CONTINUOUS-VARIABLE OPTIMIZATION WITH NEURAL NETWORK QUANTUM STATES

YABIN ZHANG[1], DAVID GORSICH[2], PARAMSOTHY JAYAKUMAR[2], AND SHRAVAN VEERAPANENI[1*]

ABSTRACT. Inspired by proposals for continuous-variable quantum approximate optimization (CV-QAOA), we investigate the utility of continuous-variable neural network quantum states (CV-NQS) for performing continuous optimization, focusing on the ground state optimization of the classical antiferromagnetic rotor model. Numerical experiments conducted using variational Monte Carlo with CV-NQS indicate that although the non-local algorithm succeeds in finding ground states competitive with the local gradient search methods, the proposal suffers from unfavorable scaling. A number of proposed extensions are put forward which may help alleviate the scaling difficulty.

Keywords. Neural Quantum States, Quantum Information, Graph Theory, Quantum Rotors

1. INTRODUCTION

It is an intriguing fact that many computationally difficult combinatorial optimization problems can be encoded in the ground states of associated physical systems. Classic examples include the *Max-Cut* and *Maximal Independent Set* problems defined on n bits, whose solution sets correspond to the minimal energy ground state configurations for the classical anti-ferromagnetic Ising model and hard-core gas model, respectively. The existence of identifications between combinatorial solution sets and ground states has motivated the pursuit of physically-inspired heuristic approximation algorithms. Perhaps the most famous example of this kind is simulated annealing (SA) [9], which attempts to interpolate, via a sequence of approximate Gibbs distributions, between the uniform distribution on n bits and the uniform distribution on the solution set. Likewise, by encoding the cost function on n bits as the endpoint of a one-parameter family of quantum Hamiltonians, quantum annealing (QA) [5] attempts to follow a trajectory in Hilbert space, which interpolates, via a sequence of approximate ground states, between the uniform superposition state of n qubits and a state in the linear span of orthonormal basis elements corresponding to valid solutions. More recently, partially inspired by the search for practical utility of noisy intermediate-scale quantum computers, variational implementations of both QA and SA have been advocated, which have been termed quantum approximate optimization algorithm (QAOA) [4] and variational neural annealing (VNA) [8]. These variational algorithms achieve the desired interpolation by optimizing over a space of trial wavefunctions (respectively, probability distributions), which are selected from a variational class by following the gradient of a stochastic objective function, estimated via Monte Carlo sampling. Both SA and VNA can thus be viewed as instances of a generalized variational Monte Carlo (VMC) with discrete classical configuration space.

Simulated annealing for continuous optimization has been investigated [11] and continuous variants of QAOA (so called CV-QAOA) have also been recently introduced [15] in the context of continuous-variable quantum computing. This paper arose from a desire to understand the contexts in which CV-QAOA might offer an advantage over classical local update algorithms such as gradient descent. Motivated by this goal, we consider a class of NP-hard continuous optimization problems from two perspectives. The first is a gradient based Riemannian trust-region method due to Burer-Monteiro-Zhang (BMZ) [2]. The second is a non-local probabilistic algorithm inspired by CV-QAOA. Due to the inability to efficiently simulate CV-QAOA at scale, we replace it by a simulation strategy based on neural network quantum states [14]. In particular, the variational wavefunction is chosen to be a rotor variant of the restricted Boltzmann machine, optimized via stochastic natural gradient descent [3, 1, 16] using a random-walk Metropolis-Hastings Markov-Chain Monte

[1]DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109 USA
[2]GROUND VEHICLE SYSTEMS CENTER, U.S. ARMY DEVCOM, WARREN, MI
[*]Corresponding author. *Email address:* shravan@umich.edu .

Carlo strategy. The use of restricted Boltzmann machines (RBMs) enables for an efficient initialization strategy using the output of the BMZ algorithm. The RBM suffers the disadvantage of depending upon an unknown normalizing constant, which precludes the possibility of utilizing an annealing schedule as discussed in [8]. In this preliminary work, we therefore focus on the limit of zero entropy regularization in which the variational Monte Carlo reduces to natural evolution strategies as discussed in [7, 16]. The second approach with neural quantum states (NQS) is quantum-inspired since the considered Hamiltonian for the optimization problem can be regarded as that for a system of quantum oscillators with zero kinetic term. This approach is based upon previous work [3] and more recent work [14] which use neural-network quantum states to simulate quantum systems.

