

Preparation circuits for matrix product states by classical variational disentanglement

Refik Mansuroglu^{1,*} and Norbert Schuch^{1,2}

¹*Faculty of Physics, University of Vienna, Boltzmanngasse 5, 1090 Wien, Austria*

²*Faculty of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria*



(Received 12 May 2025; accepted 26 March 2026; published 13 April 2026)

We study the classical compilation of quantum circuits for the preparation of matrix product states (MPS), which are quantum states of low entanglement with an efficient classical description. Our algorithm represents a near-term alternative to previous sequential approaches by reverse application of a disentangler, which can be found by minimizing bipartite entanglement measures after the application of a layer of parametrized disentangling gates. Since a successful disentangler is expected to decrease the bond dimension on average, such a layer-by-layer optimization remains classically efficient even for deep circuits. Additionally, as the Schmidt coefficients of all bonds are locally accessible through the canonical Γ - Λ form of an MPS, the optimization algorithm can be heavily parallelized. We discuss guarantees and limitations to trainability and show numerical results for ground states of one-dimensional, local Hamiltonians as well as artificially spread out entanglement among multiple qubits using error-correcting codes.

DOI: [10.1103/7x2g-twkl](https://doi.org/10.1103/7x2g-twkl)

I. INTRODUCTION

Finding quantum circuit representations of useful entangled initial states can be decisive for a successful quantum simulation of many-body physics. But also beyond quantum simulation, mapping classical solutions of parts of the problem, or educated guesses, to quantum circuits is highly relevant in ground-state problems [1] and for avoiding barren plateaus in variational algorithms and quantum machine learning [2].

Especially on one-dimensional chains, matrix product states (MPS) have been proven useful structures for investigations on ground states of gapped, local Hamiltonians on one-dimensional chains [3,4], as they are efficient classical representations of quantum states with bounded bipartite entanglement. The time evolution of such systems, however, is generally a hard task [5], which is believed to be a natural first candidate for quantum advantage [6,7]. An MPS provides classical access to physical states of quantum many-body systems and is known to be exactly preparable by a quantum circuit of depth linear in system size n and quadratic in the bond dimension D [8]. More recently, it was shown that sequential preparation could be refined for MPS with finite correlation length yielding a log-depth $O[D^4 \log(n)]$ quantum circuit for small-error approximate preparation [9]. An interesting intermediate strategy interpolating between sequential and brick-wall layouts has also been proposed [10] and benchmarked against other preparation strategies in the presence of noise. Using midcircuit measurements, the depth can

be further reduced to $O(\log \log n)$ [9] or even constant-depth circuits [11], though only for a restricted subclass of MPS, on which measurement-induced defects can be “pushed” to the physical indices.

While past approaches have been proven to have optimal asymptotic scaling [9], what is needed in near-term applications are hardware-efficient circuits with two-qubit gates on neighboring sites, allowing for approximation errors on the same scale as gate errors. In fact, a number of investigations have been built upon the sequential preparation [12,13] as well as application-specific approaches like image data encoding [14] and general classification tasks [15]. It is important to note that while sequential preparation yields exact states, or small many-body infidelity in the log-depth method, it requires a minimum number of implementable gates to begin with. This makes it poorly suited for near-term applications, where gate budgets are severely restricted and idle qubits should be avoided, especially considering state preparation as only a small subroutine in a larger quantum algorithm. Recent approaches under the paradigm of approximate quantum compiling (AQC) have therefore proposed the general task to variationally compile circuits using tensor network methods to optimize for a target fidelity [16–18] or to leverage counterdiabatic dynamics [19]. Other strategies include local infidelity minimization [20], variational quantum eigensolver (VQE) preroutines [21], and quantum circuit representations of analytical functions [22]. While these methods highlight the power of tensor-network-assisted optimization and the benefits of local cost functions for avoiding barren plateaus, fidelity-based cost function does not offer control over the bond dimension of the MPS, thus allowing a potential exponential growth of the bond dimension in circuit depth. The challenge remains to design a strategy that is both trainable and classically efficient for arbitrary circuit depths, and at the same time flexible enough to adapt to practical gate budgets on near-term devices.

*Contact author: Refik.Mansuroglu@univie.ac.at


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

In this work, we present classical variational disentanglement (CVD), a near-term focused approach to MPS preparation with parametrized brick-wall circuits optimized to disentangle the state using entanglement entropies as a cost function. At each bond, the Schmidt values can be efficiently processed to calculate von Neumann and Rényi entropies. Once a unitary disentangler U as a product of local gates is found, the preparation circuit is read off as the inverse operator $U^\dagger$. While in general, contracting quantum circuits with MPS increases the bond dimension exponentially in the circuit depth, a successful disentangler decreases the bond dimension of the MPS on average and thus ensures efficient classical scaling. We further show that the approximation error by truncating an MPS, that is close to a product state, and reverse propagating it with $U^\dagger$ is bounded by the square root of the system size $\sqrt{n}$, the circuit depth L , and the truncation error ε . Finally, we show that as a consequence of ultralocality of the cost function, i.e., being defined on only two sites of the MPS, there are no barren plateaus that would impede trainability of the model. In that, CVD enhances the idea of AQC by using the local nature of the encoding of an MPS enabling efficient, local gate-by-gate optimizations.

Although previous works have explored disentangling-based and variational approaches to prepare MPS, our work differs crucially in three respects. First, we introduce an entanglement-based cost function, using Rényi or von Neumann entropies, whose minimization is shown to provably bound the bond dimension (Lemma 1), ensuring that the algorithm remains classically efficient regardless of how deep the circuit becomes. Second, we establish trainability guarantees: together with an explicit error bound (Lemma 2), we prove absence of barren plateaus (Lemma 3). Finally, our method is designed with near-term constraints in mind, i.e., fixed gate budgets rather than exact zero-error or asymptotic behavior, and is applicable with any circuit ansatz, and not necessarily tied to a brickwall ansatz, making it more flexible than sequential preparation or global overlap-based optimization approaches. These features combined (bond-dimension bounding, trainability, fixed-budget working, and ansatz-independence) are, to our knowledge, not present together in any previous state-preparation strategy.

II. CLASSICAL VARIATIONAL DISENTANGLEMENT

We consider finite matrix product states (MPS) with open boundary conditions, i.e., we assume $D_0 = D_n = 1$:


$$|\psi\rangle = \sum_{b_1, \dots, b_n} A_{[1]}^{b_1} \cdots A_{[n]}^{b_n} |b_1 \dots b_n\rangle$$

(1)

The matrices $A_{[k]} \in \text{Mat}(D_{k-1} \times D_k, \mathbb{C})$ have a maximal bond dimension $D = \max_k D_k$, which is assumed to be constant in system size. As most of the discussion is considering local quantities, CVD can also be applied to map finite chunks of infinitely large systems to quantum circuits. In this case, we can allow for larger bond dimensions at the boundary and contract $A_{[1]}$ and $A_{[n]}$ with D_0 and D_n dimensional vectors $|\phi_{l/r}\rangle$ for the

left and right boundary, respectively. On systems with periodic boundary conditions, the concepts can be similarly applied, although the calculation of the entanglement entropies is more tedious, as there does not exist a canonical form, in general.

A. Overview

In order to minimize entanglement entropies on all bonds simultaneously, we are working with Vidal's Γ - Λ form [23,24], in which the Λ matrices are diagonal, positive-definite matrices carrying the Schmidt values of the bipartite state separated by the respective bond. A finite MPS $|\psi\rangle \in \mathcal{H}$ can always be brought into Γ - Λ form

$$|\psi\rangle = \sum_{b_1, \dots, b_n} \Gamma_1^{b_1} \Lambda_1 \Gamma_2^{b_2} \Lambda_2 \cdots \Lambda_{n-1} \Gamma_n^{b_n} |b_1 \dots b_n\rangle$$

(2)

by applying a singular value decomposition (SVD) on each site [24]. We are looking for a quantum circuit $U = \prod_i G_{i,i+1}$ built from a product of gates $G_{i,i+1} \in \text{SU}(d^2)$, that maps a given MPS $|\psi\rangle \in \mathcal{H} \cong (\mathbb{C}^d)^{\otimes n}$ as above to a product state. For the sake of concreteness, we focus on the qubit case $d = 2$, but CVD works the same way on systems with higher local dimension d if the gates $G_i \in \text{SU}(d^2)$ are adapted adequately.

Once a product state is reached, it can be mapped to $|0\rangle^{\otimes n}$ by one layer of single-qubit gates (see Fig. 1). We will absorb this last layer into U without loss of generality. This is possible since a product state, i.e., an MPS with bond dimension $D = 1$ can be prepared with a Pauli X and Pauli Z rotation with angles $\phi_{x/z}$ that can be read off of the coefficients $\Gamma^0 = \cos(\phi_x)$ and $\Gamma^1 = e^{i\phi_z} \sin(\phi_x)$. The disentangler U automatically leaves us with a state preparation circuit simply by reverse application.

