nk d \cdot c}{\varepsilon} \right)^{O(d \cdot c)}. \quad (6)$$

Here,  $c = O((3k)^{k+2}d)^k$ , and  $W$  uses  $r \cdot n$  ancilla qubits where  $r > 0$  can be chosen to be an arbitrarily small constant.

#### 3.1 The reconstruction process is robust

We first show that the final error is controlled, if we replace local inversions in the reconstruction process with approximate local inversions.

**Definition 8** ( $\varepsilon$ -approximate local inversion). *Let  $V$  be a unitary acting on  $AB$ , and denote the remaining system as  $C$  (Fig. 1b). Let  $|\psi\rangle$  be a state defined on  $ABC$ .  $V$  is an  $\varepsilon$ -approximate local inversion of the region  $A$  for  $|\psi\rangle$  if*

$$\langle 0|_A \text{Tr}_{BC} \left( V |\psi\rangle\langle\psi| V^\dagger \right) |0\rangle_A \geq 1 - \varepsilon. \quad (7)$$

**Lemma 3.** *Let  $V$  be an  $\varepsilon$ -approximate local inversion of the region  $A$  for  $|\psi\rangle$ . The corresponding  $A$ -replacement process is given by*

$$\mathcal{V}^\dagger \circ \mathcal{R}_A \circ \mathcal{V}, \quad (8)$$

where calligraphic letters denote channels, and  $\mathcal{R}_A$  is the reset channel acting on  $A$  which traces out the input and prepares  $|0\rangle\langle 0|_A$ . Then we have

$$\left\| \mathcal{V}^\dagger \circ \mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi| \right\|_1 \leq 4\sqrt{\varepsilon}. \quad (9)$$

*Proof.* First, note that due to the unitary invariance of the trace distance,

$$\begin{aligned} \left\| \mathcal{V}^\dagger \circ \mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi| \right\|_1 &= \left\| \mathcal{V}^\dagger \circ \mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - \mathcal{V}^\dagger \circ \mathcal{V}(|\psi\rangle\langle\psi|) \right\|_1 \\ &= \left\| \mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - \mathcal{V}(|\psi\rangle\langle\psi|) \right\|_1. \end{aligned} \quad (10)$$

Next, we can write  $V|\psi\rangle$  as

$$V|\psi\rangle = \sqrt{1 - \varepsilon'} |0\rangle_A |\phi\rangle_{BC} + \sqrt{\varepsilon'} |\text{else}\rangle_{ABC}, \quad (11)$$

where  $\varepsilon' \leq \varepsilon$  and  $\langle 0|_A |\text{else}\rangle_{ABC} = 0$ . Using the relationship between fidelity and trace distance,

$$\left\| \mathcal{V}(|\psi\rangle\langle\psi|) - |0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC} \right\|_1 = 2\sqrt{1 - |\langle\psi| V^\dagger |0\rangle_A |\phi\rangle_{BC}|^2} = 2\sqrt{\varepsilon'}. \quad (12)$$

Finally,

$$\begin{aligned}
& \|\mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - \mathcal{V}(|\psi\rangle\langle\psi|)\|_1 \\
& \leq \|\mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - |0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC}\|_1 + \||0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC} - \mathcal{V}(|\psi\rangle\langle\psi|)\|_1 \\
& = \|\mathcal{R}_A \circ \mathcal{V}(|\psi\rangle\langle\psi|) - \mathcal{R}_A(|0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC})\|_1 + \||0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC} - \mathcal{V}(|\psi\rangle\langle\psi|)\|_1 \\
& \leq \|\mathcal{V}(|\psi\rangle\langle\psi|) - |0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC}\|_1 + \||0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC} - \mathcal{V}(|\psi\rangle\langle\psi|)\|_1 \\
& \leq 4\sqrt{\varepsilon}.
\end{aligned} \tag{13}$$