The paper is structured as follows. In section 2 we introduce the continuous optimization problem, explaining its relationship with semi-definite program relaxations for Max-Cut and the mean-field description of quantum anti-ferromagnets. Next we outline the proposed non-local optimization algorithm based on neural network trial densities. The numerical results presented in section 4 demonstrate the ability of the algorithm to accelerate using BMZ initialization. The exploration of different regularization strategies and quantum extensions is left to future work.

2. BACKGROUND

In the computer science literature, a successful strategy for developing approximation heuristics, as well as approximation algorithms with provable approximation guarantees, involves the concept of relaxation of the search space. In the case of Max-Cut, a celebrated approximation algorithm due to Goemans and Williamson [6], proceeds by enlarging the binary search space over n bits, to the product manifold $\prod_{i=1}^{n} \mathbb{S}^{n-1}$, where $\mathbb{S}^{n-1}$ denotes the unit sphere of dimension $n-1$, isometrically embedded in n-dimensional Euclidean space $\mathbb{R}^n$. The cost function defined on the enlarged search space is equivalent to the energy of the classical antiferromagnetic $O(n)$-invariant rotor model, in which the spins are n-dimensional unit vectors. A configuration of rotors achieving the optimal ground state energy can be efficiently found using semidefinite programming and can be mapped by a randomized procedure to a classical bit string satisfying an approximation guarantee. In practice, however, heuristic approximation algorithms are found to significantly outperform the Goemans-Williamson approximation guarantee. In particular, the BMZ algorithm [2], replaces the rank-n relaxation step of Goemans and Williamson with rank-2 relaxation, resulting in the configuration space of n planar rotors $\prod_{i=1}^{n} \mathbb{S}^2$ with energy given by the antiferromagnetic $O(2)$-invariant rotor model. Although global energy optimization of the planar rotor model is not amenable to semidefinite programming, local energy optimization can be efficiently performed and BMZ produces superior cut values to GW in practice.

In order to fix notation let $G = (\mathcal{V}, \mathcal{E})$ be a simple undirected graph with vertex set $\mathcal{V} = \{1, \ldots, n\}$ for some $n \geq 2$ and edge set $\mathcal{E}$. In addition, let the weight for the edge between vertex i and j be w_{ij}. Note that $w_{ij} = 0$ for any pair of $(i,j) \notin \mathcal{E}$ and $w_{ij} = w_{ji}$ for any $i, j \in \mathcal{V}$.

2.1. **Relaxation of Max-Cut to rotor model.** The Max-Cut problem defined on G finds a bipartition of the graph such that the total weights on the edges between the two bipartite sets is maximized. The Max-Cut problem can be formulated as a binary optimization problem

(1)
$$
\begin{aligned}
&\text{minimize} \quad \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}} w_{ij}(1 - x_i x_j) \\
&\text{subject to} \quad |x_i| = 1, \quad i = 1, \ldots, n
\end{aligned}
$$

Since the Max-Cut problem is NP-hard, we follow [2] by considering a rank-2 relaxation together with a heuristic algorithm. This relaxation replaces $x_i \in \mathbb{R}$ with $\vec{v}_i \in \mathbb{R}^2$, yielding

(2)
$$
\begin{aligned}
&\text{minimize} \quad \sum_{\{i,j\} \in \mathcal{E}} w_{ij} \langle \vec{v}_i, \vec{v}_j \rangle \\
&\text{subject to} \quad \|\vec{v}_i\| = 1, \quad i = 1, \ldots, n
\end{aligned}
$$

Since v_i lives on the unit circle in $\mathbb{R}^2$, the above minimization problem is equivalent to the unconstrained minimization problem by switching to polar coordinates, where $\theta_i \in \mathbb{R}$,

$$
(3) \qquad \text{minimize} \qquad \sum_{\{i,j\}\in\mathcal{E}} w_{ij} \cos(\theta_i - \theta_j)
$$

A solution to the original Max-Cut problem can be obtained by first solving the unconstrained optimization problem for $\theta = (\theta_1, \ldots, \theta_n)$ and then applying the heuristic algorithm "Procedure-Cut" in [2] which maps the angle θ to a valid cut $x = (x_1, \ldots, x_n)$ for the graph.

