Synergy Between Quantum Circuits and Tensor Networks:
Short-cutting the Race to Practical Quantum Advantage

Manuel S. Rudolph,[1] Jacob Miller,[1] Jing Chen,[2] Atithi Acharya,[2,3] and Alejandro Perdomo-Ortiz[1,*]

[1]*Zapata Computing Canada Inc., 325 Front St W, Toronto, ON, M5V 2Y1*
[2]*Zapata Computing Inc., 100 Federal Street, Boston, MA 02110, USA*
[3]*Rutgers University, 136 Frelinghuysen Rd, Piscataway, NJ 08854, USA*
(Dated: August 29, 2022)

While recent breakthroughs have proven the ability of noisy intermediate-scale quantum (NISQ) devices to achieve quantum advantage in classically-intractable sampling tasks, the use of these devices for solving more practically relevant computational problems remains a challenge. Proposals for attaining practical quantum advantage typically involve parametrized quantum circuits (PQCs), whose parameters can be optimized to find solutions to diverse problems throughout quantum simulation and machine learning. However, training PQCs for real-world problems remains a significant practical challenge, largely due to the phenomenon of barren plateaus in the optimization landscapes of randomly-initialized quantum circuits. In this work, we introduce a scalable procedure for harnessing classical computing resources to determine task-specific initializations of PQCs, which we show significantly improves the trainability and performance of PQCs on a variety of problems. Given a specific optimization task, this method first utilizes tensor network (TN) simulations to identify a promising quantum state, which is then converted into gate parameters of a PQC by means of a high-performance decomposition procedure. We show that this task-specific initialization avoids barren plateaus, and effectively translates increases in classical resources to enhanced performance and speed in training quantum circuits. By demonstrating a means of boosting limited quantum resources using classical computers, our approach illustrates the promise of this synergy between quantum and quantum-inspired models in quantum computing, and opens up new avenues to harness the power of modern quantum hardware for realizing practical quantum advantage.

I. INTRODUCTION

In the coming years and decades, quantum computing resources will likely remain more expensive and less abundant than classical computing resources [1–3]. Despite the intrinsic theoretical advantages of quantum computers, the widespread adoption of quantum technologies will ultimately depend on the benefits they can offer for solving problems of high practical interest using these limited resources. To this end, parametrized quantum circuits (PQCs) [4–6] have been proposed as a promising formalism for leveraging near-term quantum devices for the solution of problems in quantum chemistry [7, 8], materials science [9], and quantum machine learning [10–18] applications which are difficult for classical algorithms.

However, several distinct challenges stand in the way of reaching practical advantage over classical methods using parametrized quantum algorithms, such as the existence of *barren plateaus* [19–22] and unfavorable guarantees for local minima [23–25] in the PQC optimization landscape. Such results typically apply to the setting of generic PQC ansätze and parameter initialization schemes, and much less is known about scenarios where the initial parameters of a PQC are adapted for a particular task. While this adaptation has proven useful in quantum chemistry, where circuits for computing molecular ground states have been shown to reach higher-precision results using initializations based on mean-field Hartree-Fock or more sophisticated coupled-cluster-based solutions (e.g., see Refs. [26, 27]), task-specific initializations

have seen much less use in other areas, such as *quantum machine learning* (QML).

Another difficulty for demonstrating an advantage over classical algorithms using PQCs is the increasing sophistication of classical simulation algorithms based on *tensor networks* (TNs), whose classically parametrized models can efficiently describe PQCs whose intermediate states have limited entanglement. The ability of TNs to be deployed on powerful classical hardware accelerators, such as graphical and tensor processing units (GPUs and TPUs), raises the bar for quantum hardware to overcome. This situation has led to a "zero-sum game" perspective on improvements in quantum vs. classical techologies, where advances in one domain are frequently viewed as placing additional burdens on practitioners of the other to attain relative advantage (see Ref. [28] for a representative example).

As depicted in Fig. 1, we propose a synergistic framework for boosting the performance and trainability of PQCs using a task-specific initialization strategy built on scalable TN algorithms, which leverages the complementary strengths of both technologies. This method uses TNs to first find a promising quantum state for the parametrized quantum algorithm at hand, then converts this TN state to the parameters of a PQC, where further optimization can be carried out on quantum hardware. We employ a circuit layer-efficient decomposition protocol [29] for matrix product states (MPS), whose high-performance conversion of MPS to various PQC architectures allows leveraging high-quality MPS solutions. The resulting quantum circuits can be extended with classically infeasible gates which enable better performance relative to the MPS, as well as purely quantum-optimized circuits. We empirically verify these performance improvements in various
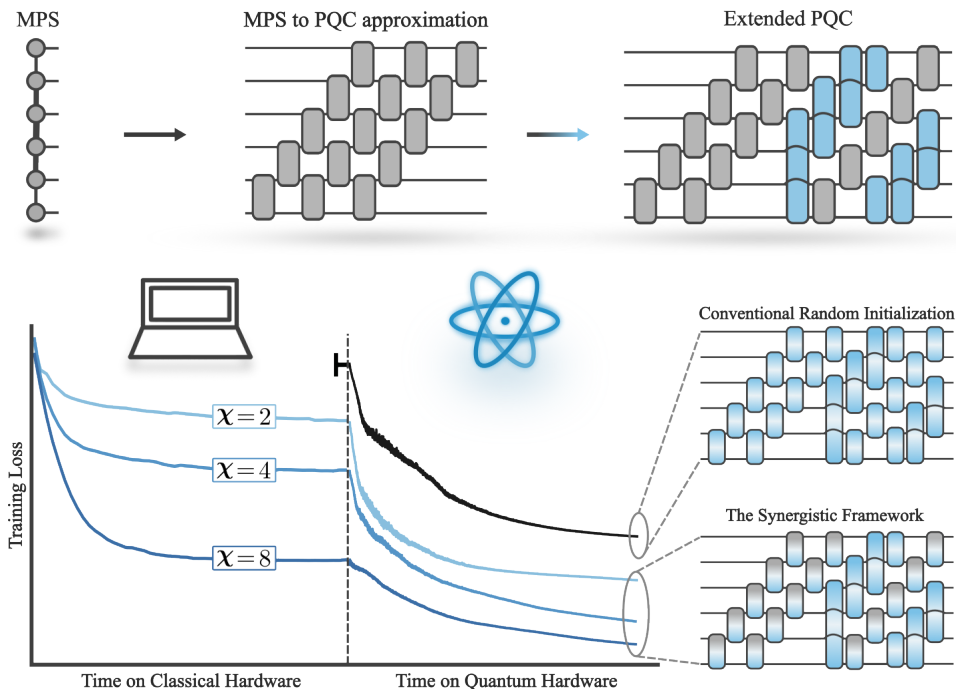

FIG. 1. Schematic depiction of the synergistic training framework utilizing TNs and PQCs. Rather than starting with a random initialization of circuit parameters (black curve), which may suffer from problems such as barren plateaus and sub-optimal local minima, we instead train a matrix product state (MPS) model on a classical simulation of the problem at hand (left half of blue curves), whose performance is bounded by the limited entanglement available via its bond dimension χ. This MPS wavefunction is then approximately transferred using a layer-efficient decomposition protocol that maps the MPS to linear layers of SU(4) gates. To improve on the classical solution, the quantum circuit is extended with additional gates (blue gates, initialized as near-identity operations) that would have been unfeasible to simulate on classical hardware. We show that quantum circuit models that leverage classically initialized circuit layers (gray & blue shaded gates) exhibit drastically improved performance over quantum circuits that were fully optimized on quantum hardware (blue shaded gates) and are likely to run into common trainability issues.

problems from generative modeling and Hamiltonian ground state search, finding that our method successfully converts deep quantum circuits from being practically untrainable to reliably converging to high-quality solutions. We further give evidence for the scalability of our synergistic framework by probing the gradient variances, i.e., the "barrenness", of PQCs with up to 20 qubits, finding gradient variances to remain stable with increasing number of qubits and circuit depth.

