

Variational Quantum Optimization with Multi-Basis Encodings

Taylor Patti (✉ taylorpatti@g.harvard.edu)

Harvard University

Jean Kossaifi

NVIDIA

Anima Anandkumar

California Institute of Technology (Caltech)

Susanne Yelin

Harvard University

Article

Keywords: quantum algorithms, nonlinear quantum optimization algorithms

Posted Date: November 12th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-654802/v2>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Variational Quantum Optimization with Multi-Basis Encodings

Taylor L. Patti,^{1, 2,*} Jean Kossaifi,² Anima Anandkumar,^{3, 2} and Susanne F. Yelin¹

¹*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

²*NVIDIA, Santa Clara, California 95051, USA*

³*Department of Computing + Mathematical Sciences (CMS), California Institute of Technology (Caltech), Pasadena, CA 91125 USA*

Despite extensive research efforts, few quantum algorithms for classical optimization demonstrate realizable quantum advantage. The utility of many quantum algorithms is limited by high requisite circuit depth and nonconvex optimization landscapes. We tackle these challenges by introducing a new variational quantum algorithm that utilizes multi-basis graph encodings and nonlinear activation functions. Our technique results in increased optimization performance, a factor of two increase in effective quantum resources, and a quadratic reduction in measurement complexity. While the classical simulation of many qubits with traditional quantum formalism is impossible due to its exponential scaling, we mitigate this limitation with exact circuit representations using factorized tensor rings. In particular, the shallow circuits permitted by our technique, combined with efficient factorized tensor-based simulation, enable us to successfully optimize the MaxCut of the nonlocal 512-vertex DIMACS library graphs on a single GPU. By improving the performance of quantum optimization algorithms while requiring fewer quantum resources and utilizing shallower, more error-resistant circuits, we offer tangible progress for variational quantum optimization.

I. INTRODUCTION

NP-hard optimization problems, such as Traveling Salesman and MaxCut, are central to a wide array of fields, such as operational research, engineering, and network design [1]. Despite the classical nature of these problems, there is immense interest in identifying variational quantum algorithms (VQAs) which can solve them faster or more precisely than any classical method, a concept known as quantum advantage [2–5].

One common approach is the variational quantum eigensolver (VQE), where energy minimization yields the ground state of a problem-encoded Hamiltonian through gradient descent update of the quantum circuit parameters [6–8]. The quantum approximate optimization algorithm (QAOA) is a related protocol in which parameterized rotations about both an initial and a problem encoded Hamiltonian are alternated in order to find a solution encoded ground state [9–13]. Novel VQA encoding strategies have also been considered in [14–16]. While the approximation ratios of VQE and QAOA can surpass those of polynomial complexity classical algorithms (e.g., Goemans-Williamson [17]) [18, 19], such guarantees require between polynomially and exponentially many gates in the number of qubits n . Such circuit depths limit the algorithms' potential to demonstrate quantum advantage, rendering them not only computationally inefficient, but also highly susceptible to quantum noise [10, 11, 20] and barren plateaus [21–26].

Furthermore, the difficulty of classically simulating large-scale quantum circuits is a central challenge to algorithm development. This is because the traditional mathematical formalism of quantum mechanics scales exponentially in the number of qubits n , with matrix operators of size 2^{2n} operating on state vectors of size 2^n .

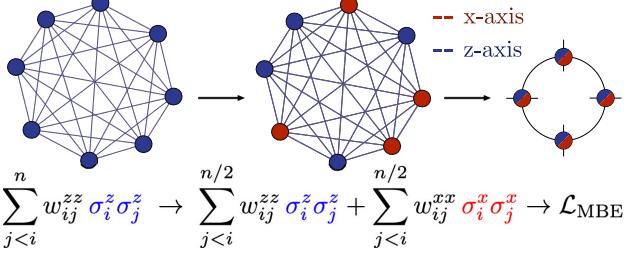
These intractable dimensions for quantum network simulation can be remediated by employing a factorized tensor formalism [27]. While many varieties of decomposed tensors exist, tensor ring has proven particularly popular in the quantum sciences due to its modularity and rank structure, which have close parallels to quantum entanglement. In tensor ring formalism, both quantum states and quantum operators are represented in factorized form by matrix product states (MPS) and matrix product operators (MPOs), respectively [28–30]. However, tensor formalism is often unsuitable for high-depth and connectivity regimes, which are most commonly used in quantum optimization, since tensor rings quickly become prohibitively large (high-rank/bond-dimension) when simulating deep or complicated circuits [31]. Moreover, they are limited to only nearest-neighbor interactions.

Due in part to these limitations, no simulation of more than ~ 100 qubits has demonstrated quantum optimization rivaling that of classical methods for non-local graph instances, even in [32] where exact representations of general tensor architectures with optimal contraction schemes are used. Other large-scale implementations have focused on more restrictive problems. For instance, QAOA MaxCut optimization with up to 210 qubits has been achieved for highly-local 3-regular graphs [33]. QAOA MaxCut optimization has also been implemented with several thousand qubits when exploring only local features of nonlocal graphs, a method which did not yield high average performance [34]. Moreover, large-scale optimization on NP-hard problems (e.g., MaxCut) have not been explored using VQE.

Quantum Computing Contribution - This manuscript introduces a novel method for quantum algorithms that not only outperforms traditional VQAs, it also requires fewer quantum resources and lower computational complexity. In particular:

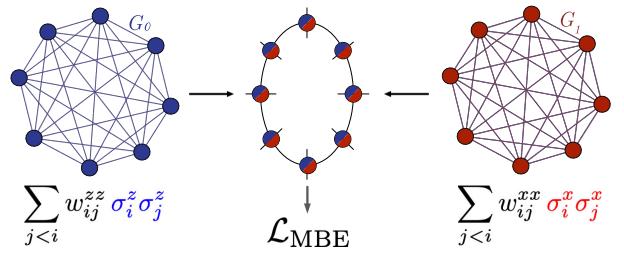
- We propose Multi-Basis Encodings (MBEs), a new quantum optimization algorithm that introduces ad-

* taylorpatti@g.harvard.edu



(a) **Multi-Basis Encoding (MBE) of a graph.**

An n -vertex graph (blue) is represented as an Ising model. We reassigned $n/2$ vertices from σ^z (blue) to σ^x (red) operators, allowing us to map the graph to just $n/2$ qubits (here a locally connected, blue/red tensor ring). The MaxCut is obtained by optimizing this state via single-qubit measurements. Although only locally connected, tensor rings effectively solve the MaxCut of highly nonlocal graphs.



(b) **Multi-Basis Encoding (MBE) with two distinct n -qubit graphs.** Each graph is mapped to the classical Ising model, with G_0 (blue) encoded along the z -basis (as is traditional) and G_1 (red) utilizing the x -basis, resulting in an n -qubit quantum state (blue/red). This encoding is similar to MBE with a single graph, except that the x and z -bases *independently* encode two separate graphs and thus no cross-terms between the z and x -bases are required.

