Efficient Matrix Product State Learning in Logarithmic Depth

Chia-Ying Lin[*] Nai-Hui Chia[†] Shih-Han Hung[‡]

Abstract

Learning the closest matrix product state (MPS) representation of a quantum state is known to enable useful tools for quantum machine learning and analysis of complex quantum systems. In this work, we study the problem of learning MPS in the following setting: given many copies of an input MPS, the task is to recover a classical description of the state. The best known polynomial-time algorithm, introduced by [LCLP10, CPF+10], requires linear circuit depth and $\widetilde{O}(n^5)$ samples, and has seen no improvement in over a decade. The combination of linear circuit depth and large sample complexity, neither known to be optimal, renders existing algorithms impractical for near-term quantum devices with limited resources.

We introduce parallel disentangling algorithms for MPS learning. For exact MPS learning, our algorithm runs in polynomial time and outputs a preparation circuit of depth $O(\log n)$, using $\widetilde{O}(n^3)$ copies, improving both the depth and the dependence on the system size n. The key idea is to exploit the bounded-rank structure of reduced states on middle blocks of an MPS and organize the disentangling operations in a tree structure.

We further extend the algorithm to closest MPS learning, improving the sample complexity dependence on n from n^9 to n^7 and complement the algorithms with an $\Omega(n)$ product-state lower bound. We also investigate MPS learning under hardware constraints, including restricted measurements and geometric connectivity. Under the Learning Parity with Noise (LPN) assumption, we show computational hardness for learning an MPS(2) family with non-adaptive single-qubit measurements. Finally, we show that our algorithm can be implemented with depth $O(qn^{1/q})$ on a q-dimensional hypercubic lattice, giving an asymptotic reduction in depth. Together, our work provides a complete characterization of the quantum resources needed for efficient MPS learning.

1 Introduction

Matrix product state is a fundamental family of quantum states defined as follows:

Definition 1 (Matrix Product State with Bond Dimension D (MPS(D)))**.** *Let $n \in \mathbb{N}$ and let $d \in \mathbb{N}$ be the local dimension of each site (for qubits, $d = 2$). An n-site* matrix product state (MPS)[4] *of bond dimension D is a pure state* $|\psi\rangle = \sum_{i_1=0}^{d-1} \sum_{i_2=0}^{d-1} \cdots \sum_{i_n=0}^{d-1} \mathrm{Tr}\left[A_{i_1}^{[1]} A_{i_2}^{[2]} \cdots A_{i_n}^{[n]}\right] |i_1 i_2 \ldots i_n\rangle$ *, where for each site $k \in \{1, \ldots, n\}$ and each physical index $i_k \in \{0, \ldots, d-1\}$, $A_{i_k}^{[k]}$ is a complex matrix of size $D \times D$. The integer D is called the* bond dimension *of the MPS.*

[*]Department of Physics and Astronomy, Rice University, USA. Email: cl207@rice.edu.

[†]Ken Kennedy Institute and Smalley-Curl Institute, Rice University, USA. Email: nc67@rice.edu.

[‡]Department of Electrical Engineering and Center for Quantum Science and Engineering, National Taiwan University, Taiwan. Email: shihhanh@ntu.edu.tw.

[4]We adopt the periodic boundary condition (PBC) convention (also known as the tensor-ring) throughout, in which all local tensors are square matrices and the physical amplitudes are obtained by taking a trace over the virtual indices. Open-boundary MPS (also known as the tensor-train) can be embedded into this framework by a standard padding construction.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Matrix Product States (MPS) exhibit several noteworthy properties. First, when the bond dimension D is small, an n-site qudit MPS requires only $O(dnD^2)$ parameters, in contrast to the $O(d^n)$ parameters needed for a general quantum state. This compressibility in description makes MPS a natural target for efficient learning, testing, and synthesis algorithms. Moreover, for an open-boundary MPS, the Schmidt rank of $|\psi\rangle$ across a bipartition induced by cutting a single virtual bond is at most D. This means the bond dimension directly characterizes the degree of entanglement: for instance, $D = 1$ corresponds to a product state with no entanglement, while if the bond dimension equals the full Hilbert space dimension, the MPS representation can express all possible quantum states.

Matrix product states (MPS) play an increasingly important role not only in quantum many-body physics, but also in machine learning and quantum computing. In machine learning, MPS and closely related tensor-train models have been used as structured, parameter-efficient representations of high-dimensional data, with applications to neural-network compression, supervised learning, high-order feature modeling, and generative modeling [NPOV15, SS16, NTO16, HWF$^+$18, SAOB22]. They have also appeared in quantum machine learning as tensor-network architectures that can be mapped to variational quantum circuits [RKR23]. From a quantum computing perspective, MPS provide a natural class of variational ansatz states for quantum algorithms, including ground-state energy estimation and combinatorial optimization [CPGSV21, WRV$^+$24, KNR$^+$25]. Their importance is also deeply rooted in quantum many-body physics. In particular, ground states of gapped one-dimensional local Hamiltonians are known to satisfy an entanglement area law and are therefore provably well approximated by matrix product states with bounded bond dimension [Has07, ALV12, AKLV13, LVV15, SW22]. This connection bridges physical Hamiltonian systems with computationally tractable descriptions, enabling efficient numerical and algorithmic methods for representing, preparing, and optimizing low-energy quantum states [Whi92, DKSV04, Vid04, VPC04, OV08]. Beyond these applications, recent developments suggest that MPS and related tensor-network structures also intersect with quantum cryptography and quantum complexity theory [ABF$^+$22]. Thus, understanding the sample complexity, circuit depth, and computational complexity of learning MPS is a timely and fundamental problem, with implications across machine learning, quantum algorithms, many-body physics, and quantum information theory.

In this work, we focus on learning the classical description of an input quantum state which is promised to be an MPS with a specified bond dimension. This problem is known as *MPS learning*. From the perspective of learning theory, MPS learning is a quantum analogue of learning a structured, low-dimensional representation of a high-dimensional object. The unknown state lives in a Hilbert space of dimension exponential in n, but the target representation has only $O(ndD^2)$ parameters. From a computational-complexity perspective, MPS provides a succinct representation for many quantum states with low entanglement; the relatively low parameter count, $O(dnD^2)$ for bond dimension D, admits the possibility of polynomial-time learning algorithms, in contrast to the exponential complexity of general quantum state learning. *MPS learning* is also of significant practical and theoretical interest: for example, efficient algorithms or provable hardness results can be leveraged to certify whether the quantum hardware prepares correct quantum states, to construct cryptographic primitives, and to analyze properties of ground states and Hamiltonians.

Foundational algorithms for efficient MPS learning were introduced in [CPF$^+$10, LCLP10]. However, these algorithms are improper: the output MPS may have a bond dimension strictly larger than that of the input state. This distinction becomes important in property testing, where the goal is to determine whether an unknown state belongs to a prescribed MPS class, rather than merely reconstruct a nearby state. In particular, improper learning does not directly yield a testing algorithm for a fixed bond dimension. The first efficient MPS testing algorithm was subsequently introduced in [SW22]. Beyond pure-state MPS learning, another line of work has investigated tomography for more general tensor-network representations, particularly matrix product operators (MPOs), which can describe mixed states. Along this direction, [QJG$^+$24] showed that an n-qudit MPO of bond dimension r can be stably recovered to Frobenius-norm error ϵ using $\widetilde{O}(n^3d^2r^2/\epsilon^2)$ samples from Haar-random projective measurements. Although this achieves polynomial sample complexity, such measurements are challenging to implement efficiently on quantum hardware. Another fundamental challenge arises when the unknown state does not belong exactly to the assumed model class, as is often the case in realistic experiments due to hardware imperfections or environmental noise. This motivates agnostic learning, which seeks an approximation whose fidelity is competitive with the best state in the target class. In this direction, [BBK$^+$25] developed an improper agnostic learning algorithm for MPS: the learned state achieves fidelity close to that of the best bond-D MPS, but its output bond dimension

Problem	Prior work	This work
Exact MPS learning	$\widetilde{O}(n^5D^2/\epsilon^4)$ samples, depth $O(n)$ [CPF$^+$10, LCLP10]	$\widetilde{O}(n^3D^6/\epsilon^4)$ samples, depth $O(\log n)$ (Theorem 1)
Improper closest-MPS learning	$\widetilde{O}(n^9D^8/\epsilon^8)$ samples [BBK$^+$25]	$\widetilde{O}(n^7D^{12}/\epsilon^{12})$ samples (Theorem 2)
2-dimensional lattice	depth $O(n)$ [CPF$^+$10, LCLP10]	depth $O(\sqrt{n})$ (Theorem 3)
q-dimensional hypercubic lattice	depth $O(n)$ [CPF$^+$10, LCLP10]	depth $O(qn^{1/q})$ (Corollary 3.1)
Product-state lower bound	–	$\Omega(nd/\epsilon)$ copies (Theorem 4)
Restricted measurements	–	LPN-hard for MPS(2) (Theorem 5 and Corollary 5.1)

Table 1: Summary of the main results and comparisons to previous works. Here, n is the number of qudits, d is the local dimension, D is the bond dimension, and ϵ is the target error. The depth entries refer to output preparation circuits for fixed d and D.

may be larger than D. This raises the further question of whether such guarantees can be achieved while keeping the learned state within the prescribed bond-dimension class. Recently, Cedillo Vayson de Pradenne and Cotler [CVdPC26] established proper agnostic learning of open-boundary MPS, where the output has bond dimension at most D and competes with the best bond-D MPS for arbitrary mixed inputs. Their copy complexity is jointly polynomial in $n, d, D, 1/\epsilon, \log(1/\delta)$, but their runtime is polynomial in n only for fixed d, D, ϵ [5]. For periodic-boundary MPS, their result is improper, returning an open-boundary MPS of bond dimension at most D^2 when competing with bond-D periodic-boundary MPS.

Despite the fundamental importance and broad applicability of MPS learning, the resource requirements of polynomial-time MPS learning remain poorly understood. The best-known algorithm uses $\widetilde{O}(n^5)$ copies of samples [LCLP10][6], while the strongest known lower bound, derived from product-state learning lower bounds, stands at only $\Omega(n)$[7]. This leaves a substantial and largely unexplored gap between the achievable upper and lower bounds. In addition, the sequential algorithms above were formulated primarily for matrix product states with open boundary conditions (OBC). Their linear circuit depth poses a severe obstacle for near-term quantum devices, where limited gate fidelity renders deep circuits prohibitively noisy. Bridging these gaps and designing shallow-depth, noise-resilient algorithms for MPS learning is therefore an open and pressing problem, carrying significant implications for both quantum algorithm design and practical deployment on emerging quantum hardware.

Along this line, we are driven by the following question:

Can we design new MPS learning algorithms whose sample complexity substantially narrows the gap to the known lower bounds, while achieving shallow circuit depth?

Our results. We answer the above questions affirmatively. We introduce a parallel disentangling algorithm for exact MPS learning, reducing the circuit depth from $O(n)$ to $O(\log n)$ and the sample complexity from $\widetilde{O}(n^5)$ to $\widetilde{O}(n^3)$. We also extend the algorithm to closest MPS learning, reducing the sample complexity from $\widetilde{O}(n^9)$ to $\widetilde{O}(n^7)$. Our algorithms improve the dependence on the system size, at the cost of a worse dependence on the bond dimension D; see Table 1. This tradeoff is meaningful in many applications where MPS and tensor-network representations are useful for succinctly representing high-dimensional data because the bond dimension is small. We complement these algorithms with an $\Omega(n)$ sample-complexity lower bound.

Finally, we study MPS learning under hardware constraints. We show that our algorithm has depth $O(qn^{1/q})$ on a q-dimensional hypercubic lattice, asymptotically improving previous approaches. We also prove, assuming LPN, that no efficient algorithm can learn the MPS(2) family using only non-adaptive single-qubit measurements. Together, these results characterize the quantum resources needed for efficient MPS learning. Table 1 summarizes our results and comparisons with prior work.

2 Logarithmic-depth algorithms for MPS learning

Definition 2 (Quantum State Learning Problem)**.** *Let $\mathscr{C}$ be a class of quantum states on n qubits. The learner has access to independent copies of an unknown state $\rho \in \mathscr{C}$. The goal is to output, with*

[5] Their Theorem 2.1 gives runtime $\mathrm{poly}(n, d, D, 1/\epsilon, \log(1/\delta))(Cn\sqrt{D}/\epsilon)^{O(D^2/\epsilon^2+dD^2)}$.

[6] The original work does not provide a detailed complexity analysis; we follow analysis in [BBK$^+$25] to get the complexity.

[7] See Theorem 4 and Section E for a formal proof.

high probability, a classical description of a hypothesis state $\hat{\rho}$ such that $\mathrm{dist}(\rho, \hat{\rho}) \leq \epsilon$*, where* $\epsilon > 0$ *specifies the allowed error in the chosen distance measure* $\mathrm{dist}(\cdot, \cdot)$*. Typical examples include the trace distance, where the goal is* $\|\rho - \hat{\rho}\|_1 \leq \epsilon$*, and the fidelity, where the goal is* $F(\rho, \hat{\rho}) \geq 1 - \epsilon$*.* [8]

We focus on the class $\mathscr{C} = \mathrm{MPS}(D)$, the set of n-qudit matrix product states with bond dimension at most D. We measure the quality of the output by fidelity. In particular, when the algorithm outputs a pure hypothesis state $|\hat{\phi}\rangle$, the fidelity guarantee takes the form $\langle\hat{\phi}| \rho |\hat{\phi}\rangle \geq 1 - \epsilon$.

Theorem 1. *For any integers* $n, d \geq 2$*, given access to copies of an* n*-qudit matrix product state* ρ *with bond dimension* D *and parameters* $\epsilon, \delta \in (0, 1)$*, Algorithm 2 outputs a classical description of a preparation circuit for a pure state* $|\hat{\phi}\rangle$*, such that, with probability* $1 - \delta$*,* $\langle\hat{\phi}| \rho |\hat{\phi}\rangle \geq 1 - \epsilon$*. The algorithm requires* $N = O(\frac{D^6 \cdot d^4 \cdot n^3 \cdot \log(n/\delta)}{(\max\{1, \log_d D\})^3 \epsilon^4})$ *copies of* ρ *and runs in time* $\mathrm{poly}(D, n, 1/\epsilon, \log(1/\delta))$ *for fixed local dimension* d*. The output circuit uses* $O(\log n)$ *layers of block unitaries.*

To prove Theorem 1, we introduce the parallel disentangling MPS learning algorithm (Algorithm 2). We provide a high-level proof sketch below and defer the complete proof to Appendix B.

The previous algorithms [LCLP10, BBK$^+$25] exploit the fact that, for an open-boundary $\mathrm{MPS}(D)$, any single cut (say between sites l and $l + 1$) induces reduced density matrices of rank at most D on the two resulting segments $\{1, \ldots, l\}$ and $\{l + 1, \ldots, n\}$. In contrast, if we isolate a middle block by making two cuts, then the reduced density matrix of that block has rank at most D^2. This allows us to go beyond disentangling only one leftmost qudit at a time. Indeed, as long as the block size k satisfies $d^{k-1} \geq D^2$, we can apply disjoint disentangling unitaries to multiple blocks in parallel, thereby disentangling the leftmost qudit of each block simultaneously.

Our parallel disentangling algorithm starts with the following two observations: (1) Any contiguous interval and its complement have Schmidt rank at most D^2 (2) Organizing the disentangling procedure into a binary tree reduces the number of sequential layers of block unitaries from $O(n)$ to $M = O(\log n)$.

To justify the rank-D^2 constraint across all layers, consider a node (j, i) of the binary tree and let $\mathscr{I}_i^j$ be the union of the first-layer blocks in its subtree. Although the surviving block $\mathscr{B}_i^j$ need not be contiguous in the original chain, $\mathscr{I}_i^j$ is a contiguous interval. By the tree-consistency condition, every successful-branch operation in the first $j - 1$ layers (including the learned unitaries and the contraction of the registers designated for removal with their all-zero bras) acts either entirely inside $\mathscr{I}_i^j$ or entirely outside it. Hence their cumulative map factorizes across the bipartition $\mathscr{I}_i^j : [n] \setminus \mathscr{I}_i^j$, mapping each side to its surviving registers. Since the original bond-D MPS has Schmidt rank at most D^2 across this bipartition and local linear maps cannot increase Schmidt rank, the resulting subnormalized vector on the surviving registers has Schmidt rank at most D^2 across $\mathscr{B}_i^j : \mathscr{B}^j \setminus \mathscr{B}_i^j$. Consequently, $\mathrm{rank}(\sigma_i^j) \leq D^2$. This holds for any realized learned unitaries, whether exact or approximate. See Lemma B.1 for a formal proof.

Due to the binary tree structure, the number of layers of block unitaries is reduced from linear in n to $M = O(\log n)$. For constant bond dimension D, each block unitary acts on a constant number of qudits, so the output circuit also has $O(\log n)$ depth in elementary one- and two-qudit gates under all-to-all connectivity. This makes the circuit significantly shallower, which is especially important for near-term quantum devices where deep circuits are highly susceptible to noise. The tree structure also improves the error accumulation. At layer j, the errors from the 2^{M-j} disjoint blocks contribute $2\sqrt{2^{M-j} \cdot 2\eta}$ to the error passed to the next layer. Therefore, the total accumulated fidelity error is bounded by $\left|\langle\phi| \rho |\phi\rangle - \langle\phi| \rho^M |\phi\rangle\right| \leq 2\sum_{j=1}^{M} \sqrt{2^{M-j} \cdot 2\eta} = O\Big(\sqrt{2^M \cdot 2\eta}\Big) = O\big(\sqrt{n\eta}\big)$. Thus, to achieve final error at most ϵ, it suffices to choose the tomography precision $\eta = O(\epsilon^2/n)$. Using optimal tomography algorithms, the sample complexity for each block is then $O(n^2/\epsilon^4)$. Since the algorithm performs tomography on $O(n)$ blocks in total, the overall sample complexity becomes $O(n^3/\epsilon^4)$.

Algorithm 1 Disentangling unitary construction

[8] Common accuracy metrics in quantum state learning include trace distance and fidelity. Fidelity $F(\rho, \sigma) = \left(\mathrm{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right)^2$ measures state overlap, with larger values indicating greater similarity.

Input: An integer D^2 and a description of a positive semidefinite operator $\hat{\sigma} \in \mathbb{C}^{d^y \times d^y}$ with $\operatorname{Tr}\hat{\sigma} \leq 1$ and $y \geq p := \max\{1, 2\lceil \log_d D \rceil\}$.
Output: A description of a Disentangling Unitary U for the state $\hat{\sigma}$.
Procedure:

1: Choose a complete orthonormal eigenbasis of $\hat{\sigma}$:

$$\hat{\sigma} = \sum_{i=1}^{d^y} a_i \left|\phi_i\right\rangle \left\langle\phi_i\right|, \qquad a_1 \geq \cdots \geq a_{d^y} \geq 0.$$

Include eigenvectors with eigenvalue zero, with arbitrary orthonormal choices within degenerate eigenspaces.
2: Select $\left|\phi_1\right\rangle, \ldots, \left|\phi_{D^2}\right\rangle$ as the leading D^2 eigenvectors; retain the rest of this basis for the unitary extension.
3: Set $p := \max\{1, 2\lceil \log_d D \rceil\}$.
4: Let $\{|a_1, \ldots, a_{y-p}\rangle\}_{a_i \in \{0,\ldots,d-1\}}$ be the standard basis for the first $y-p$ qudits, and let $\{|j\rangle\}_{j=1}^{d^p}$ be an arbitrary orthonormal basis for the remaining p qudits.
5: Define the unitary $U = \sum_{a_1,\ldots,a_{y-p}=0}^{d-1} \sum_{j=1}^{d^p} \left(|a_1, \ldots, a_{y-p}\rangle \otimes |j\rangle\right) \left\langle \phi_{\mathrm{idx}(a_1,\ldots,a_{y-p},j)} \right|$, where the index mapping is defined by $\mathrm{idx}(a_1, \ldots, a_{y-p}, j) = j + d^p \cdot \sum_{l=1}^{y-p} a_l d^{y-p-l}$.
6: **return** the unitary U.

Algorithm 1 constructs a disentangling unitary from a classical estimate of a block operator. Algorithm 2 uses this construction on disjoint blocks in parallel, layer by layer. The required block estimates are obtained by ESTIMATE; its full pseudocode and that of REPLAYSAMPLE are given in Appendix A.3.

Each ESTIMATE call uses fresh copies of ρ to replay only the previously learned layers. At each such layer, it applies the learned unitary and measures the registers designated for removal in the computational basis. A copy is retained only if every such measurement yields zero; otherwise it is discarded. We call this successful sequence the *all-zero branch*.

For fixed learned unitaries, let $(\rho^j)'$ denote the subnormalized residual operator on the surviving registers after the all-zero branch through layer j, with $(\rho^0)' = \rho$. Its trace $q_j = \operatorname{Tr}[(\rho^j)']$ is the probability of success through layer j. For each block $\mathscr{B}_i^j$ defined in Algorithm 2, the estimation target is

$$\sigma_i^j = \operatorname{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[(\rho^{j-1})'].$$

The successful outputs on this block have normalized state $\bar{\sigma}_i^j = \sigma_i^j / q_{j-1}$ when $q_{j-1} > 0$. ESTIMATE estimates this state by tomography and multiplies the estimate by the estimated success probability, as in Lemma A.2. All sample budgets count original copies, including failed trials.