To ensure classical efficiency, the bond dimensions of the MPS are truncated under a map $\mathcal{T}_D : \mathcal{H} \rightarrow \mathcal{H}$ that projects the input state to the closest bond dimension D MPS by cutting off all but the largest D Schmidt coefficients. Even if the disentangler does not create a perfect product state after L layers, we can terminate CVD by applying the truncation map $\mathcal{T}_1$ and reading off the last layer of single-qubit gates as described above. The state that would be ultimately prepared on a quantum computer, $U^\dagger|0\rangle$, carries an error that is discussed in detail in Lemma 2, where we prove an upper bound containing the truncation error

$$\varepsilon = \left\| |\psi\rangle - \left(\prod_{i=L}^1 \mathcal{T}_D \circ G_i^\dagger \right) \circ \mathcal{T}_1 \circ \left(\prod_{i=1}^L \mathcal{T}_D \circ G_i \right) |\psi\rangle \right\|$$

$$= \left\| |\psi\rangle - \left(\prod_{i=L}^1 \mathcal{T}_D \circ G_i^\dagger \right) |0\rangle \right\|, \quad (3)$$

which is calculated by forward application, truncation to the product state $|0\rangle$ (omitting the last layer of single-qubit gates without loss of generality), and finally reverse application of the disentangler while keeping a maximum bond dimension D . We relabeled the gates G_i to be ordered in the brick-wall structure as depicted in Fig. 1. The product $\prod$ orders its terms




FIG. 1. Graphical representation of a unitary disentangler U acting on an MPS $|\psi\rangle$. U is built from a product of local gates $G_{i,i+1}$ (blue) acting on neighboring sites. When the bond dimension is reduced to one, a last layer of single-site gates can be read off to finally map the MPS to the computational basis state $|0\rangle^{\otimes n}$.

from right to left, i.e., $\prod_{i=1}^L G_i$ applies G_1 first and G_L last, and $\prod_{i=L}^1$ is reverse. This error is classically accessible and can thus be used as an additional figure of merit.

B. Disentanglement by entropy minimization

In Γ - Λ form, an efficiently calculable entanglement measure on $|\psi\rangle$ is the α -Rényi entanglement entropy $S_{\mathcal{A},\alpha}(\theta) = \frac{1}{1-\alpha} \log[\text{Tr}(\rho_{\mathcal{A}}^\alpha)]$ of a bipartition into a subsystem $\mathcal{A}$ and its complement $\mathcal{A}^c$. A natural choice for $\mathcal{A}$ and $\mathcal{A}^c$ on an MPS is defined by the i th bond, whose Λ matrix contains the Schmidt values of a Schmidt decomposition with respect to the bipartition into $[i] := \{j | j \leq i\}$ and its complement. The Rényi entropy can thus be calculated without explicitly working with the reduced density matrix $\rho_{\mathcal{A}}$, it reads as

$$\begin{aligned} S_{\mathcal{A},\alpha} &= \frac{1}{1-\alpha} \log[\text{Tr}(\rho_{\mathcal{A}}^\alpha)] \\ &= \frac{1}{1-\alpha} \log\left(\sum_j \lambda_j^{2\alpha}\right), \end{aligned} \quad (4)$$

with $\lambda_j := \Lambda_{jj}$ being the j th diagonal element of the Λ matrix defining the bipartition. In order to find the disentangler, we are treating $S_{\mathcal{A},\alpha}$ as a cost function for variational optimization on the variational reduced density matrix $\rho_{\mathcal{A}}(\theta) = \text{Tr}_{\mathcal{A}^c}[U(\theta)|\psi\rangle\langle\psi|U(\theta)^\dagger]$, or in other words on the variational $\Lambda(\theta)$ matrices. As an ansatz for $U(\theta)$, we choose a brick-wall circuit

$$U(\theta) = \prod_{\mu=1}^L \left(\prod_{i \text{ odd}} G_{i,i+1}^\mu(\theta) \prod_{i \text{ even}} G_{i,i+1}^\mu(\theta) \right), \quad (5)$$

consisting of L layers of the most general disentangler [25] on two qubits, that is,

$$\begin{aligned} G_{i,j}(\theta) &= e^{-i(\theta_{i,1}X_iX_j + \theta_{i,2}Y_iY_j + \theta_{i,3}Z_iZ_j)} \\ &\quad \times (e^{-i(\theta_{i,4}X_j + \theta_{i,5}Y_j + \theta_{i,6}Z_j)} \otimes e^{-i(\theta_{i,7}X_i + \theta_{i,8}Y_i + \theta_{i,9}Z_i)}). \end{aligned} \quad (6)$$

The most general $\text{SU}(4)$ operator from [25] also includes single-qubit gates at the end. These do not need to be part of $G_{i,i+1}$, as they do not remove any entanglement. With nine parameters per gate, the total number of parameters is hence $9(n-1)L$. Although this is the most general ansatz, in some

applications it might be favorable to choose a more restricted gate set tailored to the available hardware, accepting possibly deeper circuits. In that case, the method can be straightforwardly adapted.

To ensure a classically efficient evaluation of the cost function, we perform a layer-by-layer optimization of the parameters θ since a contraction of the full circuit on nonoptimized parameters can lead to an exponential scaling of the bond dimension $D \lesssim 4^L$. The contraction of a single disentangling gate with two sites of the input MPS, on the other hand, is efficient and can be parallelized for every gate in the same layer. After contracting a gate $G_{l,r}$ on two sites l and r , the new Schmidt values Λ'_0 result from an SVD in the following way:

$$\begin{aligned} & \text{Diagram showing the contraction of a gate } G_{l,r} \text{ with two MPS tensors } \Lambda_l \text{ and } \Lambda_r. \\ & \text{The result is a single tensor } \Lambda'_0 \text{ with bond dimension } D \leq 4. \\ & \text{The decomposition is } U_l \Lambda'_0 V_r. \end{aligned} \quad (7)$$

Using a matrix product operator (MPO) form of the gate $G_{l,r}$ shows that the contraction of $\Gamma_{l/r}$ with $G_{l,r}$ can be carried out efficiently paying the prize of an increase in bond dimension of Λ_0 by a factor of 4 in the worst case. In the second step, all tensors are being contracted and isometries U_l and V_r as well as the new singular values Λ'_0 have been defined via singular value decomposition (SVD). Finally, the new Γ tensors are defined via

$$\Gamma'_l = \Lambda_l^{-1} U_l \quad \text{and} \quad \Gamma'_r = V_r \Lambda_r^{-1}. \quad (8)$$

This allows for an efficient evaluation of the new entanglement entropy by plugging Λ'_0 into Eq. (4). While the bond dimension could in principle grow by a factor of 4 in each such layer, the optimizer actively tries to reduce the rank of Λ_0 and therefore also the necessary bond dimension by minimizing the entanglement entropy. We give a rigorous explanation of this in Lemma 1.

It is clear that the Rényi index has an influence on the change in bond dimension. In the limit $\alpha \rightarrow 0$, $S_{\mathcal{A},\alpha}$ reduces to $\log(D)$, whereas for $\alpha \rightarrow \infty$, the entropy only considers the




FIG. 2. Bond dimensions for disentanglement of the ground state of a 1D Ising model with skew magnetic field of strengths $h_x = 0.5$ and $h_z = 0.05$ [see Eq. (13)]. After the target MPS is prepared with a DMRG, CVD is carried out using different values for the Rényi index α showcasing a suppression of the maximal bond dimension D for small indices α in the presence of an error threshold $p = 10^{-7}$. A bond dimension larger than one should not be interpreted as a failure of disentangling. All three algorithms converge to a per-site overlap error of approximately $\varepsilon = 10^{-4}$.

largest Schmidt value $S_{\mathcal{A},\alpha} \rightarrow -\log(\lambda_1^2)$. The Rényi index thus serves as an adaptive control parameter to balance the exploration of entangled states and classical efficiency via a maximal bond dimension. Figure 2 illustrates the application of CVD to the ground state of a transverse-field Ising model on a one-dimensional chain using different Rényi indices. While the precision of state preparation does not differ significantly, the suppression of the bond dimension by a small Rényi index becomes apparent. Note that a bond dimension larger than one does not imply a failure of the disentangling procedure or residual entanglement. While reducing the bond dimension ensures classical representability, the relevant figure of merit is the fidelity, which is high in all cases of Fig. 2. We refer to Sec. III for a more detailed discussion on numerical experiments. In the following section, we show how a successful minimization of $S_{\mathcal{A},\alpha}$ yields guarantees for classical efficiency.

C. Classical efficiency

We prove a variant of Lemma 2 from [3] to give an explicit form of the bound on the bond dimension, given an entanglement entropy value as output of the optimizer.