Compared to previous works, our method is broadly similar to the proposal of Ref. [30] to use classically trained TN models for initializing PQCs, which was predicted to yield benefits in performance and trainability within general machine learning tasks. Our findings can thus be seen as both a concrete realization of this general proposal, applicable to a diverse range of circuit architectures and learning tasks, as well as a robust experimental verification of the benefits anticipated there. Closer to our work is the pretraining method of Ref. [31], where trained MPS with bond dimension $\chi = 2$ were exactly decomposed into a staircase of two-qubit gates, which were then used to initialize quantum circuits for machine learning tasks. While this method was shown to improve the performance and trainability of PQC models, the restriction to $\chi = 2$ MPS placed a limit on the extent of classical resources which could be used to improve quantum models. By contrast, our synergistic optimization framework can

be scaled up to utilize arbitrarily large classical and quantum resources, a difference that we show gives continued returns in practice. By ensuring that PQCs are set up for success using the best solution available with today's abundant classical computing resources, we believe that parametrized quantum algorithms might finally manifest their full potential as effective methods for solving problems of deep practical interest.

II. METHODS

A. Parametrized Quantum Circuits

Parametrized quantum circuits (PQCs), and in particular the so-called *variational quantum algorithms* (VQAs), are the centerpiece of a family of quantum algorithms that aim to solve practical problems on near-term quantum devices with a limited number of qubits and noisy operations. The parameters $\boldsymbol{\theta}$ of PQCs are usually optimized (or "trained", in the context of learning from data) according to a loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \mathcal{L}\left(\psi_{\boldsymbol{\theta}}\right), \tag{1}$$

where $\psi_{\boldsymbol{\theta}}$ is the wavefunction of the quantum state prepared by the PQC. Unlike on classical hardware, one does not have

explicit access to the prepared state. Therefore, the loss $\mathcal{L}(\boldsymbol{\theta})$ needs to be estimated using quantum circuit measurements. PQCs are commonly trained via gradient descent methods, such as finite distance gradients or the *parameter shift rule* [32–34], or via gradient-free optimizers such as CMA-ES [35].

1. Quantum Circuit Born Machines

Quantum circuit Born machines (QCBMs) [14] are quantum models for solving generative learning tasks, which encode probability distributions over binary data [36] as measurement probabilities of a wavefunction prepared by a PQC. The probability assigned to a binary string $\mathbf{x}$ by a QCBM with circuit parameters $\boldsymbol{\theta}$ is given by the Born rule,

$$q_{\boldsymbol{\theta}}(\mathbf{x}) = |\langle \mathbf{x}|\psi_{\boldsymbol{\theta}}\rangle|^2, \tag{2}$$

where the parametrized wavefunction $\psi_{\boldsymbol{\theta}}$ encodes the distribution $q_{\boldsymbol{\theta}}(\mathbf{x})$. QCBMs are capable of representing complicated probability distributions [37–41], while still permitting a direct means of generating samples from any learned distribution by measuring the associated wavefunction $\psi_{\boldsymbol{\theta}}$. However, much is still unknown about the performance of QCBMs on near- to mid-term quantum devices, especially when modeling complex real-world datasets [42–44].

Many methods exist for training a QCBM to minimize a problem-specific loss function, which depends on a dataset $\mathcal{D}$ of size $|\mathcal{D}|$ and circuit parameters $\boldsymbol{\theta}$. The loss function we use here is the Kullback-Leibler (KL) divergence between the QCBM distribution $q_{\boldsymbol{\theta}}$ and the evenly weighted empirical distribution $p_{\mathcal{D}}$ associated to $\mathcal{D}$, given by

$$\begin{aligned}\mathcal{L}\left(\boldsymbol{\theta}\right) = \mathrm{KL}\left(p_{\mathcal{D}}||q_{\boldsymbol{\theta}}\right) &= \mathbb{E}_{\mathbf{x}\sim p_{\mathcal{D}}(\mathbf{x})}\left[\log\frac{p_{\mathcal{D}}(\mathbf{x})}{q_{\boldsymbol{\theta}}(\mathbf{x})}\right] \\ &= -\log\left(|\mathcal{D}|\right) - \frac{1}{|\mathcal{D}|}\sum_{\mathbf{x}\in\mathcal{D}}\log\left(q_{\boldsymbol{\theta}}(\mathbf{x})\right).\end{aligned} \tag{3}$$

For non-uniform weighting, the last line of Eq. 3 must be replaced by the appropriate expectation $\mathbb{E}_{\mathbf{x}\sim p_{\mathcal{D}}(\mathbf{x})}$.

2. Variational Quantum Eigensolver

The variational quantum eigensolver (VQE) [45] is a prototypical example of a variational quantum algorithm. The goal in VQE is to find the ground state ψ_0 or the ground state energy E_0 of a Hamiltonian H, which can be found by minimizing the variational energy function

$$\mathcal{L}(\boldsymbol{\theta}) = E(\boldsymbol{\theta}) = \langle\psi_{\boldsymbol{\theta}}|H|\psi_{\boldsymbol{\theta}}\rangle \tag{4}$$

of the parametrized trial wavefunction $\psi_{\boldsymbol{\theta}}$ on a quantum computer. The VQE algorithm can be used to calculate important properties of Hamiltonians in domains of significant practical interest, for example computing the energy of a molecule in the setting of quantum chemistry. In this setting, the qubit Hamiltonian is obtained from the fermionic Hamiltonian of the participating electrons using, for example, the Jordan-Wigner transformation [46]. Given the practical nature of the problem, and the decades of classical computational techniques towards solving such high-value problems, gave rise to highly specific quantum circuit ansätze and parameter initialization [26, 27].

B. Tensor Networks

Tensor networks (TNs) are linear-algebraic models first developed for classically simulating complex many-body quantum systems [47], which have more recently also been employed as machine learning models [30, 48–50]. TNs express many-body quantum states as the *tensor contraction* of low-order tensors, whose entries contain the parameters of the model. The manner in which the cores of a TN are contracted together is determined by an undirected graph, with different graphs determining different TN architectures. The nodes of each such graph correspond to *cores* of the TN, while the edges describe tensor contractions to be carried out between pairs of tensor cores. For applications to quantum simulation, the number of nodes N in a TN can, for example, be equal to the number of qubits in the quantum computer, and the topology of the N-node graph determines the forms of entanglement which can be faithfully reproduced in the classical TN simulation.

In this work, we utilize *matrix product states* (MPS), computationally tractable TN models whose cores are connected along a line graph, and which have a long history in the computation of ground states of quantum 1D spin chains [51, 52] via the *density matrix renormalization group* (DMRG) algorithm [53], as well as for the efficient simulation of quantum computers with limited entanglement [54]. Despite their simplicity, MPS of sufficient size can simulate any N-qubit wavefunction, making them a natural first model for many TN applications. The expressivity of an MPS is determined by its *bond dimension* χ, a quantity associated to the edges of an MPS which sets an upper bound on the amount of entanglement achievable in a simulated quantum state [55]. In cases where the entanglement of a quantum state is greater than an MPS is able to exactly reproduce, the *singular value decomposition* (SVD) can be used to find an MPS approximation of the state with near-optimal fidelity.

Although we focus on MPS, our results can be straightforwardly extended to more complicated TN architectures, allowing for different tradeoffs between expressivity and classical computational complexity [56].

Tensor Network Born Machines

One application of MPS here is as *tensor network Born machines* (TNBMs) [50], generative models which represent a probability distribution using a simulated quantum wavefunction parametrized by a TN. As such, they form the tensor net-

work analog to QCBMs described in Sec. II A 1, and we utilize TNBMs to provide the classical solutions to them.