Algorithm 2 Parallel disentangling MPS learning

Input: The parameters $n, d, D, \epsilon, \delta$ and a copy source Copy_ρ supplying fresh copies of an unknown n-qudit MPS ρ of bond dimension at most D, where $\epsilon, \delta \in (0, 1)$.
Output: A classical description of a preparation circuit for a pure hypothesis state $|\hat{\phi}\rangle$.
Procedure:

1: Set $p = \max\{1, 2\lceil \log_d D \rceil\}$, $M = \max\{0, \lceil \log_2(n/p) \rceil\}$, $\eta = (\sqrt{2}-1)^2\epsilon^2/2^{M+5}$, $\tau = \epsilon/4$, and $\zeta = \delta/(2^M + 1)$.
2: Initialize the classical layer-record list $\mathscr{H} = ()$.
3: **if** $M \geq 1$ **then**
4: Set $b = 2^{M-1}$, $r_1 = n - bp$, $\ell_1 = \lceil r_1/p \rceil$, $s_1 = r_1 - (\ell_1 - 1)p$, and $k_1 = (2\ell_1 - 1)p + s_1$.
5: Define the ordered first-layer blocks:

$$\mathscr{B}_i^1 = \begin{cases} \{2(i-1)p+1, \ldots, 2ip\}, & 1 \leq i < \ell_1, \\ \{2(\ell_1-1)p+1, \ldots, k_1\}, & i = \ell_1, \\ \{k_1 + (i-\ell_1-1)p + 1, \ldots, k_1 + (i-\ell_1)p\}, & \ell_1 < i \leq b. \end{cases} \tag{1}$$

6: Define $f(j, i) = s_1$ if $j = 1$ and $i = \ell_1$; $f(j, i) = 0$ if $j = 1$ and $i > \ell_1$; and $f(j, i) = p$ otherwise.

for $j = 1, \ldots, M$ **do**
  if $j = 1$ **then**
    Set $m_j = \ell_1$ and $\mathscr{B}^1 = [n]$. For $i > \ell_1$, set $\widetilde{\mathscr{B}}_i^1 = \mathscr{B}_i^1$.
  else
    Set $m_j = 2^{M-j}$ and $\mathscr{B}^j = \bigcup_{i=1}^{2^{M-j+1}} \widetilde{\mathscr{B}}_i^{j-1}$.
    For $i = 1, \ldots, m_j$, set $\mathscr{B}_i^j = \widetilde{\mathscr{B}}_{2i-1}^{j-1} \cup \widetilde{\mathscr{B}}_{2i}^{j-1}$.
  end if
  for $i = 1, \ldots, m_j$ **do**
    Obtain a classical matrix estimate $\hat{\sigma}_i^j \leftarrow \textsc{Estimate}(\mathsf{Copy}_\rho, \mathscr{H}, \mathscr{B}_i^j, d, D^2, \eta, \zeta)$.
    Use Algorithm 1 with inputs D^2 and $\hat{\sigma}_i^j$ to construct U_i^j.
    Let $\widetilde{\mathscr{B}}_i^j$ be the last p registers of $\mathscr{B}_i^j$, in their inherited order.
  end for
  Set $U^j = \bigotimes_{i=1}^{m_j} U_i^j$ on the processed blocks, with identity action elsewhere in $\mathscr{B}^j$.
  Set $\mathscr{D}^j = \bigcup_{i=1}^{m_j} (\mathscr{B}_i^j \setminus \widetilde{\mathscr{B}}_i^j)$.
  Append the layer record $(\mathscr{B}^j, U^j, \mathscr{D}^j)$ to $\mathscr{H}$.
end for
Set $A_{\mathrm{fin}} = \widetilde{\mathscr{B}}_1^M$.
else
  Set $A_{\mathrm{fin}} = [n]$.
end if
Obtain the classical matrix estimate $(\hat{\rho}^M)' \leftarrow \textsc{Estimate}(\mathsf{Copy}_\rho, \mathscr{H}, A_{\mathrm{fin}}, d, 1, \tau, \zeta)$.
Compute a normalized top eigenvector $|\hat{\psi}\rangle$ of $(\hat{\rho}^M)'$.
Construct an n-qudit circuit $\mathscr{C}_0$ such that $\mathscr{C}_0 |0\rangle^{\otimes n} = |\hat{\psi}\rangle_{A_{\mathrm{fin}}} \otimes |0\rangle_{[n]\setminus A_{\mathrm{fin}}}$.
Set $\mathscr{C} \leftarrow \mathscr{C}_0$.
for $j = M, M-1, \ldots, 1$ **do**
  Append $(U^j)^\dagger$ acting on $\mathscr{B}^j$ to $\mathscr{C}$.
end for
return $\mathscr{C}$, which prepares $|\hat{\phi}\rangle := \mathscr{C} |0\rangle^{\otimes n}$.

We further consider the task of learning when the target state ρ is not necessarily an exact MPS of bond dimension D, and may even be mixed. Our goal is to output a state $|\hat{\phi}\rangle$ such that its fidelity with ρ is within additive error ϵ of the best achievable fidelity achievable among all MPS of that bond dimension. This problem can be viewed as the agnostic version of MPS learning.

Definition 3 (Closest MPS Learning Problem)**.** *Let* $\mathrm{MPS}(D)$ *denote the set of* n*-qudit pure matrix product states with bond dimension at most* D*. The learner is given access to independent copies of an arbitrary* n*-qudit state* ρ*, which is not promised to belong to* $\mathrm{MPS}(D)$ *and may be mixed. Given accuracy and confidence parameters* $\epsilon, \delta \in (0,1)$*, the goal is to output, with probability at least* $1-\delta$*, a classical description of a hypothesis state* $|\hat{\phi}\rangle$ *such that* $\langle\hat{\phi}| \rho |\hat{\phi}\rangle \geq \max_{|\phi\rangle \in \mathrm{MPS}(D)} \langle\phi| \rho |\phi\rangle - \epsilon$.

Theorem 2. *For any integers* $n, d \geq 2$*, given access to copies of an* n*-qudit state* ρ*, parameters* $\epsilon, \delta \in (0,1)$*, and a bond dimension parameter* D*, Algorithm 4 in Appendix C outputs a description of a state* $|\hat{\phi}\rangle$*, such that, with probability* $1-\delta$*,* $\langle\hat{\phi}| \rho |\hat{\phi}\rangle \geq \max_{|\phi\rangle \in \mathrm{MPS(D)}} \langle\phi| \rho |\phi\rangle - \epsilon$*. The algorithm requires* $N = O\Big(\frac{D^{12} n^7 d^4 (\ln d)^7 \log\frac{n}{\delta}}{\epsilon^{12}\Big[\ln\Big(\frac{128\, n D^2 \ln d}{(\sqrt{2}-1)^2 \epsilon^2}\Big)\Big]^7}\Big)$ *copies of* ρ *and runs in time* $\mathrm{poly}(D, n, 1/\epsilon, \log(1/\delta))$ *for fixed local dimension* d*.*

Proof sketch. The proof follows the same learning framework as in Theorem 1, with one additional truncation step to handle the fact that ρ need not be an exact bond-dimension-D MPS. In the exact setting, the reduced states on the relevant blocks have rank at most D^2; in the agnostic setting, they may have full rank. Algorithm 4 therefore increases the block size to $2p$ with $d^p \geq 1/\eta$ and truncates each reduced state to its dominant eigenspace above threshold η. Since the eigenvalues of a reduced density matrix sum to one, this dominant subspace has dimension at most $1/\eta$, giving an effective low-rank substitute for the exact MPS rank constraint. With a suitable choice of η in terms of ϵ, n, and D, the truncation and estimation errors remain controlled throughout the algorithm, yielding $\langle\hat{\phi}| \rho |\hat{\phi}\rangle \geq \max_{|\phi\rangle \in \mathrm{MPS}(D)} \langle\phi| \rho |\phi\rangle - \epsilon$. The full proof is deferred to Appendix C.4. □

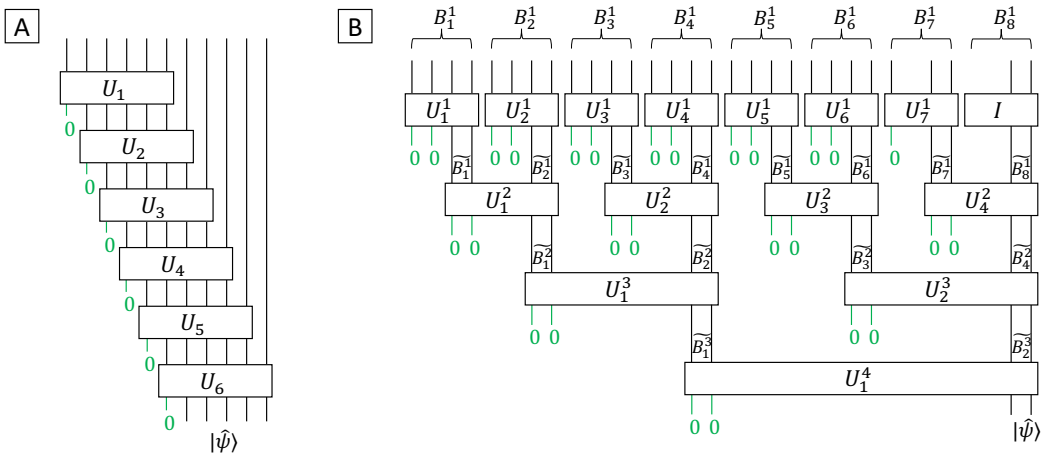


Figure 1: Comparison of circuit structures for MPS tomography. (A) Previous algorithms [LCLP10, BBK+25] use a sequential circuit structure. (B) Our Algorithm 2 produces a tree-like circuit structure, shown here for $n = 29$ and $p = 2$. Each rectangle represents a unitary U_i^j, with j denoting the layer and i the unitary index within that layer. The vertical lines are qudits ordered from left to right, and time flows downward. The brackets mark the first-layer supports $\mathscr{B}_i^1$, and the labeled wires mark the surviving sets $\widetilde{\mathscr{B}}_i^j$ carried to the next layer.

When the target state is an MPS(D), our algorithm outputs a preparation circuit containing $O(n)$ learned block unitaries. Each acts on at most $2p$ qudits and has matrix dimension at most $d^{2p} \leq d^4 D^4$. Consequently, the dense descriptions of these unitaries require at most $O(nd^{4p}) = O(nd^8 D^8)$ complex entries in total. This is an improper learning algorithm, as in [CPF+10, LCLP10]: converting the circuit description into an explicit one-dimensional MPS representation may increase the bond dimension. Nevertheless, the output remains succinct and can be used to efficiently prepare the hypothesis state.

3 Implementation on Nearest-Neighbor Lattices

The logarithmic-depth implementation of Algorithm 2 assumes an abstract circuit model in which the required qudits can be acted on jointly regardless of their geometric locations. In many realistic quantum architectures, however, gates are constrained by an underlying interaction geometry and can only be applied between neighboring qudits. It is therefore natural to ask how the depth of our construction changes under such locality constraints.

In this section, we study the implementation of the circuit output by Algorithm 2 on nearest-neighbor architectures. On a two-dimensional square lattice, the SWAP routing requires $O(\sqrt{n})$ layers. More generally, on a q-dimensional hypercubic lattice, the routing requires $O(qn^{1/q})$ layers. In both cases, the block unitaries contribute $O(\log n)$ additional layers.

Theorem 3. *On a two-dimensional square lattice, the quantum circuit output by Algorithm 2 can be implemented using $O(\sqrt{n})$ layers of nearest-neighbor SWAP gates and $O(\log n)$ layers of block unitaries, each acting on at most $2p$ qudits within adjacent patches.*

Proof sketch. View each patch of p qudits as one coarse-grained lattice site and arrange the n/p patches on a square lattice of side length $\sqrt{n/p}$. As illustrated in Figure 2, we impose a recursive 2×2 block structure on the lattice, where in each subblock the bottom-right corner is chosen as the anchor that survives to the next coarser level. At each level, two successive iterations of Algorithm 2 are implemented by first applying the disentangling unitaries horizontally and then vertically, so that within every 2×2 group only the anchor site remains active. Under nearest-neighbor connectivity, before each such disentangling step we use nearest-neighbor SWAP gates to route the qubits within each block to bring the relevant pair of active sites adjacent; if the current block has

side length L, this requires depth $O(L)$. Since all blocks at the same level are disjoint, these SWAP operations can be performed in parallel across blocks. Hence the coarse-grained depth is bounded by $\sum_{j=1}^{O(\log(n/p))} O\left(\frac{\sqrt{n/p}}{2^{j-1}}\right) = O(\sqrt{n/p})$. Since each coarse-grained site is a physical patch of p qudits occupying area $\Theta(p)$, routing within a patch contributes an additional factor $O(\sqrt{p})$. Thus the total physical depth is $O(\sqrt{n})$. We defer the full routing details to Appendix D. □

Corollary 3.1. *On a q-dimensional hypercubic lattice, the quantum circuit output by Algorithm 2 can be implemented using $O(qn^{1/q})$ layers of nearest-neighbor SWAP gates and $O(\log n)$ layers of block unitaries, each acting on at most $2p$ qudits within adjacent patches.*

Proof sketch. The proof is a direct q-dimensional analogue of Theorem 3. View each patch of p qubits as one lattice site and arrange the n/p patches on a q-dimensional hypercubic lattice of side length $(n/p)^{1/q}$. We impose a recursive 2^q-ary block structure: at each level, every block is partitioned into 2^q subblocks, and the maximal-coordinate corner of each subblock is chosen as the anchor that survives to the next coarser level. Each level is implemented by applying the disentangling operations successively along the q coordinate directions, using nearest-neighbor SWAP gates to route active sites within each block; if the current block side length is L, this costs depth $O(qL)$, while disjoint blocks are processed in parallel. Summing over the geometrically decreasing side lengths gives $\sum_{j=1}^{O(\log(n/p))} O\left(q\,\frac{(n/p)^{1/q}}{2^{j-1}} p^{1/q}\right) = O\left(q\,n^{1/q}\right)$. The full proof is deferred to Appendix D. □

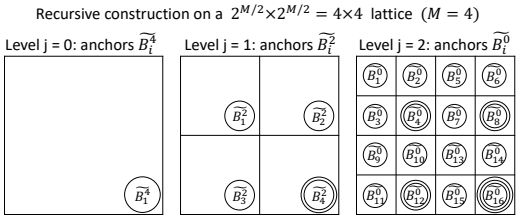


Figure 2: Illustration of the recursive block structure on a $2^{M/2} \times 2^{M/2}$ square lattice with $M = 4$. At level j, each block of size $2^{M/2-(j-1)} \times 2^{M/2-(j-1)}$ is partitioned into four subblocks, each of size $2^{M/2-j} \times 2^{M/2-j}$. In each subblock, the bottom-right site is designated as the anchor $\widetilde{\mathscr{B}}_i^{M-2j}$ (circled). These anchors form a coarser lattice for the next level.

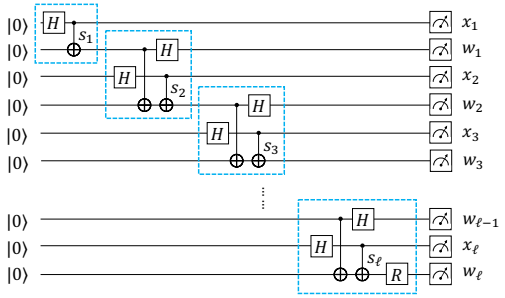


Figure 3: Quantum circuit for generating an LPN-hard MPS(2) family. The odd-indexed qubits are initialized to $|0\rangle$ and put into a uniform superposition over $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$ by Hadamard gates. The circuit then sequentially computes the prefix parities $z_i = \bigoplus_{j=1}^{i} x_j s_j$. For each $i < \ell$, a final Hadamard gate makes w_i uniformly random and independent of the other measured outputs. On the last parity register, where $z_\ell = \langle x, s\rangle \pmod 2$, a rotation $R = R_y(2\arcsin(\sqrt{\tau}))$ is applied so that the measured bit satisfies $w_\ell = \langle x, s\rangle \oplus e$ with $e \sim \text{Bernoulli}(\tau)$. Thus each measurement sample (x, w_ℓ) is distributed as a Search-LPN$_{\tau,\ell}$ sample.

4 Lower bounds and hardness for learning MPS

We prove a sample-complexity lower bound for MPS learning by considering product states, namely MPS(1). To the best of our knowledge, this is the first lower bound of this form for the problem.

Theorem 4. *For integer $n \geq 1$, $d \geq 2$, and $\epsilon \in [0, 2/3]$, consider learning a product state $\rho = \bigotimes_{i=1}^{n} \rho_i$, where each ρ_i is a state on a d-dimensional Hilbert space. Any quantum algorithm that outputs a product state $\hat{\rho}$ satisfying $F(\rho, \hat{\rho}) \geq 1 - \epsilon$ must use at least $\Omega(nd/\epsilon)$ copies of ρ.*

Proof sketch. It suffices to prove the lower bound for pure product states. We argue by reduction from single-qudit state tomography. The known tomography lower bound says that learning an arbitrary d-dimensional quantum state to fidelity at least $1-\epsilon'$ requires $\Omega(d/\epsilon')$ copies. If there were an algorithm that learned n-qudit product states to fidelity at least $1-\epsilon$ using $o(nd/\epsilon)$ copies, then by embedding a single unknown d-dimensional state into one tensor factor of an otherwise known product state, one could learn a single-qudit state to accuracy $\epsilon' = \epsilon/n$ using $o(d/\epsilon')$ copies. This would contradict the tomography lower bound, and therefore product-state learning requires $\Omega(nd/\epsilon)$ copies. We defer the formal proof to Appendix E. □

We next prove a computational hardness result for learning MPS under non-adaptive single-qubit measurements. We show that efficiently learning MPS(2) in this model would imply an efficient algorithm for Search-LPN (Definition 11) and thus falsify the LPN assumption (Assumption 1).

Theorem 5. *Assume the* $\mathrm{LPN}_{\tau,\ell}$ *assumption holds for some constant* $\tau \in (0, 1/2-\gamma)$*, where* $\gamma > 0$ *is non-negligible. Let* $\{|\psi_s\rangle : s \in \{0,1\}^\ell\}$ *be the family of MPS*(2) *states defined in Figure 3. Then there is no polynomial-time algorithm that, using only computational-basis single-qubit measurements on* $\mathrm{poly}(\ell)$ *copies of* $|\psi_s\rangle$*, outputs a classical description of a hypothesis state* $\widehat{\rho}$ *from which* $\widehat{\rho}$ *can be efficiently prepared*[9] *and* $F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1-\epsilon$ *with non-negligible probability.*

Proof sketch. By the construction in Figure 3, measuring one copy of $|\psi_s\rangle$ in the computational basis produces a uniformly random $x \in \mathbb{Z}_2^\ell$ and uniformly random $w_1, w_2, \ldots, w_{\ell-1} \in \mathbb{Z}_2$ together with a bit $w_\ell = \langle x, s\rangle \oplus e$, where $e \sim \mathrm{Bernoulli}(\tau)$. Note that, for each measured copy, the pair (x, w_ℓ) gives exactly one Search-$\mathrm{LPN}_{\tau,\ell}$ sample with secret s, while $w_1, w_2, \ldots, w_{\ell-1}$ are uniformly random binary strings independent of x and w_ℓ. Then, we show that s can be efficiently recovered from the classical description that efficiently prepares $\widehat{\rho}$. Therefore, if a polynomial-time learner could recover s from $q = \mathrm{poly}(\ell)$ such samples with non-negligible probability, then it would solve Search-$\mathrm{LPN}_{\tau,\ell}$ with non-negligible probability. This contradicts the Search-$\mathrm{LPN}_{\tau,\ell}$ assumption. Therefore no such learner exists. The formal proof is deferred to Appendix F. □

Finally, we show the following Corollary and defer the proof to Appendix F.

Corollary 5.1. *Assume the* $\mathrm{LPN}_{\tau,\ell}$ *assumption holds for* $\tau \in (0, 1/2-\gamma)$*, where* $\gamma > 0$ *is non-negligible. Then there is no polynomial-time non-adaptive algorithm that, using only single-qubit measurements on* $\mathrm{poly}(\ell)$ *copies of* $|\psi_s\rangle$*, outputs a classical description of a hypothesis state* $\widehat{\rho}$ *from which* $\widehat{\rho}$ *can be efficiently prepared and such that* $F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1-\epsilon$ *with non-negligible probability.*

4.1 Discussion and open questions

We introduce a parallel disentangling algorithm for exact MPS learning, reducing the circuit depth from $O(n)$ to $O(\log n)$ and the sample complexity from $\widetilde{O}(n^5)$ to $\widetilde{O}(n^3)$. We also extend the algorithm to closest MPS learning, reducing the sample complexity from $\widetilde{O}(n^9)$ to $\widetilde{O}(n^7)$. Our algorithms improve the dependence on the system size n [LCLP10, CPF$^+$10, BBK$^+$25], at the cost of a worse dependence on the bond dimension D (and ϵ for closest MPS learning). Our algorithms achieve the stated copy bounds using collective tomography measurements with polynomial gate complexity [PSTW25]; for comparison, an approach in [QJG$^+$24] achieves $O(n^3D^2/\epsilon^2)$ copies but relies on Haar-random measurements. We also establish a sample-complexity lower bound of $\Omega(n/\epsilon)$ for learning product states. Together, our algorithms and lower bound provide the first systematic view of the sample-complexity landscape for efficient MPS learning.

The main tradeoff in our algorithms is that the improved dependence on the system size n comes at the cost of a worse dependence on the bond dimension D. This tradeoff is meaningful from a learning-theoretic perspective: the bond dimension is a model-complexity parameter, and tensor-network methods are most useful in the low-complexity regime where D is small compared with the full Hilbert-space dimension. Quantitatively, our algorithm wins previous ones whenever $D =$

[9] To the best of our knowledge, all existing MPS learning algorithms, including those presented in this work, share this property. Specifically, they produce classical descriptions from which polynomial-sized quantum circuits can be efficiently generated to prepare the corresponding states

$o(\sqrt{n})$.[10] In particular, for any subpolynomial bond dimension $D = n^{o(1)}$, our sample complexity is asymptotically smaller in n. This is precisely the regime in which tensor-network methods serve as compact models for high-dimensional quantum states. It includes constant-bond-dimension families, such as GHZ states, W states, AKLT states, and one-dimensional cluster states, as well as one-dimensional gapped ground states that admit inverse-polynomial-accuracy MPS approximations with subpolynomial bond dimension for fixed spectral gap [AKLV13].

We also consider implementing our algorithms on quantum devices that have constraints on connectivity and measurements. In particular, we show that even with connectivity constraints, our algorithm can still outperform existing algorithms in depth and sample complexity; on the other hand, we showed that there exists no polynomial-time algorithm if the learner can only apply non-adaptive single-qubit measurement under the LPN assumption. These results provide a detailed characterization of needed quantum resources for MPS learning.

A gap still remains between our upper bound and lower bounds; closing it is the main open question. Promising directions include:

- **Closing the gap in n.** Can we design an efficient algorithm with sample complexity $O(n)$, or else prove an $\Omega(n^3)$ lower bound?
- **Closing the gap in D.** The best known lower bound is $\Omega(\sqrt{D})$ via MPS testing [SW22, ABF$^+$22]. Can we obtain a tighter lower bound specific to MPS learning?
- **Proper learning with shallow circuits.** Proper agnostic learning of OBC MPS is now established [CVdPC26], with runtime polynomial in n for fixed d, D, ϵ. Can proper learning be achieved with shallow quantum circuits or with improved dependence on D and $1/\epsilon$? Proper agnostic learning of periodic-boundary MPS also remains open: [CVdPC26] gives only an improper OBC hypothesis of bond at most D^2 for a bond-D periodic comparator.

Our algorithms have potential applications to quantum machine learning. MPS and related tensor-network models have been used as expressive but structured ansatz classes for high-dimensional data. In this light, logarithmic-depth MPS learning can be viewed as a step toward resource-efficient procedures for extracting compact tensor-network descriptions from quantum data.

Acknowledgments and Disclosure of Funding

We thank Kaden Hazzard and Tony Deng for valuable discussions on tensor-network states. Nai-Hui Chia is supported by NSF Award FET-2243659, NSF Career Award FET-2339116, Google Scholar Award, and DOE Quantum Testbed Finder Award DE-SC0024301. Chia-Ying Lin is supported by DOE Quantum Testbed Finder Award DE-SC0024301 and NSF Award FET-2243659. Shih-Han Hung is supported by the National Science and Technology Council, Taiwan, under Grant No. NSTC-114-2222-E-002-002-MY3, and by National Taiwan University under Grant Nos. 114L895503 and 114L74108.