Lemma 1. Let $S_{\mathcal{A},\alpha}(p_i) = \frac{1}{1-\alpha} \log(\sum_i p_i^\alpha)$ with $p_i := \lambda_i^2$ be the α -Rényi entropy with $0 < \alpha < 1$ given by the Λ matrix separating an MPS into left and right subsystems $\mathcal{A}$ and $\mathcal{A}^c$. Let $p = \sum_{i=D+1}^{2^n} p_i$ be the error from cutting off the squared Schmidt coefficients at bond dimension D . Then,

$$D \leq \frac{(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}}}{p^{\frac{\alpha}{1-\alpha}}} e^{S_{\mathcal{A},\alpha}}. \quad (9)$$

Lemma 1 is derived by minimization of the entanglement entropy, while considering the normalization constraint and a minimal cutoff threshold for the Schmidt values (see Appendix A 1 for details).

We observe that for $S_{\mathcal{A},\alpha} \rightarrow 0$, the bound on the rank goes to a constant that can be larger than one depending on p . Since $S_{\mathcal{A},\alpha} = 0$ implies $D = 1$, the bound becomes loose in this limit. Equation (9) supports the claim that an optimization for entropies with small α focus on lowering the bond dimension since the α -dependent prefactor becomes small for small α and diverges for $\alpha \rightarrow 1$ like $(\frac{1}{p})^{\frac{\alpha}{1-\alpha}}$.

In case of arbitrary precision $p \rightarrow 0$, the rank is unbounded. To ensure efficiency, the MPS has to be truncated regularly introducing the approximation error p at each bond. The unitary disentangling process U thus becomes $\prod_{i=1}^L \mathcal{T}_D \circ G_i$ with the truncation map $\mathcal{T}_D$ enabling the classical calculation of an error measure ε from Eq. (3). However, a state preparation on a quantum machine does not contain truncations $\mathcal{T}_D$, but prepares $U^\dagger|0\rangle$. We thus tackle the question about the total error being made in CVD when keeping a constant bond dimension D .

Lemma 2. Let $U = \prod_{i=1}^L G_i$ be a unitary disentangler resulting from a CVD with a precision $\varepsilon = \|\psi\rangle - (\prod_{i=1}^L \mathcal{T}_D \circ G_i)|0\rangle\|$ and maximal bond dimension D , then the total error of state preparation is

$$\|\psi\rangle - U^\dagger|0\rangle\| \leq \varepsilon + L\sqrt{2(n-1)p}, \quad (10)$$

with $p = \max_k p^{[k]}$ and $p^{[k]}$ denoting the truncation error at the k th bond.

See Appendix A 2 for a proof.

D. Absence of barren plateaus

The guarantees for classical efficiency rely on the performance of the optimizer finding transformations which lower the entanglement entropy. For a complete discussion on classical efficiency, we include a note on trainability of the variational ansatz due to the existence of significant gradients. Quantum variational algorithms are known to suffer from

barren plateaus [2,26], which ultimately leads to a sample problem in resolving the direction of the gradient. While in classical computations with MPS, the full state is available, an exponentially small gradient is not necessarily a sample problem, but still imposes an overhead in the convergence of the optimization. In the following, we show that variational disentangling does not suffer from barren plateaus, which becomes apparent from the locality of Eq. (7), pointing out the decomposition of the problem into gate-by-gate optimizations on two sites. In order to be able to calculate expressions of the gradient, we focus on the 2-Rényi entropy and its gradient that reads as

$$\partial_i S_{\mathcal{A},2}(\theta) = -\frac{2}{\text{Tr}[\rho_{\mathcal{A}}(\theta)^2]} \text{Tr}[\rho_{\mathcal{A}}(\theta) \partial_i \rho_{\mathcal{A}}(\theta)]$$

$$\text{with } \partial_i \rho_{\mathcal{A}}(\theta) = \text{Tr}_{\mathcal{A}^c}[(\partial_i U(\theta))|\psi\rangle\langle\psi|U(\theta)^\dagger + \text{H.c.}]. \quad (11)$$

The calculation straightforwardly generalizes to Rényi indices $\alpha > 1$, however, gradients might become unstable for $\alpha \leq 1$ in the presence of small Schmidt values. However, this is not a problem in practice since the introduction of a cutoff threshold leads to the truncation of small Schmidt values before entering a singularity. At $\theta = 0$, the elements of $\partial_i \rho_{\mathcal{A}}(\theta)$ are of the form $-i\text{Tr}_{\mathcal{A}^c}([\sigma_l \otimes \sigma_r, |\psi\rangle\langle\psi|])$ with $\sigma_{l/r}$ having support across the boundary between $\mathcal{A}$ and $\mathcal{A}^c$. An evaluation of Eq. (11) in graphical notation allows a formal identification of the extrema. Next to product states ($\Lambda_0 = 1$) being minima and maximally entangled states being maxima, states with real-valued coefficients represent saddle points. This property is discussed in detail in Appendix A3 and ensures that a gradient calculated from finite differences will thus not halt at the unstable fixed points, but move towards product states, as desired, or possibly get stuck in local minima.

In parameter spaces of large dimension, however, cost landscapes with small typical gradients, or barren plateaus, are a more serious concern than local minima would be. In order to make a statement about generic MPS with a fixed bond dimension D , we calculate the first and second moments of the gradient from Eq. (11) evaluated on randomly distributed MPS with bond dimension D . Using the isometric property of $\Lambda_l \Gamma_l$ and $\Gamma_r \Lambda_r$, we can utilize statistical tools for the Haar measure of $\text{SU}(2D)$, leading to a quantification of the variance of the gradient $\partial_i S_{\mathcal{A},2}$ in terms of the bond dimension.

Lemma 3. Let $A = \Lambda_l \Gamma_l$ and $B = \Gamma_r \Lambda_r$ be random isometries with fixed bond dimension D and Λ_0 a matrix of Schmidt values. The expected value of the gradient $\partial_i S_{\mathcal{A},2}|_{\theta=0}$ is zero and the variance is

$$\mathbb{E}[(\partial_i S_{\mathcal{A},2}|_{\theta=0})^2] = \frac{8}{(2D+1)^2} (1 - e^{2(S_{\mathcal{A},2} - S_{\mathcal{A},3})}), \quad (12)$$

for i corresponding to $\sigma_l \otimes \sigma_r \in \{X \otimes X, Y \otimes Y, Z \otimes Z\}$.

The proof consists of a straightforward, but technical, calculation in Weingarten calculus involving integrals over the Haar measure on $\text{SU}(2D)$ to account for the random isometries A and B . Technical details are provided in Appendix A3.

The statement of Lemma 3 is more general than it may first appear. The restriction to $\theta = 0$ can be lifted as long as a constant maximal bond dimension is ensured by regularly truncating the state. Doing so, an error bounded by Eq. (10) is introduced. Typically, barren plateaus are diagnosed by

randomizing over the parameters θ rather than the state that is acted upon. One can view Lemma 3 as an equivalent perspective whenever the unitary $G_{l,r}(\theta)$ is absorbed into the tensors A and B . In that case, one would have to consider a bond dimension $4D$. While this argument ensures the absence of barren plateaus, local minima can still occur. However, in practice these are mitigated by alternating between even and odd pairs after each layer, which maintains optimization momentum and reduces the chance of getting trapped.

Furthermore, it is important to distinguish the local nature of CVD from approaches based on the quantum marginal problem [27]. Starting from an MPS in canonical form allows us to locally access entanglement properties across bipartitions separating all sites to the left and right of a given bond. The corresponding local MPS tensors generally encode more information than the two-site reduced density matrix on which the optimized unitary acts.

The key distinction is that disentangling a two-site marginal addresses only strictly local correlations, whereas the MPS representation captures correlations on finite length scales beyond nearest neighbors. We refer to Sec. IIIC for illustrative examples.

III. NUMERICAL EXPERIMENTS

In order to test the performance and classical efficiency guarantees of CVD, we carried out a number of numerical experiments. We first consider ground states of one-dimensional (1D) spin Hamiltonians and an approximation to the ground state of a Fermi-Hubbard model found classically by the density matrix renormalization group (DMRG) method. We plot the ∞ -Rényi entropy as a measure for the tail weight of the Schmidt values on each bond and for each layer, as well as the per-site contribution of the overlap error ε from Eq. (3). The overlap error is calculated between the input MPS and an MPS resulting from forward propagation of the disentangler, truncation to bond dimension $D = 1$, and finally backwards propagation of the disentangler.

Since physical states of locally interacting spin systems are typically weakly correlated, we also test CVD on an artificial state of high bipartite entanglement. In particular, we disentangle a logical Bell pair, where two qubits are encoded in a quantum error-correcting code each. Although the entanglement is distributed among many qubits, we find that CVD successfully disentangles the state only considering two neighboring sites at a time. In Appendix B, we reproduce preparation circuits for known MPS structures like the Greenberger-Horne-Zeilinger (GHZ) state, the cluster state, and the Affleck-Kennedy-Lieb-Tasaki (AKLT) state exhibiting different ranges of correlation length.