While QCBMs and TNBMs are similar mathematically, with both model families using the Born rule to parametrize classical probability distributions, they nonetheless have distinct complementary benefits in real-world applications. QCBMs are fully quantum models which are able to leverage advances in quantum hardware to better reproduce the correlations present in complex datasets, but are limited by the state of current noisy intermediate-scale quantum (NISQ) devices. By contrast, TNBMs can take full advantage of recent developments in classical computing hardware, notably the development of powerful graphical/tensor processing units (GPUs/TPUs), but are fundamentally limited in their expressivity by the extent of entanglement they are able to simulate efficiently. Additionally, the analytically explicit construction of TNBMs enables exact calculation of probabilities $q_{\boldsymbol{\theta}}(\mathbf{x})$ (see Eq. (2)) and gradients of a loss function $\mathcal{L}$ with respect to the model parameters $\boldsymbol{\theta}$. The complementary strengths of both models naturally motivate the development of hybrid quantum-classical Born machine models, but this is complicated in practice by the difficulty of converting between these models.

C. MPS to PQC mapping

The parameters of PQCs and TNs can in principle be interconverted freely, with the circuit topology of a PQCs itself forming a TN via classical simulation, and with TN canonical forms [55, 57] facilitating the representation of a TN as a PQC. In practice though, there are several issues that arise with the latter conversion. The quantum "circuits" associated with a direct conversion from TNs to PQCs are composed of abstract unitary gates acting on multi-level qudits of varying size, whose compilation into gate sets of real-world quantum computers is itself a non-trivial problem (e.g., see [58]). In the general case of bond dimension of χ, an MPS will be mapped to a quantum circuit containing multi-qubit gates acting on $\lceil \log_2(\chi) \rceil + 1$ qubits per gate. Much preferred however is a decomposition into two-qubit gates. This is practical for a variety of reasons. For instance, many quantum hardware realizations natively implement two-qubit gates, removing computational overhead in applying multi-qubit operations. Additionally, two-qubit gates can be more sparingly parametrized, in contrast to the exponentially increasing number of parameters needed to fully control multi-qubit gates. Despite these challenges, we show in the following that the use of an efficient and high-performance conversion method permits MPS of increasing size and complexity to boost the performance of PQCs within several real-world applications.

We use the MPS decomposition protocol developed in Ref. [29], which augments the analytical decomposition method of Ref. [58] with intertwined constrained classical optimization steps on the circuit unitaries. Using this protocol, transferring the MPS to a PQC results in k layers of SU(4) unitaries with a next-neighbor topology, also called linear layers. The choice of an appropriate k is a hyperparameter of

the decompostion, and the quality of the decomposition for a fixed k is limited by the entanglement present in the MPS. Fortunately, the decomposition protocol used in this work allows for sequential growing of the circuit up to a desired fidelity. We refer to Appendix D for a more detailed description of the decomposition protocol used throughout this work.

One may wonder how this process is efficient on classical hardware. This is the case because the created linear quantum circuit layers tend to result in the MPS having a lower bond dimension than before. Generally speaking, if the MPS was computationally feasible beforehand, it should also be feasible to decompose it via this technique. This is opposed to alternative approaches of brute-force optimization of the linear layers. In such cases, the intermediate states reached during optimization are not guaranteed to represent an MPS with χ_{max} equal to or less than that of the original MPS.

To have a chance at improving the previously found MPS results, one needs to extend the linear layers with additional gates that would have been infeasible to simulate classically, i.e., the bond dimension χ of the MPS would need to be increased, which is likely not possible at a point where one is planning to continue optimization on a quantum computer. Extending the quantum circuit can either come in the form of increased circuit depth, more flexible entangling topologies, or both. In our work, from the k linear layers, we extend only the final layer of SU(4) gates to an all-to-all topology. Notably, we then train all existing gates in the circuit, and not just the additional gates. We refer to Appendix C 2 for details on the SU(4) gate circuit ansatz, as well as Appendix A for a brief study of the effect of adding additional gates to the mapped MPS quantum circuits.

III. RESULTS AND DISCUSSION

To assess the benefits of the synergistic training framework for real-world applications, we explore the performance of this method in a variety of probabilistic modelling tasks and Hamiltonian ground state search problems. In each case, we compare quantum circuits trained using our MPS initialization approach to those which are initialized randomly or with all gates close to the identity. The latter has been shown empirically to reduce the effects of barren plateaus and improve convergence behavior at the start of PQC optimization [59].

We study the impact of different circuit architectures on our results by designing the quantum circuit layers with either linear or all-to-all topologies of fully parametrized SU(4) gates. While all-to-all connectivity is likely not practical in a scalable manner on near- to mid-term quantum hardware, it provides a challenging use case for a initialization method leveraging a TN model with linear connectivity, while also illustrating an important advantage that quantum hardware has over classical TN simulation techniques: The flexible choice of circuit depth and entangling topology. Implementation details can be found in Appendix C.

Our results show that the use of TNs as a task-specific initialization strategy succeeds in boosting the performance of
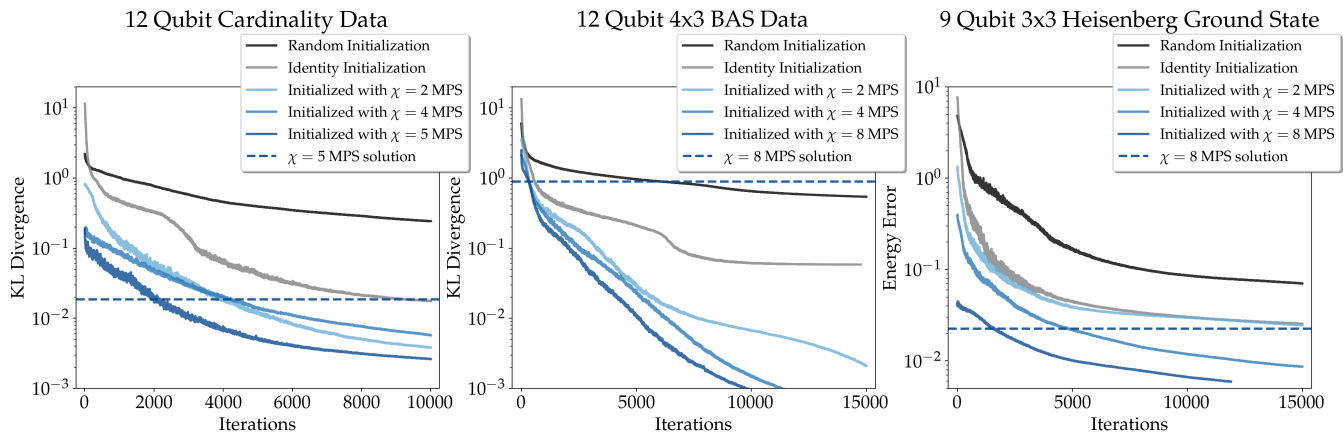
FIG. 2. Optimization results for QCBM and VQE training, where each curve represents the best performance among 6 independently initialized runs of the model. The QCBMs are trained to minimize KL divergence relative to datasets of length $N = 12$ strings of fixed cardinality $N/2$ or 4×3 bars and stripes images, whereas the VQE optimization aims to minimize the energy of a 2D Heisenberg Hamiltonian with 9 qubits, and with size 3×3. The quantum circuits for each case consist of three, four, and four layers, respectively, with the SU(4) gates of each layer arranged in a linear topology, except for the final layer whose gates are connected in an all-to-all manner. In each case, quantum circuits whose parameters are initialized randomly or close to the identity exhibit worse final losses than those initialized with a classically trained MPS model. Additionally, the use of increased classical resources (as quantified by the bond dimension χ) leads to improved performance of the trained quantum circuits. All optimization runs compared inside individual plots share exactly the same circuit layout and number of trainable parameters. Differences in training performance are only due to different initial parameters.

quantum circuits in all of these tasks, with an increase in classical computing resources (as quantified by the bond dimension χ of the MPS) in nearly every case leading to a corresponding increase in the final performance of the trained quantum circuit. This is reflected not only in the final losses in different tasks, but also through an analysis of the parameter gradients seen by the circuit at initialization. We find that although randomly initialized quantum circuits exhibit gradients of exponentially vanishing magnitude in system size, a manifestation of barren plateaus within generative modeling, the use of classically trained MPS to fix a task-specific initialization avoids this phenomenon entirely.