Here, the second line is by triangle inequality; the third line is because  $|0\rangle\langle 0|_A \otimes |\phi\rangle\langle\phi|_{BC}$  is invariant under  $\mathcal{R}_A$ ; the fourth line is by data-processing inequality.  $\square$

**Lemma 4.** *Let  $\Phi_1, \Phi_2, \dots, \Phi_T$  be arbitrary replacement processes using  $\varepsilon$ -approximate local inversions of  $|\psi\rangle$ , where each  $\Phi_i$  has the form of Eq. (8) where  $V$  is an  $\varepsilon$ -approximate local inversion for some arbitrary region. Then*

$$\|\Phi_T \circ \dots \circ \Phi_2 \circ \Phi_1(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_1 \leq 4T\sqrt{\varepsilon}. \tag{14}$$

*Proof.* Note that

$$\begin{aligned}
& \|\Phi_T \circ \dots \circ \Phi_2 \circ \Phi_1(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_1 \\
& \leq \|\Phi_T \circ \dots \circ \Phi_2 \circ \Phi_1(|\psi\rangle\langle\psi|) - \Phi_T(|\psi\rangle\langle\psi|)\|_1 + \|\Phi_T(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_1 \\
& \leq \|\Phi_{T-1} \circ \dots \circ \Phi_2 \circ \Phi_1(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_1 + 4\sqrt{\varepsilon},
\end{aligned} \tag{15}$$

where the second line is by triangle inequality, the third line is by data-processing inequality and Lemma 3. The claim follows from induction.  $\square$

### 3.2 Analysis of the learning algorithm

Next we put everything together and give a detailed analysis of the learning algorithm. The algorithm is given copies of an unknown quantum state  $|\psi\rangle$ , with the promise that  $|\psi\rangle = U|0^n\rangle$  where  $U$  is a depth- $d$  circuit acting on a  $k$ -dimensional lattice.

**Step 1: build a covering scheme.** We build a  $(k+1, c, d)$  covering scheme (here  $c = O(k^2 d)^k$  or larger, depending on the choice of  $R$ , see Remark 1) for the  $k$ -dimensional lattice according to Theorem 3. This covering scheme is divided into  $\ell = k+1$  layers. Note that there are less than  $n$  subsets in total.

**Step 2: learn reduced density matrices.** To find local inversions for the covering scheme it suffices to learn reduced density matrices on a radius- $d$  ball around each subset (dashed regions in Figs. 3b to 3d). Each reduced density matrix has size which has the same scaling as  $c$  and there are less than  $n$  of them.

Suppose we would like to learn all of them within  $\varepsilon_1$  trace distance, with  $\delta$  failure probability. Here we use a simple existing result (see e.g. [HLB+24, Lemma 23]): for each copy of  $|\psi\rangle$ , we measure each qubit in a random Pauli basis. The collection of measurement outcomes is known as classical shadows or randomized measurement dataset [HKP20; EFH+23]. The desired result can be achieved by processing this dataset, which uses

$$M = \frac{2^{O(c)}}{\varepsilon_1^2} \log \frac{n}{\delta} \tag{16}$$

copies of  $|\psi\rangle$ . The running time to process this dataset scales similarly. From now on everything is classical processing on the learned reduced density matrices, and we assume that all learned reduced density matrices are within  $\varepsilon_1$  error (this happens with probability at least  $1 - \delta$ ).

**Step 3: find approximate local inversions.** For each reduced density matrix, suppose it looks like the region  $AB$  in Fig. 1b. We will brute force search over an  $\varepsilon_1$ -net on depth- $d$  circuits acting on  $AB$  to find a depth- $d$  circuit that approximately inverts the region  $A$ . In this way we can find a  $3\varepsilon_1$ -approximate local inversion of  $A$  for  $|\psi\rangle$  (see e.g. [HLB+24, Proof of first claim of Theorem 9] for a proof). The running time scales as the size of the  $\varepsilon_1$ -net, which equals

$$\left(\frac{kd \cdot c}{\varepsilon_1}\right)^{O(d \cdot c)}. \quad (17)$$

See e.g. [HLB+24, Lemma 19] for a proof.