All numerical experiments have been carried out with using the julia library “ITensor” [28] and can be reproduced using the code available at [29].

A. Spin model ground states

An interesting quantum computing application of MPS preparation is the study of quantum many-body physics. Especially on one-dimensional spin chains, MPS have been proven faithful representations of ground states of gapped




FIG. 3. Disentanglement of the ground states of the Ising model with skew magnetic field $h_x = 0.5$, $h_z = 0.05$, the XY model with $g = 0$ and a magnetic field in X direction, $h_x = 0.1$, and the Heisenberg XXZ model with $J_z = 0.5$ and a skew magnetic field $h_z = 0.1$ and $h_x = 0.1$ defined on 50 sites. The target state has been determined using a DMRG routine and subsequently disentangled using a Rényi index $\alpha = \frac{1}{2}$ on every third layer and $\alpha = 1$ else. The color plots (a) show a gradual decrease of the ∞ -Rényi entropy starting from the boundary. As an approximation to the tail weight, it measures the distance of the MPS to a product state across the bond. Since the overlap error (b) also saturates for all three spin models, this indicates that there exist local minima with unresolved correlations. The system-size scaling (c) of the final errors after 32 layers saturate at around 10^{-4} with one outlier for the 50-site Ising model.

Hamiltonians with local interactions [3,4]. In this section, we find ground-state preparation circuits for the following spin Hamiltonians:

$$H_{\text{Ising}} = -\frac{1}{2} \sum_i Z_i Z_{i+1} + h_x \sum_i X_i + h_z \sum_i Z_i, \quad (13)$$

$$H_{XY} = -\frac{1}{2} \sum_i (X_i X_{i+1} + Y_i Y_{i+1}) + h_x \sum_i X_i, \quad (14)$$

$$H_{\text{Heisenberg}} = -\frac{1}{2} \sum_i (X_i X_{i+1} + Y_i Y_{i+1}) - \frac{J_z}{2} \sum_i Z_i Z_{i+1} - h_x \sum_i X_i - h_z \sum_i Z_i, \quad (15)$$

by first finding an MPS representation of the respective ground state using a density matrix renormalization group (DMRG) algorithm and subsequently disentangling the target MPS with CVD. There is a freedom in balancing interaction strengths in the Heisenberg model via J_z and tuning the magnetic field strengths h_x, h_z . Choosing a state deep in a phase more likely results in extremely shallow,

or even trivial, preparation circuits, as the ground states can be very close to product states for some models. We thus pick near-critical models with nontrivial MPS representations and add small magnetic field values to ensure nonintegrability.

In Fig. 3, we show results for the ground-state preparation of the three spin models above. Next to the evolution of the tail weight $S_\infty = -\log(\lambda_1^2) \approx 1 - \lambda_1^2$ of the Schmidt coefficients λ_i at any bond and after every disentangling layer, we plot the overlap error per site,

$$\epsilon_{\text{site}} = 1 - (1 - \epsilon)^{\frac{1}{n}}, \quad (16)$$

with ϵ being taken from Eq. (3). We choose this figure of merit as it reflects the error of the prepared state at the level of the individual matrices in the MPS, which contribute to the overall fidelity multiplicatively.

The ground state of the Ising model with a skew magnetic field $h_x = 0.5$ and $h_z = 0.05$ appears to have comparatively low state complexity, as it is prepared with high fidelity using around 10 brick-wall layers for a state on 50 sites.

The ground states of the XY (with $h_x = 0.1$) and Heisenberg XXZ models (with $J_z = 0.5$ and $h_x = h_z = 0.1$), on the other hand, converge to per-site overlap errors of just above 10^{-4} , which is small compared to near-term gate fidelities. From a near-term perspective, it hence makes sense to stop at this point and spend the rest of the limited gate budget on the actual computation after the initial state is prepared. Higher fidelities are achievable by tuning down the cutoff parameter p (compare Lemma 1) and allowing for higher maximal bond dimensions along the way. The tail weights shown in Fig. 3(a) display a gradual concentration towards a product state, as S_∞ decreases starting from the boundary towards the center bonds. The maximal bond dimension of $D = 100$ was never reached during CVD.

Using the MPS representation of the ground state of a local, gapped Hamiltonian defined on a one-dimensional lattice is a very convenient position to take, as there are guarantees for a faithful representation [3] leaving us with practically no restriction in system size. To showcase this fact, we plot the overlap error per site, $\varepsilon_{\text{site}}$, of state preparations of the spin model ground states after 32 layers in Fig. 3(c). Although the state complexities differ, resulting in different final error levels, the per-site error saturates. While it is to be expected that the total fidelity faces an orthogonality catastrophe, i.e., decreases exponentially in the system size n , a constant per-site error further supports the scalability of CVD.

B. 1D Fermi-Hubbard ground state

The problem of preparing low-temperature equilibrium states in the canonical ensemble of electronic structure Hamiltonians, even on one-dimensional lattices, targets highly entangled states. Tensor network approaches like DMRG are struggling to find scalable approximations, in general [30,31]. This is why quantum algorithms are being investigated for the quantum simulation of electronic structure problems, but (like most quantum algorithms) they suffer from scalability problems in sample [2] and gate complexity [32]. An educated initial guess with sufficient overlap with the target state eases the ground-state problem possibly even by an exponential factor [33]. This is where CVD comes into play. A (potentially faulty) DMRG and subsequent CVD yield a preparation circuit that produces a state in the low-energy subspace, for which efficient ground-state algorithms have been investigated, for instance via digital cooling [1] or with variational algorithms [34]. In this section, we illustrate this classical preprocessing step on the Fermi-Hubbard model with Hamiltonian

$$H_{\text{FH}} = -t \sum_{i,\sigma} (c_{i,\sigma}^\dagger c_{i+1,\sigma} + c_{i,\sigma} c_{i+1,\sigma}^\dagger) + U \sum_i n_{i,\uparrow} n_{i,\downarrow}, \quad (17)$$

constructed from creation and annihilation operators $c_{i,\sigma}^\dagger$ and $c_{i,\sigma}$ with site indices $i \in \{1, \dots, n\}$ and spin indices $\sigma \in \{\uparrow, \downarrow\}$. The operators $n_{i,\sigma} := c_{i,\sigma}^\dagger c_{i,\sigma}$ are the number operators on site i and spin σ . Running a DMRG within a particle-number-conserving subspace yields an MPS with local dimension 4. This state can be mapped to qubits yielding an MPS with local dimension 2, which is suitable for CVD.

Each electronic site is separated into two neighboring qubit sites representing the spin-up and -down degrees of freedom, respectively. Although the DMRG with heavily restricted bond dimension might not be quite successful in finding the ground state of H_{FH} , it provides a low-energy state as an educated guess for a potential quantum computing application. For illustration, we probe the energy of an MPS approximation with bond dimension $D = 5$ of the Fermi-Hubbard ground state on 25 electronic sites or 50 qubit sites, respectively. We compare the energy of the state prepared by local unitary gates using CVD to a benchmark of the ground-state energy using DMRG on an MPS with bond dimension $D = 1000$ as well as the energy of the target state ($D = 5$) in Fig. 4. Even though the prepared quantum state lies far above the actual ground state in energy, it still exemplifies how CVD can generate structured initial states that may reduce the complexity of a subsequent quantum solution to the ground-state problem. Importantly, our example is meant as a preparatory step rather than a proposal for solving ground-state problems in 1D, where DMRG is already highly efficient. The task becomes much harder and more interesting in 2D, where the Fermi-Hubbard model is QMA-hard [30], but can be reduced to BQP when a ground-state guess with inverse-polynomial precision is available [33]. Cooling algorithms such as [1] only require that the initial state lies in a sufficiently low-energy subspace; in some instances this can lead to high overlap with the ground state, but this is certainly not guaranteed in general, assuming no collapse of complexity classes. With an appropriate snaking of the 2D lattice, a faulty MPS approximation of the ground state can thus be transferred to the quantum processor.

C. Entangled logical qubits

Going beyond ground states of local Hamiltonians, in this section we challenge CVD on states whose entanglement is distributed such that it is not removable by single local operations. Consider the example of a Bell pair distributed on far away sites. Its state complexity is proportional to the distance of the entangled qubits, as the state can only be disentangled by first swapping the pair to neighboring sites and subsequently applying a disentangling gate, a CNOT gate for instance. Swapping the Bell pair closer together only decreases the entanglement entropy of the leftmost and rightmost bonds. Thus, there is hope that CVD, which only looks at pairs of two neighboring sites at a time, can find a good approximation of this state.