2.2. **Heisenberg model.** As additional motivation for the continuous optimization problem (3), consider the following Heisenberg quantum Hamiltonian formulated on the graph G,

$$
(4) \qquad H = \sum_{\{i,j\}\in\mathcal{E}} w_{ij}(X_i X_j + Z_i Z_j) \ .
$$

and define the following variational wavefunction parametrized by $\theta = (\theta_1, \ldots, \theta_n)$,

$$
(5) \qquad \rho = \bigotimes_{i=1}^{n} \frac{1}{2}(I + \sin\theta_i X_i + \cos\theta_i Z_i) \ .
$$

Evaluating the quantum expectation value of H in the state ρ gives the variational upper bound $\operatorname{tr}(H\rho) \geq \lambda_{\min}(H)$ for all θ. Optimizing the variational bound over θ is equivalent to problem (3).

3. Neural quantum state relaxation

In this section we pass from deterministic rotor configurations $\theta \in [0, 2\pi)^n$ to parametrized probability densities over $[0, 2\pi)^n$ satisfying periodic boundary conditions. In particular, we introduce a set of m hidden rotor variables $\varphi = (\varphi_1, \ldots, \varphi_m) \in [0, 2\pi)^m$ and define a wavefunction by integrating a Boltzmann factor with respect to the flat measure,

$$
(6) \qquad \psi(\theta) = \int_{[0,2\pi]^m} \mathrm{d}^m\varphi \, \exp\left(-E(\theta, \varphi)\right)
$$

where the energy is of the following restricted Boltzmann form

$$
(7) \qquad E(\theta, \varphi) = -\sum_{i=1}^{m}\sum_{j=1}^{n} a_{ij}\langle \vec{z}_i, \vec{v}_j \rangle - \sum_{i=1}^{m}\langle \vec{b}_i, \vec{z}_i \rangle - \sum_{j=1}^{n}\langle \vec{c}_j, \vec{v}_j \rangle \ .
$$

Here we have passed back to Cartesian coordinates $\vec{v}_i = (\cos\theta_i, \sin\theta_i)$ and $\vec{z}_i = (\cos\varphi_i, \sin\varphi_i)$. The wavefunction is parametrized by weights $a_{ij} \in \mathbb{R}$ and biases $\vec{b}_i \in \mathbb{R}^2$ and $\vec{c}_j \in \mathbb{R}^2$ for all $(i,j) \in [m] \times [n]$. It is equally possible to consider holomorphic parametrization in the spirit of [3], although we only consider real parametrization in this work. The wavefunction defines a probability density given by the Born rule,

$$
(8) \qquad \pi(\theta) := \frac{|\psi(\theta)|^2}{Z}
$$

which is normalized by the partition function

$$
(9) \qquad Z := \int_{[0,2\pi]^n} \mathrm{d}^n\theta \, |\psi(\theta)|^2 \ .
$$

Although the integral defining the partition function is intractable, the integrals over φ can be computed in closed form owing to the bipartite structure of the energy function. It follows that the probability density $\pi(\theta)$ over visible rotor configurations is known up to an overall normalizing constant. The stochastic optimization problem is thus defined as

$$
(10) \qquad \text{minimize} \qquad \mathbb{E}_{\theta \sim \pi}\left[\sum_{\{i,j\}\in\mathcal{E}} w_{ij} \cos(\theta_i - \theta_j)\right]
$$

where the unconstrained minimization is over the density class parameterized by $\mathbb{R}^{nm+2(n+m)}$. In practice, this minimization is performed using stochastic reconfiguration [10, 3, 1]. For details on the stochastic configuration, please refer to the supplemental materials of [3].

FIGURE 1. Flow chart for solving the Max-Cut problem: the (black) solid arrows represent the NQS algorithm proposed in this manuscript; the (blue) dashed line arrow represent an deterministic approach given in [2].