FIG. 1

ditional constraints (regularization) that are *beneficial* to the algorithm's performance, reducing its susceptibility to local minima in the training landscape.

- By doubling the amount of optimization features encoded into a single qubit, MBEs halve the number of qubits required for a given optimization task, a valuable asset for a developing field which has invested millions of dollars and spent multiple decades to achieve ~ 50 -qubit registers and where additional coherence limitations emerge at scale [35]. Moreover, by utilizing single-qubit measurements, these algorithms yield up to a quadratic reduction in runtime.
- By combining our MBEs with non-linear activation functions and an exact factorized tensor network approach, we solve nonlocal MaxCut graph optimization problems with exceptionally shallow quantum circuits. Furthermore, sampling ~ 5 initializations of our MBE experiments on shallow circuits (depth $L = 7$ for 100-vertex graphs, such that L approximately logarithmic in the number of vertices) leads to optimal cut convergence with near unit probability. This shallow-circuit, multi-shot procedure is both more coherent and time-efficient than deterministic convergence with deep circuits, which require up to an exponential number of parameters.

Large-Scale Simulation Contribution - This work utilizes tensor networks, developing new software in order to simulate practical quantum algorithms at unprecedented scale. Specifically:

- We develop TensorLy-Quantum [36], a new software package for simulating efficient quantum circuits with decomposed tensors on CPU and GPU. TensorLy-Quantum is based on the TensorLy software family [37].

- Using TensorLy-Quantum on a single NVIDIA A100 GPU, we simulate solving a 512-vertex MaxCut problem using MBE, which demonstrates superior performance than comparable classical algorithms. This sets a new record for the large-scale simulation of a successful, nonlocal quantum optimization algorithm.

By introducing a new variety of algorithms that improve optimization performance, require fewer quantum resources, and operate on shallower, more error-resistant circuits, we offer tools to increase the utility of variational quantum algorithms.

A. MaxCut Optimization Problems

The Maximum Cut problem, most commonly referred to as *MaxCut*, is a partitioning problem on undirected graphs $G = (V, E)$, where V is a set of vertices (blue orbs in Fig. 2, left) connected by edges E (black lines connecting orbs) [38]. The objective is to optimally assign all vertices $v_i, v_j \in \{-1, 1\}$, so as to maximize the edge weights $w_{ij} \in E$, where any such assignment is referred to as a “cut”. In this work, we will consider a generalized form of the problem known as *weighted* MaxCut, in which w_{ij} take arbitrary real values.

Two formulations of MaxCut exist: the NP-complete decision problem and the NP-hard optimization problem [39]. The former seeks to determine if a cut of size c or greater exists for a given graph G , whereas the latter attempts to identify the largest cut of G possible. We here focus on the more general optimization problem formulation, the ground truth of which we denote $\text{MaxCut}(G)$. It is common practice to express the objective function in its binary quadratic form [38]:

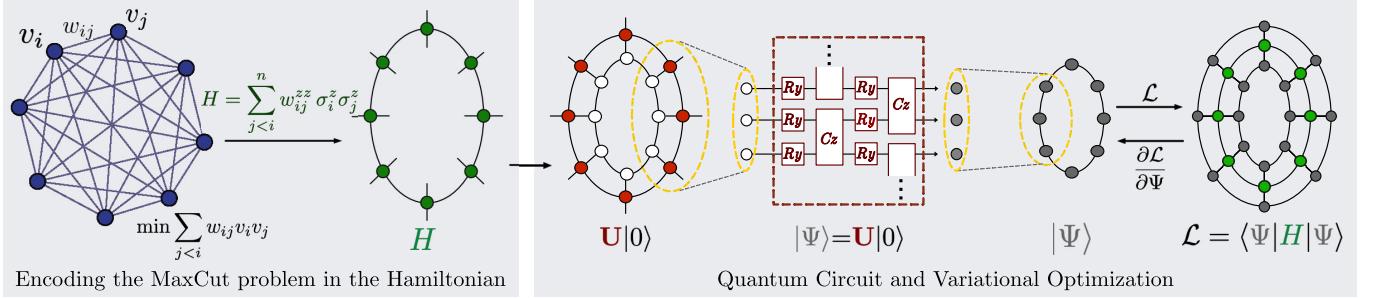


FIG. 2: **Overview of traditional MaxCut encoding and VQE using tensor ring factorizations**, which are tensor train networks with periodic boundary conditions. (Left) A graph G with n vertices v_i, v_j and weights w_{ij} is mapped into an n -qubit Hamiltonian H in MPO form. The MPS ground state $|\psi_g\rangle$ of H encodes the solution to $\text{MaxCut}(G)$. (Right) To find $\text{MaxCut}(G)$ variationally, the null input state $|0\rangle$ (an MPS) is evolved under a parameterized quantum circuit U (an MPO), producing an output state $|\psi\rangle$. U encodes a circuit of depth L (here $L = 4$, red box) in this manuscript’s layer (block) pattern: one layer (block) of single-qubit y -axis rotations R_y followed by a layer of control-Z gates which alternate between even and odd qubits. The energy expectation value $\mathcal{L} = E$ is minimized via gradient descent. The global minimum of \mathcal{L} corresponds to $|\psi\rangle = |\psi_g\rangle$.

$$\text{maximize} \quad \frac{1}{2} \sum_{j < i} w_{ij} (1 - v_i v_j). \quad (1)$$

B. VQE Framework and Tensor Network Formalism

To find the MaxCut of a given graph on a quantum computer, it is convenient to minimize the equivalent summation, $\sum_{j < i} w_{ij} v_i v_j$. For a graph with n vertices v_i , this reduces the problem to finding the n -qubit wavefunction $|\psi\rangle$ that minimizes the energy expectation value $E = \langle\psi|H|\psi\rangle$ of the classical Ising Model Hamiltonian:

$$H = \sum_{j < i}^n w_{ij}^{zz} \sigma_i^z \sigma_j^z. \quad (2)$$

H is obtained by substituting vertices v_i for the Pauli-Z spin operators σ_i^z , as depicted in Fig. 2, and $w_{ij}^{zz} = w_{ij}$ is a relabeling to specify the zz -spin interactions. As H contains only terms in the z -basis, its eigenvectors are classical (zero-entanglement product states), such that $|\psi_i\rangle = \bigotimes_s |s\rangle$, where $|s\rangle \in \{|0\rangle, |1\rangle\}$. We here denote the lowest eigenvalue or “ground state” solution as $|\psi_g\rangle$, the qubits of which form a bijection with the optimal v_i of $\text{MaxCut}(G)$.