In order to make the example more interesting, we additionally smear out the entanglement among many qubits by encoding the Bell pair in a stabilizer code. There, local operations should not alter the state, if stabilizer measurements are carried out fast enough. We thus expect that a single layer of unitaries cannot significantly disentangle the state alone, but there will be a minimal necessary depth to be able to concentrate the entanglement entropy into two neighboring qubits. We consider two examples of stabilizer codes with distances three and five, in particular the $[5,1,3]$ code and the $[[11,1,5]]$ code defined by the symmetric subspace of the stabilizers displayed in Table I. Our goal is to disentangle the logical Bell pair $\frac{1}{\sqrt{2}}(|01\rangle_L - |10\rangle_L)$ that defines




FIG. 4. Disentanglement of an MPS approximation to the ground state of the Fermi-Hubbard model with hopping parameter $t = 1$ and Coulomb interaction strength $U = 2$ on 25 electronic sites at half-filling (12 spins up and 12 spins down) using $\alpha = 1$. These are encoded into 50 qubits with each qubit the occupation of a spin-up or -down particle, respectively. Next to a plot of the tail weight (a), the energy of the prepared low-rank state is plotted against the layer count (b). The ground-state energy is approximated by a DMRG calculation with bond dimension $D = 1000$ (orange) and the target state ($D = 5$) energy is plotted in green.

a 10-qubit or 22-qubit quantum state, respectively, shown in Fig. 5(d).

We plot the results for both the [5,1,3] and the [11,1,5] codes for separated and interlaced logical Bell pairs in Fig. 5. While the separated states can be prepared exactly (up to numerical precision) with approximately $\frac{n}{2}$ layers, disentangling the interlaced logical qubits needs a deeper circuit to properly swap the entanglement to neighboring sites. Although the optimizer eventually reaches overlap errors of below 10^{-3} per

site, the interlaced [11,1,5] Bell pair is not exactly prepared, but seems to be stuck in a local minimum. The evolution of the tail weights behaves similar in both codes, while the decrease of the tail weights in the [11,1,5] code is evidently subject to fluctuations. It appears that CVD, being a strictly local disentangling strategy, is having trouble finding the correct order of swapping back the interlaced qubits before removing the nonlocal entanglement as in the separated case. The attempt to directly remove entanglement from the start, i.e., without swapping qubits around first, leads the optimizer into a convergence barrier.

TABLE I. Stabilizers for the [5,1,3] stabilizer code [35] (a) and the [11,1,5] stabilizer code [36] (b) used for the disentangling test case shown in Fig. 5.

(a)										
$S_1^{[5]}$	I	X	Z	Z	X					
$S_2^{[5]}$	X	Z	Z	X						
$S_3^{[5]}$	Z	Z	X	I						
$S_4^{[5]}$	Z	X	I	X						
(b)										
$S_1^{[11]}$	Z	Z	Z	Z	Z	I	I	I	I	I
$S_2^{[11]}$	X	X	X	X	X	I	I	I	I	I
$S_3^{[11]}$	I	I	I	Z	X	Y	Y	Y	X	Z
$S_4^{[11]}$	I	I	I	X	Y	Z	Z	Z	Y	X
$S_5^{[11]}$	Z	Y	X	I	I	I	Z	Y	X	I
$S_6^{[11]}$	X	Z	Y	I	I	I	X	Z	Y	I
$S_7^{[11]}$	I	I	I	Z	Y	X	X	Y	Z	I
$S_8^{[11]}$	I	I	I	X	Z	Y	Z	X	Y	I
$S_9^{[11]}$	Z	X	Y	I	I	I	Z	Z	Z	X
$S_{10}^{[11]}$	Y	Z	X	I	I	I	Y	Y	Y	Z

IV. CONCLUSION

We present classical variational disentanglement, or CVD, a classical mapping of MPS representations to quantum circuits for state preparation. The classically efficient representation of MPS using local information is utilized to optimize a parametrized circuit to minimize the entanglement entropy for bipartitions defined by each bond of the MPS. We prove guarantees for classical efficiency ensuring trainability and bounds on the maximal bond dimension and the error of the output state and showcase the algorithm on a number of examples.

We expect CVD to be a flexible preprocessing routine for the simulation of quench dynamics, where the initial state is taken as the ground state of a gapped, local Hamiltonian. Since the initial-state preparation is not the quantumly hard part, it should be processed classically. Also, generic ground-state algorithms like the variational quantum eigensolver, imaginary-time evolution or cooling protocols benefit from an educated guess, which CVD provides. CVD also improves the access to entangled bases for sample efficient




FIG. 5. Disentanglement of logical Bell pair in a [5,1,3] and [11,1,5] stabilizer code using $\alpha = 1$ in both the separated and the interlaced versions as depicted in (d). In (a) and (b), the tail weights on each bond and for each layer in the disentangling process are shown and in (c), the overlap errors from backwards preparation of the initial MPS in the [11,1,5] code are shown. We find that in all cases a linear circuit depth is required to fully disentangle the state. Although the bond dimension is decreasing along the way, the overlap error undergoes an abrupt jump.

measurements [37], as well as for tomography of low-rank states [38] and processes [39], which are further known ways of connecting classical and quantum information processing. CVD is not restricted to disentangling only local correlations, but can in practice become trapped in local minima that yield good approximations but neglect longer-range correlations. Such cases may be alleviated by blocking sites before optimization, which offers a natural extension of the method to more strongly correlated systems.

The idea of disentangling tensor networks with local gates reaches beyond state preparation. Interesting generalizations include classical compilation of circuits for time evolution by disentanglement of matrix product operators (MPOs) going beyond Trotterization-based classical compilation schemes like [40–42] or sequential preparation [43]. Similarly, matrix product functions (MPFs) have been successfully applied for machine-learning tasks [44,45] and thus their compilation into circuits (potentially including nonunitary structures) provide classical guesses for quantum machine-learning models boosting trainability with CVD as a systematic segue past barren plateaus and towards narrow gorges.

ACKNOWLEDGMENTS

This research was funded in parts by the Austrian Science Fund (FWF) via the Cluster of Excellence “quantA” (Grant No. 10.55776/COE1) and via the Special Research Programme “BeyondC: Quantum Information Systems Beyond Classical Capabilities” (Grant No. 10.55776/F71), by the European Union–NextGenerationEU, and by the European Union’s Horizon 2020 research and innovation programme through Grant No. 863476 (ERC-CoG SEQUAM).

DATA AVAILABILITY

The data that support the findings of this article are openly available [29].

APPENDIX A: CLASSICAL EFFICIENCY OF CVD

In this appendix, we provide technical details for the proofs of classical efficiency of CVD including a bound on the bond dimension for a given entanglement entropy value, an error

bound for the approximation of the prepared state by truncation of MPS and finally a proof of absence of barren plateaus in optimization. Together these results yield good evidence for efficient disentanglement.

1. Bounds on bond dimension: Proof of Lemma 1

Just as the entanglement entropy is bounded by the bond dimension D , vice versa for a given entropy value $S_{\mathcal{A},\alpha}$, the bond dimension is bounded. We derive a relation for the minimal entropy defined by the squared Schmidt coefficients $\lambda_i^2 =: p_i$, which also satisfy the normalization constraint $\sum_{i=1}^D p_i = 1 - p$. The entropy constraint can be reformulated to $\sum_{i=1}^D p_i^\alpha = e^{S_{\mathcal{A},\alpha}(1-\alpha)}$. We observe that the gradient of the Lagrangian $\sum_i p_i^\alpha - \mu(\sum_i p_i - 1)$ becomes zero only at the maximum where $p_i = p_j \forall i, j$. The concavity of the Rényi entropy implies that the minimum is taken at the boundary where $p_1 = 1 - p - (D-1)h$ and $p_2 = \dots = p_D = h$, such that

$$\sum_{i=1}^{2^n} p_i^\alpha \geq [1 - p - (D-1)h]^\alpha + (D-1)h^\alpha + \frac{p}{h} h^\alpha, \quad (\text{A1})$$

using $p_i \leq h$ for all $i > D$ and therefore $\sum_{i>D} p_i^{\alpha-1+1} \geq \frac{p}{h} h^\alpha$. The normalization constraint and $p_i \geq h$ for all $i \leq D$ further implies $Dh \leq 1 - p$, which we can use to make the estimate

$$\begin{aligned} \sum_i p_i^\alpha &\geq Dh^\alpha + ph^{\alpha-1} \\ &\geq D^{1-\alpha} p^\alpha \left(\left(\frac{1-\alpha}{\alpha} \right)^\alpha + \left(\frac{1-\alpha}{\alpha} \right)^{\alpha-1} \right) \\ &= \frac{D^{1-\alpha} p^\alpha}{\alpha^\alpha (1-\alpha)^{1-\alpha}}. \end{aligned} \quad (\text{A2})$$