Note that this can be significantly improved if we assume a discrete gate set, which we do not discuss here.

**Step 4: find approximate reconstruction circuit.** We have found a  $3\varepsilon_1$ -approximate local inversion for every subset in the covering scheme. Construct a reconstruction process (Definition 5) using these approximate local inversions, denote as  $\Phi$ . By Lemma 4, we have

$$\|\Phi(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_1 \leq 4n\sqrt{3\varepsilon_1}. \quad (18)$$

By Theorem 2, the covering scheme guarantees that the backward lightcone  $C$  of the output of  $\Phi$  is contained entirely within  $\Phi$  (in fact, the output state of  $\Phi$  is independent of its input state).  $C$  is a depth- $(2k+1)d$  circuit. Our final learned state  $\hat{\rho}$  is thus equal to  $\Phi(|\psi\rangle\langle\psi|)$ , which can be prepared by running  $C$  on all- $|0\rangle$  input and trace out those qubits that do not belong to the output of  $\Phi$ . To achieve  $\varepsilon$ -closeness in trace distance, it suffices to choose  $\varepsilon_1 = \frac{\varepsilon^2}{48n^2}$ .

This concludes the proof of the complexity statements in the main result.

### 3.3 Optimizing the number of ancilla qubits

Here we explicitly calculate (a rough upper bound of) the number of ancilla qubits needed for reconstructing  $|\psi\rangle$ , which corresponds to the red qubits in Fig. 2c. In Figs. 3b to 3d, the ancilla qubits live in the colored regions outside solid black boxes.

We start from a lattice coloring (e.g. Fig. 3a). Note that each small colored region is contained in a  $k$ -dimensional box of length scale  $2kR$ ; meanwhile each small colored region at least contains a  $k$ -dimensional box of length scale  $R$  (see e.g. [HLB+24, below Definition 17]). The number of colored regions is at most  $n/R^k$ .

Next we upper bound the number of ancilla qubits that can be associated with a colored region. The length scale of a subset in a covering scheme is at most  $2kR + 2 \times (2k+1)d \leq 3kR$ , since we will choose  $R$  to be at least  $(4k+5)d$ . The ancilla qubits live at  $2k$  different  $k-1$  dimensional surfaces with thickness  $(2k+1)d \leq 3kd$ ; the total volume is at most  $(2k) \times (3kR)^{k-1} \times (3kd)$ . Overall, the total number of ancilla qubits in the entire circuit is at most

$$\frac{n}{R^k} \times (2k) \times (3kR)^{k-1} \times (3kd) \leq \frac{n}{R} \times (3k)^{k+1}d. \quad (19)$$

Choosing  $R = L \cdot (3k)^{k+1}d$  for some sufficiently large constant  $L$ , the total number of ancilla qubits is an arbitrarily small constant times  $n$ . In fact we could also afford to choose  $L = \omega(1)$  to be some

small function (e.g.  $\log \log \log n$ ) so that the total number of ancilla qubits is sublinear. Below we just consider  $L$  as being a constant.

The complexity scales with  $c$  as shown in [Section 3.2](#), which is given by

$$c \leq (3kR)^k = O((3k)^{k+2}d)^k. \quad (20)$$

## 4 Testing quantum circuit complexity

The following result is a stronger version than stated in [Corollary 1](#).

**Theorem 5.** *Given copies of an unknown state  $|\psi\rangle$  on a  $k$ -dimensional lattice, with the promise that one of the following two cases hold:*

- **Case 1: Low complexity.**  $|\psi\rangle = U|0^n\rangle$  where the depth of  $U$  is at most  $d$ ;
- **Case 2: High complexity.** Let  $\rho$  be any  $n$ -qubit state prepared by a depth at most  $(2k+1)d$  circuit using  $r \cdot n$  ancilla qubits, where  $r$  is some small constant. Then  $\|\rho - |\psi\rangle\langle\psi|\|_1 > \varepsilon$ .