Fig. 2 (right) depicts the VQE framework [6–8]. Eq. 1 is optimized by defining the loss function $\mathcal{L} = E$ and varying the parameters $\hat{\theta}$ of a quantum circuit with unitary $U(\hat{\theta})$, which acts on the input quantum state (Fig. 2, right). Without loss of generality, we define the input state as the n -qubit zero state $|0\rangle = \bigotimes_n |0\rangle$, such that

$$|\psi\rangle = U(\hat{\theta})|0\rangle. \quad (3)$$

We decompose this unitary matrix U as Λ subunitaries $U(\hat{\theta}) = \prod_k^\Lambda U_k(\hat{\theta}_k)$, where $\hat{\theta}_k$ is the corresponding sub-

set of $\hat{\theta}$ and $U_k(\hat{\theta}_k) = \prod_{j=1}^n \exp(-i\hat{\theta}_j W_j) M_k$ for generic Hermitian operators W_j and unitary matrices M_k . Thus, the gradient $g_l(\hat{O}) = \frac{\partial \langle \hat{O} \rangle}{\partial \theta_l}$ of operator \hat{O} with respect to any parameter $\theta_l \in \hat{\theta}$ is

$$g_l(\hat{O}) = i \langle 0 | U_R^\dagger \left[W_l, U_L^\dagger \hat{O} U_L \right] U_R | 0 \rangle, \quad (4)$$

where U_L and U_R are the compositions of unitaries U_k with $k \geq l$ and $k < l$, respectively. Rather than using circuits with extensive connectivity, we instead focus on 1D tensor ring circuits of n qubits. In particular, tensor rings have periodic boundary conditions such that qubit $n-1$ is connected to qubit 0. Such local connectivity makes the circuit amenable to both near-term quantum hardware [10, 12] and simulation via decomposed tensors. We accomplish this simulation with TensorLy-Quantum [36]. A nascent and expanding software package, TensorLy-Quantum strives to leverage the structure of decomposed tensors in order to simulate quantum machine learning in the most efficient, non-approximate manner possible. While tensor ring-based tensor networks are typically used for approximate inference and obtained by applying tensor decomposition to dense state vectors and operators, we build a low-rank but exact factorized representation of the simulated quantum circuits. When judiciously constructed, tensor simulations yield a low-rank quantum formalism that permits enormous compression of state and operator spaces. Although in the quantum sciences tensor methods are most frequently associated with state approximations and truncations, like the density matrix renormalization group [40], we here advocate for their use in exact quantum simulation. Similarly, due to their local connectivity, tensor ring factorizations in quantum computing have traditionally been employed for local optimization problems, such as 3-regular MaxCut [41], however we here emphasize their utility for nonlocal optimization tasks.

To analyze VQE with tensor formalism, the Hamiltonian of Eq. 2 is represented as an MPO $H^{\{\beta, \gamma\}}$, with physical indices β and γ . The energy $\mathcal{L} = E$ is then calculated with a single large contraction (Fig. 2, right)

$$E = \sum_{\{\beta, \gamma, \delta, \epsilon\}} \Psi^{\{\beta\}} U^{\{\beta, \gamma\}} H^{\{\gamma, \delta\}} U^{\{\delta, \epsilon\}} \Psi^{\{\epsilon\}}, \quad (5)$$

where

$$\Psi^{\{\beta\}} = \Psi^{\beta_0, \dots, \beta_{m-1}} = \sum_{\{\alpha\}} \psi_{\alpha_0 \alpha_1, \dots, \alpha_{m-1} \alpha_0}^{\beta_0}, \dots, \psi_{\alpha_0 \alpha_1, \dots, \alpha_{m-1} \alpha_0}^{\beta_{m-1}}$$

is an n -qubit MPS of m cores and

$$U^{\{\beta, \gamma\}} = U^{\beta_0, \gamma_0, \dots, \beta_{m-1}, \gamma_{m-1}} = \sum_{\{\alpha\}} u_{\alpha_0, \alpha_1, \dots, \alpha_{m-1}, \alpha_0}^{\beta_0, \gamma_0}, \dots, u_{\alpha_0, \alpha_1, \dots, \alpha_{m-1}, \alpha_0}^{\beta_{m-1}, \gamma_{m-1}}$$

is the corresponding MPO unitary.

As we work in the absence of quantum noise, states $|\psi\rangle$ display time-reversal symmetry and can be fully expressed with real numbers [42]. We thus restrict our rotations to those of the Pauli-Y generator σ^y and implement a simple, repeating subunitary pattern of two layers, also known as blocks. The pattern is illustrated in Fig. 2 (right): a row of parameterized single-qubit rotations $R_y(\theta)$ ($W = \sigma^y$) is followed by a row of control-z (CZ) gates, with the latter alternating control between even and odd qubits. As each single qubit rotation is a 2×2 dense matrix and each two-qubit control-z gate is a rank-2 MPO of two, eight-element cores, the memory requirements of the uncontracted circuit representation scale only linearly in both n and L , an exponential reduction in resources compared to circuits described in traditional quantum formalism. Likewise, a factorized representation of the input state $|0\rangle$ in tensor ring form requires exponentially fewer terms, as it is represented by a rank- $\prod_{i=0}^n 1$ MPS with just n , two-element cores.

II. MULTI-BASIS ENCODING (MBE)

Our MBE protocol uses a loss function which is inspired by the long-range, ZX Hamiltonian

$$H_{zx} = \sum_{j < i} w_{ij}^{zz} \sigma_i^z \sigma_j^z + \sum_{j < i} w_{ij}^{xx} \sigma_i^x \sigma_j^x + \sum_{i,j} w_{ij}^{zx} \sigma_i^z \sigma_j^x, \quad (6)$$

but which encodes separate vertices into the z and x -bases and utilizes single-qubit measurements and non-linear activation functions. The utilization of two, rather than a single, quantum basis has proven useful in other quantum machine learning algorithms [43]. MBE for weighted graphs is depicted in Fig. 1a. An n -vertex graph G is expressed similarly to the Ising model Hamiltonian in Eq. 2, save that half of the vertices are mapped to the x -axis (red) rather than the z -axis (blue), thus enabling

n vertices to be encoded into only $n/2$ qubits. Specifically, MBE halves the number of qubits required for a given optimization, providing a meaningful decrease in quantum hardware overhead.

In order to optimize both axes as independent vertices, we must make several alterations to standard VQE. To begin, $\langle H_{zx} \rangle$ itself is an unsuitable loss function, as the quantum ground state it encodes does not correspond to classical MaxCut of G . We instead focus on the products of single-qubit measurements $\langle \sigma_i^x \rangle$ and $\langle \sigma_i^z \rangle$, such that σ_i^x and σ_i^z operators are simultaneously optimized. This yields the MBE loss function

$$\begin{aligned} \mathcal{L}_{\text{MBE}} = & \sum_{j < i}^{n/2} w_{ij}^{zz} \tanh(\langle \sigma_i^z \rangle) \tanh(\langle \sigma_j^z \rangle) \\ & + \sum_{j < i}^{n/2} w_{ij}^{xx} \tanh(\langle \sigma_i^x \rangle) \tanh(\langle \sigma_j^x \rangle) \\ & + \sum_{i,j}^{n/2} w_{ij}^{zx} \tanh(\langle \sigma_i^z \rangle) \tanh(\langle \sigma_j^x \rangle), \end{aligned} \quad (7)$$