The second inequality represents the global minimum of $Dh^\alpha + ph^{\alpha-1}$ as a function of h . Plugging in $\sum_i p_i^\alpha = e^{S_{\mathcal{A},\alpha}(1-\alpha)}$ yields the result. ■

2. Approximation error bound: Proof of Lemma 2

We show the error bound $\| |\psi\rangle - U^\dagger |0\rangle \| \leq \varepsilon + L\sqrt{2(n-1)p}$ by separating the error into two parts. The first part is the accuracy of the optimizer $\varepsilon = \| |\psi\rangle - (\prod_{i=L}^1 \mathcal{T}_D \circ G_i^\dagger) |0\rangle \|$, which can be efficiently calculated classically. The second part is the truncation error arising for the sake of a bounded bond dimension. By the triangle inequality, we can extract ε via

$$\left\| |\psi\rangle - \prod_{i=L}^1 G_i^\dagger |0\rangle \right\| \leq \varepsilon + \left\| \left(\prod_{i=L}^1 (\mathcal{T}_D \circ G_i^\dagger) - \prod_{i=L}^1 G_i^\dagger \right) |0\rangle \right\|. \quad (\text{A3})$$

We show that every layer in the second term contributes at most linearly to the worst-case error. To do this, we prove the following estimation by induction:

$$\begin{aligned} &\left\| \left(\prod_{i=L}^1 (\mathcal{T}_D \circ G_i^\dagger) - \prod_{i=L}^1 G_i^\dagger \right) |\phi\rangle \right\| \\ &\leq \sum_{i=1}^L \left\| (\mathcal{T}_D \circ G_i^\dagger - G_i^\dagger) \prod_{j=L}^{i+1} (\mathcal{T}_D \circ G_j^\dagger) |\phi\rangle \right\| \forall |\phi\rangle. \end{aligned} \quad (\text{A4})$$

The statement is trivially correct for $L = 1$ and the induction step is done via

$$\begin{aligned} &\left\| \left(\prod_{i=L+1}^1 (\mathcal{T}_D \circ G_i^\dagger) - \prod_{i=L+1}^1 G_i^\dagger \right) |\phi\rangle \right\| \\ &= \left\| \left(\prod_{i=L+1}^1 (\mathcal{T}_D \circ G_i^\dagger) - \left(\prod_{i=L}^1 G_i^\dagger \right) (\mathcal{T}_D \circ G_{L+1}^\dagger) + \left(\prod_{i=L}^1 G_i^\dagger \right) (\mathcal{T}_D \circ G_{L+1}^\dagger) - \prod_{i=L+1}^1 G_i^\dagger \right) |\phi\rangle \right\| \\ &\leq \sum_{i=1}^L \left\| (\mathcal{T}_D \circ G_i^\dagger - G_i^\dagger) \left(\prod_{j=L}^{i+1} (\mathcal{T}_D \circ G_j^\dagger) \right) (\mathcal{T}_D \circ G_{L+1}^\dagger) |\phi\rangle \right\| + \| (\mathcal{T}_D \circ G_{L+1}^\dagger - G_{L+1}^\dagger) |\phi\rangle \|, \end{aligned} \quad (\text{A5})$$

where we used the triangle inequality in the last step together with the induction assumption on the state $(\mathcal{T}_D \circ G_{L+1}^\dagger) |\phi\rangle$ for the first term and unitary invariance of the Euclidean norm for the second. It has been shown [3,46] that the truncation error of a state $|\psi\rangle$ to a bond dimension D state $\mathcal{T}_D |\psi\rangle$ is bounded by $\| |\psi\rangle - \mathcal{T}_D |\psi\rangle \|^2 \leq 2 \sum_k \sum_{i=D+1}^{2^n} (\lambda_i^{[k]})^2 =: 2 \sum_k p^{[k]}$ with $\lambda_i^{[k]}$ being the i th Schmidt value on the k th bond. We can estimate the right-hand side of Eq. (A4) by the truncation errors of the bond dimension D MPS $|\phi_i\rangle := \prod_{j=L}^{i+1} (\mathcal{T}_D \circ G_j^\dagger) |\psi\rangle$. This reduces to L truncation errors, which are in the worst case equal to $\sqrt{2(n-1)p}$. ■

Although, we can find cases in which the triangle inequalities and the bounds by maximizing all tail Schmidt values are being exhausted, we expect the “typical error” of the state preparation to be below the derived bound.

3. Absence of barren plateaus in CVD: Proof of Lemma 3

At $\theta = 0$, the gradient from Eq. (11) has a relatively simple form and can be calculated in graphical notation to be

$$\begin{aligned}
 & -i \text{Tr}_{\mathcal{A}} (\text{Tr}_{\mathcal{A}^c} (|\psi\rangle\langle\psi|) \text{Tr}_{\mathcal{A}^c} ([\sigma_j^{(l)} \otimes \sigma_j^{(r)}] |\psi\rangle\langle\psi|)) \\
 &= \text{Diagram 1} - \text{Diagram 2} \\
 &= \text{Diagram 3} - \text{Diagram 4} = 2i \text{Im} \left(\text{Diagram 5} \right), \tag{A6}
 \end{aligned}$$

The diagrams are graphical representations of tensor networks. Diagram 1 and 2 show the difference of two terms. Diagram 3 and 4 show the same terms with different internal indices. Diagram 5 is a single term inside an imaginary part and a factor of 2i.

using the isometry property of $\Gamma_r \Lambda_r$ and $\Lambda_l \Gamma_l$. It becomes evident that for a product state, i.e., $\Lambda_0 = 1$, the gradient vanishes. As the Rényi entropy is always non-negative, product states are global minima of the cost function. The gradient also vanishes exactly, whenever the Schmidt values are evenly distributed, $\Lambda_0 \propto \mathbb{1}$, or when all transition amplitudes between Schmidt vectors over $\sigma_l \otimes \sigma_r$ are real. In these cases, we are facing an unstable fixed point.

Consider the example of two qubits in the state $|\psi\rangle = \alpha|00\rangle + \beta e^{i\phi}|11\rangle$. The gradient for the 2-Rényi entanglement entropy reads as

$$\begin{aligned}
 -i \text{Tr}(\rho \partial_j \rho) &= 2 \text{Im}(e^{-i\phi} \alpha \beta^3 \langle 1|\sigma_l|0\rangle \langle 1|\sigma_r|0\rangle \\
 &+ e^{i\phi} \alpha^3 \beta \langle 0|\sigma_l|1\rangle \langle 0|\sigma_r|1\rangle), \tag{A7}
 \end{aligned}$$

with two nonvanishing elements for $\sigma_l = \sigma_r = X$: $-i \text{Tr}(\rho \partial_X \rho) = 2 \sin(\phi)(\beta \alpha^3 - \alpha \beta^3)$ and similarly $\sigma_l = \sigma_r = Y$. One can quickly see that the gradient vanishes for the product state $\alpha = 0$ or $\beta = 0$, as well as for the Bell state $\alpha = \beta$. As the Bell state is maximally entangled, it can only be a maximum of the entanglement entropy and thus an unstable fixed point. Indeed, adding a slight imbalance in the amplitudes $\alpha^2 \rightarrow \frac{1}{2} + \delta$, $\beta^2 \rightarrow \frac{1}{2} - \delta$ results in an increasing gradient. Another case, in which the gradient vanishes, is $\phi = 0$, where all transition amplitudes with respect to $\sigma_l \otimes \sigma_r$ are real valued. Also here, the gradient increases whenever $\phi \rightarrow \pm\delta$ is perturbed around the fixed point $\phi = 0$. The gradient then pulls toward $\alpha = 1$ or $\beta = 1$ depending on the sign of ϕ .

More generally, a state $|\psi\rangle = \sum_j \lambda_j |\phi_j^{(l)}\rangle \otimes |\phi_j^{(r)}\rangle$ with only real transitions $\langle \phi_j^{(l)} | \sigma^{(l)} | \phi_k^{(l)} \rangle \langle \phi_j^{(r)} | \sigma^{(r)} | \phi_k^{(r)} \rangle$ will have a vanishing gradient. However, that fixed point is easily escaped by tilting the relative phase towards one of the $|\phi_j^{(r)}\rangle$, for instance.