Competing Interests. The authors declare no competing interests.

References

[ABF$^+$22] Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum pseudoentanglement. *arXiv preprint arXiv:2211.00747*, 2022.

[AKLV13] Itai Arad, Alexei Kitaev, Zeph Landau, and Umesh Vazirani. An area law and sub-exponential algorithm for 1d systems. *arXiv preprint arXiv:1301.1162*, 2013.

[ALV12] Itai Arad, Zeph Landau, and Umesh Vazirani. Improved one-dimensional area law for frustration-free systems. *Physical Review B—Condensed Matter and Materials Physics*, 85(19):195145, 2012.

[10] Prior work uses $O(n^5 D^2/\epsilon^4)$ samples, whereas our algorithm uses $O(n^3 D^6/\epsilon^4)$ samples. Thus, comparing the leading polynomial dependencies, our algorithm has better sample complexity whenever $D = o(\sqrt{n})$.

[BBK+25] Ainesh Bakshi, John Bostanci, William Kretschmer, Zeph Landau, Jerry Li, Allen Liu, Ryan O'Donnell, and Ewin Tang. Learning the closest product state. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1212–1221, 2025.

[CPF+10] 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):149, 2010.

[CPGSV21] J. Ignacio Cirac, David Perez-Garcia, Norbert Schuch, and Frank Verstraete. Matrix product states and projected entangled pair states: Concepts, symmetries, theorems. *Reviews of Modern Physics*, 93(4):045003, 2021.

[CVdPC26] Constantin Cedillo Vayson de Pradenne and Jordan Cotler. Proper Agnostic Learning of Matrix Product States and Tree Tensor Networks. arXiv preprint arXiv:2609.30148, 2026.

[DKSV04] Andrew John Daley, Corinna Kollath, Ulrich Schollwöck, and Guifré Vidal. Time-dependent density-matrix renormalization-group using adaptive effective hilbertspaces. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(04):P04005, 2004.

[Has07] Matthew B. Hastings. An area law for one-dimensional quantum systems. *Journal of statistical mechanics: theory and experiment*, 2007(08):P08024, 2007.

[HWF+18] Zhao-Yu Han, Jun Wang, Heng Fan, Lei Wang, and Pan Zhang. Unsupervised generative modeling using matrix product states. *Physical Review X*, 8(3):031012, 2018.

[KNR+25] Andrew D King, Alberto Nocera, Marek M Rams, Jacek Dziarmaga, Roeland Wiersema, William Bernoudy, Jack Raymond, Nitin Kaushal, Niclas Heinsdorf, Richard Harris, et al. Beyond-classical computation in quantum simulation. *Science*, 388(6743):199–204, 2025.

[LCLP10] Olivier Landon-Cardinal, Yi-Kai Liu, and David Poulin. Efficient Direct Tomography for Matrix Product States, 2010.

[LVV15] Zeph Landau, Umesh Vazirani, and Thomas Vidick. A polynomial time algorithm for the ground state of one-dimensional gapped local hamiltonians. *Nature Physics*, 11(7):566–569, 2015.

[NPOV15] Alexander Novikov, Dmitrii Podoprikhin, Anton Osokin, and Dmitry P Vetrov. Tensorizing neural networks. *Advances in neural information processing systems*, 28, 2015.

[NTO16] Alexander Novikov, Mikhail Trofimov, and Ivan Oseledets. Exponential machines. *arXiv preprint arXiv:1605.03795*, 2016.

[OV08] Roman Orus and Guifre Vidal. Infinite time-evolving block decimation algorithm beyond unitary evolution. *Physical Review B—Condensed Matter and Materials Physics*, 78(15):155117, 2008.

[Pie12] Krzysztof Pietrzak. Cryptography from learning parity with noise. In *International Conference on Current Trends in Theory and Practice of Computer Science*, pages 99–114. Springer, 2012.

[PSTW25] Angelos Pelecanos, Jack Spilecki, Ewin Tang, and John Wright. Mixed state tomography reduces to pure state tomography, 2025.

[QJG+24] Zhen Qin, Casey Jameson, Zhexuan Gong, Michael B Wakin, and Zhihui Zhu. Quantum state tomography for matrix product density operators. *IEEE Transactions on Information Theory*, 70(7):5030–5056, 2024.

[RKR23] Hans-Martin Rieser, Frank Köster, and Arne Peter Raulf. Tensor networks for quantum machine learning. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 479(2275), 2023.

[SAOB22] Richik Sengupta, Soumik Adhikary, Ivan Oseledets, and Jacob Biamonte. Tensor networks in machine learning. *European Mathematical Society Magazine*, (126):4–12, 2022.

[SS16] Edwin Stoudenmire and David J Schwab. Supervised learning with tensor networks. *Advances in neural information processing systems*, 29, 2016.

[SW22] Mehdi Soleimanifar and John Wright. Testing matrix product states. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1679–1701. SIAM, 2022.

[Vid04] Guifré Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Physical review letters*, 93(4):040502, 2004.

[VPC04] Frank Verstraete, Diego Porras, and J Ignacio Cirac. Density matrix renormalization group and periodic boundary conditions: A quantum information perspective. *Physical review letters*, 93(22):227205, 2004.

[Whi92] Steven R White. Density matrix formulation for quantum renormalization groups. *Physical review letters*, 69(19):2863, 1992.

[WRV+24] Dian Wu, Riccardo Rossi, Filippo Vicentini, Nikita Astrakhantsev, Federico Becca, Xiaodong Cao, Juan Carrasquilla, Francesco Ferrari, Antoine Georges, Mohamed Hibat-Allah, et al. Variational benchmarks for quantum many-body problems. *Science*, 386(6719):296–301, 2024.

[Yue23] Henry Yuen. An improved sample complexity lower bound for (fidelity) quantum state tomography. *Quantum*, 7:890, January 2023.

Appendix A: Preliminaries and Tomography

A.1 MPS Rank Properties

Definition 4. *Consider a quantum system consisting of n sites, each associated with a local Hilbert space $\mathbb{C}^d$. A generic state in the global Hilbert space $(\mathbb{C}^d)^{\otimes n}$ can be written as*

$$|\psi\rangle = \sum_{i_1,i_2,\ldots,i_n=1}^{d} c_{i_1 i_2 \cdots i_n} \; |i_1\rangle \otimes |i_2\rangle \otimes \cdots \otimes |i_n\rangle \,,$$

where the coefficients $c_{i_1 i_2 \cdots i_n}$ are complex amplitudes. In the MPS representation, these amplitudes are expressed as products of matrices. Specifically, a state $|\psi\rangle$ is an MPS with bond dimension D if

$$c_{i_1 i_2 \cdots i_n} = \mathrm{Tr}\Big(A_{i_1}^{[1]} A_{i_2}^{[2]} \cdots A_{i_n}^{[n]}\Big) \,,$$

where for each site j, the tensors $\{A_i^{[j]}\}_{i=1}^d$ are $D \times D$ complex matrices (with the first and last possibly rectangular if open boundary conditions are used). The parameter D is referred to as the bond dimension of the MPS, and controls the amount of entanglement that the state can carry.

Lemma A.1. *Let $|\psi\rangle$ be a Matrix Product State (MPS) on n sites with periodic boundary conditions and bond dimension D. For any block $\mathscr{B}_j$ of k contiguous sites, the reduced density matrix*

$$\rho_j = \mathrm{Tr}_{[n]\setminus\mathscr{B}_j}(|\psi\rangle\langle\psi|)$$

has rank at most D^2.

Proof. Let $\mathscr{C} := [n] \setminus \mathscr{B}_j$ denote the complement of the block. Represent the periodic MPS $|\psi\rangle$ by contracting local tensors around the ring. Cutting the ring at the two virtual bonds that border the contiguous block $\mathscr{B}_j$ exposes two virtual indices (one on each cut), denoted by $\alpha, \beta \in \{1, \ldots, D\}$. By contracting all physical indices inside $\mathscr{B}_j$ while leaving the two boundary virtual indices open, we obtain vectors

$$\{\, |\Phi_{\alpha\beta}\rangle_{\mathscr{B}_j} \in \mathbb{H}_{\mathscr{B}_j} : \alpha, \beta \in [D] \,\}.$$

Similarly, contracting all tensors in the complement $\mathscr{C}$ (with the two virtual indices open) gives vectors

$$\{\, |E_{\beta\alpha}\rangle_{\mathscr{C}} \in \mathbb{H}_{\mathscr{C}} : \alpha, \beta \in [D] \,\}.$$

Thus the global state can be written as

$$|\psi\rangle \;=\; \sum_{\alpha,\beta=1}^{D} |\Phi_{\alpha\beta}\rangle_{\mathscr{B}_j} \otimes |E_{\beta\alpha}\rangle_{\mathscr{C}}\,.$$

Taking the partial trace over $\mathscr{C}$ yields

$$\rho_j = \mathrm{Tr}_{\mathscr{C}}(|\psi\rangle\langle\psi|) = \sum_{\alpha,\beta,\alpha',\beta'=1}^{D} \langle E_{\beta'\alpha'}|E_{\beta\alpha}\rangle \;|\Phi_{\alpha\beta}\rangle_{\mathscr{B}_j}\langle\Phi_{\alpha'\beta'}|\,.$$

Therefore, the support of ρ_j is contained in $\mathrm{span}\{|\Phi_{\alpha\beta}\rangle : \alpha, \beta \in [D]\}$, which has dimension at most D^2. Hence

$$\mathrm{rank}(\rho_j) \le D^2,$$

as claimed. □

A.2 Subnormalized State Tomography

Lemma A.2 (Rank-constrained subnormalized tomography)**.** *Let ρ be a normalized input state, and let $\mathscr{T}$ be a known, physically implementable, completely positive trace-nonincreasing map describing a specified successful measurement branch, followed by discarding any registers not retained in the output. Write*

$$\sigma = \mathscr{T}(\rho), \qquad \mu = \mathrm{Tr}\,\sigma, \qquad \bar{\sigma} = \sigma/\mu \quad (\mu > 0).$$

Suppose that the output space has dimension h and $\mathrm{rank}(\sigma) \le R$, where $1 \le R \le h$. For $\eta, \delta \in (0,1)$, there exists an algorithm using at most

$$O\!\left(\frac{Rh}{\eta^2}\log\frac{2}{\delta}\right)$$

original copies of ρ, including all failed postselection trials, that outputs a classical description of a positive semidefinite operator $\hat{\sigma}$ with $\mathrm{Tr}\,\hat{\sigma} \le 1$ and $\|\sigma - \hat{\sigma}\|_1 \le \eta$ with probability at least $1-\delta$. The algorithm does not require prior knowledge of μ and remains well-defined when $\mu = 0$. The tomography measurements and classical reconstruction can be implemented in time polynomial in $h, R, 1/\eta, \log(2/\delta)$ using an efficient collective-measurement tomography procedure [PSTW25]. The cost of implementing $\mathscr{T}$ on each input copy is additional; no shallow-depth bound is claimed for the collective tomography measurements. In particular, for an $(r-i)$-qudit output with rank at most D^2, taking $R = \min\{D^2, d^{r-i}\}$ gives the sample bound

$$O\!\left(\frac{D^2 d^{r-i}}{\eta^2}\log\frac{2}{\delta}\right).$$

Proof. Set $L = \log(8/\delta)$. A trial is successful when the specified measurement branch occurs; its probability is μ. For $\mu > 0$, the retained quantum output conditioned on success has state $\bar{\sigma}$. First use $m_0 = \lceil C_0 \eta^{-2} L \rceil$ fresh input copies to estimate μ by the observed success frequency $\hat{\mu}$. Hoeffding's inequality gives, for sufficiently large C_0,

$$\Pr\big[|\hat{\mu} - \mu| > \eta/4\big] \le \delta/4.$$

If $\mu = 0$, every trial fails, so $\hat{\mu} = 0$ with probability one. The procedure then outputs zero, which equals σ because $\sigma \succeq 0$ and $\mathrm{Tr}\,\sigma = 0$. Henceforth assume $\mu > 0$ and consider the event $|\hat{\mu} - \mu| \le \eta/4$. If $\hat{\mu} \le \eta/2$, output $\hat{\sigma} = 0$. Then $\|\sigma - \hat{\sigma}\|_1 = \mu \le 3\eta/4$.

Otherwise, $\hat{\mu} > \eta/2$, which implies

$$\hat{\mu}/2 \le \mu \le 2\hat{\mu}.$$

Set $a = \eta/(4\hat{\mu})$. By the gate-efficient mixed-state tomography theorem of [PSTW25, Theorem 1.3], normalized-state tomography with rank bound R and trace-norm accuracy a can be implemented with polynomial quantum-gate and classical complexity. Indeed, running that algorithm with infidelity

tolerance $a^2/4$ gives trace-norm error at most a by the Fuchs–van de Graaf inequalities, using $O((Rh + \log(4/\delta))/a^2)$ successful copies. Since $Rh \geq 1$ and $L = \log(8/\delta)$, a sufficiently large universal constant C_1 allows us to take

$$k = \lceil C_1 Rh\, a^{-2} L \rceil = \lceil 16 C_1 Rh\, \hat{\mu}^2 \eta^{-2} L \rceil$$

independent copies of $\bar{\sigma}$ to produce a normalized positive semidefinite estimate $\widehat{\widehat{\sigma}}$ satisfying $\|\bar{\sigma} - \widehat{\widehat{\sigma}}\|_1 \leq a$, except with probability $\delta/4$.

To obtain these copies with a fixed input budget, apply the successful branch to a second, independent batch of $m_1 = \lceil C_2 Rh\eta^{-2} L \rceil$ fresh input copies. Conditioned on any pilot outcome satisfying $|\hat{\mu} - \mu| \leq \eta/4$ and $\hat{\mu} > \eta/2$, the second-batch success count S is binomial with mean μm_1. Indeed, since $\mu \geq \hat{\mu}/2$, $\hat{\mu} \leq 1$, and $Rh\eta^{-2} L\, \hat{\mu} > 1$, we have

$$\mu m_1 \geq \frac{C_2}{2} Rh\hat{\mu}\eta^{-2} L, \qquad 2k \leq (32C_1 + 2) Rh\hat{\mu}\eta^{-2} L.$$

Also, $\mu > \eta/4$ implies

$$\mu m_1 > \frac{C_2 Rh}{4\eta} L \geq \frac{C_2}{4} L.$$

Thus choosing $C_2 \geq \max\{64C_1 + 4, 32\}$ ensures

$$\mu m_1 \geq 2k, \qquad \mu m_1 \geq 8L.$$

A multiplicative Chernoff bound therefore gives, for every such pilot outcome,

$$\Pr[S < k \mid \hat{\mu}] \leq \exp(-\mu m_1/8) \leq \delta/8.$$

If $S < k$, output zero and count this branch as a failure event. If $S \geq k$, perform normalized-state tomography on the quantum outputs of the first k successful trials, obtaining $\widehat{\widehat{\sigma}}$, and set $\hat{\sigma} = \hat{\mu}\widehat{\widehat{\sigma}}$. Since the trials use independent input copies, conditioning on the pilot data and on which trials succeeded leaves the retained quantum outputs as independent copies of $\bar{\sigma}$.

Whenever the pilot frequency estimate, the second-batch success count, and the normalized-state tomography satisfy the above guarantees,

$$\begin{aligned} \|\sigma - \hat{\sigma}\|_1 &\leq \mu \|\bar{\sigma} - \widehat{\widehat{\sigma}}\|_1 + |\mu - \hat{\mu}| \|\widehat{\widehat{\sigma}}\|_1 \\ &\leq 2\hat{\mu}\, \frac{\eta}{4\hat{\mu}} + \frac{\eta}{4} \leq \eta. \end{aligned}$$

Moreover, $\hat{\sigma} \succeq 0$ and $\operatorname{Tr} \hat{\sigma} = \hat{\mu} \leq 1$. A union bound over the pilot estimation, insufficient-success, and tomography failure events gives total failure probability at most $\delta/4 + \delta/8 + \delta/4 < \delta$. The deterministic input budget is $m_0 + m_1 = O(Rh\eta^{-2} \log(2/\delta))$, since $Rh \geq 1$; it includes all pilot copies and all failed second-batch trials. □

Lemma A.3 (Subnormalized tomography)**.** *Under the same input and successful-branch model as Lemma A.2, suppose only that the output space has dimension h. For $\eta, \delta \in (0, 1)$, an algorithm using at most*

$$O\left(\frac{h^2}{\eta^2} \log \frac{2}{\delta}\right)$$

original copies of ρ, including failed postselection trials, outputs a classical description of a positive semidefinite operator $\hat{\sigma}$ with $\operatorname{Tr} \hat{\sigma} \leq 1$ and $\|\mathscr{T}(\rho) - \hat{\sigma}\|_1 \leq \eta$ with probability at least $1 - \delta$. No prior knowledge of the success probability is required, and the algorithm remains well-defined when it is zero. For an $(r - i)$-qudit output, the sample bound becomes

$$O\left(\frac{d^{2(r-i)}}{\eta^2} \log \frac{2}{\delta}\right).$$

As in Lemma A.2, the tomography measurements and classical postprocessing have polynomial implementation cost in $h, 1/\eta, \log(2/\delta)$, although no shallow-depth guarantee is asserted for the measurements. The cost of implementing $\mathscr{T}$ is additional.

Proof. Apply Lemma A.2 with $R = h$. □

A.3 Physical sampling and tomography for Algorithm 2

We first specify the inputs and notation used by REPLAYSAMPLE and ESTIMATE.

- Copy_ρ: A source that supplies one fresh independent copy of ρ on the registers $[n]$ per invocation.
- $\mathscr{H}$: An ordered list of previously completed layer records,

$$\mathscr{H} = \big((\mathscr{B}^1, U^1, \mathscr{D}^1), \ldots, (\mathscr{B}^a, U^a, \mathscr{D}^a)\big),$$

  where each record consists of:
  - $\mathscr{B}^t$: The registers present before layer t.
  - U^t: The learned unitary of layer t, acting as the identity on unprocessed registers.
  - $\mathscr{D}^t \subseteq \mathscr{B}^t$: The registers measured and removed at layer t.

  The empty history is denoted by $()$. When processing layer j, $\mathscr{H}$ contains only layers $1, \ldots, j-1$. The record for layer j is appended only after all its block estimates have been computed.
- A: The target register set, satisfying $A \subseteq [n] \setminus \bigcup_{t=1}^{a} \mathscr{D}^t$.

REPLAYSAMPLE returns either a failure flag or actual quantum registers on A. ESTIMATE also takes the local dimension d, an output rank bound R, the trace-norm accuracy α, and a failure probability ζ, and returns a classical positive semidefinite (PSD) matrix of trace at most one. The true success probability is not an input. All variables inside these functions are local to their invocation. Fix universal constants C_0, C_1, C_2 as in the proof of Lemma A.2.

function REPLAYSAMPLE($\mathsf{Copy}_\rho, \mathscr{H}, A$)
  Invoke Copy_ρ to obtain one fresh input copy.
  for each $(\mathscr{B}^t, U^t, \mathscr{D}^t)$ in $\mathscr{H}$, in order **do**
    Apply U^t on $\mathscr{B}^t$.
    Measure $\mathscr{D}^t$ in the computational basis.
    if any measured outcome is nonzero **then**
      Discard all remaining registers of this copy.
      return FAIL.
    end if
    Discard the measured registers $\mathscr{D}^t$.
  end for
  Discard all remaining registers outside A.
  return the quantum registers in A.
end function

function ESTIMATE($\mathsf{Copy}_\rho, \mathscr{H}, A, d, R, \alpha, \zeta$)
  Set $h = d^{|A|}$, $R = \min\{R, h\}$, and $L = \log(8/\zeta)$.
  if $\mathscr{H} = ()$ **then**
    Set $\hat{q} = 1$.
  else
    Set $m_0 = \lceil C_0 \alpha^{-2} L \rceil$.
    Run REPLAYSAMPLE($\mathsf{Copy}_\rho, \mathscr{H}, A$) independently m_0 times, using fresh input copies.
    Let S_0 be the number of calls that do not return FAIL; discard their quantum outputs.
    Set $\hat{q} = S_0/m_0$.
  end if
  if $\hat{q} \le \alpha/2$ **then**
    return the classical zero matrix $0_{h\times h}$.
  end if
  Set $e = \alpha/(4\hat{q})$, $k = \lceil C_1 Rh\, e^{-2} L \rceil$, and $m_1 = \lceil C_2 Rh\, \alpha^{-2} L \rceil$.
  Run REPLAYSAMPLE($\mathsf{Copy}_\rho, \mathscr{H}, A$) independently m_1 times, using fresh input copies.
  Let S_1 be the number of calls that do not return FAIL.
  Store the quantum outputs of the first $\min\{S_1, k\}$ successful calls and discard any additional outputs.

if $S_1 < k$ **then**
    Discard the stored outputs.
    return the classical zero matrix $0_{h\times h}$.
end if
Perform normalized-state tomography on the k stored output samples, with rank bound R, trace-norm accuracy e, and failure probability $\zeta/4$, obtaining a classical PSD matrix $\hat{\omega}$ with $\operatorname{Tr}\hat{\omega} = 1$.
return the classical matrix $\hat{q}\hat{\omega}$.
end function

Operators used in the analysis. For fixed learned unitaries, set $(\rho^0)' = \rho$ and define

$$Z^j = |0^{|\mathscr{D}^j|}\rangle\langle 0^{|\mathscr{D}^j|}|\,, \qquad Q^j = Z^j \otimes I_{\mathscr{B}^j\setminus\mathscr{D}^j},$$
$$(\rho^j)' = \operatorname{Tr}_{\mathscr{D}^j}\left[Q^j U^j (\rho^{j-1})' (U^j)^\dagger Q^j\right],$$

Here Z^j acts only on the discarded registers $\mathscr{D}^j$, whereas Q^j is the corresponding full-space postselection projector on $\mathscr{B}^j$. All tensor products are understood with registers in their inherited order. Later unitaries act only on surviving registers, so sequential postselection is equivalent to applying all learned unitaries first and then jointly projecting the discarded registers onto their all-zero state. For a history of length a, let $q_a = \operatorname{Tr}[(\rho^a)']$. When $q_a > 0$, successful outputs on A have normalized state

$$\omega_A = \frac{\operatorname{Tr}_{\mathscr{B}^{a+1}\setminus A}[(\rho^a)']}{q_a},$$

where $\mathscr{B}^{a+1} = [n] \setminus \bigcup_{t=1}^{a} \mathscr{D}^t$ denotes the registers surviving that history, including the final registers when $a = M$. The target of ESTIMATE is the subnormalized operator $q_a\omega_A$. If $q_a = 0$, that target is zero and ω_A is left undefined. Thus block calls estimate σ_i^j and the final call estimates $(\rho^M)'$, while all physical tomography samples are normalized.