where $\tanh(x)$ is trivially implemented on the classical computer controlling gradient descent. We again emphasize that, as Eq. 7 is comprised of distinct Pauli strings that are independently measured on separate circuit preparations, the uncertainty principle is not violated for w_{ij}^{zx} with $j = i$. The projection of high-dimensional quantum data into a lower-dimensional representation has also been explored in [44, 45]. The inclusion of the non-linear activation function $\tanh(x)$ disincentives the extremization of one basis at the expense of another, which could otherwise occur because the optimal values of both σ_i^x and σ_i^z cannot be linearly encoded by a single quantum state due to the normalization condition of the Bloch sphere of each qubit i

$$\langle \sigma_i^z \rangle^2 + \langle \sigma_i^x \rangle^2 \leq 1, \quad (8)$$

where equality holds for real-valued pure states. As the gradient of $\tanh(x)$ reduces near the ± 1 poles (inset Fig. 3a), full optimization of one axis at the expense of the other is discouraged and optimal cuts for are deduced by a rounding procedure (detailed below) despite the normalization condition of Eq. 8. In this manner, MBE is a dual-axis quantum analog to linear programming relaxations [46]. Furthermore, the normalization constraint of Eq. 8 means that \mathcal{L}_{MBE} can only ever partially descend into local minima and is better equipped to escape their regions of attraction. Finally, we note that we have for simplicity neglected both external fields and y -basis interactions in Eq. 7, however the addition of y -basis terms could immediately be used to both improve the algorithm's performance, as well as to simultaneously optimize three (rather than two) graph vertices.

As minimizing Eq. 7 under the constraints of Eq. 8 cannot yield classical solutions to Eq. 1, we define a rounding

TABLE I: **Comparison of single-graph MBE and traditional VQE** for $n = 100$ vertex graphs for circuits of depth $L = 7$. MBE requires half the number of qubits and parameters as VQE, yet produces significantly better solutions (higher cut \mathcal{C}), both on average and with higher probability.

Method	Depth	#Vertices	#Qubits	#Param	Mean(\mathcal{C})/MaxCut(G)	$P(\mathcal{C} > T)$
VQE	$L = 7$	100	100	400	0.921	12.5%
MBE [Ours]	$L = 7$	100	50	200	0.971	50.0%

TABLE II: **Comparison of single-graph MBE with circuits of depth $L = 13$ and the leading single-shot classical relaxation heuristic** [47], with comparable parameters, for $n = 512$ vertex graphs. MBE produces improved solutions (higher cut \mathcal{C}), both on average and in the most successful run.

Method	Mean(\mathcal{C})/MaxCut(G)	Max(\mathcal{C})/MaxCut(G)
Classical Relaxation	0.939	0.969
MBE [Ours] ($L = 13$)	0.948	0.978

procedure for the classification and scoring of a cut \mathcal{C} for a graph G :

$$\begin{aligned} \mathcal{C}_{\text{MBE}}(\hat{\theta}; G) = & \sum_{j < i}^{n/2} \frac{w_{ij}^{zz}}{2} [1 - R(\langle \sigma_i^z \rangle)R(\langle \sigma_j^z \rangle)] + \\ & \sum_{j < i}^{n/2} \frac{w_{ij}^{xx}}{2} [1 - R(\langle \sigma_i^x \rangle)R(\langle \sigma_j^x \rangle)] + \\ & \sum_{i,j}^{n/2} \frac{w_{ij}^{zx}}{2} [1 - R(\langle \sigma_i^z \rangle)R(\langle \sigma_j^x \rangle)], \end{aligned} \quad (9)$$

where the classically implemented function R rounds the measured expectation values to ± 1 . We note that this scoring is our true, or *computational* MaxCut estimate, as it is the MaxCut assignment which results from projecting the qubit measurements of our quantum state from the $[-0.76, 0.76]$ codomain of our linear programming relaxation ($\tanh(x)$ activation function) back into the ± 1 codomain of MaxCut nodes.

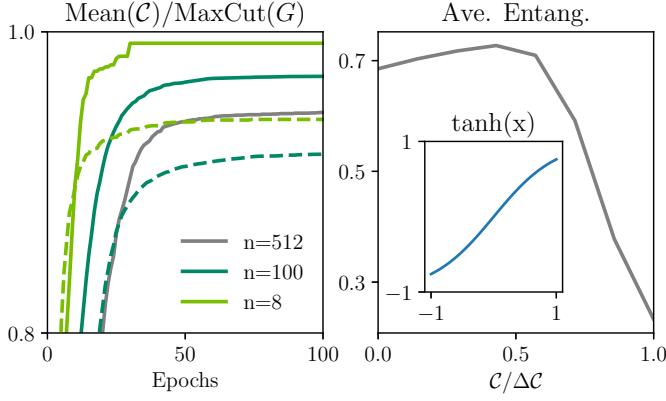
III. RESULTS

In this section, we empirically validate our approach's performance by solving the MaxCut problem on a diverse set of highly nonlocal graphs with up to 512 vertices. We first introduce the experimental settings and implementation details before presenting the results for two scenarios: i) using MBE to solve n -vertex MaxCut problems with only $n/2$ qubits, and ii) using MBE to encode two separate MaxCut graph instances in a single circuit. In addition to having an inherently lower quantum hardware overhead and measurement complexity, both implementations of MBE demonstrate superior optimization performance.

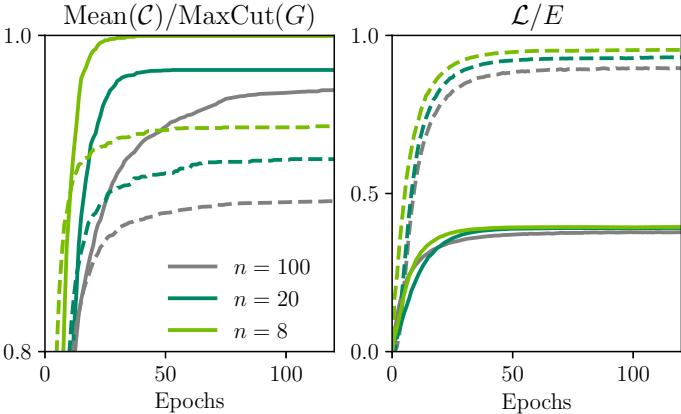
Fig. 3a illustrates the average performance (ratio of cut obtained with largest known solution) of both MBE and VQE circuits for graphs of $n = 8, 100$ vertices and the

MBE circuit alone for $n = 512$. The $n = 512$ graph with traditional VQE was too memory inefficient for evaluation on a single NVIDIA A100 GPU. The simulations were completed using PyTorch [48] and the tensor contractions implemented with Opt-Einsum [49]. The $n = 8$ instances are complete (all-to-all, $n(n-1)/2$ -edge) graphs for which we calculated the exact ground truth through brute force computation, the $n = 100$ graphs are the first three 0.9 density weighted (4455-edge) MaxCut graphs (cataloged as the w09-100 instances) from the extensively studied Biq Mac library [50], and the $n = 512$ graph is the pm3-8-50 instance of the DIMACS library [51]. While the pm3-8-50 graph is relatively sparse (1536 edges), it is nonlocally connected. Like other recent works [22, 52], we implement simple entanglement-based pre-training prior to the MBE algorithm (details in the Supplementary Information [53]). Shallow circuits of depth $L = 7$ ($n = 8$ and $n = 100$ graphs) and $L = 13$ ($n = 512$ graph) are selected in order to adopt a protocol suitable for near-term quantum devices, however the performance of the larger graphs ($n = 100, 512$) increases with moderately deeper circuits.