We now prove Lemma 3 using the result of Eq. (A6). To do so, we consider uniformly distributed isometries $A = \Lambda_l \Gamma_l$ and $B = \Gamma_r \Lambda_r$ and calculate the first and second moments of the gradient $\partial_i S_{\mathcal{A},2}|_{\theta=0}$. Using the isometric property of A and B , we can formulate the calculation in terms of unitaries U, V defined via

$$\begin{aligned}
 \text{Diagram A} &=: \text{Diagram U} \\
 \text{Diagram B} &=: \text{Diagram V}
 \end{aligned} \tag{A8}$$

The diagrams show the mapping from isometry blocks A and B to unitary blocks U and V, with a reference state |0>.

and consider them distributed via the Haar measure on $\text{SU}(2D)$. With this, we can integrate polynomials of U and $U^\dagger$ or V and $V^\dagger$, respectively, using Weingarten calculus, in particular we make use of the formulas

$$\int_{\text{SU}(2D)} dU U_{ab} \overline{U_{cd}} = \frac{\delta_{ac} \delta_{bd}}{2D}, \tag{A9}$$

$$\begin{aligned}
 & \int_{\text{SU}(2D)} dU U_{ab} U_{cd} \overline{U_{ef}} \overline{U_{gh}} \\
 &= \frac{1}{4D^2 - 1} (\delta_{ae} \delta_{bf} \delta_{cg} \delta_{dh} + \delta_{ag} \delta_{bh} \delta_{ce} \delta_{df}) \\
 &- \frac{1}{2D(4D^2 - 1)} (\delta_{ae} \delta_{bh} \delta_{cg} \delta_{df} + \delta_{ag} \delta_{bf} \delta_{ce} \delta_{dh}), \tag{A10}
 \end{aligned}$$

which also can be written in graphical notation

$$\int_{\text{SU}(2D)} dU \begin{array}{c} a \\ | \\ \boxed{U} \\ | \\ b \end{array} \begin{array}{c} c \\ | \\ \boxed{\bar{U}} \\ | \\ d \end{array} = \frac{1}{2D} \begin{array}{cc} a & c \\ \text{---} & \text{---} \\ b & d \end{array} \quad (\text{A11})$$

$$\begin{aligned} \int_{\text{SU}(2D)} dU \begin{array}{c} a \\ | \\ \boxed{U} \\ | \\ b \end{array} \begin{array}{c} c \\ | \\ \boxed{U} \\ | \\ d \end{array} \begin{array}{c} e \\ | \\ \boxed{\bar{U}} \\ | \\ f \end{array} \begin{array}{c} g \\ | \\ \boxed{\bar{U}} \\ | \\ h \end{array} &= \frac{1}{4D^2 - 1} \left(\begin{array}{cc} a & e \\ \text{---} & \text{---} \\ c & g \\ \text{---} & \text{---} \\ b & f \\ \text{---} & \text{---} \\ d & h \end{array} + \begin{array}{cc} a & e \\ \text{---} & \text{---} \\ c & g \\ \text{---} & \text{---} \\ b & f \\ \text{---} & \text{---} \\ d & h \end{array} \right) \\ &- \frac{1}{2D(4D^2 - 1)} \left(\begin{array}{cc} a & e \\ \text{---} & \text{---} \\ c & g \\ \text{---} & \text{---} \\ b & f \\ \text{---} & \text{---} \\ d & h \end{array} + \begin{array}{cc} a & e \\ \text{---} & \text{---} \\ c & g \\ \text{---} & \text{---} \\ b & f \\ \text{---} & \text{---} \\ d & h \end{array} \right). \end{aligned} \quad (\text{A12})$$

Being ultimately a trace of Pauli operators, the expected value of the gradient thus vanishes. The first term of Eq. (A6) reads as, for instance,

$$\mathbb{E}_A \mathbb{E}_B \left[\begin{array}{cc} A & B \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{A} & \bar{B} \end{array} \right] = \int_{\text{SU}(2D)} dU \int_{\text{SU}(2D)} dV \left[\begin{array}{cc} U^{(0)} & V^{(0)} \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{U}^{(0)} & \bar{V}^{(0)} \end{array} \right] = \frac{1}{4} \text{Tr}(\sigma_l) \text{Tr}(\sigma_r) \text{Tr}(\Lambda_0) \text{Tr}(\Lambda_0^3) = 0, \quad (\text{A13})$$

and similarly the second term vanishes. The second moment is calculated using Eq. (A10), whereas two of the four contractions vanish as they separate into terms of the form of Eq. (A13). The remaining terms are integrals over squares of $\left[\begin{array}{cc} A & B \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{A} & \bar{B} \end{array} \right]$, once with Λ_0^3 on the upper contraction and once on the lower contraction, as well as the mixed term between both. The squared terms read as

$$\begin{aligned} \mathbb{E}_A \mathbb{E}_B \left[\begin{array}{cc} A & B \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{A} & \bar{B} \end{array} \right]^2 &= \int_{\text{SU}(2D)} dU \int_{\text{SU}(2D)} dV \left[\begin{array}{cc} U^{(0)} & V^{(0)} \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{U}^{(0)} & \bar{V}^{(0)} \end{array} \right] \left[\begin{array}{cc} U^{(0)} & V^{(0)} \\ \boxed{\sigma_l} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{U}^{(0)} & \bar{V}^{(0)} \end{array} \right] \\ &= \frac{1}{2D+1} \int_{\text{SU}(2D)} dV \left[\begin{array}{cc} \Lambda_0 & \Lambda_0^3 \\ \boxed{\sigma_r} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{V}^{(0)} & \bar{V}^{(0)} \end{array} \right] \left[\begin{array}{cc} \Lambda_0 & \Lambda_0^3 \\ \boxed{\sigma_r} & \boxed{\sigma_r} \\ \Lambda_0 & \Lambda_0^3 \\ \bar{V}^{(0)} & \bar{V}^{(0)} \end{array} \right] \\ &= \frac{\text{Tr}(\Lambda_0^4)^2}{(2D+1)^2} \end{aligned} \quad (\text{A14})$$

The second squared term yields the same outcome and the mixed term can be calculated similarly by just swapping Λ_0 and Λ_0^3 in one of the terms. This yields $\frac{\text{Tr}(\Lambda_0^6)}{(2D+1)^2}$ instead. Plugging in everything into the formula of the gradient of the Rényi entropy, Eq. (11), finally yields the second moment as claimed:

$$\mathbb{E}_A \mathbb{E}_B [(\partial_i S_{\mathcal{A},2} |_{\theta=0})^2] = \left(-\frac{2}{\text{Tr}(\Lambda_0^4)} \right)^2 \frac{2}{(2D+1)^2} (\text{Tr}(\Lambda_0^4)^2 - \text{Tr}(\Lambda_0^6)) = \frac{8}{(2D+1)^2} \left(1 - \frac{\text{Tr}(\Lambda_0^6)}{\text{Tr}(\Lambda_0^4)^2} \right). \quad (\text{A15})$$

■

APPENDIX B: CVD ON THE GHZ, CLUSTER, AND AKLT STATE

Some matrix product states are of conceptual interest in quantum information theory and quantum many-body physics. In the following, we want to consider the GHZ state $|\psi\rangle = \frac{1}{\sqrt{2}}(\bigotimes_{i=1}^n |0\rangle + \bigotimes_{i=1}^n |1\rangle)$, the cluster state, which is the ground state of the




FIG. 6. Disentanglement of the GHZ, cluster, and AKLT state on 50 sites using $\alpha = 1$. (a) A plot of the tail weights reveals that exact preparation circuits for the GHZ and the cluster state can be found with the expected circuit depth, that is $\frac{n}{2}$ for the GHZ state and 2 for the cluster state. (b) The AKLT state on 25 spin-1 sites (50 qubit sites restricted to the triplet subspaces) is approximately prepared in 10–13 layers, as the overlap error comes close to $\varepsilon_{\text{site}} = 10^{-3}$. After that, some tail weights can be exactly set to 0 as can be seen in the third color plot in (a). (c) The AKLT state preparation is scalable to larger system sizes, as $\varepsilon_{\text{site}}$ saturates.

cluster Ising Hamiltonian $H = \sum_i Z_{i-1} X_i Z_{i+1}$ and the (unnormalized) AKLT state,

$$\bigotimes_{\text{odd}} P^{(1)} \bigotimes_{\text{even}} |\psi^-\rangle = \left(\bullet \bullet \right) - \left(\bullet \bullet \right) - \left(\bullet \bullet \right) - \left(\bullet \bullet \right) - \left(\bullet \bullet \right), \quad (\text{B1})$$

which is built from neighboring singlet states $|\psi^-\rangle$ with each an end point of two different singlets projected onto the triplet subspace. In the literature, one often finds the pictorial representation shown in Eq. (B1) of the AKLT state, where the dots are spin-1/2 sites, straight lines denote singlet states and the ellipse depicts a projection onto the triplet subspace spanned by $\{|00\rangle, |11\rangle, \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)\}$.

Our results to these three states are illustrated in Fig. 6. It is well known that the GHZ state exhibits topological order and therefore cannot be prepared with a local circuit of sub-linear depth. The cluster state, on the other hand, is known for having particularly low-state complexity being exactly preparable with a depth-2 circuit. Our results are in agreement with these expectations and show exact preparation of

the GHZ state in circuit depth $\frac{n}{2}$ and the cluster state with depth 2.