Failure probability and sample counting. The low-frequency zero-output branch is accurate on the pilot concentration event; the insufficient-success branch is instead counted in the lemma's failure budget. With the nonempty history held fixed, the proof of Lemma A.2 applies directly. For an empty history the success probability is known to be one, so its pilot batch can be omitted without changing the guarantee. Each call uses at most $O(Rh\alpha^{-2}\log(2/\zeta))$ original copies, including failed trials. Lemma B.1 supplies the rank bound for block calls; the final residual has rank at most one. Conditioned on the previous learning history, each call uses fresh independent inputs and a fixed replay procedure. There are at most 2^M calls for $M \geq 1$ and one for $M = 0$; therefore the choice $\zeta = \delta/(2^M + 1)$ yields total failure probability at most δ by a conditional union bound.

Appendix B: Proof of Theorem 1: Exact MPS Learning

In this section, we prove the correctness of our algorithm. Let ρ be an n-qudit unknown state that is guaranteed to be a matrix product state of bond dimension D. Taking ρ as input, Algorithm 2 aims to learn and output a description of a state $|\hat{\phi}\rangle$ that is sufficiently close to ρ.

A key structural property required by Algorithm 2 is that every reduced operator σ_i^j on which tomography is performed has rank at most D^2. This cross-layer rank invariant is established in Lemma B.1, which shows that $\operatorname{rank}(\sigma_i^j) \leq D^2$ for every pair (j, i) for which Algorithm 2 performs tomography. Accordingly, we set

$$p := \max\{1, 2\lceil \log_d D\rceil\}$$

so that $d^p \geq D^2$. This ensures that the surviving p-qudit subsystem has sufficient dimension to support the rank-D^2 subspace identified by tomography. In the agnostic setting, where the input need not be an MPS and may be mixed, we instead use a spectral-thresholding construction.

B.1 Reconstructed States and the Cross-Layer Rank Invariant

Definition 5 (Step-wise Reconstructed State)**.** *Let $(\rho^0)' = \rho$ be the initial state before any iteration and set $\rho^0 = \rho$ and $E^0 = I$. For $j = 1, \ldots, M$, define*

$$E^j := U^j U^{j-1} \cdots U^1,$$

where each U^a is extended by the identity on registers discarded in earlier layers, so that E^j acts on all n registers. Use the discarded-register projector Z^a and full-space projector $Q^a = Z^a \otimes I_{\mathscr{B}^a \setminus \mathscr{D}^a}$ defined in Appendix A.3. Let $(\rho^j)'$ denote the subnormalized residual operator after successful postselection through layer j. Define the reconstructed operator on the original n registers by

$$\rho^j \coloneqq (E^j)^\dagger \left[\left(\bigotimes_{a=1}^{j} Z^a \right) \otimes (\rho^j)' \right] E^j.$$

Tensor factors are placed on their corresponding physical registers, with identities implicit where required by the register ordering.

We call ρ^j the step-wise reconstructed state *corresponding to the j-th layer of the circuit, representing the state that would have been reconstructed had the algorithm stopped after the j-th iteration.*

For $j = 0, \ldots, M$, define

$$q_j \coloneqq \mathrm{Tr}[(\rho^j)'], \qquad \bar{\rho}^j \coloneqq (\rho^j)'/q_j \quad (q_j > 0).$$

Here $q_0 = 1$, and q_j is the probability that all projections through layer j succeed on one original copy of ρ, conditioned on the fixed learned unitaries. The operators $(\rho^j)'$ and ρ^j are subnormalized operators used in the analysis, whereas $\bar{\rho}^j$ is the normalized physical state in the successful branch. If $q_j = 0$, then $(\rho^j)' = 0$ and no normalized conditional state is required. The tomography procedure is defined in this case by Lemma A.2.

Lemma B.1 (Cross-layer rank invariant)**.** *Let $\rho = |\psi\rangle\langle\psi|$ be an n-qudit matrix product state with periodic boundary conditions and bond dimension D. Let $\mathscr{B}_i^j$ and $\widetilde{\mathscr{B}}_i^j$ be the blocks constructed by Algorithm 2. For every pair (j, i) for which Algorithm 2 performs tomography, define*

$$\sigma_i^j \coloneqq \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j} \left[(\rho^{j-1})' \right].$$

Then

$$\mathrm{rank}(\sigma_i^j) \leq D^2.$$

This conclusion holds for every realized collection of learned disentangling unitaries, whether exact or approximate.

Proof. Fix a pair (j, i) for which Algorithm 2 performs tomography. Define the indices of the first-layer blocks contained in the subtree rooted at (j, i) by

$$S_i^j \coloneqq \left\{ (i-1)2^{j-1} + 1, \ldots, i2^{j-1} \right\}.$$

The corresponding ancestor interval in the original chain is

$$\mathscr{I}_i^j \coloneqq \bigcup_{k \in S_i^j} \mathscr{B}_k^1.$$

Since the first-layer blocks $\mathscr{B}_k^1$ are consecutive and ordered along the original chain, $\mathscr{I}_i^j$ is a contiguous interval. Although the surviving block $\mathscr{B}_i^j$ need not be contiguous in the original chain, it consists precisely of the registers in $\mathscr{I}_i^j$ that remain immediately before the node (j, i) is processed.

We next describe the successful-branch map associated with each earlier node. For a processed node (a, k), let

$$\mathscr{D}_k^a \coloneqq \mathscr{B}_k^a \setminus \widetilde{\mathscr{B}}_k^a$$

be the registers discarded at that node, and define

$$F_k^a \coloneqq \left(\langle 0^{f(a,k)}|_{\mathscr{D}_k^a} \otimes I_{\widetilde{\mathscr{B}}_k^a} \right) U_k^a.$$

Thus, F_k^a combines the learned unitary U_k^a, projection onto the successful all-zero outcome, and removal of the projected registers.

At the first layer, the blocks with $k > \ell_1$ are not processed. For these blocks, we use the conventions

$$\widetilde{\mathscr{B}}_k^1 = \mathscr{B}_k^1, \qquad F_k^1 = I_{\mathscr{B}_k^1}.$$

Define the complete successful-branch map at the first layer by

$$F^1 \coloneqq \bigotimes_{k=1}^{2^{M-1}} F_k^1,$$

and, for $a \geq 2$, define

$$F^a \coloneqq \bigotimes_{k=1}^{2^{M-a}} F_k^a.$$

The cumulative successful-branch map through layer $j-1$ is

$$F^{<j} \coloneqq F^{j-1} \circ \cdots \circ F^1,$$

with the convention $F^{<1} = I$.

By the tree-consistent construction in Algorithm 2, no block from an earlier layer is split between the subtree rooted at (j,i) and its complement. Therefore, every map applied before (j,i) is processed acts either entirely on registers descended from $\mathscr{I}_i^j$ or entirely on registers descended from $[n] \setminus \mathscr{I}_i^j$. Consequently, the cumulative map factorizes as

$$F^{<j} = K_i^{<j} \otimes L_i^{<j}$$

across the original bipartition

$$\mathscr{I}_i^j : [n] \setminus \mathscr{I}_i^j,$$

where $K_i^{<j}$ and $L_i^{<j}$ are linear maps acting on the two respective sides.

Because $\mathscr{I}_i^j$ is a contiguous interval, Lemma A.1 implies that the original MPS $|\psi\rangle$ has Schmidt rank at most D^2 across this bipartition. Hence there exist vectors $|x_\alpha\rangle_{\mathscr{I}_i^j}$ and $|y_\alpha\rangle_{[n]\setminus\mathscr{I}_i^j}$ such that

$$|\psi\rangle = \sum_{\alpha=1}^{R} |x_\alpha\rangle_{\mathscr{I}_i^j} \otimes |y_\alpha\rangle_{[n]\setminus\mathscr{I}_i^j}, \qquad R \leq D^2.$$

After all successful-branch operations through layer $j-1$, the subnormalized residual vector is

$$|\Psi^{j-1}\rangle \coloneqq F^{<j} |\psi\rangle .$$

Using the factorization of $F^{<j}$, we obtain

$$|\Psi^{j-1}\rangle = \sum_{\alpha=1}^{R} K_i^{<j} |x_\alpha\rangle \otimes L_i^{<j} |y_\alpha\rangle .$$

This decomposition contains at most $R \leq D^2$ product terms. Therefore, local linear maps cannot increase the Schmidt rank, and $|\Psi^{j-1}\rangle$ has Schmidt rank at most D^2 across the bipartition

$$\mathscr{B}_i^j : \mathscr{B}^j \setminus \mathscr{B}_i^j.$$

Since the input state is pure and we have conditioned on a fixed sequence of successful projection outcomes, the subnormalized residual operator is

$$(\rho^{j-1})' = |\Psi^{j-1}\rangle\langle\Psi^{j-1}| .$$

For a possibly unnormalized pure state, the rank of its reduced operator equals its Schmidt rank across the corresponding bipartition. It follows that

$$\begin{aligned} \operatorname{rank}(\sigma_i^j) &= \operatorname{rank}\Big(\operatorname{Tr}_{\mathscr{B}^j\setminus\mathscr{B}_i^j}\big[(\rho^{j-1})'\big]\Big) \\ &\leq D^2. \end{aligned}$$

The proof uses only the locality of the maps F_k^a and the tree-consistent block construction. It does not require the learned unitaries U_k^a to be exact. Therefore, the invariant holds for both exact and approximate learned disentangling unitaries. □

B.2 Projection Identities and Local Estimation Bounds

Definition 6 (Level-j transformed and projected states)**.** *For any unit vector $|\phi\rangle$ on the original n-qudit system, set $|\phi^j\rangle = E^j |\phi\rangle$ and define*

$$\Pi^{(j)} := \left(\bigotimes_{a=1}^{j} Z^a\right) \otimes I_{\mathscr{B}^{j+1}}, \qquad \Pi^{(0)} = I,$$

where $\mathscr{B}^{j+1}$ is the set of surviving registers after layer j (and $\mathscr{B}^1 = [n]$). The level-j projected vector has the form

$$\Pi^{(j)} |\phi^j\rangle = \left(\bigotimes_{a=1}^{j} |0^{|\mathscr{D}^a|}\rangle_{\mathscr{D}^a}\right) \otimes |\psi^j\rangle ,$$

which defines the possibly subnormalized residual vector $|\psi^j\rangle$ on $\mathscr{B}^{j+1}$. For $j = 0$, set $|\psi^0\rangle = |\phi\rangle$.

Fact B.1. *For any state $|\phi\rangle$ and for each $j = 0, \ldots, M$, we have $\langle\phi| \rho^j |\phi\rangle = \langle\psi^j| (\rho^j)' |\psi^j\rangle$.*

Definition 7 (Top-Eigenspace W_i^j of the Estimated Reduced State)**.** *Fix a block (j, i) processed by Algorithm 2. Define the subnormalized reduced operator*

$$\sigma_i^j = \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[(\rho^{j-1})'].$$

When $q_{j-1} > 0$, define its normalized conditional counterpart by

$$\bar{\sigma}_i^j = \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[\bar{\rho}^{j-1}], \qquad \sigma_i^j = q_{j-1} \bar{\sigma}_i^j.$$

Let $\hat{\sigma}_i^j$ be the PSD, trace-at-most-one estimate produced by the subnormalized tomography procedure. It combines conditional-state tomography with an independently observed frequency of the entire successful history, or returns zero in the branches specified by the lemma. All input copies are fresh original copies of ρ.

For every processed block, define W_i^j as the span of the leading d^p eigenvectors of $\hat{\sigma}_i^j$, and let $\Pi_{W_i^j}$ be its orthogonal projector. For unprocessed first-layer blocks $i > \ell_1$, define instead

$$W_i^1 = \mathbb{H}_{\mathscr{B}_i^1}, \qquad \Pi_{W_i^1} = I_{\mathscr{B}_i^1}.$$

These blocks are retained unchanged, and their leakage is zero.

Fact B.2. *With Q^j as defined above and with the identity convention for unprocessed first-layer blocks,*

$$(U^j)^\dagger Q^j U^j = \bigotimes_{i=1}^{2^{M-j}} \Pi_{W_i^j}.$$

Fact B.3. *For each $j = 1, \ldots, M$, the orthogonal complement of $\Pi_{W_1^j} \otimes \cdots \otimes \Pi_{W_{2^{M-j}}^j}$ can be decomposed as*

$$I - \bigotimes_{k=1}^{2^{M-j}} \Pi_{W_k^j} = \sum_{i=1}^{2^{M-j}} \left(\bigotimes_{k=1}^{i-1} \Pi_{W_k^j}\right) \otimes \left(I_{\mathscr{B}_i^j} - \Pi_{W_i^j}\right) \otimes \left(\bigotimes_{k=i+1}^{2^{M-j}} I_{\mathscr{B}_k^j}\right),$$

where $I_{\mathscr{B}_i^j}$ denotes the identity operator on the Hilbert space of the block $\mathscr{B}_i^j$.

Here, we adopt the convention that the tensor product over an empty index set (i.e., when the upper limit is less than the lower limit) is defined to be the identity operator on the trivial (one-dimensional) Hilbert space.

Lemma B.2. *Let $\sigma, \hat{\sigma}$ be positive semidefinite operators on an h-dimensional space, each with trace at most one, where $h \geq D^2$ and $\mathrm{rank}(\sigma) \leq D^2$. Choose a complete orthonormal eigenbasis of $\hat{\sigma}$, including zero eigenvalues. Let W be the span of its leading D^2 eigenvectors, and let Π_W be the orthogonal projector onto W. If $\|\sigma - \hat{\sigma}\|_1 \leq \eta$, then*

$$\mathrm{Tr}\big[(I - \Pi_W)\, \sigma\big] \leq 2\eta.$$

Proof. Write the spectral decomposition of $\hat{\sigma}$ as

$$\hat{\sigma} = \sum_{i=1}^{h} \hat{\sigma}_i \left|v_i\right\rangle \left\langle v_i\right|,$$

with eigenvalues sorted in descending order $\hat{\sigma}_1 \geq \hat{\sigma}_2 \geq \cdots \geq \hat{\sigma}_h \geq 0$. Define the best rank-$D^2$ approximation of $\hat{\sigma}$ as

$$\hat{\sigma}_{\mathrm{tr}} = \sum_{i=1}^{D^2} \hat{\sigma}_i \left|v_i\right\rangle \left\langle v_i\right|.$$

By the Eckart–Young–Mirsky theorem, $\hat{\sigma}_{\mathrm{tr}}$ minimizes the trace norm error among all rank-D^2 matrices. Since σ has rank at most D^2,

$$\|\hat{\sigma}_{\mathrm{tr}} - \hat{\sigma}\|_1 \leq \|\sigma - \hat{\sigma}\|_1 \leq \eta.$$

Note that $\|I - \Pi_W\|_\infty = 1$. We compute

$$\mathrm{Tr}[(I - \Pi_W)\sigma] = \mathrm{Tr}\left[(I - \Pi_W)(\sigma - \hat{\sigma})\right] + \mathrm{Tr}\left[(I - \Pi_W)\hat{\sigma}\right].$$

For the first term, by Hölder (duality of Schatten norms),

$$\left|\,\mathrm{Tr}\left[(I - \Pi_W)(\sigma - \hat{\sigma})\right]\right| \leq \|I - \Pi_W\|_\infty \, \|\sigma - \hat{\sigma}\|_1 \leq \eta.$$

For the second term we use positivity of $\hat{\sigma} - \hat{\sigma}_{\mathrm{tr}}$ (indeed $\hat{\sigma} - \hat{\sigma}_{\mathrm{tr}} = \sum_{i>D^2} \hat{\sigma}_i \left|v_i\right\rangle \left\langle v_i\right| \succeq 0$) to get

$$\mathrm{Tr}[(I - \Pi_W)\hat{\sigma}] = \mathrm{Tr}\left[\hat{\sigma} - \hat{\sigma}_{\mathrm{tr}}\right] = \|\hat{\sigma} - \hat{\sigma}_{\mathrm{tr}}\|_1 \leq \eta.$$

Combining the two estimates yields

$$\mathrm{Tr}[(I - \Pi_W)\sigma] \leq \eta + \eta = 2\eta,$$

as desired. □

Lemma B.3. *Fix $j \in \{1, \ldots, M\}$. For each $i = 1, \ldots, 2^{M-j}$, let $\mathbb{H}_{\mathscr{B}_i^j}$ be the Hilbert space associated with $\mathscr{B}_i^j$, where $\mathscr{B}_i^j$ is defined in Algorithm 2. Let $W_i^j \subseteq \mathbb{H}_{\mathscr{B}_i^j}$ be a subspace defined in Definition 7, and let $\Pi_{W_i^j}$ denote the orthogonal projector onto W_i^j. Let $\mathscr{B}^j \coloneqq \bigcup_{k=1}^{2^{M-j}} \mathscr{B}_k^j$ be the collection of all qudits under consideration at step j. Let $(\rho^{j-1})'$ denote the state on $\mathbb{H}_{\mathscr{B}^j}$ defined in Algorithm 2, and let $\left|\varphi\right\rangle$ be a unit vector in the same Hilbert space. Define the product projector $\Pi_{W^j} \coloneqq \bigotimes_{i=1}^{2^{M-j}} \Pi_{W_i^j}$. Suppose that*

$$(I - \Pi_{W^j}) \left|\varphi\right\rangle = \left|\varphi\right\rangle,$$

and that for each $i = 1, \ldots, 2^{M-j}$,

$$\mathrm{Tr}\Big[(I - \Pi_{W_i^j}) \,\mathrm{Tr}_{(\mathscr{B}^j \setminus \mathscr{B}_i^j)}\big[(\rho^{j-1})'\big]\Big] \leq 2\eta.$$

Then

$$\left\langle\varphi\right| (\rho^{j-1})' \left|\varphi\right\rangle \leq 2\eta \, 2^{M-j}.$$

Proof. Let $X = (\rho^{j-1})'$ and $V = \Pi_{W^j}$. Since $V \left|\varphi\right\rangle = 0$ and $\|\varphi\| = 1$, we have $0 \preceq \left|\varphi\right\rangle\!\left\langle\varphi\right| \preceq I - V$. By Fact B.3, the orthogonal first-failure projectors F_i sum to $I - V$, where

$$F_i = \left(\bigotimes_{k<i} \Pi_{W_k^j}\right) \otimes (I_{\mathscr{B}_i^j} - \Pi_{W_i^j}) \otimes \left(\bigotimes_{k>i} I_{\mathscr{B}_k^j}\right).$$

Each $F_i \preceq (I_{\mathscr{B}_i^j} - \Pi_{W_i^j}) \otimes I_{\mathscr{B}^j \setminus \mathscr{B}_i^j}$. Since $X \succeq 0$, the assumptions give

$$\begin{aligned}\left\langle\varphi\right| X \left|\varphi\right\rangle &\leq \mathrm{Tr}[(I - V)X] = \sum_i \mathrm{Tr}[F_i X] \\ &\leq \sum_i \mathrm{Tr}\Big[(I_{\mathscr{B}_i^j} - \Pi_{W_i^j}) \;\mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j} X\Big] \leq 2\eta \, 2^{M-j}.\end{aligned}$$

This argument also covers zero-weight components and unprocessed first-layer blocks, without normalizing any zero vector. □

B.3 Layerwise Error Propagation

Lemma B.4. *Fix $j \in \{1, \dots, M\}$. For each $i = 1, \dots, 2^{M-j}$, let $\mathbb{H}_{\mathscr{B}_i^j}$ denote the Hilbert space associated with $\mathscr{B}_i^j$, where $\mathscr{B}_i^j$ is defined in Algorithm 2. Let $W_i^j \subseteq \mathbb{H}_{\mathscr{B}_i^j}$ be a subspace defined in Definition 7, and denote by $\Pi_{W_i^j}$ the orthogonal projector onto W_i^j. Let $\mathscr{B}^j := \bigcup_{k=1}^{2^{M-j}} \mathscr{B}_k^j$ be the set of all qudits under consideration at step j. Let $(\rho^{j-1})'$ denote the state on $\mathbb{H}_{\mathscr{B}^j}$ defined in Algorithm 2, and let ρ^j be the state given in Definition 5.*

Suppose that, for each $i = 1, \dots, 2^{M-j}$, the tomography procedure succeeds in the sense that

$$\operatorname{Tr}\Big[(I - \Pi_{W_i^j}) \operatorname{Tr}_{(\mathscr{B}^j \setminus \mathscr{B}_i^j)}\big[(\rho^{j-1})'\big]\Big] \le 2\eta.$$

Then, for any pure state $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$,

$$\big|\, \langle\phi|\, \rho^{j-1} \,|\phi\rangle - \langle\phi|\, \rho^j \,|\phi\rangle \,\big| \;\le\; 2\sqrt{2\eta\; 2^{M-j}}\,.$$

Proof. Write $X = (\rho^{j-1})'$, $V = \Pi_{W^j}$, and $|v\rangle = |\psi^{j-1}\rangle$ for the level-$(j-1)$ residual vector of the unit input $|\phi\rangle$. Thus $\|v\| \le 1$ and $0 \preceq X \preceq I$, since X is subnormalized. The well-typed postselection identity on the current registers is

$$Z^j \otimes (\rho^j)' = Q^j U^j X (U^j)^\dagger Q^j. \tag{2}$$

Using the reconstruction definition, the residual-vector representation in Fact B.1, and Fact B.2, we obtain

$$\langle\phi|\, \rho^{j-1} \,|\phi\rangle = \langle v|\, X \,|v\rangle\,, \qquad \langle\phi|\, \rho^j \,|\phi\rangle = \langle v|\, VXV \,|v\rangle\,.$$

The latter identity follows by applying U^j only to the surviving registers in the level-$(j-1)$ projected vector, inserting the all-zero outcomes for $\mathscr{D}^j$, and using $(U^j)^\dagger Q^j U^j = V$. Set $|w\rangle = (I - V)\,|v\rangle$. By Cauchy–Schwarz,

$$\begin{aligned} |\, \langle v|\, X \,|v\rangle - \langle v|\, VXV \,|v\rangle \,| &\le |\, \langle v|\, X \,|w\rangle \,| + |\, \langle w|\, XV \,|v\rangle \,| \\ &\le 2\|X \,|w\rangle\,\|. \end{aligned}$$

If $w = 0$, the difference is zero. Otherwise put $|\varphi\rangle = |w\rangle\,/\|w\|$; then $(I - V)\,|\varphi\rangle = |\varphi\rangle$. As $X^2 \preceq X$ and $\|w\| \le 1$, Lemma B.3 gives

$$2\|X \,|w\rangle\,\| \le 2\sqrt{\langle\varphi|\, X^2 \,|\varphi\rangle} \le 2\sqrt{\langle\varphi|\, X \,|\varphi\rangle} \le 2\sqrt{2\eta\, 2^{M-j}}.$$

□

B.4 Proof of Theorem 1 and Sample Complexity

Theorem (Restatement of Theorem 1)**.** *For any integers $n, d \ge 2$, given access to copies of an n-qudit matrix product state ρ with bond dimension D and parameters $\epsilon, \delta \in (0, 1)$, Algorithm 2 outputs a classical description of a preparation circuit for a pure state $|\hat{\phi}\rangle$, such that, with probability $1 - \delta$, $\langle\hat{\phi}|\, \rho \,|\hat{\phi}\rangle \;\ge\; 1 - \epsilon$. The algorithm requires $N = O(\frac{D^6 \cdot d^4 \cdot n^3 \cdot \log(n/\delta)}{(\max\{1, \log_d D\})^3 \epsilon^4})$ copies of ρ and runs in time $\operatorname{poly}(D, n, 1/\epsilon, \log(1/\delta))$ for fixed local dimension d. The output circuit uses $O(\log n)$ layers of block unitaries.*