MBE consistently demonstrates a 5%-7% average performance increase across all n , as seen in Fig. 3a. We emphasize that not only is the MBE algorithm more accurate than traditional VQE, it simultaneously solves MaxCut(G) with *half* the required qubits and parameters, as summarized in Table I. As quantum state space scales exponentially in n , this factor of two reduction in required qubits remains significant for quantum computing at scale. Even with very shallow circuit-depth (L increasing only sublogarithmically in n), MBE outperforms the leading single-shot classical algorithm (Table II) for the 512-vertex DIMACS graph, achieving an average cut of $\sim 95\%$ of the largest known solution [54]. MBE also outperforms the classical algorithm in terms of the largest cut obtained for any given run, with $\sim 98\%$ accuracy from just thirty total runs compared to $\sim 97\%$ accuracy from one-hundred total runs. These performance increases would be even greater for deeper circuits, however our current contraction algorithm yields a maximum



(a) Right: Average cut \mathcal{C} convergence (left) for both MBE (solid lines) and traditional VQE (dashed) with $L = 7$ ($n = 8, 100$) and $L = 13$ ($n = 512$). We note the significantly increased performance for $n = 8, 100$ with MBE over VQE. While VQE with $n = 512$ was prohibitively memory inefficient to simulate for comparison, MBE with $n = 512$ outperforms VQE with $n = 8$, a system 1/64th of its size, as well as the leading single-shot classical algorithm (Table II). Left: Average entanglement entropy for two-qubit subpartitions (maximum value per qubit is 1) vs fraction of calculated MaxCut convergence for nonlinear loss functions. Product state formation occurs because minimizing \mathcal{L}_{MBE} maximizes $\langle \sigma_i^z \rangle^2 + \langle \sigma_i^x \rangle^2$. Inset: $\tanh(x)$ nonlinear activation function further disincentivizes the maximization of one axis at the expense of the other.



(b) Average cut \mathcal{C} convergence (left) and raw loss function \mathcal{L} (right) with both two-graph MBE (solid lines) and traditional VQE (dashed) for $n = 8, 20$, and 100 . MBE improves *calculated* MaxCut convergence \mathcal{C} , although its ability to satisfy by the two encoded Ising models is limited by the normalization condition of Eq. 8. This is remedied by the rounding procedure of Eq. 13.

FIG. 3

MBE circuit depth of $L = 13$ for 512-vertex graphs on a single GPU. The simulation of deeper circuits could be provided by tensor contraction backends with improved memory management, such as the cuTensor library, while

implementations of this scale on quantum hardware is consistent with the projections for moderate-term quantum devices. Although computational benchmarking for optimization problems has been demonstrated for thousands of qubits [34], to our knowledge, MBE with $n = 512$ is the largest simulation of successful quantum optimization algorithms on highly nonlocal graphs yet conducted.

MBE's improved performance on optimization problems is due to the two-axis constraint on each qubit, which only permits convergence to local minima that are *bistable* points for both the z and x -axes. This is in contrast with the monostable condition of traditional VQE. Convergence to a local minima with bistability requires the concurrence of a zero gradient for both independently parametrized axes at a single, non-optimal point in parameter space. As \mathcal{L}_{MBE} is best extremized by larger $\langle \sigma_i^\zeta \rangle$, the circuit will tend towards satisfying the equality in Eq. 8. As this corresponds to entanglement-free qubits, there is a systematic disentanglement of the circuit into product states throughout training (Fig. 3a, right). To understand this process, note that for the general wavefunction

$$|\phi\rangle = \alpha|0_i0_r\rangle + \beta|0_i1_r\rangle + \gamma|1_i0_r\rangle + \delta|1_i1_r\rangle$$

describing any two qubits i and r , the lefthand side of Eq. 8 for qubit i can be written as

$$\langle \sigma_i^z \rangle^2 + \langle \sigma_i^x \rangle^2 = [(\beta + \gamma)^2 + (\alpha - \delta)^2] [(\beta - \gamma)^2 + (\alpha + \delta)^2]. \quad (10)$$

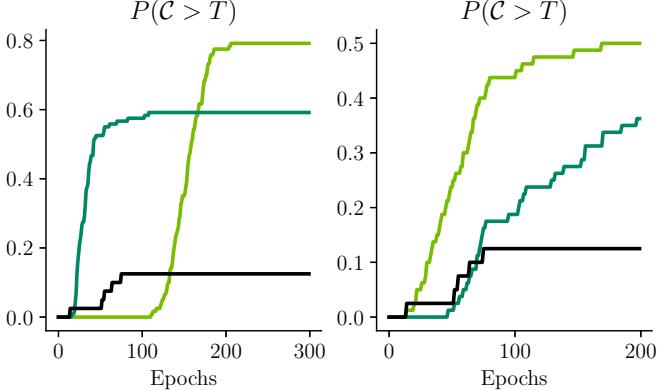
In this form, we note that Eq. 8 is maximized when the concurrence (entanglement [55, 56]) is minimized and vice versa, driving the wave function towards product states as training progresses. Once disentanglement nears completion, the equality in Eq. 8 begins to hold and for any θ_t and qubit i , such that

$$\langle \sigma_i^z \rangle g_t(\sigma_i^z) = -\langle \sigma_i^x \rangle g_t(\sigma_i^x), \quad (11)$$

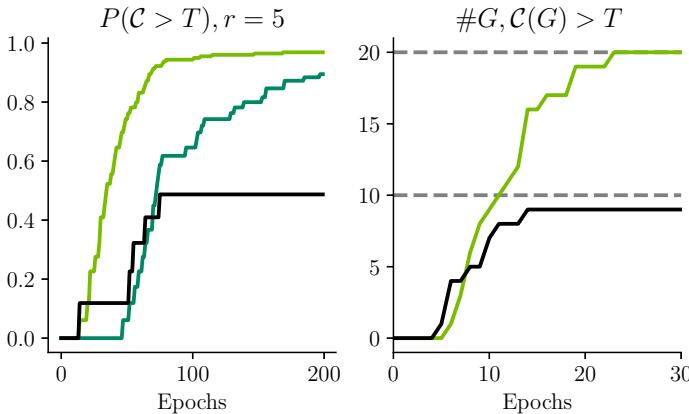
where g_t are the gradients as given by Eq. 4. As $\langle \sigma_q^\zeta \rangle = 0$ is unfavorable for the optimization of \mathcal{L}_{MBE} , both axes of each qubit i must be bistable with respect to each angle θ_t in order for update of that parameter to halt.