4. NUMERICAL RESULTS

We investigated the performance of the NQS solver on two types of graphs, characterized by having either randomized or uniform edge weights, respectively. Previous studies on Kuramoto oscillators show that the continuous optimization problem (3) formulated with uniform weights and high edge density can be solved via local algorithms since all local extremas are also global extremas [13]. Since the BMZ algorithm uses only local information whereas the NQS algorithm updates a global density, the latter potentially holds promise in producing improved solutions in the regime where local extremas are not always global extremas. Thus, we included some low edge density uniform weight graphs in our tests in addition to general random weight graphs.

We will consider the following two graphs in particular:

- `RWN50E619`: a graph with 50 nodes and 619 edges with random weight for each edge in the range (0,15). The edge density is approximately 50%.
- `UWN100E519`: a graph with 100 nodes and 519 edges with uniform weight for each edge. The edge density is about 10%.

For comparison with NQS we use an implementation of the BMZ algorithm in [2]. The main idea of the BMZ approach is to relax the original binary minimization problem of Max-cut to a unconstrained continuous variable minimization problem defined on $\mathbb{R}^n$, where n is the number of nodes on the graph. The Hessian

Graph	Energy value		
	Mean	Standard deviation	Minimum
RWN50E619	-1449.35	10.5	-1459.35
UWN100E519	-168.45	1.52	-170.65

TABLE 1. Observed solution energy for applying the BMZ algorithm to the two tested graphs.

matrix of the relaxed problem is easy to compute and thus the relaxed problem can be solved via standard classical minimization algorithms such as the trust-region method used in our implementation. The solution results for the two tested graphs are given in Table 1.

The proposed NQS algorithm has several parameters whose values were found to control the solution quality. Some of these parameters are held fixed throughout all testing. In particular, the hidden unit density, i.e., the ratio between the number of hidden nodes and the visible nodes $\alpha = m/n$, was set to 1, the descent direction was solved via MINRES with TOL=10^{-10}, and the learning rate was set to 0.01. Parameters that were varied amongst the tests include

- N_{samp} and N_{warm}: the number of Metropolis samples and warm samples (early samples that are discarded when calculating the expectation) used in the Markov chain.
- N_{iter}: the number of stochastic reconfiguration steps
- λ_{reg}: the regularization parameter used in solving the linear system for each stochastic reconfiguration step

Section 4.1 describes how these parameters affect the solution for the RWN50E619 graph. Both the BMZ algorithm and the proposed NQS algorithm were executed 10 times with 10 random seeds, and the resulting mean energy, standard deviation, and minimum energy are reported for all tests. We focus on the solution quality rather than computational cost in wall-clock time when comparing the results from the BMZ and the NQS approaches. As for running time, our implementation of BMZ solves the test problems on the order of 1 second while the NQS may require running time on the order of 10^3 seconds. Such gap is expected as the sampling step of the NQS approach is costly and similar gap is also observed in previous work [16].

Section 4.2 analyzes the NQS algorithm performance for uniform weight graphs with low edge density. While in theory we expect the proposed algorithm to outperform the BMZ algorithm due to the latter's local nature, in practice sub-optimal solutions are found. The quality of the solutions is expected to improve with improved parameter tuning or simply by increasing N_{samp}, N_{warm}, and N_{iter}.

Finally, section 4.3 discusses an initialization technique for the NQS algorithm, which exploits the output of the BMZ algorithm to initialize the weights and biases of the restricted Boltzmann machine.

All tests are run on the Great Lakes cluster provided by Advanced Research Computing at the University of Michigan, with four 2x 3.0 GHz Intel Xeon Gold 6154 CPUs and 1GB memory per CPU requested per task. Although this section is dedicated to test problems which require nontrivial optimization effort, a collection of much smaller and/or simpler graphs are also included in the Appendix, which verifies that the proposed approach produce the optimal solution for simple graphs.

4.1. **Performance on random weight graph.** For the RWN50E619 graph, the BMZ algorithm with 10 different random seeds produced an mean energy of -1449.35, standard deviation of 10.5, and minimum of -1459.35. In this section, we study how the VMC parameters in the NQS algorithm affect the solution quality when applied to this graph.