Proof. By Lemma B.1, every reduced operator

$$\sigma_i^j = \operatorname{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}\big[(\rho^{j-1})'\big]$$

defined in Algorithm 2 satisfies

$$\operatorname{rank}(\sigma_i^j) \le D^2.$$

Condition on the preceding learning history at each tomography call. The successful-branch map is then fixed, and Lemma A.2 applies with the output-rank bound above. There are at most 2^M block and final tomography calls. Assigning failure probability $\zeta = \delta/(2^M + 1)$ to each and using a union

bound, all estimates satisfy their stated bounds simultaneously with probability at least $1-\delta$. On this event, each block estimate $\hat{\sigma}_i^j$ satisfies

$$\|\sigma_i^j - \hat{\sigma}_i^j\|_1 \;\le\; \frac{(\sqrt{2}-1)^2\epsilon^2}{2^{M+5}}.$$

In the next step, Algorithm 2 invokes Algorithm 1 to select the D^2 largest eigenvectors from the estimated subnormalized PSD operator $\hat{\sigma}_i^j$ and constructs a disentangling unitary accordingly. Let $\widetilde{W}_i^j$ denote the subspace spanned by the D^2 largest eigenvectors. From Lemma B.2, we know that

$$\mathrm{Tr}\Big[(I - \Pi_{\widetilde{W}_i^j})\,\sigma_i^j\Big] \le \frac{(\sqrt{2}-1)^2\epsilon^2}{2^{M+4}}.$$

Since $\widetilde{W}_i^j$ is contained in the space spanned by the d^p largest eigenvectors of $\hat{\sigma}_i^j$, we have $\widetilde{W}_i^j \subseteq W_i^j$, where W_i^j is defined in Definition 7. Equivalently, their orthogonal complements satisfy

$$(W_i^j)^\perp \subseteq (\widetilde{W}_i^j)^\perp$$

Therefore, it follows that

$$\mathrm{Tr}\Big[(I - \Pi_{W_i^j})\,\sigma_i^j\Big] \le \frac{(\sqrt{2}-1)^2\epsilon^2}{2^{M+4}}.$$

Let $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$ be a pure state on an n-qudit system. Applying Lemma B.4, we get

$$\big|\,\langle\phi|\,\rho^{j-1}\,|\phi\rangle - \langle\phi|\,\rho^{j}\,|\phi\rangle\,\big| \;\le\; \frac{(\sqrt{2}-1)\epsilon}{2\sqrt{2^j}}\,.$$

Using the triangular inequality, we derive

$$\big|\,\langle\phi|\,\rho^{0}\,|\phi\rangle - \langle\phi|\,\rho^{M}\,|\phi\rangle\,\big| \;\le\; \sum_{j=1}^{M}\frac{(\sqrt{2}-1)\epsilon}{2\sqrt{2^j}} \le \frac{\epsilon}{2},$$

which implies

$$\langle\phi|\,\rho^M\,|\phi\rangle \ge \langle\phi|\,\rho\,|\phi\rangle - \frac{\epsilon}{2}. \tag{3}$$

Recall that $\rho^0 \coloneqq \rho$ by Definition 5. For any normalized state $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$, we define its (possibly subnormalized) level-j projected vectors $|\psi^j\rangle \in \mathbb{H}_{\widetilde{\mathscr{B}}_1^j} \otimes \cdots \otimes \mathbb{H}_{\widetilde{\mathscr{B}}_{2^{M-j}}^j}$ through Definition 6. By Fact B.1, we have

$$\langle\phi|\,\rho^M\,|\phi\rangle = \langle\psi^M|\,(\rho^M)'\,|\psi^M\rangle\,. \tag{4}$$

On the other hand, for the reconstructed state $|\hat{\phi}\rangle = (E^M)^\dagger\left(\bigotimes_{j=1}^{M}\bigotimes_{i=1}^{2^{M-j}}|0^{f(j,i)}\rangle_{\mathscr{B}_i^j\setminus\widetilde{\mathscr{B}}_i^j}\otimes|\hat{\psi}\rangle\right)$, since $(\rho^M)' = \left(\bigotimes_{j=1}^{M}\bigotimes_{i=1}^{2^{M-j}}\langle 0^{f(j,i)}|_{\mathscr{B}_i^j\setminus\widetilde{\mathscr{B}}_i^j}\right)E^M\rho(E^M)^\dagger\left(\bigotimes_{j=1}^{M}\bigotimes_{i=1}^{2^{M-j}}|0^{f(j,i)}\rangle_{\mathscr{B}_i^j\setminus\widetilde{\mathscr{B}}_i^j}\right)$, we have

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle = \langle\hat{\psi}|\,(\rho^M)'\,|\hat{\psi}\rangle \tag{5}$$

If the final subnormalized tomography procedure succeeds, combining conditional tomography on $\bar{\rho}^M$ with the observed frequency of all M projections succeeding (or returning zero as specified by the lemma), it yields a PSD estimate $(\hat{\rho}^M)'$ such that $\|(\rho^M)' - (\hat{\rho}^M)'\|_1 \le \epsilon/4$. By Hölder's inequality for Schatten norms, for any vector $|\varphi\rangle \in \mathbb{H}_{\widetilde{\mathscr{B}}_1^M}$ with $\|\varphi\| \le 1$,

$$\big|\langle\varphi|\,(\rho^M)'\,|\varphi\rangle - \langle\varphi|\,(\hat{\rho}^M)'\,|\varphi\rangle\big| \le \epsilon/4. \tag{6}$$

Applying this bound gives

$$\langle\hat{\psi}|\,(\rho^M)'\,|\hat{\psi}\rangle \ge \langle\hat{\psi}|\,(\hat{\rho}^M)'\,|\hat{\psi}\rangle - \epsilon/4 \ge \langle\psi^M|\,(\hat{\rho}^M)'\,|\psi^M\rangle - \epsilon/4 \ge \langle\psi^M|\,(\rho^M)'\,|\psi^M\rangle - \epsilon/2. \tag{7}$$

Equation (6) gives the first and last inequalities. For the middle inequality, $|\psi^M\rangle$ is a projected, possibly subnormalized vector with $\|\psi^M\| \le 1$, whereas $(\hat{\rho}^M)' \succeq 0$ and $|\hat{\psi}\rangle$ is its normalized top eigenvector. Hence its expectation on $|\psi^M\rangle$ cannot exceed its largest eigenvalue. Combining with Equations (4) and (5), we obtain

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \ge \langle\phi|\,\rho^M\,|\phi\rangle - \frac{\epsilon}{2}. \tag{8}$$

Together with Equation (3), the desired inequality is established:

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \geq \langle\phi|\,\rho\,|\phi\rangle - \epsilon \tag{9}$$

Notice that this inequality holds for any pure state $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$. In particular, by choosing $|\phi\rangle$ to be our target MPS(D) state, we obtain the maximal fidelity $\langle\phi|\,\rho\,|\phi\rangle = 1$.

Next, we compute the sample complexity of Algorithm 2. At each block call, the successful branch consists of replaying all previously learned layers on a fresh original copy of ρ and retaining the current block only if every preceding projection succeeds. The output dimension is at most d^{2p}, and Lemma B.1 bounds its rank by D^2. Lemma A.2, with $\eta = (\sqrt{2}-1)^2\epsilon^2/2^{M+5}$ and $\zeta = \delta/(2^M+1)$, therefore uses at most

$$O\left(\frac{D^2 d^{2p}}{\eta^2}\log\frac{2}{\zeta}\right) = O\left(\frac{D^2 d^{2p} 2^{2M}}{\epsilon^4}\log\frac{n}{\delta}\right)$$

original copies per block, including frequency-estimation copies and all failed postselection trials. The number of processed blocks is $\ell_1 + \sum_{j=2}^{M} 2^{M-j} \leq 2^M - 1$. Thus the total block-tomography budget is

$$O\left(\frac{D^2 d^{2p} 2^{3M}}{\epsilon^4}\log\frac{n}{\delta}\right).$$

For $M \geq 1$, the final residual output has dimension d^p and rank at most one, so the same lemma with accuracy $\tau = \epsilon/4$ uses $O(d^p\epsilon^{-2}\log(n/\delta))$ original copies. This is dominated by the block-tomography budget. Using $d^{2p} \leq d^4D^4$ and $2^M \leq 2n/p$ for $M \geq 1$ yields the desired bound

$$O\left(\frac{D^6 d^4 n^3 \log(n/\delta)}{(\max\{1,\log_d D\})^3\epsilon^4}\right).$$

If $M = 0$, equivalently $n \leq p$, there are no learned layers and the only call is final tomography of the original pure state on n qudits. Lemma A.2 with rank bound one gives $O(d^n\epsilon^{-2}\log(2/\delta))$ copies. Here $d^n \leq d^p \leq d^2D^2$, and $(\max\{1,\log_d D\})^3 = O(D^4)$ for $D \geq 1$ and $d \geq 2$. Thus the $M = 0$ cost is also bounded by the displayed theorem bound, using $n \geq 1$, $\epsilon < 1$, and $\log(2/\delta) \leq \log(n/\delta)$ when $n \geq 2$. No additional factor $1/q_{j-1}$ is needed: the fixed-budget procedure in Lemma A.2 already includes unsuccessful trials and relaxes the conditional-state accuracy in proportion to η/q_{j-1} when the success probability is not negligible. When it is small, the lemma controls the error by returning zero. For every block, $h \leq d^{2p} \leq d^4D^4$ and $R \leq D^2$. By Lemma A.2 and the gate-efficient normalized-state tomography of [PSTW25], each sampling and tomography call uses polynomially many quantum gates and classical operations in $D, d, n, 1/\epsilon, \log(1/\delta)$, excluding the cost of supplying fresh input copies. Replaying at most $M = O(\log n)$ layers and diagonalizing the polynomial-size block matrices also take polynomial time. Thus the entire learning procedure has polynomial runtime (for fixed local dimension d), assuming standard polynomial-precision implementations of the learned block unitaries and classical linear-algebra operations. The $M = O(\log n)$ depth counts layers of disjoint block unitaries in the circuit output by the algorithm, not the depth of collective tomography measurements. For fixed d, D, the block unitaries and the final p-register seed preparation have bounded dimension, so this also yields $O(1+\log n)$ elementary-gate depth; no uniform elementary-gate depth claim for growing D follows solely from the layer count. □

Appendix C: Closest-MPS Learning

C.1 Algorithms and Spectral Thresholding

In this section, we consider the learning task where the target state ρ is not necessarily an exact matrix product state (MPS) of bond dimension D, as assumed in the previous sections, and may even be a mixed state. We define the maximum fidelity between ρ and and the set MPS(D) of MPS of bond dimension D as

$$\max_{|\phi\rangle\in\text{MPS}(D)} \langle\phi|\,\rho\,|\phi\rangle\,.$$

Our goal is to output a description of state $|\hat{\phi}\rangle$ such that, given an error parameter $\epsilon > 0$ and access to copies of ρ, this state is close enough to the closest matrix product state of bond dimension D in terms of fidelity

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \geq \max_{|\phi\rangle\in\text{MPS}(D)} \langle\phi|\,\rho\,|\phi\rangle - \epsilon.$$

This is an agnostic fidelity benchmark: the input need not belong to the target class, and no lower bound on the optimal achievable fidelity is promised. We use the term "closest-MPS learning" to emphasize the comparator in the displayed inequality.

Algorithm 3 Disentangling unitary construction (threshold η)

Input: A threshold $\eta > 0$ and a description of a positive semidefinite operator $\hat{\sigma} \in \mathbb{C}^{d^y \times d^y}$ with $\operatorname{Tr}\hat{\sigma} \leq 1$.
Output: A description of a Disentangling Unitary U for the state $\hat{\sigma}$ (constructed from eigenvectors whose eigenvalues exceed η).
Procedure:

1: Perform the spectral decomposition of $\hat{\sigma}$:

$$\hat{\sigma} = \sum_{i=1}^{r} a_i \left|\hat{\phi}_i\right\rangle \left\langle\hat{\phi}_i\right|,$$

where $r \leq d^y$.
2: Define the index set of selected eigenvalues:

$$S \coloneqq \{\, i \in \{1, \ldots, r\} \mid a_i > \eta \,\}.$$

3: Let $m = |S|$. If $m > 0$, then $m\eta < \sum_{i \in S} a_i \leq \operatorname{Tr}\hat{\sigma} \leq 1$, so $m < 1/\eta$.
4: Reindex the selected eigenvectors as $\{|\phi_i\rangle\}_{i=1}^{m}$, where $|\phi_i\rangle \coloneqq |\hat{\phi}_{s_i}\rangle$ for $s_i \in S$.
5: Extend $\{|\phi_i\rangle\}_{i=1}^{m}$ to a full orthonormal basis $\{|\phi_i\rangle\}_{i=1}^{d^y}$ of $(\mathbb{C}^d)^{\otimes y}$ by choosing arbitrary orthonormal vectors for indices $i = m+1, \ldots, d^y$.
6: Set $t = \lceil \log_d \max\{1, m\} \rceil$. If $m = 0$, choose an arbitrary complete orthonormal basis in the preceding extension step; this makes the unitary construction well-defined.
7: Let $\{|a_1, \ldots, a_{y-t}\rangle\}_{a_i \in \{0,\ldots,d-1\}}$ be the standard basis for the first $y - t$ qudits, and let $\{|j\rangle\}_{j=1}^{d^t}$ be an arbitrary orthonormal basis for the remaining t qudits. Define the unitary

$$U = \sum_{a_1,\ldots,a_{y-t}=0}^{d-1} \sum_{j=1}^{d^t} \left(|a_1, \ldots, a_{y-t}\rangle \otimes |j\rangle\right) \left\langle \phi_{\mathrm{idx}(a_1,\ldots,a_{y-t},j)} \right|,$$

where the index mapping is defined by $\mathrm{idx}(a_1, \ldots, a_{y-t}, j) = j + d^t \cdot \sum_{l=1}^{y-t} a_l d^{y-t-l}$.
8: **return** the unitary U.

The normalized conditional states, success probabilities, fresh-copy replay convention, and zero-success convention are the same as for Algorithm 2. Here all tomography calls use Lemma A.3, since the input need not be an MPS and the successful-branch reduced output can have full rank.

Algorithm 4 Learning the closest matrix product states

Input: Fresh copies of an arbitrary n-qudit state ρ, bond-dimension parameter D, and accuracy and failure probabilities $\epsilon, \delta \in (0, 1)$.
Output: A description of a quantum state $|\hat{\phi}\rangle \in (\mathbb{C})^{\otimes n}$.
Procedure:

1: Set $B = 128nD^2/((\sqrt{2}-1)^2\epsilon^2)$, $p = \lceil W(B \ln d)/\ln d \rceil$, and $\eta = p/B$ (see Appendix C.2.1).
2: **if** $p \geq n$ **then**
3: Estimate ρ directly by full-state tomography (Lemma A.3, $\mathscr{T} = \mathrm{id}$, $h = d^n$, accuracy $\epsilon/2$, failure probability δ).
4: **return** a normalized leading eigenvector of this estimate.
5: **end if**
6: Set $M = \lceil \log_2(n/p) \rceil$.
7: Set $r_1 = n - 2^{M-1}p$, $\ell_1 = \lceil r_1/p \rceil$, $s_1 = r_1 - (\ell_1 - 1)p$, and $k_1 = 2\ell_1 p - p + s_1$.
8: Define the support $\mathscr{B}_i^1$ set for the i-th disentangling unitaries in the first layer:

$$\mathscr{B}_i^1 \coloneqq \begin{cases} \{2(i-1)p+1, \ldots, 2ip\} & \text{if } 1 \leq i < \ell_1 \\ \{2(\ell_1-1)p+1, \ldots, k_1\} & \text{if } i = \ell_1 \\ \{k_1+1+(i-\ell_1-1)p, \ldots, k_1+(i-\ell_1)p\} & \text{if } \ell_1 < i \leq 2^{M-1} \end{cases}. \tag{10}$$

Define a function

$$f(j,i) = \begin{cases} s_1 & \text{when } j = 1, i = \ell_1 \\ 0 & \text{when } j = 1, i > \ell_1 \\ p & \text{otherwise} \end{cases}.$$

Set $(\rho^0)' = \bar{\rho}^0 = \rho$, $q_0 = 1$, and $\zeta = \delta/(2^M + 1)$.

Let $\mathscr{B}^1 = \bigcup_{i=1}^{2^{M-1}} \mathscr{B}_i^1$.

for i from 1 to ℓ_1 **do**

Define $\bar{\sigma}_i^1 = \sigma_i^1 = \mathrm{Tr}_{\mathscr{B}^1 \setminus \mathscr{B}_i^1}[\rho]$.

Apply Lemma A.3 with output dimension $h = d^{|\mathscr{B}_i^1|} \le d^{2p}$, accuracy η, and failure probability ζ. Using at most $O(d^{4p}\eta^{-2}\log(2/\zeta))$ fresh original copies of ρ, obtain a PSD estimate $\hat{\sigma}_i^1$ with trace at most one.

Use Algorithm 3 to generate the disentangling unitary U_i^1 from a description of a $(p + f(1,i))$-qudit state $\hat{\sigma}_i^1$ and a parameter η.

Let $\widetilde{\mathscr{B}}_i^1$ be the last p qudits in $\mathscr{B}_i^1$.

end for

Let $U^1 = \bigotimes_{i=1}^{\ell_1} U_i^1$.

Let $P^1 = \bigotimes_{i=1}^{\ell_1} \left(|0^{f(1,i)}\rangle\langle 0^{f(1,i)}|\right)_{\mathscr{B}_i^1 \setminus \widetilde{\mathscr{B}}_i^1} \otimes I_{\widetilde{\mathscr{B}}_i^1}$.

Extend U^1 and P^1 by the identity on all unprocessed first-layer blocks $\mathscr{B}_i^1$ with $i > \ell_1$.

Define the successful-branch operator for the analysis:

$$(\rho^1)' = \mathrm{Tr}_{\bigcup_{i=1}^{\ell_1}(\mathscr{B}_i^1 \setminus \widetilde{\mathscr{B}}_i^1)}\left[P^1 U^1 \rho (U^1)^\dagger P^1\right].$$

Set $q_1 = \mathrm{Tr}[(\rho^1)']$ and, when $q_1 > 0$, $\bar{\rho}^1 = (\rho^1)'/q_1$. Future conditional samples replay the learned layer on fresh original copies and retain only successful trials.

Let $\widetilde{\mathscr{B}}_i^1 = \mathscr{B}_i^1$ for $\ell_1 < i \le 2^{M-1}$.

for j from 2 to M **do**

Let $\mathscr{B}^j = \bigcup_{i=1}^{2^{M-j+1}} \widetilde{\mathscr{B}}_i^{j-1}$.

for i from 1 to 2^{M-j} **do**

Let $\mathscr{B}_i^j = \widetilde{\mathscr{B}}_{2i-1}^{j-1} \cup \widetilde{\mathscr{B}}_{2i}^{j-1}$.

Define $\sigma_i^j = \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[(\rho^{j-1})']$ and, when $q_{j-1} > 0$, $\bar{\sigma}_i^j = \sigma_i^j / q_{j-1}$.

On fresh original copies of ρ, replay the learned layers through $j-1$, postselect on all preceding projections succeeding, and discard registers outside $\mathscr{B}_i^j$. Apply Lemma A.3 with $h = d^{2p}$, accuracy η, and failure probability ζ to obtain $\hat{\sigma}_i^j$, using at most $O(d^{4p}\eta^{-2}\log(2/\zeta))$ original copies, including failed trials.

Use Algorithm 3 to generate the disentangling unitary U_i^j from a description of $2p$-qudit state $\hat{\sigma}_i^j$ and a parameter η.

Let $\widetilde{\mathscr{B}}_i^j$ be the last p qudits in $\mathscr{B}_i^j$.

end for

Let $U^j = \bigotimes_{i=1}^{2^{M-j}} U_i^j$.

Let $P^j = \bigotimes_{i=1}^{2^{M-j}} \left(|0^{f(j,i)}\rangle\langle 0^{f(j,i)}|\right)_{\mathscr{B}_i^j \setminus \widetilde{\mathscr{B}}_i^j} \otimes I_{\widetilde{\mathscr{B}}_i^j}$.

Define the successful-branch operator for the analysis:

$$(\rho^j)' = \mathrm{Tr}_{\bigcup_{i=1}^{2^{M-j}}(\mathscr{B}_i^j \setminus \widetilde{\mathscr{B}}_i^j)}\left[P^j U^j (\rho^{j-1})' (U^j)^\dagger P^j\right].$$

Set $q_j = \mathrm{Tr}[(\rho^j)']$ and, when $q_j > 0$, $\bar{\rho}^j = (\rho^j)'/q_j$. Future conditional samples replay all learned layers through j on fresh original copies and retain only successful trials.

end for

On fresh original copies of ρ, replay all M learned layers and retain the final p qudits only when every projection succeeds. Apply Lemma A.3 with $h = d^p$, accuracy $\tau = \epsilon/4$, and failure probability ζ to estimate $(\rho^M)'$ (equal to $q_M \bar{\rho}^M$ when $q_M > 0$) using at most $O(d^{2p}\tau^{-2}\log(2/\zeta))$ original copies, including failed trials.

Let $|\hat{\psi}\rangle$ be a normalized top eigenvector of the PSD estimate $(\hat{\rho}^M)'$, chosen arbitrarily if the estimate is zero.

return the state $|\hat{\phi}\rangle := (U^1)^\dagger \cdots (U^M)^\dagger \left(\bigotimes_{j=1}^{M} \bigotimes_{i=1}^{2^{M-j}} |0^{f(j,i)}\rangle_{\mathscr{B}_i^j \setminus \widetilde{\mathscr{B}}_i^j} \otimes |\hat{\psi}\rangle\right)$.

Remark 1. *The block size is determined by the minimal feasibility inequality $pd^p \geq B$, not by the generally unsolvable integer fixed point $p = \lceil \log_d(B/p) \rceil$. It therefore applies to arbitrary $\epsilon \in (0, 1)$. If $p \geq n$, the direct tomography branch avoids defining a tree with blocks larger than the entire system.*

C.2 Block-Size Parameter and Integer Feasibility

C.2.1 A Block-Size Choice Valid for Every Accuracy

Let $d \geq 2$, $n, D \geq 1$, and $\epsilon \in (0, 1)$, and set

$$B := \frac{128nD^2}{(\sqrt{2}-1)^2\epsilon^2}.$$

The algorithm needs enough room in its surviving p-qudit register for all eigenvectors of a trace-at-most-one estimate whose eigenvalues exceed the threshold

$$\eta := \frac{(\sqrt{2}-1)^2\epsilon^2 p}{128D^2n} = \frac{p}{B}.$$

It is sufficient that $d^p \geq 1/\eta = B/p$, or equivalently, $pd^p \geq B$. We therefore choose the *smallest positive integer* p satisfying

$$pd^p \geq B. \tag{11}$$

Since $x \mapsto xd^x$ is strictly increasing for $x > 0$, its positive real threshold is $a = W(B \ln d)/\ln d$, where W is the principal Lambert function. Thus the above integer choice is exactly

$$p = \left\lceil \frac{W(B \ln d)}{\ln d} \right\rceil. \tag{12}$$

Unlike the earlier self-consistent equality $p = \lceil \log_d(B/p) \rceil$, the condition (11) has a solution for *every* $\epsilon \in (0, 1)$; equality of those two integer expressions is neither assumed nor needed.