In this manner, MBE is a sort of quantum analog to alternating minimization in classical algorithms [57], but which uses both quantum superposition and classical nonlinearity to minimize two cost functions simultaneously, rather than one sequentially. Alternating minimization has also proven useful in QAOA protocols [15, 58–60], as has other perturbations, such as filtered measurements [61]. Because \mathcal{L}_{MBE} is calculated from single-qubit measurements, it is a form of measurement-based quantum computation [62–64]. Moreover, as the number of possible single-qubit measurements scales linearly with circuit width, \mathcal{L}_{MBE} represents up to a quadratic

reduction in the number of observables required to solve complete graphs from $\sim n^2$ (specifically $n(n - 1)/2$ two-operator Pauli strings) to $\sim 2n$ (two single-qubit measurements per qubit), lowering the measurement complexity and runtime of the algorithm on real quantum hardware [65, 66].



(a) (Left) The probability $P(\mathcal{C} > T)$ that cut \mathcal{C} of an $n = 100$ graph is optimal using: MBE with $L = 13$ (light green), MBE with $L = 7$ (dark green), and VQE with $L = 7$ (black). Increasing depth from $L = 7$ to $L = 13$, while still shallow for $n = 100$, markedly improves performance. (Right) $P(\mathcal{C} > T)$ of $n = 100$ graphs using: two-graph MBE with $L = 7$ (light green), two-graph MBE with $L = 1$ (dark green), and VQE with $L = 7$ (black). While the $L = 1$ case is entanglement-free, it benefits from MBE's two-axis constraints.



(b) (Left) The probability of achieving an optimal cut ($\mathcal{C} > T$) of an $n = 100$ graph with $r = 5$ repeats using: two-graph MBE with $L = 7$ (light green), two-graph MBE with $L = 1$ (dark green), and VQE with $L = 7$ (black). For the shallow $L = 7$ MBE circuit, five repetitions produces nearly deterministic results with less than 200 epochs. (Right) Number of $n = 20$ graphs with identified optimal cuts from set of ten instances and $r = 10$ repeats using: two-graph MBE (green), and VQE (black). MBE not only successfully optimizes all (vs 90%) of G , it solves twice as many graphs in the same number of epochs.

FIG. 4

MBE can also encode two *distinct* n -vertex graphs into a single register of n -qubits and solve their two MaxCuts in parallel. This is equivalent to the simplified case of $w_{ij}^{zx} = 0 \forall i, j \leq n$ in Eqs. 7 and 9 using n qubits, yielding

$$\begin{aligned} \mathcal{L}_{\text{MBE}} = & \sum_{j < i}^n w_{ij}^{zz} \tanh(\langle \sigma_i^z \rangle) \tanh(\langle \sigma_j^z \rangle) \\ & + \sum_{j < i}^n w_{ij}^{xx} \tanh(\langle \sigma_i^x \rangle) \tanh(\langle \sigma_j^x \rangle), \end{aligned} \quad (12)$$

and

$$\begin{aligned} \mathcal{C}_{\text{MBE}}(\hat{\theta}; G) = & \sum_{j < i}^n \frac{w_{ij}^{zz}}{2} [1 - R(\langle \sigma_i^z \rangle)R(\langle \sigma_j^z \rangle)] \\ & + \sum_{j < i}^n \frac{w_{ij}^{xx}}{2} [1 - R(\langle \sigma_i^x \rangle)R(\langle \sigma_j^x \rangle)]. \end{aligned} \quad (13)$$

The average performance of MBE for solving two n -vertex graphs in parallel vs that of traditional VQE with a single graph is displayed in Fig. 3b for graphs of $n = 8$, 20, and 100 vertices with $L = 7$. For $n = 8$ and $n = 20$, we generate exact solutions to complete (all-to-all) graphs through brute force computation, whereas the $n = 100$ graphs are again the first three 0.9 density weighted MaxCut graphs from the Biq Mac library [50]. While for this fixed L , both VQE and two-graph MBE suffer decreasing performance with increasing n , two-graph MBE consistently demonstrates a 5%-7% average performance increase across n . We again note that the performance for large- n graphs increases with greater L . Finally, we emphasize that not only is the MBE algorithm more accurate than traditional VQE, it simultaneously solves MaxCut(G) for *two* graphs G , rather than only one as with traditional VQE.

Although much emphasis is placed on the development of quantum algorithms that deterministically obtain optimal cuts, studies have indicated that this requires up to an exponential number of parameters with traditional VQE [8]. This is an unfeasible quantity, reaching $\sim 2^{99}$ ($\sim 2^{511}$) parameters for the $n = 100$ ($n = 512$) graphs considered here. Conversely, the cumulative effects of probabilistic sampling (that is, running the randomly initialized circuit multiple times) lead to high-confidence convergence with markedly few repetitions r . In what follows, we reason that a probabilistic sampling of various shallow MBE circuit initializations is a more efficient alternative. As larger values of \mathcal{C} are a direct certificate of superior optimization, there should be no preference for less efficient single-shot techniques. Furthermore, shallow implementations are particularly important for near-term quantum devices, which are prohibitively susceptible to noise at even moderate circuit-depth.

Fig. 4a displays the probability that an optimal cut, which we define as $\mathcal{C} > T = 0.97 \times \text{MaxCut}(G)$, will be found for $n = 100$ graphs with both MBE and VQE. For depth $L = 7$, MBE produces an optimal cut with

upwards of 50% probability for both the single-graph (n vertices in $n/2$ qubits, Fig. 4a left) and double-graph (two n vertex graphs in n qubits, Fig. 4a right) protocols. In contrast, traditional VQE with $L = 7$ produces optimal cuts with just 12.5% probability. Furthermore, the likelihood of obtaining an optimal cut with MBE increases considerably with moderate circuit depth, rising to approximately 80% for $L = 13$ (left). We note that $L = 1$ circuits (right) obtain optimal cuts with probability 0.36, tripling the convergence rate of standard VQE with 1/7th the resources. As circuits with $L = 1$ are comprised of only local rotations without control gates, the totality of the performance is due to mutual constraints on multi-basis *superpositions*, and not due to quantum entanglement. Like other entanglement-free formulations [67–69], this renders the circuit efficient for classical simulation and indicates that algorithms for simulated superposition with multi-basis constraints may hold promise as “quantum inspired” classical algorithms. However, we note that quantum implementations are still of interest, because other entanglement-free relaxations are known to suffer decreased performance with increasing circuit width n [8]. Furthermore, MBE with even modest entanglement and circuit-depth markedly increases the probability of optimal convergence.

Fig. 4b (left) shows the probability of obtaining at least one optimal cut for $n = 100$ graphs with $L = 7$ and $r = 5$, which nears 97% in fewer than 100 training steps for two-graph MBE circuits. For $r = 10$, convergence is greater than 99.9% and the $4nr = 4000$ parameters utilized for ten repetitions still pale in comparison to the exponentially many required by deep-circuit techniques. As traditional VQE with $L = 7$ and $n = 100$ produces optimal cuts only 12.5% of the time, MBE is four times more effective than VQE for probabilistic optimization.