In our first experiments, we explore the optimization performance of QCBM and VQE, i.e., the progression of the loss function values (Eqs. (3) & (4), respectively). For the QCBM, we study two distinct datasets of bitstrings of length $N = 12$. The first is the *cardinality* dataset that is the set of all strings having a cardinality (i.e., Hamming weight, or number of 1s) of $N/2$. The second QCBM dataset is the dataset of bars and stripes (BAS) images [60] containing horizontal or vertical lines. The VQE optimization problem uses $N = 9$ qubits and minimizes the energy of the 2D Heisenberg model Hamiltonian on a 3×3 rectangular lattice:

$$H = \frac{1}{4} \sum_{\langle i,j \rangle} \sigma_X^{(i)} \sigma_X^{(j)} + \sigma_Y^{(i)} \sigma_Y^{(j)} + \sigma_Z^{(i)} \sigma_Z^{(j)}. \quad (5)$$

$\langle i,j \rangle$ indexes all nearest-neighbor spins in a 2D rectangular grid with open boundary conditions, and $\sigma_\mu^{(i)}$, $\mu = X, Y, Z$ denote the Pauli operators acting on the i'th spin. We measure the energy error $\Delta E(\boldsymbol{\theta}) = E(\boldsymbol{\theta}) - E_0$ relative to the exact ground state energy E_0

In all cases, we compare the performance of PQCs initial-

ized with random SU(4) or near-identity unitaries to those initialized with previously found MPS solutions. Transferring the MPS state is done via the decomposition protocol described in Ref. [29]. As described in Sec. II C, the topologies utilized here contain k layers of gates, which are arranged linearly in the first $k - 1$ layers and in an all-to-all topology in the last layer. For the cardinality dataset we utilize $k = 3$ layers, and for the BAS dataset, as well as for the 2D Heisenberg Hamiltonian, $k = 4$ layers. The parameters of the quantum circuits are optimized using the *CMA-ES* algorithm [35, 61], a gradient-free optimizer that is based on an adaptive evolutionary strategy.

Our optimization results in Fig. 2 depict the best optimization runs out of 6 repetitions. It becomes evident that the models without the MPS initialization do not converge to high-quality solutions. In fact, we have observed that, while all-to-all layers clearly increase the expressive capabilities of a PQC as compared to linear layers, the presence of a single all-to-all entangling layer has detrimental effect on their trainability (see also Appendix A). By choosing an initialization which makes use of the parameters of a classically trained MPS model however, all models exhibit a drastic increase in performance on all tasks. This behavior is enhanced even more by the use of MPS with larger bond dimension χ. We emphasize that all PQC models compared inside one plot have precisely the same circuit layout and number of parameters. Differences in performance are only due to different choices of parameter initialization.

One may expect that the informed initialization merely affects the number of quantum circuit evaluations needed to achieve a target training loss, however, we show that it qualitatively changes convergence behaviour. For example, in the case of the cardinality dataset, QCBMs initialized with $\chi = 4$
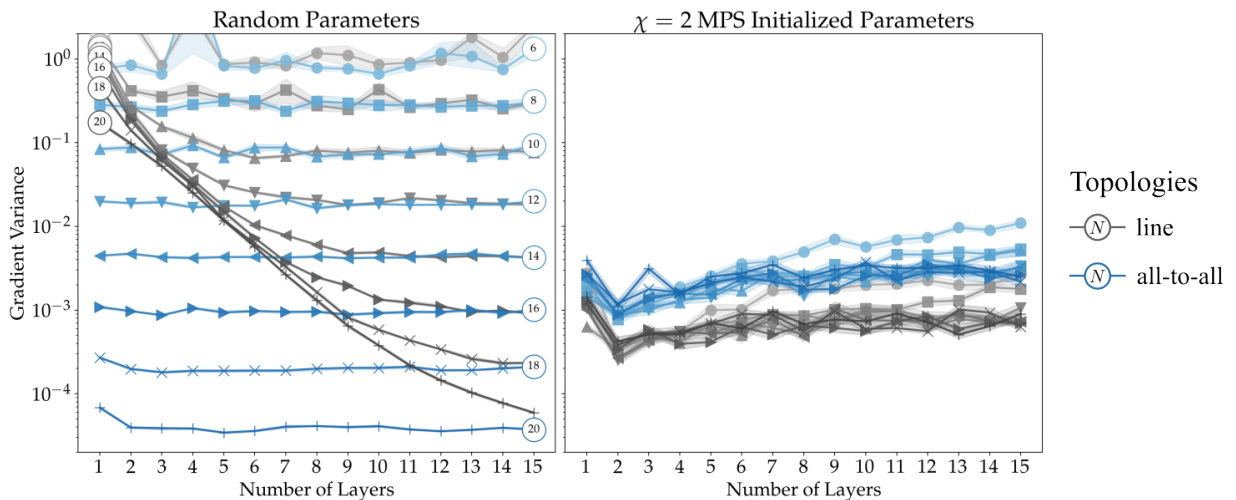
FIG. 3. Removal of barren plateaus in QCBMs inside our synergistic optimization framework as demonstrated by the gradient variances with respect to the KL divergence loss in Eq. (3). The grey lines indicate the gradient variances of linear topology circuits, whereas the blue lines indicate gradient variances of all-to-all topology circuits. The numbers at the beginning or the end of the lines denote the number of qubits. We record the median gradient variances over 1000 repetitions for the randomly initialized circuits (left), and 100 repetitions for the $\chi = 2$ MPS initialized circuits (right), as well as bootstrapped 25-75 percentile confidence intervals of the median inside the shaded areas. We study the Cardinality $N/2$ dataset for the respective number of qubits N. The gradients are measured with respect to the YY-entangling gate contribution of the first SU(4) gate in the circuit between qubit 1 and 2. For random parameter initializations, the gradient variance decays exponentially in the number of qubits, and also the circuit depth up until a certain limit. This is clear indication for the existence of barren plateaus. One all-to-all layer of SU(4) gates appears to fully maximize the degree of barrenness. In contrast, the gradient variances of MPS initialized circuits neither decay significantly in the number of qubits nor with increasing circuit depth.

or $\chi = 5$ begin training at approximately the same KL values, but the PQC initialized with the larger χ MPS rapidly achieves better loss values shortly after. A similar behavior can be seen in the case of the more challenging BAS dataset. In this case, the MPS solutions achieve relatively high KL divergence values, which consequently leads to high initial KL values for the MPS initialized QCBMs. While the randomly initialized circuits generally reach the KLs at which the MPS initialized models start out, the latter are able to converge fully, while the former appear to plateau at much higher KL values.

Crucially, we observe in the VQE optimization example that initializing with $\chi = 2$ MPS does not suffice to reliably improve over a naive near-identity initialization of the circuit unitaries. Only MPS with larger bond dimension χ, facilitated by the layer-efficient decomposition in Ref. [29], enable significant enhancements. This is also highlighted by the depicted loss values achieved by the MPS solutions with highest χ that were used to initialize the respective PQC models. In the VQE simulation, initializing the PQC with $\chi = 2$ is not sufficient for the PQC to outperform what the MPS on classical hardware may have been capable of. The particular case of $\chi = 2$, where the MPS readily maps to two-qubit gates, was studied in Ref. [31], and does not require the layer-efficient decomposition scheme used in this work which enables arbitrary χ.