Dimension and asymptotic bounds. The choice gives $1/\eta \leq d^p$ and consequently the number m of eigenvalues exceeding η is strictly less than d^p. For $p \geq 2$, minimality implies $(p-1)d^{p-1} < B$, and therefore

$$d^p < \frac{dB}{p-1} = \frac{dp}{(p-1)\eta} \leq \frac{2d}{\eta}.$$

The same bound holds for $p = 1$ because $B > 1$, so $\eta < 1$. Consequently,

$$\frac{1}{\eta} \leq d^p \leq \frac{2d}{\eta}. \tag{13}$$

In particular, $p = \Theta(1 + W(B \ln d)/\ln d)$ and for fixed $d \geq 2$, $p = \Theta(\log_d B)$ as $B \to \infty$. The previously used inequality $d^p < d/\eta$ need not hold for arbitrary accuracy; (13) supplies the uniform bound needed for all big-O sample estimates.

C.2.2 Why the Original Integer Fixed Point Is Unnecessary

The stronger equation

$$p = \lceil \log_d(B/p) \rceil \tag{14}$$

holds if and only if $pd^{p-1} < B \leq pd^p$. These intervals leave gaps: for example, $d = 2$, $B = 9$ has no positive integer solution. Nevertheless, the minimal-feasibility rule (11) always defines p, even within those gaps. The analysis uses only $d^p \geq B/p$ and the constant-factor upper bound (13); no restriction to specially selected accuracy parameters is necessary.

C.3 Reconstruction and Auxiliary Bounds for Algorithm 4

Definition 8 (Step-wise Reconstructed State)**.** *Let $(\rho^0)' = \rho$ be the initial state before any iteration. Define*

$$\rho^0 \coloneqq (\rho^0)' = \rho.$$

For each $j = 1, \dots, M$, define the cumulative unitary

$$E^j \coloneqq U^j U^{j-1} \cdots U^1.$$

For each layer a, let $\mathscr{D}^a \coloneqq \bigcup_i (\mathscr{B}_i^a \setminus \widetilde{\mathscr{B}}_i^a)$ be its discarded registers and write $Q^a = |0_{\mathscr{D}^a}\rangle\langle 0_{\mathscr{D}^a}|$ for the all-zero projector acting only on those registers. Thus the full-layer projector in Algorithm 4 is $P^a = Q^a \otimes I$ (with the identity on the surviving registers). Let $(\rho^j)'$ denote the subnormalized post-measurement operator on the surviving registers after j successful layers. Define

$$\rho^j \coloneqq (E^j)^\dagger \left[\left(\bigotimes_{a=1}^{j} Q^a \right) \otimes (\rho^j)' \right] E^j,$$

where tensor factors follow the physical register order and the unitaries are extended by identity on previously discarded registers.

We call ρ^j the step-wise reconstructed state *corresponding to the j-th layer of the circuit, representing the state that would have been reconstructed had the algorithm stopped after the j-th iteration.*

For $j = 0, \dots, M$, define

$$q_j \coloneqq \mathrm{Tr}[(\rho^j)'], \qquad \bar{\rho}^j \coloneqq (\rho^j)'/q_j \quad (q_j > 0).$$

Here $q_0 = 1$, and q_j is the probability that all projections through layer j succeed on one original copy of ρ, conditioned on the fixed learned unitaries. The operators $(\rho^j)'$ and ρ^j are subnormalized operators used in the analysis, whereas $\bar{\rho}^j$ is the normalized physical state in the successful branch. If $q_j = 0$, then $(\rho^j)' = 0$ and no normalized conditional state is required. The tomography procedure is defined in this case by Lemma A.2.

Definition 9 (Level-j transformed and projected states)**.** *Fix $j \in \{1, \dots, M\}$. For any pure state $|\phi\rangle$ on an n-qudit system define the* level-j transformed state

$$|\phi^j\rangle \coloneqq E^j |\phi\rangle,$$

where E^j is the cumulative unitary introduced above. Define the level-j projector

$$\Pi^{(j)} \coloneqq \left(\bigotimes_{a=1}^{j} Q^a \right) \otimes I_{\mathrm{surv},j},$$

where $I_{\mathrm{surv},j}$ is the identity on the registers still present after layer j. The level-j projection *of $|\phi^j\rangle$ is $\Pi^{(j)} |\phi^j\rangle$.*

$\Pi^{(j)} |\phi^j\rangle$ has the product form

$$\Pi^{(j)} |\phi^j\rangle = \left(\bigotimes_{a=1}^{j} \bigotimes_{b=1}^{2^{M-a}} |0^{f(a,b)}\rangle_{\mathscr{B}_b^a \setminus \widetilde{\mathscr{B}}_b^a} \right) \otimes |\psi^j\rangle,$$

we call $|\psi^j\rangle \in \mathbb{H}_{\widetilde{\mathscr{B}}_1^j} \otimes \cdots \otimes \mathbb{H}_{\widetilde{\mathscr{B}}_{2^{M-j}}^j}$ the projected (residual) state of $|\phi\rangle$ at level j. *By convention, we define the projected state at level 0 to be the input state itself, $|\psi^0\rangle \coloneqq |\phi\rangle$.*

Fact C.1. *For any state $|\phi\rangle$ and for each $j = 0, \dots, M$, we have $\langle\phi| \rho^j |\phi\rangle = \langle\psi^j| (\rho^j)' |\psi^j\rangle$.*

Definition 10 (Retained Subspace W_i^j of the Estimated Reduced State)**.** *Fix a block (j, i) processed by Algorithm 4. Define the subnormalized reduced operator*

$$\sigma_i^j = \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[(\rho^{j-1})'].$$

When $q_{j-1} > 0$, define its normalized conditional counterpart by

$$\bar{\sigma}_i^j = \mathrm{Tr}_{\mathscr{B}^j \setminus \mathscr{B}_i^j}[\bar{\rho}^{j-1}], \qquad \sigma_i^j = q_{j-1} \bar{\sigma}_i^j.$$

Let $\hat{\sigma}_i^j$ be the PSD, trace-at-most-one estimate produced by the subnormalized tomography procedure. It combines conditional-state tomography with an independently observed frequency of the entire successful history, or returns zero in the branches specified by the lemma. All input copies are fresh original copies of ρ.

Let V_i^j denote the span of the eigenvectors of $\hat{\sigma}_i^j$ with eigenvalues exceeding η. Let W_i^j be the actual retained subspace *of dimension d^p selected by Algorithm 3: it contains V_i^j and is completed using the extended eigenbasis in that algorithm. Thus $V_i^j \subseteq W_i^j$, and we write $\Pi_{W_i^j}$ for its orthogonal projector. If a processed block contains only p registers, the retained subspace is the full block Hilbert space. For an unprocessed first-layer block ($j = 1$, $i > \ell_1$), set $W_i^1 = \mathbb{H}_{\mathscr{B}_i^1}$ and $\Pi_{W_i^1} = I$, corresponding to no projection.*

Fact C.2. *Let U^j and P^j as operators from Algorithm 4. By Definition 10 of W_i^j and $\Pi_{W_i^j}$, we have*

$$(U^j)^\dagger P^j U^j = \Pi_{W_1^j} \otimes \cdots \otimes \Pi_{W_{2^{M-j}}^j}.$$

Fact C.3. *For each $j = 1, \ldots, M$, the orthogonal complement of $\Pi_{W_1^j} \otimes \cdots \otimes \Pi_{W_{2^{M-j}}^j}$ can be decomposed as*

$$I - \bigotimes_{k=1}^{2^{M-j}} \Pi_{W_k^j} = \sum_{i=1}^{2^{M-j}} \left(\bigotimes_{k=1}^{i-1} \Pi_{W_k^j} \right) \otimes \left(I_{\mathscr{B}_i^j} - \Pi_{W_i^j} \right) \otimes \left(\bigotimes_{k=i+1}^{2^{M-j}} I_{\mathscr{B}_k^j} \right),$$

where $I_{\mathscr{B}_i^j}$ denotes the identity operator on the qudit set $\mathscr{B}_i^j$.

Here, we adopt the convention that the tensor product over an empty index set (i.e., when the upper limit is less than the lower limit) is defined to be the identity operator on the trivial (one-dimensional) Hilbert space.

Lemma C.1 (Lemma B.9 in [BBK+25])**.** *Let ρ, σ be unnormalized mixed states satisfying $\|\rho - \sigma\|_1 \leq \eta$. Let W be the span of all eigenvectors of σ with eigenvalues exceeding η. Then, letting Π_W denote orthogonal projection onto W, we have that $\|(I - \Pi_W)^\dagger \rho (I - \Pi_W)\|_\infty \leq 2\eta$.*

Lemma C.2. *Fix $j \in \{1, \ldots, M\}$. For each $i = 1, \ldots, 2^{M-j}$, let $\mathbb{H}_{\mathscr{B}_i^j}$ be the Hilbert space associated with $\mathscr{B}_i^j$. Let $(\rho^{j-1})'$ be a positive semidefinite operator with trace at most one on the Hilbert space associated with $\mathscr{B}^j$. For each i, let $W_i^j \subseteq \mathbb{H}_{\mathscr{B}_i^j}$ be a subspace with orthogonal projector $\Pi_{W_i^j}$. Let $|\chi\rangle$ be a (possibly subnormalized) vector with Schmidt rank at most D^2 across every bipartition $\mathscr{B}_i^j : (\mathscr{B}^j \setminus \mathscr{B}_i^j)$. Define $\Pi_{W^j} \coloneqq \bigotimes_{i=1}^{2^{M-j}} \Pi_{W_i^j}$. Assume $(I - \Pi_{W^j}) |\chi\rangle \neq 0$ and set*

$$|\varphi\rangle = \frac{(I - \Pi_{W^j}) |\chi\rangle}{\|(I - \Pi_{W^j}) |\chi\rangle\|}.$$

Suppose further that, for each $i = 1, \ldots, 2^{M-j}$,

$$\left\| (I - \Pi_{W_i^j}) \operatorname{Tr}_{(\mathscr{B}^j \setminus \mathscr{B}_i^j)} \left[(\rho^{j-1})' \right] (I - \Pi_{W_i^j}) \right\|_\infty \leq 2\eta.$$

Then

$$\langle \varphi | (\rho^{j-1})' | \varphi \rangle \leq 2\eta D^2 \, 2^{M-j}.$$

Proof. Let $m = 2^{M-j}$ and, for each $i = 1, \ldots, m$, define

$$Q_i^{\mathrm{leak}} = \left(\bigotimes_{k<i} \Pi_{W_k^j} \right) \otimes (I - \Pi_{W_i^j}) \otimes \left(\bigotimes_{k>i} I_{\mathscr{B}_k^j} \right).$$

By Fact C.3, these are mutually orthogonal projectors satisfying $\sum_i Q_i^{\mathrm{leak}} = I - \Pi_{W^j}$. Writing $|\varphi_i\rangle = Q_i^{\mathrm{leak}} |\varphi\rangle$ and $p_i = \|\varphi_i\|^2$, we have $|\varphi\rangle = \sum_i |\varphi_i\rangle$ and $\sum_i p_i = 1$. Terms with $p_i = 0$ are omitted below.

Since $Q_i^{\mathrm{leak}} \Pi_{W^j} = 0$, the vector $|\varphi_i\rangle$ is proportional to $Q_i^{\mathrm{leak}} |\chi\rangle$. This projector factorizes across $\mathscr{B}_i^j : (\mathscr{B}^j \setminus \mathscr{B}_i^j)$, so the normalized $|v_i\rangle = |\varphi_i\rangle / \sqrt{p_i}$ has Schmidt rank $r_i \leq D^2$ on

that cut, with its $\mathscr{B}_i^j$ Schmidt vectors in $(W_i^j)^\perp$. Write its Schmidt decomposition as $|v_i\rangle = \sum_{l=1}^{r_i} \sqrt{q_l}\,|x_l\rangle_{\mathscr{B}_i^j}\,|y_l\rangle_{\mathscr{B}^j\setminus\mathscr{B}_i^j}$, where $\sum_l q_l = 1$. With $\omega_i = \mathrm{Tr}_{\mathscr{B}^j\setminus\mathscr{B}_i^j}[(\rho^{j-1})']$, the assumed local operator bound gives

$$\langle x_l|\,\omega_i\,|x_l\rangle \le 2\eta, \qquad \langle x_l, y_l|\,(\rho^{j-1})'\,|x_l, y_l\rangle \le 2\eta.$$

Positivity, Cauchy–Schwarz, and $\sum_l q_l = 1$ imply

$$\langle v_i|\,(\rho^{j-1})'\,|v_i\rangle \le \left(\sum_{l=1}^{r_i} \sqrt{q_l}\sqrt{2\eta}\right)^2 \le 2\eta r_i \le 2\eta D^2.$$

Finally, applying Cauchy–Schwarz to the cross terms for $(\rho^{j-1})' \succeq 0$ yields

$$\begin{aligned}\langle\varphi|\,(\rho^{j-1})'\,|\varphi\rangle &\le 2\eta D^2 \left(\sum_{i=1}^{m} \sqrt{p_i}\right)^2 \\ &\le 2\eta D^2 m,\end{aligned}$$

as claimed. □

Lemma C.3. *Fix $j \in \{1,\dots,M\}$. For each $i = 1,\dots,2^{M-j}$, let $\mathbb{H}_{\mathscr{B}_i^j}$ denote the Hilbert space associated with $\mathscr{B}_i^j$, where $\mathscr{B}_i^j$ is defined in Algorithm 4. Let $W_i^j \subseteq \mathbb{H}_{\mathscr{B}_i^j}$ be a subspace defined in Definition 10, and denote by $\Pi_{W_i^j}$ the orthogonal projection onto W_i^j. Let $\mathscr{B}^j \coloneqq \bigcup_{k=1}^{2^{M-j}} \mathscr{B}_k^j$ be the set of all qudits under consideration at step j. Let $(\rho^{j-1})'$ be a positive semidefinite operator with trace at most one on $\mathbb{H}_{\mathscr{B}^j}$, as defined in Algorithm 4, and let ρ^j be the subnormalized reconstructed operator defined in Definition 8.*

Suppose that, for each $i = 1,\dots,2^{M-j}$, the tomography procedure succeeds in the sense that

$$\left\|(I - \Pi_{W_i^j})\,\mathrm{Tr}_{\mathscr{B}^j\setminus\mathscr{B}_i^j}\left[(\rho^{j-1})'\right](I - \Pi_{W_i^j})\right\|_\infty \;\le\; 2\eta.$$

Then, for any n-qudit matrix product state $|\phi\rangle$ with bond dimension D,

$$\left|\,\langle\phi|\,\rho^{j-1}\,|\phi\rangle - \langle\phi|\,\rho^j\,|\phi\rangle\,\right| \;\le\; 2\sqrt{4\eta\,D^2\,2^{M-j}}\,.$$

Proof. According to Definition 8 and 9, we can express the fidelity $\langle\phi|\,\rho^j\,|\phi\rangle$ as

$$\begin{aligned}\langle\phi|\,\rho^j\,|\phi\rangle &= \langle\phi|\,(E^j)^\dagger \left[\left(\bigotimes_{a=1}^{j} Q^a\right) \otimes (\rho^j)'\right] E^j\,|\phi\rangle \\ &= \langle\phi^{j-1}|\,(U^j)^\dagger \left[\left(\bigotimes_{a=1}^{j} Q^a\right) \otimes (\rho^j)'\right] U^j\,|\phi^{j-1}\rangle \\ &= \langle\phi^{j-1}|\left(\bigotimes_{a=1}^{j-1} Q^a\right)(U^j)^\dagger\left[Q^j \otimes (\rho^j)'\right]U^j\left(\bigotimes_{a=1}^{j-1} Q^a\right)|\phi^{j-1}\rangle \\ &= \langle\psi^{j-1}|\,(U^j)^\dagger\left[Q^j \otimes (\rho^j)'\right]U^j\,|\psi^{j-1}\rangle \qquad (15)\end{aligned}$$

Using the Q^j defined above, the successful-branch residual satisfies

$$Q^j \otimes (\rho^j)' = P^j U^j (\rho^{j-1})' (U^j)^\dagger P^j, \qquad (16)$$

where the tensor factors are ordered by discarded and surviving registers. Using this identity and Fact C.2, we can rewrite the fidelity $\langle\phi|\,\rho^j\,|\phi\rangle$ as

$$\langle\phi|\,\rho^j\,|\phi\rangle = \langle\psi^{j-1}|\,\Pi_{W_1^j} \otimes \cdots \otimes \Pi_{W_{2^{M-j}}^j}\,(\rho^{j-1})'\,\Pi_{W_1^j} \otimes \cdots \otimes \Pi_{W_{2^{M-j}}^j}\,|\psi^{j-1}\rangle \qquad (17)$$

For brevity, we denote $\Pi_{W^j} \coloneqq \bigotimes_{i=1}^{2^{M-j}} \Pi_{W_i^j}$. We use $\|\psi^{j-1}\| \le 1$ and $0 \preceq (\rho^{j-1})' \preceq I$; the latter follows because the residual operator is PSD with trace at most one. Together with Fact C.1, this

implies

$$\begin{aligned}
\left|\langle\phi|\,\rho^{j-1}\,|\phi\rangle - \langle\phi|\,\rho^{j}\,|\phi\rangle\right| &= \left|\langle\psi^{j-1}|\,(\rho^{j-1})'\,|\psi^{j-1}\rangle - \langle\psi^{j-1}|\,\Pi_{W^j}(\rho^{j-1})'\,\Pi_{W^j}\,|\psi^{j-1}\rangle\right| \\
&\le \left|\langle\psi^{j-1}|\,(\rho^{j-1})'\,(I-\Pi_{W^j})\,|\psi^{j-1}\rangle\right| + \left|\langle\psi^{j-1}|\,(I-\Pi_{W^j})(\rho^{j-1})'\,\Pi_{W^j}\,|\psi^{j-1}\rangle\right| \\
&\le \left\||\psi^{j-1}\rangle\right\|\left\|(\rho^{j-1})'\,(I-\Pi_{W^j})\,|\psi^{j-1}\rangle\right\| + \left\|\Pi_{W^j}\,|\psi^{j-1}\rangle\right\|\left\|(\rho^{j-1})'\,(I-\Pi_{W^j})\,|\psi^{j-1}\rangle\right\| \\
&\le 2\left\|(\rho^{j-1})'\,(I-\Pi_{W^j})\,|\psi^{j-1}\rangle\right\| \qquad (18)
\end{aligned}$$

Let $|\widetilde{\varphi}\rangle = (I-\Pi_{W^j})\,|\psi^{j-1}\rangle$. If $|\widetilde{\varphi}\rangle = 0$, the expectation difference is zero and the bound follows. Otherwise define the normalized vector $|\varphi\rangle \coloneqq |\widetilde{\varphi}\rangle\,/\|\widetilde{\varphi}\|$. Then

$$\begin{aligned}
2\left\|(\rho^{j-1})'\,(I-\Pi_{W^j})\,|\psi^{j-1}\rangle\right\| &= 2\left\|(\rho^{j-1})'\;|\varphi\rangle\right\|\;\left\|\widetilde{\varphi}\right\| \\
&\le 2\left\|\widetilde{\varphi}\right\|\sqrt{\langle\varphi|\,[(\rho^{j-1})']^2\,|\varphi\rangle} \\
&\le 2\left\|\widetilde{\varphi}\right\|\sqrt{\langle\varphi|\,(\rho^{j-1})'\,|\varphi\rangle} \\
&\le 2\sqrt{\langle\varphi|\,(\rho^{j-1})'\,|\varphi\rangle}. \qquad (19)
\end{aligned}$$

For $j \in \{1,\dots,M\}$ and $i \in \{1,\dots,2^{M-j}\}$, let the first-layer ancestor interval of $\mathscr{B}_i^j$ be

$$S(\mathscr{B}_i^j) \coloneqq \bigcup_{a=(i-1)2^{j-1}+1}^{i2^{j-1}} \mathscr{B}_a^1.$$

The first-layer blocks partition the original chain in order, so $S(\mathscr{B}_i^j)$ is a contiguous interval. An MPS of bond D (with the periodic-boundary convention) has Schmidt rank at most D^2 across this interval and its complement. For all layers $1,\dots,j-1$, the binary-tree organization ensures that each learned unitary, all-zero projection, and removal of a projected register acts entirely inside $S(\mathscr{B}_i^j)$ or entirely outside it. Thus, up to the ordering of surviving registers, the map producing the unnormalized residual vector $|\psi^{j-1}\rangle$ from $|\phi\rangle$ factorizes into linear maps on the two sides of this bipartition. Such local maps cannot increase Schmidt rank. After removing the projected registers, $|\psi^{j-1}\rangle$ therefore has Schmidt rank at most D^2 across $\mathscr{B}_i^j : (\mathscr{B}^j \setminus \mathscr{B}_i^j)$. This holds for each i, including unprocessed first-layer blocks.