MBE also offers superior performance over traditional VQE in terms of the diversity of tenable graphs (Fig. 4b, right). For $r = 10$, not only does two-graph MBE find optimal solutions for *all* of the complete $n = 20$ graphs tested (compared to 90% for VQE), its parallel implementation doubles the number of MaxCut instances optimized.

Simulation Considerations - Numerically, \mathcal{L}_{MBE} is more compact for large or dense graphs, where the MPO H quickly becomes cumbersome. However, for the single-qubit measurements required for \mathcal{L}_{MBE} , contraction with a simple, single-qubit operator needs to occur n times. In order to efficiently compute n single-qubit measurements on large, exact tensor networks without either reconstructing an exponentially large ($2^{n/2}$) space or contracting over the full network $\sim n$ times, we use an efficient partial trace-based contraction scheme in which we construct k distinct reduced density matrix operators

$$\rho_k = \sum_{\{\beta, \gamma, \delta \notin K\}} \Psi^{\{\beta\}} U^{\{\beta, \gamma\}} U^{\{\gamma, \delta\}} \Psi^{\{\delta\}}, \quad (14)$$

where K is the k th set of kept indices. K should be sufficiently small so that the $2^{|K|}$ elements of ρ_k remain

numerically tractable. For each ρ_k , $|K|$ smaller partial traces are done to isolate single-qubit density matrices ρ_q , with which we take the single-qubit expectation values of Eq. 12

$$\langle \sigma_q^\zeta \rangle = \text{Tr} [\sigma_q^\zeta \rho_q], \quad (15)$$

where $\zeta = z, x$.

IV. DISCUSSION

In this manuscript, we introduced Multi-Basis Encoding (MBE), a novel technique for quantum optimization algorithms. With shallow circuit-depths, MBE’s performance on a diverse set of graphs exceeds that of traditional VQAs. MBE also provides meaningful efficiency improvements over similar VQAs, and thus lowers the overhead of demonstrating quantum advantage for classical optimization problems. These efficiency improvements include up to a quadratic reduction in circuit measurements, as well as a factor of two decrease in required qubits, which could readily be extended to a factor of three with the inclusion of the y -basis. This reduction is particularly useful as the up to exponential complexity of quantum algorithms scales specifically n , such that MBEs have meaningful efficiency advantages at-scale. Furthermore, when we extend our definition of accuracy to encompass probabilistic sampling of various circuit initializations, we find that remarkably few quantum resources are requisite for classical optimization problems.

MBE could be expanded to a broad framework of multi-axis qubit encodings, which would include any non-linear quantum loss function that permits the optimization of multiple, mutually regularizing observables on a single qubit. These findings are likely to spur additional research in efficient qubit encodings and the application of our techniques to related algorithms. Since deeper circuits would be attainable with more efficient tensor contraction methods or distributed computing efforts, this work encourages further development of large-scale quantum simulation with tensor methods. Most critically, as these simulations are ultimately memory-bound, the implementation of MBE at-scale constitutes a strong and novel candidate for quantum advantage.

We also leverage the powerful tensor techniques packaged in TensorLy-Quantum to complete large-scale simulations of effective optimization algorithms on a single, consumer-grade GPU. To our knowledge, we have produced the largest to-date simulation of a quantum algorithm for a nonlocal optimization problem that rivals classical performance. Such a successful and large-scale implementation demonstrates that exceedingly simple and low-rank tensor representations are sufficient to model diverse techniques in quantum machine learning, and to do so without truncation or approximation. Finally, through the use of large-scale nonlocal graphs, we

demonstrate that the local connectivity and low entanglement capacity of both the MPS formalism and linearly connected near-term quantum devices do not preclude nonlocal quantum optimization routines.

V. ACKNOWLEDGEMENTS

This work was done during T.L.P.’s internship at NVIDIA. At CalTech, A.A. is supported in part by the Bren endowed chair, and Microsoft, Google, Adobe faculty fellowships. S.F.Y. thanks the AFOSR and the NSF for funding. The authors would like to thank Brucek Khailany, Johnnie Gray, Garnet Chan, Andreas Hehn, and Adam Jedrych for conversations.

-
- [1] W. Li, Y. Ding, Y. Yang, R. S. Sherratt, J. H. Park, and J. Wang, *Human-centric Computing and Information Sciences* **10**, 29 (2020).
 - [2] A. Lucas, *Frontiers in Physics* **2**, 5 (2014).
 - [3] D. Wecker, M. B. Hastings, and M. Troyer, *Phys. Rev. A* **92**, 042303 (2015).
 - [4] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New Journal of Physics* **18**, 023023 (2016).
 - [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,” (2020), arXiv:2012.09265 [quant-ph].
 - [6] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nature Communications* **5**, 4213 (2014).
 - [7] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature Publishing Group* **549**, 242 (2017), arXiv:1704.05018.
 - [8] J. Lee, A. B. Magann, H. A. Rabitz, and C. Arenz, “Towards favorable landscapes in quantum combinatorial optimization,” (2021), arXiv:2105.01114 [quant-ph].
 - [9] E. Farhi, J. Goldstone, and S. Gutmann, arXiv (2014), arXiv:1411.4028.
 - [10] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, *et al.*, *Nature Physics* **3**, 1745 (2021).
 - [11] G. G. Guerreschi and A. Y. Matsurra, *Scientific Reports* **9**, 6903 (2019).
 - [12] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyriyanidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, *Proceedings of the National Academy of Sciences* **117**, 25396 (2020), <https://www.pnas.org/content/117/41/25396.full.pdf>.
 - [13] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, *Phys. Rev. X* **10**, 021067 (2020).
 - [14] I. H. Kim and B. Swingle, “Robust entanglement renormalization on a noisy quantum computer,” (2017), arXiv:1711.07500 [quant-ph].
 - [15] Z. Wang, N. C. Rubin, J. M. Dominy, and E. G. Rieffel, *Phys. Rev. A* **101**, 012320 (2020).
 - [16] F. G. Fuchs, H. Kolden, A. Øie, N. Henrik, and G. Sartor, *SN Computer Science* **2**, 2661 (2021).
 - [17] M. X. Goemans and D. P. Williamson, *J. of the ACM* **42**, 1115 (1995).
 - [18] J. Håstad, *J. ACM* **48**, 798–859 (2001).
 - [19] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell, *SIAM* **37**, 319 (2005).
 - [20] J. Preskill, *Quantum* **2**, 79 (2018).
 - [21] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nature Communications* **9**, 4812 (2018).
 - [22] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, “Entanglement devised barren plateau mitigation,” (2020), arXiv:2012.12658 [quant-ph].
 - [23] C. O. Marrero, M. Kieferová, and N. Wiebe, “Entanglement induced barren plateaus,” (2020), arXiv:2010.15968.
 - [24] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, *PRX Quantum* **1**, 020319 (2020).
 - [25] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, “Connecting ansatz expressibility to gradient magnitudes and barren plateaus,” (2021), arXiv:2101.02138 [quant-ph].
 - [26] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, “Cost-function-dependent barren plateaus in shallow quantum neural networks,” (2020), arXiv:2001.00550.
 - [27] M. Fishman, S. R. White, and E. M. Stoudenmire, “The itensor software library for tensor network calculations,” (2020), arXiv:2007.14822 [cs.MS].
 - [28] R. Orús, *Annals of Physics* **349**, 117 (2014).
 - [29] J. C. Bridgeman and C. T. Chubb, *Journal of Physics A: Mathematical and Theoretical* **50**, 223001 (2017).
 - [30] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, *Quantum Science and Technology* **4**, 024001 (2019).
 - [31] Y. Zhou, E. M. Stoudenmire, and X. Waintal, *Phys. Rev. X* **10**, 041038 (2020).
 - [32] E. S. Fried, N. P. D. Sawaya, Y. Cao, I. D. Kivlichan, J. Romero, and A. Aspuru-Guzik, *PLOS ONE* **13**, 1 (2018).
 - [33] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, “Tensor network quantum simulator with step-dependent parallelization,” (2020), arXiv:2012.02430 [quant-ph].
 - [34] C. Huang, M. Szegedy, F. Zhang, X. Gao, J. Chen, and Y. Shi, “Alibaba cloud quantum development platform: Applications to quantum algorithm design,” (2019), arXiv:1909.02559 [quant-ph].
 - [35] N. P. de Leon, K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. S. Palmer, N. Samarth, S. Sangtawesin, and D. W. Steuerman, *Science* **372**, eabb2823 (2021), <https://www.science.org/doi/pdf/10.1126/science.abb2823>.
 - [36] T. L. Patti, J. Kossaifi, and A. Anandkumar, “Tensorly-quantum: Tensor-based quantum machine learning,” (2021).
 - [37] J. Kossaifi, Y. Panagakis, A. Anandkumar, and M. Pantic, *J. Mach. Learn. Res.* **20**, 925–930 (2019).
 - [38] C. W. Commander, “Maximum cut problem, max-cutmaximum cut problem, max-cut,” in *Encyclopedia of Optimization*, edited by C. A. Floudas and P. M. Pardalos