However, the final MPS losses (indicated by the dashed lines) also showcase how the PQC solutions can improve on solutions attained on classical hardware by leveraging the more flexible capabilities of quantum hardware and initializing with strong classical models. The gaps between the fi-

nal MPS losses and the respective PQC initializations stem from imperfect decomposition of the MPS into a low number of two-qubit gate layers, as well as the close-to-identity extension of the quantum circuits into the all-to-all topologies, and the initial exploration step size of the CMA-ES optimizer. While initializing of the QCBM with $\chi = 8$ on the BAS dataset here achieves the best result, we note that it does not clearly outperform $\chi = 4$ on average (see Fig. 5 in Appendix B) The likely reason is that the BAS dataset is 2d-correlated and thus the MPS with growing bond dimension χ increasingly biases the quantum circuit to a 1d-correlated solution. In other words, there is a bias miss-match between the TN architecture used and the task at hand. Depending on the number of additional free parameters that the PQC is given access to, this can lead to saddle points and local minima, because the PQC needs to correct the unsuitable bias. In such cases, one may try to train another TN model which is adapted for more general correlation structures, and then, if needed, map this TN to a quantum circuit.

To assess whether the synergistic framework is expected to be effective at improving the trainability of PQCs as the number of qubits increases, we now assess the variance of parameter gradients, i.e. the *barrenness*, of QCBMs training on the cardinality dataset. The results are shown in Fig. 3. We probe the gradient of the KL divergence loss with respect to the parameter controlling the YY-entangling component (according to the KAK-decomposition [62]) of the first SU(4) gate between qubits 1 and 2 (see Appendix C 2 for details). Gradient magnitudes for that parameter are recorded 1000 times per data point in the case of random parameters, and 100 times

per data point in the case of MPS $\chi = 2$ initialization. The latter case contains the training of the MPS, as well as the mapping to a quantum circuit, and the (potential) extension of the linear layers to all-to-all topologies. We note that our results are robust to different choices of the parameter for which the gradients are estimated.

In the case of randomly initialized parameters ($\boldsymbol{\theta} \in [0, 2\pi]$), we observe a clear exponential decay in gradient variances with increasing circuit depth and number of qubits. To the best of our knowledge, we provide the first experimental demonstration of barren plateaus for QCBMs trained with the KL divergence loss function. The nature of this decay depends on the quantum circuit topology used, with a single all-to-all layer being sufficient to saturate the barreness for a specific number of qubits up until $N = 18 - 20$. In contrast, we observe that QCBMs initialized with a classically trained MPS avoid this exponential decay - something which likely contributes to the increased trainability observed in Fig. 2. Fascinatingly, the gradients in this initialization can actually exhibit an *increase* in variance with circuit depth, a trend which is more visible in the all-to-all extended circuit. Overall, the gradients of the all-to-all topologies have a larger gradient variance, indicating that the circuit extension after transferring the MPS solution is crucial to the success of the PQC. These findings show that the use of trained MPS places the quantum circuit model in a region of the parameter space without evident barren plateaus, but where the additional flexibility provided by increased connectivity in the quantum circuit enables it to effectively improve on the classical MPS solution. With a more sparse set of measurements, we can confirm a very similar trend when utilizing $\chi = 4$ MPS solutions.

The removal of barren plateaus, as indicated in Fig. 3, is vital to ensuring the trainability of PQCs and their viability on quantum hardware. Vanishing gradient variances imply that gradient magnitudes also vanish [19], which leads the estimation of parameter gradients on quantum hardware to require a number of measurements which grows exponentially in the number of qubits. Additionally, barren plateaus have been linked to the phenomena of *cost concentration* and *narrow gorges* [25], which hinder the ability of gradient-based and gradient-free optimizers to find high-quality solutions, as well as the existence of large numbers of low-quality local minima [24], which present further difficulties in learning. Aside from improving the training performance in practice (as seen in Fig. 2), stable gradient variances (such as those in Fig. 3) hint that a finite (or at worst, non-exponential) number of quantum circuit evaluations may be sufficient to estimate parameter gradients and perform PQC optimization on quantum hardware in a scalable manner.

IV. CONCLUSIONS

This work introduces a synergistic training framework for quantum algorithms, which employs classical tensor network simulations towards boosting the performance of PQCs. Our framework allows a problem of interest to be attacked first with the aid of abundant classical resources (e.g. GPUs and TPUs), before being transitioned onto quantum hardware to find a solution with further improved performance. By moving the work of quantum computers to improving on promising classical solutions, rather than finding such solutions *de novo*, we ensure that scarce quantum resources are allocated where they are most effective, setting up parametrized quantum algorithms for success.

Assessing the performance of our methods on generative modeling and Hamiltonian minimization problems, we found that PQCs initialized with this synergistic training framework not only obtained better training losses using identical quantum resources, but also exhibited qualitatively improved optimization behavior, with deep quantum circuits transformed from being practically untrainable to reliably converging on high-quality solutions. A study of gradient variances shows the promise of this method for avoiding barren plateaus and related worst-case guarantees, which for randomly initialized PQCs lead to gradients which decay exponentially in both the number of qubits and the depth of the circuit. For PQCs initialized using classically obtained TN solutions however, we observed gradient variances which remain essentially constant with respect to size, even showing a slight increase in deeper circuits. These results point towards the promise of this framework for enabling PQCs to scale to large number of qubits, thereby unlocking the latent capabilities of quantum computers for optimization and learning problems which remain out of reach for purely classical methods.

These findings naturally open up several related questions. We have employed MPS as our classical TN ansatz, whose bond dimension χ determines the classical resources allocated for PQC initialization, but have yet to characterize the performance of our method on problem sizes requiring very large values of χ. While our findings show that larger bond dimensions in our synergistic framework lead to increased PQC performance, we also anticipate an eventual need to employ more sophisticated TN models whose topology is better adapted to the connectivity of the circuit architecture at hand. To this end, initializing PQCs using tree tensor networks is a natural next area of study, as the simple canonical forms available to such models permit a straightforward extension of the decomposition procedure used here [29]. We anticipate the use of more flexible TN models to lead to further improvements in the performance of quantum algorithms, complementary to those identified for the use of larger values of χ.

As a final remark, we note that our synergistic framework highlights the benefits of moving beyond the adversarial mindset of "classical vs. quantum" which is typica l of discussions surrounding quantum supremacy. By embracing the rich connections between classical TN algorithms and PQCs, we show that good use can be made of the complementary strengths of both. Moving forward, we believe that the existence of powerful classical simulation methods should not be seen as an obstacle on the path to demonstrating practical quantum advantage, but rather as a guide to help quantum methods find the way [11].

[1] S. Aaronson, Read the fine print, Nature Physics **11**, 291 (2015), commentary.

[2] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[3] National Academies of Sciences, Engineering, and Medicine and others, *Quantum computing: progress and prospects* (National Academies Press, 2019).

[4] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parame-terized quantum circuits as machine learning models, Quantum Science and Technology **4**, 043001 (2019).

[5] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, Nature Reviews Physics **3**, 625 (2021).

[6] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, Rev. Mod. Phys. **94**, 015004 (2022).

[7] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. John-son, M. Kieferová, I. D. Kivlichan, T. Menke, B. Per-opadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum comput-ing, Chemical Reviews **119**, 10856 (2019), pMID: 31469277, https://doi.org/10.1021/acs.chemrev.8b00803.

[8] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, Reviews of Mod-ern Physics **92**, 015003 (2020).