The impact of number of stochastic reconfiguration steps N_{iter} holding fixed $N_{\text{samp}} = 40, N_{\text{warm}} = 0, \lambda_{\text{reg}} = 10^{-9}$ is shown in Table 2. As N_{iter} increases, the energy approaches but is less than the value given by the BMZ algorithm.

We then fix $\lambda_{\text{reg}} = 10^{-9}, N_{\text{iter}} = 12000$, and explore the changes in the energy value for different choices of N_{samp} and N_{warm}. Since it is computationally costly to explore a wide range of pair values $(N_{\text{samp}}, N_{\text{warm}})$, we picked four choices as representatives: $(40, 0)$, $(100, 10)$, $(200, 20)$, and $(400, 40)$. The results are visualized in Figure 2. The results show that as N_{samp} and N_{warm} both increase, the resulting energy increases and gets closer to the results from the BMZ algorithm.

N_{iter}	Mean	Energy value Standard deviation	Minimum
8000	-1359.29	36.23	-1405.88
12000	-1388.47	24.81	-1416.89
16000	-1404.48	18.64	-1431.56
20000	-1407.50	22.08	-1435.02

TABLE 2. Observed solution energy for the `RWN50E619` graph with respect to N_{iter}, the number of stochastic reconfiguration steps. $N_{\text{samp}} = 40$, $N_{\text{warm}} = 0$, and $\lambda_{\text{reg}} = 10^{-9}$ for all test cases. See Table 1 for the corresponding results obtained by the BMZ algorithm.

FIGURE 2. Optimal energy found for the `RWN50E619` graph with respect to N_{samp} and N_{warm} the number of samples and warm samples in Markov chain. Each vertical line corresponds to one choice of the parameter pair $(N_{\text{samp}}, N_{\text{warm}})$ or the BMZ algorithm, and the three markers for each line give the observed minimum, average, and maximum of the energy approximations for 10 runs. For the NQS approach, $N_{\text{iter}} = 12000$ and $\lambda_{\text{reg}} = 10^{-9}$ for all test cases.

λ_{reg}	Mean	Energy value Standard deviation	Minimum
10^{-9}	-1388.47	24.81	-1416.89
10^{-7}	-1389.31	26.85	-1418.32
10^{-5}	-1375.55	36.13	-1418.18

TABLE 3. Observed solution energy for the `RWN50E619` graph with respect to λ_{reg}, the regularization parameter in solving the descent direction. $N_{\text{samp}} = 40$, $N_{\text{warm}} = 0$, and $N_{\text{iter}} = 12000$ for all test cases. See Table 1 for the corresponding results obtained by the BMZ algorithm.

We also studied the impact of the regularization parameter, λ_{reg}, holding fixed $N_{\text{samp}} = 40, N_{\text{warm}} = 0, N_{\text{iter}} = 12000$. Results are given in Table 3. Based on the results, we decided to use $\lambda_{\text{reg}} = 10^{-6}$ for larger size problems.

Table 4 reports the wall-clock time required for the NQS with selected parameter values. As we expected, the cost roughly doubles when the value of the parameters including N_{iter}, N_{samp} and N_{warm} doubles.

N_{iter}	$(N_{\text{samp}}, N_{\text{warm}})$	Time in seconds
800	(40,0)	1.5×10^3
1600	(40,0)	2.9×10^3
1200	(40,0)	2.2×10^3
1200	(100,10)	5.2×10^3
1200	(200,20)	1.1×10^4
1200	(400,40)	2.5×10^4

TABLE 4. Observed running time in seconds for applying the NQS algorithm to the RWN50E619 graph. $\lambda_{\text{reg}} = 10^{-9}$ for all cases.