- los (Springer US, Boston, MA, 2009) pp. 1991–1999.
- [39] R. M. Karp, *Reducibility among Combinatorial Problems*, edited by B. J. Miller R.E., Thatcher J.W. (Springer US, Boston, MA, 1972).
- [40] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
- [41] J. Wurtz and P. Love, Phys. Rev. A **103**, 042612 (2021).
- [42] H. Shen, P. Zhang, Y.-Z. You, and H. Zhai, Phys. Rev. Lett. **124**, 200504 (2020).
- [43] X. Gao, E. R. Anschuetz, S.-T. Wang, J. I. Cirac, and M. D. Lukin, “Enhancing generative models via quantum correlations,” (2021), arXiv:2101.08354 [quant-ph].
- [44] A. Eddins, M. Motta, T. P. Gujarati, S. Bravyi, A. Mezzacapo, C. Hadfield, and S. Sheldon, “Doubling the size of quantum simulators by entanglement forging,” (2021), arXiv:2104.10220 [quant-ph].
- [45] M. Matty, Y. Zhang, T. Senthil, and E.-A. Kim, Phys. Rev. Research **3**, 023212 (2021).
- [46] K. Aardal and R. Weismantel, *Polyhedral combinatorics : an annotated bibliography*, Universiteit Utrecht. UU-CS, Department of Computer Science (Utrecht University, Netherlands, 1996).
- [47] S. Burer, R. D. C. Monteiro, and Y. Zhang, SIAM J. Optim. (2001), 10.1137/S1052623400382467.
- [48] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, in *Advances in Neural Information Processing Systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019).
- [49] D. G. a. Smith and J. Gray, Journal of Open Source Software **3**, 753 (2018).
- [50] A. Wiegele, *Big Mac library – a collection of Max-Cut and quadratic 0-1 programming instances of medium size*, Tech. Rep. (2007).
- [51] “he dimacs library of mixed semidefinite-quadratic-linear programs,” .
- [52] J. Dborin, F. Barratt, V. Wimalaweera, L. Wright, and A. G. Green, “Matrix product state pre-training for quantum machine learning,” (2021), arXiv:2106.05742 [quant-ph].
- [53] L. to be Inserted by Publisher,.
- [54] P. Festa, P. Pardalos, M. Resende, and C. Ribeiro, Optimization Methods and Software **17**, 1033 (2002), <https://doi.org/10.1080/1055678021000090033>.
- [55] W. K. Wootters, Quantum Info. Comput. **1**, 27–44 (2001).
- [56] X. Gao, A. Sergio, K. Chen, S. Fei, and X. Li-Jost, Front. Comput. Sci-Chi **2**, 114 (2008).
- [57] P. Jain and P. Kar, Foundations and Trends® in Machine Learning **10**, 142 (2017).
- [58] S. Hadfield, Z. Wang, B. O’Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, Algorithms **12** (2019), 10.3390/a12020034.
- [59] L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, “An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer,” (2020), arXiv:2005.10258 [quant-ph].
- [60] J. Cook, S. Eidenbenz, and A. Bärtschi, “The quantum alternating operator ansatz on maximum k-vertex cover,” (2020), arXiv:1910.13483 [quant-ph].
- [61] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, “Filtering variational quantum algorithms for combinatorial optimization,” (2021), arXiv:2106.10055 [quant-ph].
- [62] R. Raussendorf and H. J. Briegel, Phys. Rev. Lett. **86**, 5188 (2001).
- [63] R. Raussendorf, D. E. Browne, and H. J. Briegel, Phys. Rev. A **68**, 022312 (2003).
- [64] R. R. Ferguson, L. Dellantonio, A. A. Balushi, K. Jansen, W. Dür, and C. A. Muschik, Phys. Rev. Lett. **126**, 220501 (2021).
- [65] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, J. Chem. Phys. **152** (2020), 10.1063/1.5141458.
- [66] O. Shehab, I. H. Kim, N. H. Nguyen, K. Landsman, C. H. Alderete, D. Zhu, C. Monroe, and N. M. Linke, “Noise reduction using past causal cones in variational quantum algorithms,” (2019), arXiv:1906.00476 [quant-ph].
- [67] T. Wang and J. Roychowdhury, “Oim: Oscillator-based ising machines for solving combinatorial optimisation problems,” (2019), arXiv:1903.07163 [cs.ET].
- [68] H. Goto, K. Tatsumura, and A. R. Dixon, Science Advances **5** (2019), 10.1126/sciadv.aav2372, <https://advances.sciencemag.org/content/5/4/eaav2372.full.pdf>.
- [69] E. Crosson and A. W. Harrow, 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS) (2016), 10.1109/focs.2016.81.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SI.tex](#)