[9] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, Quantum algorithms for quantum chemistry and quantum materials sci-ence, Chemical Reviews **120**, 12685 (2020).

[10] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, Nature **549**, 195 (2017).

[11] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gómez, and R. Biswas, Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers, Quantum Science and Technology **3**, 030502 (2018).

[12] J. Tian, X. Sun, Y. Du, S. Zhao, Q. Liu, K. Zhang, W. Yi, W. Huang, C. Wang, X. Wu, M.-H. Hsieh, T. Liu, W. Yang, and D. Tao, Recent advances for quantum neural networks in generative learning, arXiv:2206.03066 (2022).

[13] E. Farhi and H. Neven, Classification with quantum neural net-works on near term processors, arXiv:1802.06002 (2018).

[14] M. Benedetti, D. Garcia-Pintos, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, npj Quantum Information **5** (2018).

[15] M. Schuld and N. Killoran, Quantum machine learning in fea-ture hilbert spaces, Physical Review Letters **122**, 10.1103/phys-revlett.122.040504 (2019).

[16] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kan-dala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, Nature **567**, 209 (2019).

[17] M. S. Rudolph, N. B. Toussaint, A. Katabarwa, S. Johri, B. Per-opadre, and A. Perdomo-Ortiz, Generation of high-resolution handwritten digits with an ion-trap quantum computer, Phys. Rev. X **12**, 031010 (2022).

[18] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, Science **376**, 1182 (2022), https://www.science.org/doi/pdf/10.1126/science.abn7293.

[19] J. Mcclean, S. Boixo, V. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nature Communications **9** (2018).

[20] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, Nature communications **12**, 1 (2021).

[21] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connect-ing ansatz expressibility to gradient magnitudes and barren plateaus, PRX Quantum **3**, 010313 (2022).

[22] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cin-cio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, Nature communications **12**, 1 (2021).

[23] E. R. Anschuetz and B. T. Kiani, Beyond barren plateaus: Quantum variational algorithms are swamped with traps, arXiv preprint arXiv:2205.05786 (2022).

[24] E. R. Anschuetz, Critical points in quantum generative mod-els, in *International Conference on Learning Representations* (2021).

[25] A. Arrasmith, Z. Holmes, M. Cerezo, and P. J. Coles, Equiva-lence of quantum barren plateaus to cost concentration and nar-row gorges, arXiv preprint arXiv:2104.05868 (2021).

[26] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. J. Love, and A. Aspuru-Guzik, Strategies for quantum comput-ing molecular energies using the unitary coupled cluster ansatz, Quantum Science and Technology **4**, 014008 (2018).

[27] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, Physical Review A **92**, 042303 (2015).

[28] A. Cho, Ordinary computer matches Google's quantum com-puter, Science **377**, 563 (2022).

[29] M. S. Rudolph, J. Chen, J. Miller, A. Acharya, and A. Perdomo-Ortiz, Decomposition of matrix product states into trainable quantum circuits, (2022).

[30] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, Towards quantum machine learning with tensor networks, Quantum Science and Technology **4**, 024001 (2019).

[31] J. Dborin, F. Barratt, V. Wimalaweera, L. Wright, and A. Green, Matrix product state pre-training for quantum machine learn-ing, Quantum Science and Technology (2022).

[32] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid quantum-classical approach to quantum optimal control, Physical review letters **118**, 150503 (2017).

[33] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, Phys. Rev. A **98**, 032309 (2018).

[34] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, Physical

Review A **99**, 032331 (2019).

[35] N. Hansen and A. Ostermeier, Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation, in *Proceedings of IEEE international conference on evolutionary computation* (IEEE, 1996) pp. 312–317.

[36] QCBMs can model probability distributions over arbitrary discrete data, but we restrict our attention to the case of binary random variables. Given the possibility of encoding discrete random variables as random bitstrings, this entails no loss of generality.

[37] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, Expressive power of parametrized quantum circuits, Physical Review Research **2**, 033125 (2018).

[38] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and I. Cirac, Expressive power of tensor-network factorizations for probabilistic modeling, Advances in neural information processing systems **32** (2019).

[39] R. Sweke, J.-P. Seifert, D. Hangleiter, and J. Eisert, On the quantum versus classical learnability of discrete distributions, Quantum **5**, 417 (2021).

[40] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The born supremacy: quantum advantage and training of an ising born machine, npj Quantum Information **6** (2020).

[41] M. Hinsche, M. Ioannou, A. Nietner, J. Haferkamp, Y. Quek, D. Hangleiter, J.-P. Seifert, J. Eisert, and R. Sweke, Learnability of the output distributions of local quantum circuits (2021), arXiv:2110.05517 [quant-ph].

[42] M. Benedetti, B. Coyle, M. Fiorentini, M. Lubasch, and M. Rosenkranz, Variational inference with a quantum computer, Phys. Rev. Applied **16**, 044057 (2021).

[43] K. Gili, M. Hibat-Allah, M. Mauri, C. Ballance, and A. Perdomo-Ortiz, Do quantum circuit born machines generalize? (2022), arXiv:2207.13645.

[44] M. Hinsche, M. Ioannou, A. Nietner, J. Haferkamp, D. Hangleiter, J.-P. Seifert, J. Eisert, and R. Sweke, A single t-gate makes distribution learning hard (2022), arXiv:2207.03140.

[45] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'brien, A variational eigenvalue solver on a photonic quantum processor, Nature communications **5**, 1 (2014).

[46] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Simulating physical phenomena by quantum networks, Physical Review A **65**, 042323 (2002).

[47] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, Annals of physics **349**, 117 (2014).

[48] E. Stoudenmire and D. J. Schwab, Supervised learning with tensor networks, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016) pp. 4799–4807.

[49] A. Novikov, M. Trofimov, and I. Oseledets, Exponential machines (2016).

[50] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised generative modeling using matrix product states, PRX **8**, 031012 (2018).

[51] I. Affleck, T. Kennedy, E. Lieb, and H. Tasaki, Rigorous results on valence-bond ground states in antiferromagnets, Physical Review Letters **59**, 799 (1987).

[52] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, Communications in mathematical physics **144**, 443 (1992).

[53] S. R. White, Density matrix formulation for quantum renormalization groups, Physical review letters **69**, 2863 (1992).

[54] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, Physical review letters **91**, 147902 (2003).

[55] D. Perez-Garcia, F. Verstraete, M. Wolf, and J. Cirac, Matrix product state representations, Quantum Information & Computation **7**, 401 (2007).

[56] I. L. Markov and Y. Shi, Simulating quantum computation by contracting tensor networks, SIAM Journal on Computing **38**, 963 (2008).

[57] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, Physical review A **74**, 022320 (2006).

[58] S.-J. Ran, Encoding of matrix product states into quantum circuits of one-and two-qubit gates, Physical Review A **101**, 032310 (2020).

[59] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, Quantum **3**, 214 (2019).

[60] D. J. C. MacKay, *Information Theory, Inference & Learning Algorithms* (Cambridge University Press, New York, NY, USA, 2002).

[61] N. Hansen, Y. Akimoto, and P. Baudis, CMA-ES/pycma on Github, Zenodo, DOI:10.5281/zenodo.2559634 (2019).

[62] R. R. Tucci, An introduction to Cartan's KAK decomposition for QC programmers, arXiv preprint arXiv:0507171 (2005).

[63] M. Fishman, S. R. White, and E. M. Stoudenmire, The itensor software library for tensor network calculations (2020).

[64] Y. Suzuki, Y. Kawase, Y. Masumura, Y. Hiraga, M. Nakadai, J. Chen, K. M. Nakanishi, K. Mitarai, R. Imai, S. Tamiya, T. Yamamoto, T. Yan, T. Kawakubo, Y. O. Nakagawa, Y. Ibe, Y. Zhang, H. Yamashita, H. Yoshimura, A. Hayashi, and K. Fujii, Qulacs: a fast and versatile quantum circuit simulator for research purpose, Quantum **5**, 559 (2021).