The normalized leaked vector $|\varphi\rangle$ has the form required by Lemma C.2, with $|\chi\rangle = |\psi^{j-1}\rangle$; the Schmidt-rank condition on $|\chi\rangle$ holds across every block cut by the preceding argument. Hence

$$\left|\langle\phi|\,\rho^{j-1}\,|\phi\rangle - \langle\phi|\,\rho^{j}\,|\phi\rangle\right| \le 2\sqrt{\langle\varphi|\,(\rho^{j-1})'\,|\varphi\rangle} \le 2\sqrt{4\eta D^2 2^{M-j}}. \qquad (20)$$

□

C.4 Proof of Theorem 2 and Sample Complexity

Theorem (Restatement of Theorem 2)**.** *For any integers $n, d \ge 2$, given access to copies of an n-qudit state ρ, parameters $\epsilon, \delta \in (0,1)$, and a bond dimension parameter D, Algorithm 4 outputs a description of a state $|\hat{\phi}\rangle$, such that, with probability $1-\delta$, $\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \ge \max_{|\phi\rangle\in MPS(D)}\langle\phi|\,\rho\,|\phi\rangle - \epsilon$. The algorithm requires $N = O\Big(\frac{D^{12}n^7\,d^4(\ln d)^7\,\log\frac{n}{\delta}}{\epsilon^{12}\left[\ln\left(\frac{128\,nD^2\ln d}{(\sqrt{2}-1)^2\epsilon^2}\right)\right]^7}\Big)$ copies of ρ and runs in time $\mathrm{poly}(D, n, 1/\epsilon, \log(1/\delta))$ for fixed local dimension d.*

Proof. First consider the direct-tomography branch $p \ge n$. Let $\hat{\rho}$ be the trace-norm-$\epsilon/2$ estimate supplied by the rank-agnostic tomography lemma. If $|\hat{\phi}\rangle$ is a leading eigenvector of $\hat{\rho}$, then for any normalized $|\phi\rangle$,

$$\langle\hat{\phi}|\rho|\hat{\phi}\rangle \ge \langle\hat{\phi}|\hat{\rho}|\hat{\phi}\rangle - \epsilon/2 \ge \langle\phi|\hat{\rho}|\phi\rangle - \epsilon/2 \ge \langle\phi|\rho|\phi\rangle - \epsilon.$$

This branch uses $O(d^{2n}\epsilon^{-2}\log(2/\delta))$ copies; its bound is absorbed by the final stated bound in the regime $p \ge n$ (see Appendix C.5.1). Henceforth suppose $p < n$. Conditioned on the preceding learning history, each replay map is fixed. Lemma A.3 estimates each subnormalized target using fresh original input copies. Assign failure probability $\zeta = \delta/(2^M+1)$ to every call. There are at

most 2^M block and final calls, so a union bound ensures simultaneous success with probability at least $1-\delta$. On this event, each block estimate $\hat{\sigma}_i^j$ satisfies

$$\|\sigma_i^j - \hat{\sigma}_i^j\|_1 \;\le\; \eta.$$

In the next step, Algorithm 4 invokes Algorithm 3 to select eigenvectors corresponding to eigenvalues exceeding η from the estimated subnormalized PSD operator $\hat{\sigma}_i^j$ to construct a disentangling unitary accordingly. We denote $\widetilde{W}_i^j$ to be the subspace spanned by the eigenvectors of $\hat{\sigma}_i^j$ with eigenvalues exceeding η. From Lemma C.1, we know that

$$\|(I-\Pi_{\widetilde{W}_i^j})^\dagger \sigma_i^j (I-\Pi_{\widetilde{W}_i^j})\|_\infty \le \frac{(\sqrt{2}-1)^2\epsilon^2 p}{64D^2 n}.$$

By Algorithm 3 and the dimension bound $d^p \ge 1/\eta$, all eigenvectors with eigenvalues exceeding η belong to the retained subspace W_i^j. Hence $\widetilde{W}_i^j = V_i^j \subseteq W_i^j$, and $(W_i^j)^\perp \subseteq (V_i^j)^\perp$. Therefore, it follows that

$$\|(I-\Pi_{W_i^j})^\dagger \sigma_i^j (I-\Pi_{W_i^j})\|_\infty \le \frac{(\sqrt{2}-1)^2\epsilon^2 p}{64D^2 n}.$$

Let $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$ be a matrix product state with bond dimension D on an n-qudit system. Applying Lemma C.3, we get

$$\left|\, \langle\phi|\, \rho^{j-1}\, |\phi\rangle - \langle\phi|\, \rho^{j}\, |\phi\rangle \,\right| \;\le\; \frac{(\sqrt{2}-1)\epsilon}{2}\sqrt{\frac{p}{2n}\, 2^{M-j}}\,.$$

Using the triangular inequality, we derive

$$\left|\, \langle\phi|\, \rho^{0}\, |\phi\rangle - \langle\phi|\, \rho^{M}\, |\phi\rangle \,\right| \;\le\; \sum_{j=1}^{M} \frac{(\sqrt{2}-1)\epsilon}{2}\sqrt{\frac{p}{2n}\, 2^{M-j}} \le \frac{\epsilon}{2}\sqrt{\frac{p}{2n}2^M} \;\le\; \frac{\epsilon}{2},$$

where in the last inequality we use the fact that $n > p2^{M-1}$ (see line 3 in Algorithm 4). This implies

$$\langle\phi|\, \rho^{M}\, |\phi\rangle \ge \langle\phi|\, \rho^{0}\, |\phi\rangle - \frac{\epsilon}{2}. \tag{21}$$

For any matrix product state $|\phi\rangle \in (\mathbb{C}^d)^{\otimes n}$ with bond dimension D, we define its level-j projected states $|\psi^j\rangle \in \mathbb{H}_{\widetilde{\mathscr{B}}_1^j} \otimes \cdots \otimes \mathbb{H}_{\widetilde{\mathscr{B}}_{2^{M-j}}^j}$ through Definition 9. By Fact C.1, we have

$$\langle\phi|\, \rho^{M}\, |\phi\rangle = \langle\psi^M|\, (\rho^M)'\, |\psi^M\rangle\,. \tag{22}$$

With $\mathscr{D}^j$ as defined above, let $\mathscr{D} = \bigsqcup_{j=1}^M \mathscr{D}^j$ be the set of registers discarded through the M layers and write $|0_{\mathscr{D}}\rangle$ for their joint all-zero state. Later unitaries act as the identity on previously discarded registers. Therefore the sequential successful-branch construction gives

$$(\rho^M)' = (\langle 0_{\mathscr{D}}| \otimes I) E^M \rho (E^M)^\dagger (|0_{\mathscr{D}}\rangle \otimes I),$$

with the tensor factors ordered by the discarded and surviving registers. Algorithm 4 outputs $|\hat{\phi}\rangle = (E^M)^\dagger(|0_{\mathscr{D}}\rangle \otimes |\hat{\psi}\rangle)$. Consequently, by the above identity and Definition 8,

$$\langle\hat{\phi}|\, \rho\, |\hat{\phi}\rangle = \langle\hat{\psi}|\, (\rho^M)'\, |\hat{\psi}\rangle = \langle\hat{\phi}|\, \rho^M\, |\hat{\phi}\rangle\,. \tag{23}$$

If the final subnormalized tomography procedure succeeds, combining conditional tomography on $\bar{\rho}^M$ with the observed frequency of all M projections succeeding (or returning zero as specified by the lemma), it yields a PSD estimate $(\hat{\rho}^M)'$ such that $\|(\rho^M)' - (\hat{\rho}^M)'\|_1 \le \epsilon/4$. By Hölder's inequality for Schatten p-norms, for any state $|\varphi\rangle \in \mathbb{H}_{\widetilde{\mathscr{B}}_1^M}$,

$$\left|\langle\varphi|\, (\rho^M)'\, |\varphi\rangle - \langle\varphi|\, (\hat{\rho}^M)'\, |\varphi\rangle\right| \le \epsilon/4. \tag{24}$$

Applying this bound gives

$$\langle\hat{\psi}|\, (\rho^M)'\, |\hat{\psi}\rangle \ge \langle\hat{\psi}|\, (\hat{\rho}^M)'\, |\hat{\psi}\rangle - \epsilon/4 \ge \langle\psi^M|\, (\hat{\rho}^M)'\, |\psi^M\rangle - \epsilon/4 \ge \langle\psi^M|\, (\rho^M)'\, |\psi^M\rangle - \epsilon/2. \tag{25}$$

Here, the first and last inequalities use the trace-norm tomography bound. This bound applies to the possibly subnormalized $|\psi^M\rangle$ because $\|\psi^M\| \leq 1$. The middle inequality holds since $(\hat{\rho}^M)' \succeq 0$ and its normalized leading eigenvector maximizes the expectation over all vectors of norm at most one. Combining with Equations (22) and (23), we obtain

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \geq \langle\phi|\,\rho^M\,|\phi\rangle - \frac{\epsilon}{2}. \tag{26}$$

Combining this with Equation (21) gives

$$\langle\hat{\phi}|\,\rho\,|\hat{\phi}\rangle \geq \langle\phi|\,\rho\,|\phi\rangle - \epsilon. \tag{27}$$

Since this holds for every bond-D MPS comparator $|\phi\rangle$, the desired closest-MPS guarantee follows. No operator monotonicity of ρ^j is used.

Next, we compute the sample and time complexity of Algorithm 4. Each successful-branch map replays the previously learned layers on a fresh original copy and retains the current block only if all preceding projections succeed. Lemma A.3, with output dimension at most $h = d^{2p}$, accuracy η, and failure probability $\zeta = \delta/(2^M + 1)$, uses at most

$$O\left(\frac{d^{4p}}{\eta^2}\log\frac{n}{\delta}\right)$$

original copies per block, including frequency-estimation copies and all failed trials. There are at most $\ell_1 + \sum_{j=2}^{M} 2^{M-j} \leq 2^M - 1 = O(n/p)$ block calls. Hence the total block budget is

$$N = O\left(\frac{nd^{4p}}{p\eta^2}\log\frac{n}{\delta}\right).$$

The lemma's fixed-budget implementation already accounts for the success probability; no additional $1/q_{j-1}$ factor is required. By Equation (13), $d^p \leq 2d/\eta$, where $\eta = (\sqrt{2}-1)^2\epsilon^2 p/(128D^2 n)$. the uniform block-size bound in Appendix C.2.1 yields

$$N = O\left(\frac{nd^4}{p\eta^6}\log\frac{n}{\delta}\right) = O\left(\frac{D^{12}n^7\,d^4(\ln d)^7\,\log\frac{n}{\delta}}{\epsilon^{12}\left[\ln(\ln d) + \ln\frac{128\,n\,D^2}{(\sqrt{2}-1)^2\epsilon^2}\right]^7}\right).$$

See Appendix C.5.1 for a more detailed derivation. The final successful-branch output has dimension $h = d^p$. Lemma A.3, with accuracy $\tau = \epsilon/4$ and failure probability ζ, uses

$$K = O\left(\frac{d^{2p}}{\epsilon^2}\log\frac{n}{\delta}\right)$$

original copies of ρ, including failed trials. We show in Appendix C.5.2 that the complexity bound K is asymptotically dominated by N. Therefore, the overall sample complexity is

$$N = O\left(\frac{D^{12}n^7\,d^4(\ln d)^7\,\log\frac{n}{\delta}}{\epsilon^{12}\left[\ln(\ln d) + \ln\frac{128\,n\,D^2}{(\sqrt{2}-1)^2\epsilon^2}\right]^7}\right),$$

and the complete procedure runs in polynomial time. Indeed, the block parameter obeys $p = O(\log_d B + 1)$ for $B = 128nD^2/((\sqrt{2}-1)^2\epsilon^2)$, so each tomography output has dimension at most $d^{2p} = \mathrm{poly}(d, n, D, 1/\epsilon)$. In the direct-tomography branch $p \geq n$, its dimension $h = d^n \leq d^p$ is polynomial in the same parameters. The gate-efficient tomography of [PSTW25], postselection replay, matrix diagonalization and construction of the learned unitaries consequently all require polynomial time (for fixed local dimension d). This does not require the collective tomography measurements to have shallow circuit depth. □

C.5 Detailed Sample-Complexity Calculations

C.5.1 Substituting the Block-Size Bound

We start from

$$N = O\left(\frac{nd^{4p}}{p\eta^2}\,\log\frac{n}{\delta}\right), \qquad \eta = \frac{(\sqrt{2}-1)^2\,\epsilon^2\,p}{128\,D^2 n}.$$

Step 1. Rewriting in terms of η. Using $d^p \le 2d/\eta$ from Equation (13) and absorbing a universal constant, one finds

$$N = O\left(\frac{nd^4}{p\eta^6}\,\log\frac{n}{\delta}\right).$$

Step 2. Expanding η. Since

$$\eta = c_1\frac{\epsilon^2 p}{D^2 n}, \qquad c_1 := \frac{(\sqrt{2}-1)^2}{128},$$

we obtain

$$\frac{1}{p\eta^6} = \frac{1}{c_1^6}\,\frac{D^{12}n^6}{\epsilon^{12}}\,\frac{1}{p^7}.$$

Step 3. Eliminating p. Recall that the block-size choice in Appendix C.2.1 satisfies $p \ge W(z)/\ln d$, where

$$z := (\ln d)B, \qquad B := \frac{128nD^2}{(\sqrt{2}-1)^2\epsilon^2}, \qquad L := \ln z.$$

Since $z > e^2$, the standard Lambert-W bound $W(z) > L - \ln L \ge L/2$ gives the uniform estimate

$$\frac{1}{p^7} = O\left(\frac{(\ln d)^7}{L^7}\right).$$

For fixed $d \ge 2$ and growing B, this becomes a Θ bound. Substituting yields

$$\frac{1}{p\eta^6} = O\left(\frac{D^{12}n^6}{\epsilon^{12}}\,\frac{(\ln d)^7}{L^7}\right).$$

Step 4. Final expression. Collecting all terms, we obtain

$$N = O\left(\frac{D^{12}n^7}{\epsilon^{12}}\,\frac{d^4(\ln d)^7}{L^7}\,\log\frac{n}{\delta}\right),$$

where

$$L = \ln\Big(\ln d \cdot \tfrac{128\,n\,D^2}{(\sqrt{2}-1)^2\epsilon^2}\Big).$$

Step 5. Simplified $\widetilde{O}$ form. Up to polylogarithmic factors, this can be expressed as

$$N = \widetilde{O}\left(\frac{D^{12}n^7}{\epsilon^{12}}d^4(\ln d)^7\log\frac{n}{\delta}\right).$$

C.5.2 Comparing Block and Final Tomography Budgets

Consider the two complexity expressions:

$$K = O\left(\frac{d^{2p}}{\epsilon^2}\,\log\frac{n}{\delta}\right), \qquad N = O\left(\frac{nd^{4p}}{p\eta^2}\,\log\frac{n}{\delta}\right), \qquad \eta = \frac{(\sqrt{2}-1)^2\,\epsilon^2 p}{128D^2 n}.$$

To compare the displayed *upper-bound expressions*, define

$$K_{\mathrm{bd}} := d^{2p}\epsilon^{-2}\log(n/\delta), \qquad N_{\mathrm{bd}} := \frac{nd^{4p}}{p\eta^2}\log(n/\delta).$$

Their ratio, using $\eta = c_1\epsilon^2 p/(D^2 n)$, is

$$\frac{N_{\mathrm{bd}}}{K_{\mathrm{bd}}} = \frac{n\epsilon^2 d^{2p}}{p\eta^2} = \frac{D^4 n^3 d^{2p}}{c_1^2 p^3\epsilon^2}.$$

Since $B = p/\eta$ and $d^p \ge B/p = 1/\eta$, the above ratio is at least 1 whenever $p < n$, since $0 < c_1 < 1$, $D \ge 1$, $n > p$, $d \ge 2$ and $0 < \epsilon < 1$. Thus the final-state tomography term is absorbed in the stated big-O sample bound. For $p \ge n$, Algorithm 4 uses direct full-state tomography instead. Indeed, writing $c_1 = (\sqrt{2}-1)^2/128$, the ratio of the upper-bound expressions in this regime is

$$\frac{nd^{4p}/(p\eta^2)}{d^{2n}/\epsilon^2} = \frac{D^4 n^3 d^{4p-2n}}{c_1^2\epsilon^2 p^3} \ge \frac{D^4 n^3 d^{2p}}{c_1^2\epsilon^2 p^3} \ge 1,$$

where the last inequality uses $d \ge 2$, $4^p \ge p^3$ for integers $p \ge 1$, and $c_1, \epsilon < 1$. Hence the direct-tomography copy budget is also absorbed.

Appendix D: Nearest-Neighbor Implementations

Theorem (Restatement of Theorem 3)**.** *On a two-dimensional square lattice, the quantum circuit output by Algorithm 2 can be implemented using $O(\sqrt{n})$ layers of nearest-neighbor SWAP gates and $O(\log n)$ layers of block unitaries, each acting on at most $2p$ qudits within adjacent patches.*

Proof. Without loss of generality, we assume that $\sqrt{n/p} = 2^{M/2}$ and that M is even [11]. Note that we may arrange the qubits on a square lattice in any convenient way. To prove the lemma, we view a bunch of p qubits as a single lattice point and adopt the following construction to form a $2^{M/2} \times 2^{M/2}$ square lattice. For simplicity, we assume that the circuit depth M is even.

First, we place $\widetilde{\mathscr{B}}_1^M$ at the bottom-right lattice site of the entire square lattice, which serves as the anchor for the 0-th level.

Next, we partition the lattice into 4 blocks, each of size $2^{M/2-1} \times 2^{M/2-1}$. In each block, we place $\widetilde{\mathscr{B}}_i^{M-2}$ at its bottom-right lattice site, which we define as anchors for the 1-st level. The blocks are indexed in row-major order, i.e., from left to right and then from top to bottom, with $i = 1, \ldots, 4$.

We then iterate this procedure. At level $j \geq 1$, for each block indexed by i (whose bottom-right anchor is $\widetilde{\mathscr{B}}_i^{M-2(j-1)}$), we partition the block into 2×2 subblocks, each of size $2^{M/2-j} \times 2^{M/2-j}$. In each subblock we place $\widetilde{\mathscr{B}}_{k+4(i-1)}^{M-2j}$ at the bottom-right lattice site as the next-level anchor, with $k = 1, \ldots, 4$. The four subblocks are again indexed in row-major order. At the $M/2$-th level, we obtain a $2^{M/2} \times 2^{M/2}$ square lattice (see Figure 2). Each lattice site is associated with a patch of p qubits, which we label by $\widetilde{\mathscr{B}}_\ell^0$, where $\ell = 1, \ldots, 2^M$.

Now, we have placed all the qubits on a $2^{M/2} \times 2^{M/2}$ square lattice. Using the level structure described above, we then implement the algorithm iteratively over $M/2$ levels, from $j = M/2$ down to $j = 1$.

Each level j consists of two steps. First, we apply the disentangling unitaries horizontally on pairs of p-qudit blocks $(\widetilde{\mathscr{B}}_{k-1}^{M-2j}, \widetilde{\mathscr{B}}_k^{M-2j})$, where k ranges over even integers in $\{1, \ldots, 2^{2j}\}$. These horizontal unitaries disentangle the qubits at the left lattice sites (i.e., $\widetilde{\mathscr{B}}_{k-1}^{M-2j}$ for even k) to the all-zero state. In the second step, we apply the disentangling unitaries vertically on pairs of p-qudit blocks $(\widetilde{\mathscr{B}}_{k-2}^{M-2j}, \widetilde{\mathscr{B}}_k^{M-2j})$ where k ranges over multiples of 4 in $\{1, \ldots, 2^{2j}\}$. These vertical unitaries disentangle the qubits at the top lattice site to the all-zero state. As a result, only the sites $\widetilde{\mathscr{B}}_k^{M-2j}$ where k are multiples of 4 in $\{1, \ldots, 2^{2j}\}$ remain. Notice that the site $\widetilde{\mathscr{B}}_k^{M-2j}$ is exactly the same site with the label $\widetilde{\mathscr{B}}_{k/4}^{M-2(j-1)}$ in the next level. Each level corresponds to a pair of successive iterations in Algorithm 2.

The collection of these surviving lattice sites then forms the input to the next level. After executing $M/2$ levels, only a single lattice site containing p qubits remains. We then perform full state tomography on this site to complete the two-dimensional version of Algorithm 2.

Now, we analyze the circuit depth under the restriction to nearest-neighbor quantum gates. At level j, before applying the horizontal disentangling unitaries, we need to move the qubits from the left lattice sites (i.e., $\widetilde{\mathscr{B}}_{k-1}^{M-2j}$ for even k) to the lattice sites immediately to the left of their right partners. Similarly, before applying the vertical disentangling unitaries, we need to move the qubits from the top lattice sites (i.e., $\widetilde{\mathscr{B}}_{k-2}^{M-2j}$ for k being multiples of 4) to the lattice sites immediately above their bottom partners. Therefore, at level j, this requires $O(2^{M/2-(j-1)})$ SWAP gates within each block $\widetilde{\mathscr{B}}_k^{M-2(j-1)}$, where $k = 1, \ldots, 2^{2(j-1)}$. Since these blocks are disjoint, the SWAP gates acting on different blocks at the same level can be executed in parallel. This gives a coarse-grained SWAP routing depth of $\sum_{j=1}^{M/2} 2^{M/2-(j-1)} = O(2^{M/2}) = O(\sqrt{n/p})$. Each coarse-grained site corresponds to a patch of p qudits. Exchanging two adjacent patches can be implemented using $O(\sqrt{p})$ layers of nearest-neighbor two-qudit SWAP gates. Hence, the total nearest-neighbor SWAP depth is $O(\sqrt{n})$. The learned block unitaries contribute $O(\log n)$ additional layers, where each block unitary is counted

[11] If this is not the case, we can replace n by the smallest $n' \geq n$ such that $\sqrt{n'/p} = 2^{M/2}$. The following argument then applies to n' without any change.

as a single operation, as in Theorem 1. The elementary-gate decomposition of these block unitaries is not included in this depth bound. □

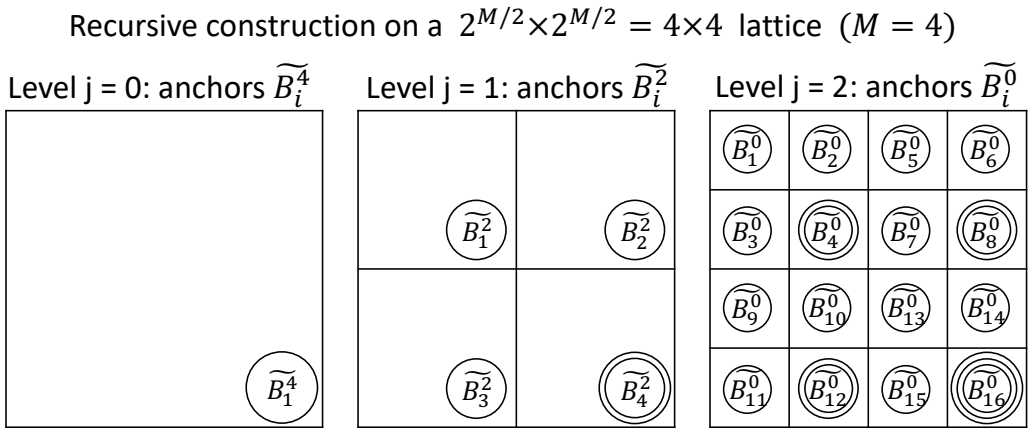


Figure 4: Illustration of the recursive block structure on a $2^{M/2} \times 2^{M/2}$ square lattice with $M = 4$. At level j, each block of size $2^{M/2-(j-1)} \times 2^{M/2-(j-1)}$ is partitioned into four subblocks, each of size $2^{M/2-j} \times 2^{M/2-j}$. In each subblock, the bottom-right site (corresponding to the maximal-coordinate corner in each block) is designated as the anchor $\widetilde{\mathscr{B}}_i^{M-2j}$ (circled). These anchors form a coarser lattice for the next level.

Corollary (Restatement of Corollary 3.1)**.** *On a q-dimensional hypercubic lattice, the quantum circuit output by Algorithm 2 can be implemented using $O(qn^{1/q})$ layers of nearest-neighbor SWAP gates and $O(\log n)$ layers of block unitaries, each acting on at most $2p$ qudits within adjacent patches.*

Proof. The proof is a direct generalization of Theorem 3. Without loss of generality, assume that $n/p = 2^M$ and M is divisible by q. [12] We view each patch of p qubits as a single lattice point and arrange the n/p patches on a q-dimensional hypercubic lattice of side length $2^{M/q}$ in each coordinate direction.