*There is an algorithm that decides which is the case, with success probability at least  $1 - \delta$ , where the sample complexity and running time is the same as in [Theorem 4](#).*

*Proof.* We perform Step 1 and 2 as prescribed by [Section 3.2](#) (note that  $\varepsilon_1 = \frac{\varepsilon^2}{48n^2}$ ). Then, we declare “Case 1: Low complexity” if the search procedure to find all desired approximate local inversions in Step 3 all succeed. Otherwise, declare “Case 2: High complexity”.

We assume that the reduced density matrices in Step 2 are all within  $\varepsilon_1$  error, which happens with probability at least  $1 - \delta$ . Conditioned on this event, the algorithm is always correct, because of the following reasoning.

- In Case 1, all search procedures to find the desired approximate local inversions are guaranteed to succeed, so the algorithm must output “Case 1”.
- In Case 2, suppose by contradiction that all search procedures succeed. Then, by the arguments of the main result, this actually gives a *proof* that  $|\psi\rangle$  can be prepared within  $\varepsilon$  error, using a quantum circuit of depth  $(2k+1)d$  and a small amount of ancilla qubits, which contradicts the assumption of Case 2. Therefore, some of the search procedures must fail, and the algorithm must output “Case 2”.

□

## Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. Additional support is acknowledged from DOE Grant No. DE-SC0024124 and NSF Grant No. 2311733.

## References

- [AA24] Anurag Anshu and Srinivasan Arunachalam. “A survey on the complexity of learning quantum states”. In: *Nature Reviews Physics* 6.1 (Jan. 2024), pp. 59–69. DOI: [10.1038/s42254-023-00662-4](https://doi.org/10.1038/s42254-023-00662-4).