[65] ORQUESTRA®, https://www.orquestra.io/.

[66] T. Shirakawa, H. Ueda, and S. Yunoki, Automatic quantum circuit encoding of a given arbitrary quantum state, arXiv preprint arXiv:2112.14524 (2021).

Appendix A: Performance enhancement when extending linear layers

When initializing PQCs for model optimization, the approach throughout this work has been to decompose the classically trained MPS into k linear layers of two-qubit gates, and extend only the final layer into an all-to-all topology. Fig. 2 demonstrates how this approach can reliably improve on the final loss values that the MPS could achieve. In contrast, Fig. 4 depicts an example case of an $N = 10$ qubit QCBM training on the Cardinality dataset where we do not extend the linear layers after transferring the MPS solution to the quantum circuit. Interestingly, with random initial parameters, the linearly arranged circuit reaches better KL divergence values (black, left panel) than when the final layer is extended to an all-to-all topology (black, right panel). This highlights the deterioration of trainability with increasing circuit depth that can also seen in Fig. 3. The near-identity initialization is notably less hampered by this. When classically initializing the QCBMs, the models can only aim to recover the MPS performance when the circuit is not extended with additional gates
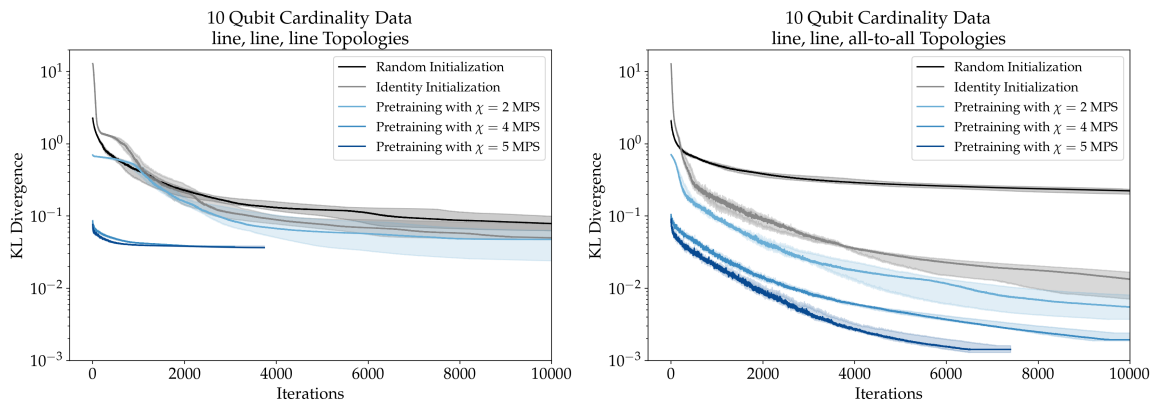
FIG. 4. Optimization results for QCBMs training on the Cardinality dataset with $N = 10$ qubits and three linear layers. We showcase the enhancement of the mapped MPS solutions when the final linear layer is extended to a all-to-all topology. While trainability suffers for random initial parameters, the classically initialized PQCs are able to converge to significantly better solutions.

and free parameters. In contrast, the all-to-all topology allows the QCBM to achieve great performance and, as also seen in Fig. 2, QCBM initialized with larger χ MPS perform better even when starting at similar loss values.

Appendix B: Median Optimization Performance from Sec. III

In addition to the best of 6 optimization runs shown in Fig. 2, we here report the bootstrapped median and 25-75 percentile of the losses. Fig. 5 shows that initializing the PQC circuits with classically found MPS solutions is beneficial in all cases and significantly enhances the trainability of the models. In the cases of the Cardinality dataset in the generative modeling task, and the Heisenberg ground state minimization, the synergistic optimization framework shows continuous improvements with increasing MPS bond dimension χ. The case of the BAS dataset is different. As discussed in Sec. III, the BAS dataset is a 2d-correlated dataset, whereas the MPS solutions are 1d-correlated. This introduces an unbeneficial bias miss-match which the PQC after circuit extension needs to unlearn. In such cases, initializing the PQC with large χ MPS may not deliver the desired robust improvements. Still, the synergistic approach with any χ greatly outperforms the uninformed random initialization and also the near-identity initialization, which was empirically shown to deal with barren plateaus [59].

Appendix C: Implementation Details of the MPS and PQCs

1. MPS optimization

The MPS for the generative modelling task, i.e., the TNBMs, were implemented and trained according to the method outlined in Ref. [50]. The TNBMs are trained via gradient descent using gradients of the KL divergence loss function (see Eq. (3)) that can be calculated efficiently using a *density matrix renormalization group* (DMRG) method. This ap-

proach allows for adaptive bond dimensions χ for each MPS bond up to a maximum of χ_{max}, or a singular value threshold of $5 \cdot 10^{-5}$ in the SVD truncation. The learning rate η for the gradient descent updates was chosen to be $\eta = 0.01$, with 50 optimization sweeps for the results in Fig. 2 for complete convergence, and 30 optimization sweeps for the gradient results in Fig. 3.

The training of the MPS for Hamiltonian ground state minimization uses an analogous DMRG-based gradient calculation approach for the energy loss function described Eq. (4). The MPS calculations utilize the ITensor library [63], were $\chi = 16$ is used to find the exact ground state without any approximation. The exact MPS was then truncated via SVD to $\chi = 2, 4, 8$ for the results in Fig. 2. This is viable because fidelity with the ground state (which is what SVD empirically optimizes) and the energy in Eq. (4) are proportional, whereas fidelity is not proportional to the KL divergence in Eq. (3).

2. SU(4) gates

The PQCs throughout this work used SU(4) gates between qubit pairs with 15 parameters per gate. Up to a global phase, SU(4) gates represent fully parametrized two-qubit interactions. By the KAK-decomposition [62], an SU(4) rotation can be decomposed into four single-qubit U(3) rotations, and the entangling gates, i.e., XX, YY, and ZZ. Concretely,

$$\begin{aligned} \text{SU}(4)_{i,j}(\boldsymbol{\theta}) = \text{U}(3)_i(\theta_{1:3}) \times \text{U}(3)_j(\theta_{4:6}) \times \\ \text{XX}_{i,j}(\theta_7) \times \text{YY}_{i,j}(\theta_8) \times \text{ZZ}_{i,j}(\theta_9) \times \quad \text{(C1)} \\ \text{U}(3)_i(\theta_{9:12}) \times \text{U}(3)_j(\theta_{13:15}), \end{aligned}$$