The AKLT state does not have a straightforward circuit representation except the trivial approach of sequential preparation. There has been some effort to find preparation circuits [47, 48] with imaginary-time evolution on the AKLT Hamiltonian. We manage to find a shallow quantum circuit to prepare the AKLT state on 25 spin-1 sites, where each spin-1 site is encoded into the triplet subspace of two qubits. We reach a per-site overlap error of 1.8×10^{-3} with 13 layers of disentangling gates. Although this is a non-negligible remaining error, it is below the typical gate error rate of near-term quantum hardware. Better fidelities can be achieved using a lower cutoff for the Schmidt values and deeper circuits.

- [1] L. Marti, R. Mansuroglu, and M. J. Hartmann, Efficient quantum cooling algorithm for fermionic systems, *Quantum* **9**, 1635 (2025).
- [2] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [3] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
- [4] M. B. Hastings, An area law for one-dimensional quantum systems, *J. Stat. Mech.* (2007) P08024.
- [5] N. Schuch, M. M. Wolf, K. G. H. Vollbrecht, and J. I. Cirac, On entropy growth and the hardness of simulating time evolution, *New J. Phys.* **10**, 033032 (2008).
- [6] M. Will, T. A. Cochran, E. Rosenberg, B. Jobst, N. M. Eassa, P. Roushan, M. Knap, A. Gammon-Smith, and F. Pollmann, Probing non-equilibrium topological order on a quantum processor, *Nature (London)* **645**, 348 (2025).
- [7] S. J. Evered, M. Kalinowski, A. A. Geim, T. Manovitz, D. Bluvstein, S. H. Li, N. Maskara, H. Zhou, S. Ebadi, M. Xu, J. Campo, M. Cain, S. Ostermann, S. F. Yelin, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological matter and fermion dynamics on a neutral-atom quantum computer, *Nature (London)* **645**, 341 (2025).
- [8] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, Sequential generation of entangled multiqubit states, *Phys. Rev. Lett.* **95**, 110503 (2005).
- [9] D. Malz, G. Styliaris, Z.-Y. Wei, and J. I. Cirac, Preparation of matrix product states with log-depth quantum circuits, *Phys. Rev. Lett.* **132**, 040404 (2024).
- [10] Z.-Y. Wei and D. Malz, State preparation with parallel-sequential circuits, *arXiv:2503.14645*.
- [11] K. C. Smith, A. Khan, B. K. Clark, S. M. Girvin, and T.-C. Wei, Constant-depth preparation of matrix product states with adaptive quantum circuits, *PRX Quantum* **5**, 030344 (2024).
- [12] S.-J. Ran, Encoding of matrix product states into quantum circuits of one- and two-qubit gates, *Phys. Rev. A* **101**, 032310 (2020).
- [13] M. S. Rudolph, J. Chen, J. Miller, A. Acharya, and A. Perdomo-Ortiz, Decomposition of matrix product states into shallow quantum circuits, *Quantum Sci. Technol.* **9**, 015012 (2024).
- [14] B. Jobst, K. Shen, C. A. Riofrío, E. Shishenina, and F. Pollmann, Efficient MPS representations and quantum circuits from the Fourier modes of classical image data, *Quantum* **8**, 1544 (2024).
- [15] K. Murota, Adiabatic encoding of pre-trained MPS classifiers into quantum circuits, *arXiv:2504.09250*.
- [16] B. Jaderberg, G. Pennington, K. V. Marshall, L. W. Anderson, A. Agarwal, L. P. Lindoy, I. Rungger, S. Mensa, and J. Crain, Variational preparation of normal matrix product states on quantum computers, *arXiv:2503.09683*.
- [17] N. F. Robertson, A. Akhrie, J. Vala, and S. Zhuk, Approximate quantum compiling for quantum simulation: A tensor network based approach, *ACM Trans. Quantum Comput.* **6**, 1 (2025).
- [18] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, Variational quantum algorithms for nonlinear problems, *Phys. Rev. A* **101**, 010301(R) (2020).
- [19] C. Mc Keever and M. Lubasch, Towards adiabatic quantum computing using compressed quantum circuits, *PRX Quantum* **5**, 020362 (2024).
- [20] M. Ben-Dov, D. Shnaiderov, A. Makmal, and E. G. Dalla Torre, Approximate encoding of quantum states using shallow circuits, *npj Quantum Inf.* **10**, 65 (2024).
- [21] A. Khan, B. K. Clark, and N. M. Tubman, Pre-optimizing variational quantum eigensolvers with tensor networks, *arXiv:2310.12965*.
- [22] A. A. Melnikov, A. A. Termanova, S. V. Dolgov, F. Neukart, and M. R. Perelshtein, Quantum state preparation using tensor networks, *Quantum Sci. Technol.* **8**, 035027 (2023).
- [23] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [24] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (NY)* **326**, 96 (2011).
- [25] B. Kraus and J. I. Cirac, Optimal creation of entanglement using a two-qubit gate, *Phys. Rev. A* **63**, 062309 (2001).
- [26] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, E. R. Anschuetz, and Z. Holmes, Does provable absence of barren plateaus imply classical simulability? *Nat. Commun.* **16**, 7907 (2025).
- [27] L. Cigna and C. Schilling (private communication).
- [28] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* **4** (2022).
- [29] R. Mansuroglu and N. Schuch, Classical variational disentanglement, Software on Zenodo (2025), <https://doi.org/10.5281/zenodo.15058443>.
- [30] N. Schuch and F. Verstraete, Computational complexity of interacting electrons and fundamental limitations of density functional theory, *Nat. Phys.* **5**, 732 (2009).
- [31] S. R. White and E. M. Stoudenmire, Multisliced gausslet basis sets for electronic structure, *Phys. Rev. B* **99**, 081110(R) (2019).
- [32] G. González-García, R. Trivedi, and J. I. Cirac, Error propagation in NISQ devices for solving classical optimization problems, *PRX Quantum* **3**, 040326 (2022).
- [33] S. Gharibian and F. Le Gall, Dequantizing the quantum singular value transformation: Hardness and applications to quantum chemistry and the quantum PCP conjecture, *SIAM J. Comput.* **52**, 1009 (2023).
- [34] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, Strategies for solving the Fermi-Hubbard model on near-term quantum computers, *Phys. Rev. B* **102**, 235122 (2020).
- [35] R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, Perfect quantum error correction code, *Phys. Rev. Lett.* **77**, 198 (1996).
- [36] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane, Quantum error correction via codes over GF(4), in *Proceedings of IEEE International Symposium on Information Theory* (IEEE, Ulm, Germany, 2002).
- [37] T. Eckstein, R. Mansuroglu, S. Wolf, L. Nützel, S. Tasler, M. Kliesch, and M. J. Hartmann, Shot-noise reduction for lattice Hamiltonians, *PRX Quantum* **7**, 020303 (2026).
- [38] H.-Y. Huang, Y. Liu, M. Broughton, I. Kim, A. Anshu, Z. Landau, and J. R. McClean, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC '24 (ACM, New York, 2024), pp. 1343–1351.
- [39] R. Mansuroglu, A. Adil, M. J. Hartmann, Z. Holmes, and A. T. Sornborger, Quantum tensor-product decomposition

- from choi-state tomography, [PRX Quantum](#) **5**, 030306 (2024).
- [40] C. M. Keever and M. Lubasch, Classically optimized Hamiltonian simulation, [Phys. Rev. Res.](#) **5**, 023146 (2023).
- [41] R. Mansuroglu, T. Eckstein, L. Nützel, S. A. Wilkinson, and M. J. Hartmann, Variational Hamiltonian simulation for translational invariant systems via classical pre-processing, [Quantum Sci. Technol.](#) **8**, 025006 (2023).
- [42] R. Mansuroglu, F. Fischer, and M. J. Hartmann, Problem-specific classical optimization of Hamiltonian simulation, [Phys. Rev. Res.](#) **5**, 043035 (2023).
- [43] M. Nibbi and C. B. Mendl, Block encoding of matrix product operators, [Phys. Rev. A](#) **110**, 042427 (2024).
- [44] J. I. Latorre, Image compression and entanglement, [arXiv:quant-ph/0510031](#).
- [45] E. M. Stoudenmire and D. J. Schwab, Supervised learning with quantum-inspired tensor networks, [Adv. Neural Inf. Proces. Syst.](#) **29**, 4799 (2016).
- [46] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Entropy scaling and simulability by matrix product states, [Phys. Rev. Lett.](#) **100**, 030504 (2008).
- [47] T. Chen, R. Shen, C. H. Lee, and B. Yang, [SciPost Phys.](#) **15**, 170 (2023).
- [48] T. Chen and T. Byrnes, Efficient preparation of the AKLT state with measurement-based imaginary time evolution, [Quantum](#) **8**, 1557 (2024).