- [AK22] Eric R Anschuetz and Bobak T Kiani. “Quantum variational algorithms are swamped with traps”. In: *Nature Communications* 13.1 (2022), p. 7760.
- [AT18] Dorit Aharonov and Yonathan Touati. *Quantum Circuit Depth Lower Bounds For Homological Codes*. 2018. arXiv: [1810.03912](https://arxiv.org/abs/1810.03912) [quant-ph].
- [BCK+22] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, et al. “Noisy intermediate-scale quantum algorithms”. In: *Rev. Mod. Phys.* 94 (1 Feb. 2022), p. 015004. DOI: [10.1103/RevModPhys.94.015004](https://doi.org/10.1103/RevModPhys.94.015004).
- [BHS+18] Juan Bermejo-Vega, Dominik Hangleiter, Martin Schwarz, Robert Raussendorf, and Jens Eisert. “Architectures for Quantum Simulation Showing a Quantum Speedup”. In: *Phys. Rev. X* 8 (2 Apr. 2018), p. 021010. DOI: [10.1103/PhysRevX.8.021010](https://doi.org/10.1103/PhysRevX.8.021010).
- [BHV06] S. Bravyi, M. B. Hastings, and F. Verstraete. “Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order”. In: *Phys. Rev. Lett.* 97 (5 July 2006), p. 050401. DOI: [10.1103/PhysRevLett.97.050401](https://doi.org/10.1103/PhysRevLett.97.050401).
- [BK19] Fernando G. S. L. Brandão and Michael J. Kastoryano. “Finite Correlation Length Implies Efficient Preparation of Quantum Thermal States”. In: *Communications in Mathematical Physics* 365.1 (Jan. 2019), pp. 1–16. DOI: [10.1007/s00220-018-3150-8](https://doi.org/10.1007/s00220-018-3150-8).
- [CGW10] Xie Chen, Zheng-Cheng Gu, and Xiao-Gang Wen. “Local unitary transformation, long-range quantum entanglement, wave function renormalization, and topological order”. In: *Phys. Rev. B* 82 (15 Oct. 2010), p. 155138. DOI: [10.1103/PhysRevB.82.155138](https://doi.org/10.1103/PhysRevB.82.155138).
- [CIKK16] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. “Learning Algorithms from Natural Proofs”. In: *31st Conference on Computational Complexity (CCC 2016)*. Ed. by Ran Raz. Vol. 50. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016, 10:1–10:24. DOI: [10.4230/LIPIcs.CCC.2016.10](https://doi.org/10.4230/LIPIcs.CCC.2016.10).
- [EFH+23] Andreas Elben, Steven T. Flammia, Hsin-Yuan Huang, Richard Kueng, John Preskill, Benoît Vermersch, and Peter Zoller. “The randomized measurement toolbox”. In: *Nature Reviews Physics* 5.1 (Jan. 2023), pp. 9–24. DOI: [10.1038/s42254-022-00535-2](https://doi.org/10.1038/s42254-022-00535-2).
- [GWD17] Xun Gao, Sheng-Tao Wang, and L.-M. Duan. “Quantum Supremacy for Simulating a Translation-Invariant Ising Spin Model”. In: *Phys. Rev. Lett.* 118 (4 Jan. 2017), p. 040502. DOI: [10.1103/PhysRevLett.118.040502](https://doi.org/10.1103/PhysRevLett.118.040502).
- [HE23] Dominik Hangleiter and Jens Eisert. “Computational advantage of quantum random sampling”. In: *Rev. Mod. Phys.* 95 (3 July 2023), p. 035001. DOI: [10.1103/RevModPhys.95.035001](https://doi.org/10.1103/RevModPhys.95.035001).
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Predicting many properties of a quantum system from very few measurements”. In: *Nature Physics* 16.10 (Oct. 2020), pp. 1050–1057. DOI: [10.1038/s41567-020-0932-7](https://doi.org/10.1038/s41567-020-0932-7).
- [HLB+24] Hsin-Yuan Huang, Yunchao Liu, Michael Broughton, Isaac Kim, Anurag Anshu, Zeph Landau, and Jarrod R. McClean. *Learning shallow quantum circuits*. 2024. arXiv: [2401.10095](https://arxiv.org/abs/2401.10095) [quant-ph].
- [KKR24] Hyun-Soo Kim, Isaac Kim, and Daniel Ranard. *Private communication*. 2024.
- [LMN93] Nathan Linial, Yishay Mansour, and Noam Nisan. “Constant depth circuits, Fourier transform, and learnability”. In: *Journal of the ACM (JACM)* 40.3 (1993), pp. 607–620.

- [MOS03] Elchanan Mossel, Ryan O’Donnell, and Rocco P. Servedio. “Learning Juntas”. In: *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*. STOC ’03. San Diego, CA, USA: Association for Computing Machinery, 2003, pp. 206–212. DOI: [10.1145/780542.780574](https://doi.org/10.1145/780542.780574).
- [PSC21] Lorenzo Piroli, Georgios Styliaris, and J. Ignacio Cirac. “Quantum Circuits Assisted by Local Operations and Classical Communication: Transformations and Phases of Matter”. In: *Phys. Rev. Lett.* 127 (22 Nov. 2021), p. 220503. DOI: [10.1103/PhysRevLett.127.220503](https://doi.org/10.1103/PhysRevLett.127.220503).
- [RF21] Cambyse Rouzé and Daniel Stilck França. *Learning quantum many-body systems from a few copies*. 2021. arXiv: [2107.03333](https://arxiv.org/abs/2107.03333) [[quant-ph](#)].
- [TD04] Barbara M. Terhal and David P. DiVincenzo. *Adaptive Quantum Computation, Constant Depth Quantum Circuits and Arthur-Merlin Games*. 2004. arXiv: [quant-ph/0205133](https://arxiv.org/abs/quant-ph/0205133) [[quant-ph](#)].
- [YW23] Nengkun Yu and Tzu-Chieh Wei. *Learning marginals suffices!* 2023. arXiv: [2303.08938](https://arxiv.org/abs/2303.08938) [[quant-ph](#)].