with a parameter vector $\boldsymbol{\theta}$ of length 15. The notation $\theta_{l:l+2}$ for $l = 1; 4; 9; 13$, refers to the three elements in the $\boldsymbol{\theta}$ vector that go each U3 gate, e.g., $\theta_{1:3} = [\theta_1, \theta_2, \theta_3]$. We note that PQCs with only SU(4) gates are redundantly parametrized because of the application of consecutive U(3) gates to the same qubit. When implementing on hardware, one may first com-
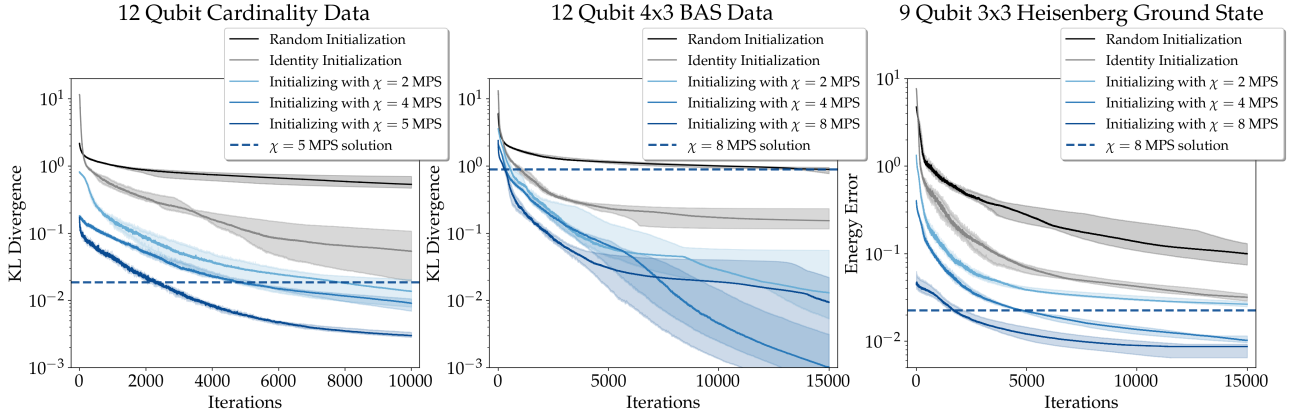
FIG. 5. Optimization results for QCBM and VQE training, where each curve represents the bootstrapped median and 25-75 percentiles of 6 independently initialized runs of the model. The data is complementary to the best of 6 runs shown in Fig. 2. Initializing the PQCs with classically found MPS solutions appears to be generally beneficial as compared to randomly or near-identity initialized parameters. While continuous improvements can be attained with larger bond dimension χ in the Cardinality dataset, as well as the Heisenberg ground state minimization, the BAS dataset - which is a 2D correlated dataset - exhibits an unfavorable bias miss-match with the MPS that the PQC needs to un-learn. Even so, the synergistic framework with small χ greatly improves on the uninformed initializations.

pile the circuit into hardware-native gate sets and then combine trivially redundant gates.

In the case of the MPS initialized PQCs, the decomposed MPS tensors map to a U(4) rotation, which includes a global phase $e^{-i\phi}$ with $\phi \in [0, 2\pi]$. The U(4) were then decomposed into a trainable SU(4) rotation with a fixed global phase via the KAK-decomposition:

$$U(4) \xrightarrow{\text{KAK}} SU(4)(\boldsymbol{\theta}) \times e^{-i\phi} \qquad (C2)$$

The SU(4) gates in the MPS initialized simulations that were not part of the linear layers, i.e., the additional gates that were used to extend the PQC layer to an all-to-all topology, were initialized by sampling a normal distribution $\mathcal{N}$ with zero mean $\mu = 0$ and a small standard deviation $\sigma = 0.01$. This was also done for all gates in the near-identity initializations in Fig. 2. The only exceptions were the additional gates in the VQE simulation where $\sigma = 0$ was chosen, because of the sensitivity of the energy loss in Eq. (4) at these small scales.

3. PQC simulation

All PQCs in this work were simulated using the Qulacs [64] quantum circuit simulator through the interface provided by the ORQUESTRA® platform [65]. All results utilized exact statevector simulation of the quantum states $\psi_{\boldsymbol{\theta}}$ and the corresponding probabilities $q_{\boldsymbol{\theta}}(\mathbf{x})$ described in Secs. II A 1 & II A 2.

4. PQC optimization

The parameters of the PQC models in Sec. III were optimized using a Python implementation of the CMA-ES [35, 61] optimizer through the interface provided by the

ORQUESTRA® platform [65]. In the case of the QCBM training, the initial step size was chosen as $\sigma_{\text{cma}} = 10^{-2}$ in all depicted cases. For the VQE optimization, we found the energy error at the scales $10^{-1} - 10^{-2}$ to be very sensitive to changes in parametrization. Therefore, we chose $\sigma_{\text{cma}} = 10^{-2}$ for the random and near-identity initializations, and $\sigma_{\text{cma}} = 7.5 \cdot 10^{-3}, \sigma_{\text{cma}} = 5.0 \cdot 10^{-3}$, and $\sigma_{\text{cma}} = 2.5 \cdot 10^{-3}$ for the MPS initialized modes with $\chi = 2, \chi = 4$, and $\chi = 8$, respectively. The population sizes λ_{cma} were always chosen to be $\lambda_{\text{cma}} = 20$, meaning that each iteration in Fig. 3 corresponds to 20 quantum circuit simulations. In addition to a limit to the number of optimization steps, i.e., 10000, 15000, and 15000 for the PQCs in Fig. 2 (for Cardinality, BAS, and the Heisenberg model respectively), we also set a loss tolerance of $5 \cdot 10^{-4}$ which may stop the optimization if differences of loss values between steps stay below this threshold. This can be observed in the $\chi = 8$ example on the right tile of Fig. 2.

5. Gradients

The gradients with respect to the KL divergence loss (see Eq. (3)) for Fig. 3 were calculated using a finite-distance gradient estimator

$$\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_l} = \frac{\mathcal{L}(\boldsymbol{\theta} + \epsilon \cdot \theta_l) - \mathcal{L}(\boldsymbol{\theta} - \epsilon \cdot \theta_l)}{2\epsilon} \qquad (C3)$$

with $\epsilon = 10^{-8}$ and for parameter index $l = 8$ in Eq. (C1).

Appendix D: MPS decomposition

Our MPS decomposition protocol used throughout this work was proposed and demonstrated in Ref. [29], and combines the analytical decomposition technique in Ref. [58] with

intertwined optimization steps of the created unitaries using TN techniques on classical computers. Concretely, for a target MPS wavefunction $|\psi_{\chi_{max}}\rangle$, the goal of the decomposition protocol is to output a sequence of circuit layers, i.e., $\prod_{i=k}^{1} U^{(i)}$, where $U^{(i)}$ represent the two-qubit unitaries in the linear quantum circuit layer i, such that the fidelity

$$
\begin{aligned}
f\big(\{U^{(i)}\}_i\big) &= \big|\langle 0^{\otimes N}|\prod_{i=k}^{1} U^{(i)\dagger}|\psi_{\chi_{max}}\rangle\big| \\
&= \big|\langle 0^{\otimes N}|U^{(k)\dagger}\cdots U^{(1)\dagger}|\psi_{\chi_{max}}\rangle\big|
\end{aligned}
\tag{D1}
$$

is maximized. The number of layers $k \leq K$ is chosen to be less or equal than a predetermined circuit depth limit K, or until a target fidelity $\hat{f}$ is achieved via $f\big(\{U^{(i)})\}\big) \geq \hat{f}$. The layer indexing convention from k to 1 is intentional as the decomposition protocol creates circuit layers "from the MPS backwards", such that the newest layer is implemented first in the quantum circuit.

The decomposition protocol sequentially creates circuit layers $U^{(i)}$ via the truncation and disentangling technique described in Ref. [58]. All existing circuit layers $\prod_{i=k'}^{1} U^{(i)}$ for a given decomposition iteration $k' < k$ are optimized via the constrained optimization algorithm described in Ref. [66]. That algorithm leverages the *singular value decomposition* (SVD) and the fact that, for any matrix $\mathcal{M}$ that is decomposed via SVD, i.e., $\mathcal{M} = \mathcal{U} \cdot \mathcal{S} \cdot \mathcal{V}$, the product $\mathcal{U} \cdot \mathcal{V}$ represents the unitary matrix that best approximates $\mathcal{M}$. Ref. [29] shows that this protocol is very effective at decomposing MPS into a low number k of unitary circuit layers, especially with a limited computational budget for the optimization steps. The precision of the decomposition, i.e., the fidelity $f\big(\{U^{(i)}\}_i\big)$ in Eq. (D1), can be sequentially improved with additional circuit layers and optimization steps. Additionally, unlike the method proposed in Ref. [31], the protocol is not hindered by cumulative approximation error build-up when only a limited number of circuit layers are employed.