We recursively define a level structure on this lattice. At level 0, the entire hypercube has a distinguished anchor site, chosen to be the corner $(2^{M/q}, \ldots, 2^{M/q})$. At level $j \geq 1$, each block from level $j-1$ is partitioned into 2^q equal subblocks of side length $2^{M/q-j}$. In each subblock, we designate as its anchor the corner with maximal coordinates in all directions. After M/q levels, every lattice site is identified with one patch $\widetilde{\mathscr{B}}_\ell^0$, where $\ell = 1, \ldots, 2^M$.

We now implement the algorithm level by level, from level M/q down to level 1. Each level consists of q successive steps, one for each coordinate direction. In the r-th step, where $r \in \{1, \ldots, q\}$, we apply disentangling unitaries to pairs of neighboring subblocks that differ only in the r-th coordinate, so that after this step only one half of the sites in that coordinate direction remain active. After performing all q steps, only one corner site remains in each block, which becomes the anchor of the corresponding block at the next coarser level. Thus, each level corresponds to q successive iterations of Algorithm 2.

It remains to bound the depth under the nearest-neighbor constraint. Consider level j. Inside each block labeled by the level $j-1$ anchor, the active sites form a hypercube of side length $L_j = 2^{M/q-(j-1)}$. Before applying the disentangling unitaries in a given coordinate direction, we move the qubits from one half of the active sites to the neighboring sites adjacent to their partner blocks along that direction. This routing can be implemented by nearest-neighbor SWAP gates with

[12] For general n and p, we may pad the lattice with dummy sites so that these conditions hold. This changes the side length by at most a constant factor and does not affect the asymptotic depth.

depth $O(L_j)$ inside each block. Since there are q such directions in each level, the total routing depth per level is $O(qL_j)$.

Because distinct blocks at the same level are disjoint, the routing operations can be performed routing depth is $O\left(\sum_{j=1}^{M/q} q\, 2^{M/q-(j-1)}\right) = O\left(q(n/p)^{1/q}\right)$. Each coarse-grained site corresponds to a patch of p qudits, which can be embedded in a q-dimensional tile of side length $\lceil p^{1/q} \rceil$. Exchanging two adjacent patches requires $O(p^{1/q})$ layers of nearest-neighbor SWAP gates. Hence, the total SWAP routing depth is $O\left(q(n/p)^{1/q} p^{1/q}\right) = O(qn^{1/q})$. The learned block unitaries contribute $O(\log n)$ additional layers, with each block unitary counted as a single operation, as in Theorem 1. The elementary-gate decomposition of these block unitaries is not included in this depth bound. The final state preparation contributes one additional block operation. □

Appendix E: Product-State Learning Lower Bound

To establish a lower bound, we prove that for product state learning, the sampling complexity lower bound for producing an approximate description of the given quantum state to fidelity ϵ-close to optimum is $\Omega(nd/\epsilon)$, where n is the number of d-dimensional systems.

Theorem (Restatement of Theorem 4)**.** *For integer $n \geq 1$, $d \geq 2$ and $\epsilon \in [0, 2/3]$, consider learning a product state $\rho = \bigotimes_{i=1}^{n} \rho_i$, where each ρ_i is a state on a d-dimensional Hilbert space. Any quantum algorithm that outputs a product state $\hat{\rho}$ satisfying $F(\rho, \hat{\rho}) \geq 1-\epsilon$ must use at least $\Omega(nd/\epsilon)$ copies of ρ.*

Proof. It suffices to consider the special case in which ρ is a pure product state.

The sample complexity for approximating a single-qudit Haar random state to fidelity at least $1-\epsilon$ requires $\Omega(d/\epsilon)$ samples. This follows from a direct worst-case to average-case reduction. Suppose there exists a learning algorithm $\mathscr{B}$ which, given $o(d/\epsilon)$ samples of a single-qudit Haar random state, outputs a estimate with fidelity at least $1-\epsilon$. Then we can construct an algorithm $\mathscr{B}^*$ that learns *arbitrary* single-qudit states using $o(d/\epsilon)$ samples, contradicting the $\Omega(d/\epsilon)$ lower bound for learning d-dimensional states to fidelity $1-\epsilon$ due to Yuen [Yue23].

$\mathscr{B}^*$ can be described as follows: given an arbitrary single-qudit state $|\phi\rangle$, one first samples a single-qudit Haar-random unitary U_1 and apply U_1 on $|\phi\rangle$ to prepare a single-qudit Haar-random state $U_1 |\phi\rangle$, runs $\mathscr{B}$ on $U|\phi\rangle$ using $o(d/\epsilon)$ samples to obtain a state $\hat{\rho}$ such that $F(\hat{\rho}, U|\phi\rangle) \geq 1-\epsilon$, and then outputs $U^\dagger |\varphi\rangle U$. By unitary invariance of fidelity, this estimate achieves fidelity at least $1-\epsilon$ with $|\phi\rangle$, completing the reduction and yielding the desired contradiction to the sample complexity lower bound.

Now, consider a product state of n single-qudit Haar random state sampled independently, $|\psi\rangle = \otimes_{i=1}^{n} |\psi_i\rangle_{Haar}$. Suppose there exists a quantum algorithm $\mathscr{A}$ that takes $o(\frac{nd}{\epsilon})$ copies of state $|\psi\rangle$ and outputs another product state $\hat{\rho}$ such that $F(\hat{\rho}, |\psi\rangle) \geq 1-\epsilon$. Then, we can obtain an algorithm $\mathscr{A}^*$ to learn a single-qudit Haar random state with fidelity at least $1-\epsilon/n$ using at most $o(nd/\epsilon)$ samples.

Note that $F(\hat{\rho}, |\psi\rangle) \geq 1-\epsilon$ implies that $F_1 \cdot F_2 \cdots F_n \geq 1-\epsilon$, where $F_i = \langle\psi_i|_{Haar} \operatorname{Tr}_{\bar{i}}(\hat{\rho}) |\psi_i\rangle_{Haar}$ for $i \in [n]$ and $\bar{i} = \{1, 2, \ldots, i-1, i+1, \ldots, n\}$. Let i^* be the position that has the largest F_{i^*}, by an averaging argument, $F_{i^*}^n \geq F_1 \cdot F_2 \cdots F_n \geq 1-\epsilon$. Then, for $\epsilon \leq 2/3$, we can obtain that

$$F_{i^*} \geq (1-\epsilon)^{1/n} \geq 1 - 2\epsilon/n. \tag{28}$$

$\mathscr{A}^*$ is described as follows: Given a single-qudit Haar-random input state $|\phi\rangle$, the algorithm first samples n independent single-qudit Haar-random states, preparing $|\psi\rangle = \bigotimes_{i=1}^{n} |\psi_i\rangle_{\text{Haar}}$. It then runs $\mathscr{A}$ on copies of $|\psi\rangle$, evaluates the per-qudit fidelities, and identifies an index i^* achieving the largest fidelity so that $F_{i^*} \geq 1 - 2\epsilon/n$ by Equation 28. (This step will take many copies of $|\psi\rangle$ to obtain desired precision of the per-qudit fidelities; however, it does not require any copy of the input $|\phi\rangle$.) Then $\mathscr{A}^*$ replaces the state $|\psi_{i^*}\rangle_{\text{Haar}}$ with $|\phi\rangle$ and applies $\mathscr{A}$ to learn a state $\hat{\rho}$ that approximates

$$|\psi^*\rangle = \left(\bigotimes_{i<i^*} |\psi_i\rangle_{\text{Haar}}\right) \otimes |\phi\rangle \otimes \left(\bigotimes_{i>i^*} |\psi_i\rangle_{\text{Haar}}\right)$$

to fidelity at least $1-\epsilon$. Finally, it outputs the reduced state $\rho^* = \operatorname{Tr}_{\overline{i^*}}(\hat{\rho})$. Since $F_{i^*} \geq 1 - 2\epsilon/n$ and $\mathscr{A}^*$ uses only $o(nd/\epsilon)$ copies, this contradicts the $\Omega(nd/\epsilon)$ sample-complexity lower bound for

learning a single-qudit Haar-random state to fidelity at least $1 - 2\epsilon/n$ (setting the $\Omega(d/\epsilon')$ bound with $\epsilon' = \epsilon/n$). Therefore, any algorithm that learns $|\psi\rangle = \bigotimes_{i=1}^{n} |\psi_i\rangle_{\mathrm{Haar}}$ to fidelity $\geq 1 - \epsilon$ must use $\Omega(nd/\epsilon)$ copies. Consequently, learning an arbitrary n-qudit product state to fidelity at least $1 - \epsilon$ also requires $\Omega(nd/\epsilon)$ copies, since the average-case lower bound extends to the worst case.

□

Appendix F: Hardness Under Single-Qubit Measurements

Notation. Let $\mathbb{Z}_2 = \{0, 1\}$ denote the field with two elements. For $\ell \in \mathbb{N}$, let $\mathbb{Z}_2^\ell$ denote the set of binary vectors of length ℓ. For vectors $a, s \in \mathbb{Z}_2^\ell$, we denote their inner product over $\mathbb{Z}_2$ by

$$\langle a, s \rangle = \sum_{i=1}^{\ell} a_i s_i \bmod 2.$$

We write $x \xleftarrow{\$} \mathscr{X}$ to denote that x is sampled uniformly at random from the set $\mathscr{X}$. For $\tau \in (0, 1/2)$, we denote by Ber_τ the Bernoulli distribution over $\mathbb{Z}_2$ satisfying

$$\Pr_{e \sim \mathrm{Ber}_\tau}[e = 1] = \tau.$$

Definition 11 (Search $\mathrm{LPN}_{\tau,\ell}$ [Pie12])**.** *Let $\ell \in \mathbb{N}$ and $\tau \in (0, 1/2)$. A secret vector $s \xleftarrow{\$} \mathbb{Z}_2^\ell$ is chosen uniformly at random. An adversary is given q independent samples of the form*

$$(a_i, b_i) \in \mathbb{Z}_2^\ell \times \mathbb{Z}_2,$$

from the distribution $(a_i, b_i) \sim D_{LPN}^{r,\ell}$ where

$$a_i \xleftarrow{\$} \mathbb{Z}_2^\ell, \quad b_i = \langle a_i, s \rangle \oplus e_i, \quad e_i \sim \mathrm{Ber}_\tau \quad for \quad i = 1, \ldots, q.$$

The goal is to recover the secret s.

Assumption 1 (Search $\mathrm{LPN}_{\tau,\ell}$ Assumption [Pie12])**.** *Let $\ell \in \mathbb{N}$ and $\tau \in (0, 1/2)$. For every quantum polynomial-time adversary $\mathscr{A}$, given $q = \mathrm{poly}(\ell)$ samples (a_i, b_i) generated as in Definition 11, the probability that $\mathscr{A}$ outputs the secret s is negligible in ℓ, i.e.,*

$$\Pr\left[\mathscr{A}((a_1, b_1), \ldots, (a_q, b_q)) = s\right] \leq \mathrm{negl}(\ell),$$

where the probability is taken over the distribution $(a_i, b_i) \sim D_{LPN}^{r,\ell}$ and the quantum algorithm $\mathscr{A}$.

Lemma F.1. *For every $s \in \mathbb{Z}_2^\ell$, the state $|\psi_s\rangle$ generated by the circuit in Figure 5 is a matrix product state with bond dimension at most 2.*

Proof. The prefix variable satisfies $z_i = z_{i-1} \oplus x_i s_i$. Thus the evolution of the parity register depends only on the previous parity bit z_{i-1} and the current input bit x_i. This induces a sequential tensor-network structure where the virtual bond stores the current parity value, which takes only two possible values: 0 or 1. Therefore the bond dimension is at most 2. □

Theorem 6. *Assume the $\mathrm{LPN}_{\tau,\ell}$ assumption holds for some constant $\tau \in (0, 1/2 - \gamma)$, where $\gamma > 0$ is a constant. Let $\{|\psi_s\rangle : s \in \{0,1\}^\ell\}$ be the family of MPS(2) states defined in Figure 5. Then there is no polynomial-time algorithm that, using only computational-basis single-qubit measurements on $\mathrm{poly}(\ell)$ copies of $|\psi_s\rangle$, outputs a classical description of a hypothesis state $\widehat{\rho}$ from which $\widehat{\rho}$ can be efficiently prepared and $F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1 - \epsilon$.*

Proof. First, the string $x \in \mathbb{Z}_2^n$ is uniformly random; indeed, each of the corresponding qubits is first prepared in the state $H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Subsequent CNOT gates, if any, with these qubits as controls and $|0\rangle$ ancillas as targets do not change the marginal distribution of the control qubits. For each input $x \in \mathbb{Z}_2^\ell$, the circuit computes $z_\ell(x) = \langle x, s \rangle \bmod 2$. After applying $R_y(2\theta)$ to the final parity register and measuring in the computational basis, the observed bit w_ℓ satisfies

$$w_\ell = z_\ell(x) \oplus e, \qquad e \sim \mathrm{Bernoulli}(\sin^2\theta).$$

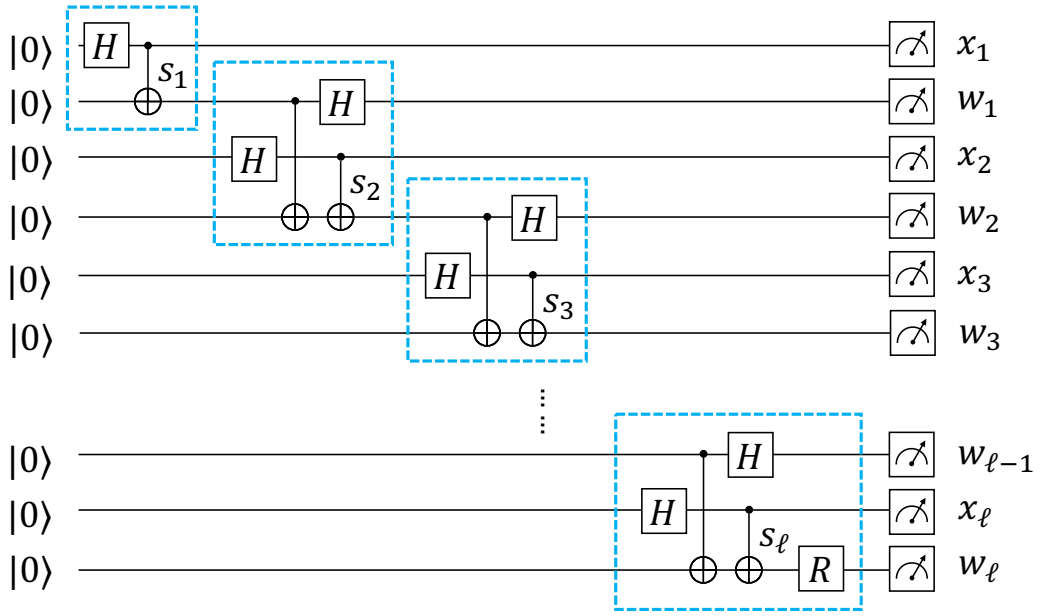


Figure 5: Quantum circuit for generating an LPN-hard MPS(2) family. The odd-indexed qubits are initialized to $|0\rangle$ and put into a uniform superposition over $x = (x_1, \ldots, x_\ell) \in \{0,1\}^\ell$ by Hadamard gates. The circuit then sequentially computes the prefix parities $z_i = \bigoplus_{j=1}^{i} x_j s_j$. For each $i < \ell$, a final Hadamard gate makes w_i uniformly random and independent of the other measured outputs. On the last parity register, where $z_\ell = \langle x, s\rangle \pmod 2$, a rotation $R = R_y(2\arcsin(\sqrt{\tau}))$ is applied so that the measured bit satisfies $w_\ell = \langle x, s\rangle \oplus e$ with $e \sim \mathrm{Bernoulli}(\tau)$. Thus each measurement sample (x, w_ℓ) is distributed as a Search-LPN$_{\tau,\ell}$ sample.

Choosing θ such that $\sin^2\theta = \tau$, we obtain

$$w_\ell = \langle x, s\rangle \oplus e, \qquad e \sim \mathrm{Bernoulli}(\tau).$$

Therefore, each observed pair (x, w_ℓ) is distributed exactly as an LPN$_{\tau,\ell}$ sample.

For each $i \in \{1, 2, \ldots, \ell-1\}$, immediately before the final Hadamard gate on the w_i-register, the value of w_i is determined by the previously generated variables. More precisely, there exists a Boolean function f_i such that $w_i = f_i(x_1, \ldots, x_i, w_1, \ldots, w_{i-1})$ at this point in the circuit. Thus, without loss of generality, the state can be written as

$$\sum_{x,w} a_{x,w} \left|x_1, \ldots, x_i, w_1, \ldots, w_{i-1}\right\rangle \left|\sum_{j=1}^{i} x_i s_i\right\rangle \left|x_{i+1}, \ldots, x_\ell, w_{i+1}, \ldots, w_\ell\right\rangle,$$

where the middle register is the w_i-register.

Applying the Hadamard gate to this middle register gives

$$\sum_{x,w} \frac{a_{x,w}}{\sqrt{2}} \left|x_1, \ldots, x_i, w_1, \ldots, w_{i-1}\right\rangle \Big(|0\rangle + (-1)^{\sum_{j=1}^{i} x_i s_i} |1\rangle \Big) \otimes \left|x_{i+1}, \ldots, x_\ell, w_{i+1}, \ldots, w_\ell\right\rangle.$$

Therefore, conditioned on any fixed values of the other registers, measuring w_i in the computational basis gives 0 and 1 with equal probability. Hence each w_i, for $i \in \{1, \ldots, \ell-1\}$, is uniformly random after the final Hadamard gate and thus is independent of the other computational-basis measurement outcomes.

Suppose that we obtain a classical description of a state $\widehat{\rho}$ from which $\widehat{\rho}$ can be efficiently prepared, and suppose that

$$F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1 - \epsilon.$$

We show that this allows us to recover the secret s in polynomial time with non-negligible probability.

Consider $|\psi_s\rangle$. By applying Hadamard gates H to the qubits $w_1, w_2, \ldots, w_{\ell-1}$, we obtain the values $(x_1 s_1, x_1 s_1 + x_2 s_2, \ldots, \sum_{i=1}^{\ell-1} x_i s_i)$. By measuring these qubits and x, we can directly recover $s_1, s_2, \ldots, s_{\ell-1}$. For s_ℓ, we can also uncompute $R_y(2\theta)$ to obtain $\sum_{i=1}^{\ell} x_i s_i$ and then identify s_ℓ after computing x and $s_1, s_2, \ldots, s_{\ell-1}$. Therefore, if we can prepare $|\psi_s\rangle$, we can recover s with probability 1.

Now consider the learned state $\widehat{\rho}$. Since

$$F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1 - \epsilon,$$

we can obtain that $D\big(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|\big) \leq \sqrt{\epsilon}$, where $D(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1$ denotes trace distance. Thus, for any measurement, the total variation distance between the outcome distributions obtained from $\widehat{\rho}$ and from $|\psi_s\rangle\langle\psi_s|$ is at most $\sqrt{\epsilon}$.

Applying the algorithm above to $\widehat{\rho}$ therefore recovers s with probability at least $1 - \sqrt{\epsilon}$, which is non-negligible. Hence a high-fidelity efficiently preparable $\widehat{\rho}$ allows us to recover the secret s in polynomial time with non-negligible probability. This completes the proof.

□

Corollary 6.1. *Assume the* $\mathrm{LPN}_{\tau,\ell}$ *assumption holds for* $\tau \in (0, 1/2 - \gamma)$*, where* $\gamma > 0$ *is non-negligible. Then there is no polynomial-time non-adaptive algorithm that, using only single-qubit measurements on* $\mathrm{poly}(\ell)$ *copies of* $|\psi_s\rangle$*, outputs a classical description of a hypothesis state* $\widehat{\rho}$ *from which* $\widehat{\rho}$ *can be efficiently prepared and such that* $F(\widehat{\rho}, |\psi_s\rangle\langle\psi_s|) \geq 1 - \epsilon$ *with non-negligible probability.*

Proof. Consider any non-adaptive single-qubit measurement strategy specified by local unitaries $U_1, \ldots, U_{2\ell}$, where the learner measures the i-th qubit in the basis defined by U_i. Equivalently, this measurement procedure can be viewed as first applying $U_1^\dagger \otimes \cdots \otimes U_{2\ell}^\dagger$ to the state and then performing computational-basis measurements. Now define a modified state family

$$|\phi_s\rangle = \left(U_1^\dagger \otimes \cdots \otimes U_{2\ell}^\dagger\right) |\psi_s\rangle .$$

Since local unitaries do not increase MPS bond dimension, Lemma F.1 implies that $|\phi_s\rangle \in \mathrm{MPS}(2)$. Measuring $|\phi_s\rangle$ in the computational basis is statistically equivalent to measuring $|\psi_s\rangle$ using the original non-adaptive single-qubit measurement strategy. Therefore, if a polynomial-time learner could efficiently learn $\mathrm{MPS}(2)$ under arbitrary non-adaptive single-qubit measurements, then it could also recover the hidden secret s for the family $\{|\phi_s\rangle\}$. This contradicts Theorem 5. Hence learning $\mathrm{MPS}(2)$ remains hard in the worst case under non-adaptive single-qubit measurements. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the main theoretical contributions of the paper, including the learning guarantees, sample-complexity bounds, and the assumptions on the target state and measurement access. The detailed algorithm and its analysis are presented in the main body, and the claims are formalized in the main theorem statements and restricted to the settings covered by those theorems.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses limitations and scope in the discussion section (see Section 4.1) comparing parameter regimes. In particular, it identifies that the improved dependence on n comes at the cost of worse dependence on D, so the sample-complexity improvement is most meaningful when $D = o(\sqrt{n})$. It also discusses remaining gaps between upper and lower bounds, limitations of proper learning, and open questions for the closest-MPS setting.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The assumptions for each theoretical result are stated in the corresponding theorem, lemma, or corollary statement. Full proofs are provided in the technical appendix, and the main text points to the relevant appendix sections.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: This paper is theoretical and does not include experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A]

   Justification: This paper is theoretical and does not include experiments requiring code or data.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: The paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: The paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: The paper does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The authors have reviewed the NeurIPS Code of Ethics, and the research conducted in this paper conforms with it.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

Justification: This work is theoretical and focuses on the study of quantum state learning algorithms. It does not involve deployment, data, or applications that directly impact individuals or society. As such, we do not identify immediate societal impacts of the work.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: This work is theoretical and does not involve the release of data, models, or systems that could pose risks of misuse. Therefore, no safeguards are required.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [N/A]

Justification: This paper does not use external assets such as datasets, code, or pre-trained models.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.

- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: This paper does not introduce or release new assets such as datasets, code, or models.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This work does not involve crowdsourcing or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification: This work does not involve crowdsourcing or research with human subjects, and therefore IRB approval is not applicable.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core methods and results of this paper do not involve the use of large language models. Any use of LLMs was limited to writing or editing assistance and does not affect the scientific content.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.