4.2. **Performance on uniform weight graph.** Inspired by the discussion in [13], we also applied the proposed algorithm to uniformly weighted graphs with low edge density. The idea being that, in this regime, local extrema may not correspond to global extrema, and thus non-local search heuristics may be advantageous compared to local search. The results for graph UWN100E246 are visualized in Figure 3, using $\lambda_{\text{reg}} = 10^{-6}$ and varying the parameters $N_{\text{samp}}, N_{\text{warm}}, N_{\text{iter}}$. The BMZ algorithm produces energy mean of -168.45, standard deviation of 1.52, and minimum of -170.65. Somewhat surprisingly, the NQS approach fared comparatively worse compared with BMZ in this regime. Presumably this could be overcome by investing more time and effort into tweaking the parameters and running tests with larger value for N_{samp}, N_{warm}, and N_{iter}. For example, the current results show consistent improvement in solution optimality as $(N_{\text{samp}}, N_{\text{warm}})$ increases from $(40, 0)$ to $(400, 40)$. This suggests that simple parameter tweaking, such as increasing N_{samp} and N_{warm}, could bring the the NQS results on par with BMZ.


FIGURE 3. Optimal energy for the graph UWN100E246 with respect to N_{samp} and N_{warm} the number of samples and warm samples in Markov chain and N_{iter} the number of reconfiguration steps. Each vertical line corresponds to one choice of the parameters or the BMZ algorithm, and the three markers for each line give the observed minimum, average, and maximum of the energy approximations for 10 runs. The runs with $N_{\text{iter}} = 12000$ are in black and those with $N_{\text{iter}} = 16000$ are in blue. For the NQS approach, $\lambda_{\text{reg}} = 10^{-6}$ for all test cases.

4.3. **Pretrained initialization.** The cost of the NQS algorithm can be significantly reduced by initializing the weights and biases of the RBM with an educated guess. If the rotor configuration output from the BMZ algorithm is available, this can be used to initialize the RBM weights and biases as follows. Given the output $(\theta_1^*, \ldots, \theta_n^*)$ of BMZ, we initialize the bias on the visible units as $\vec{c}_i = (r\cos(\theta_i^*), r\sin(\theta_i^*))$, where $r \geq 0$ is a parameter. The initial hidden bias was chosen to be zero and the weights a_{ij} are randomly drawn from

a normal distribution with zero mean and 0.1 standard deviation. Figure 4 confirms the intuition that the algorithm converges faster using the pretrained initialization.


FIGURE 4. Performance of the proposed NQS algorithm with smart initialization from solution given by the BMZ algorithm applied to the `RWN50E619` graph with $N_{\text{samp}} = 40$, $N_{\text{warm}} = 0$, $N_{\text{iter}} = 12000$, and $\lambda_{\text{reg}} = 10^{-6}$.

5. CONCLUSIONS AND FUTURE WORK

We demonstrated the feasibility of continuous optimization using a quantum-inspired variational algorithm, which we interpret in two ways: either as simulating the continuous-variable QAOA for an efficiently simulable neural-network wavefunction, or alternatively as a realization of continuous-variable natural evolution strategies. Performance was assessed using a challenging non-convex optimization problem and it was found that, despite the non-local nature of the algorithm, a purely local heuristic out-performed the quantum-inspired algorithm in all cases. It is an interesting open problem to find continuous optimization tasks for which variational probabilistic/quantum algorithms offer an advantage compared to local search heuristics. It is also interesting to consider improvements to the choice of classically simulable wavefunction. For example, we have focused on unnormalized densities where samples are generated using the Metropolis-Hastings MCMC algorithm inspired by [3]. A more sophisticated proposal distribution based on Hamiltonian Monte Carlo is expected to deliver significantly improved approximations of Monte Carlo samples. Since the algorithm we described can be also be understood as the endpoint of variational neural annealing (VNA) flow, it will be natural to consider the continuous-variable extension of VNA using normalized densities, and corresponding initialization strategies.

6. ACKNOWLEDGEMENTS

REFERENCES

[1] Shun-ichi Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276, 02 1998.
[2] Samuel Burer, Renato D. C. Monteiro, and Yin Zhang. Rank-two relaxation heuristics for max-cut and other binary quadratic programs. *SIAM Journal on Optimization*, 12(2):503–521, 2002.
[3] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.

[4] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

[5] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000.

[6] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

[7] Joseph Gomes, Keri A McKiernan, Peter Eastman, and Vijay S Pande. Classical quantum optimization with neural network quantum states. *arXiv preprint arXiv:1910.10675*, 2019.

[8] Mohamed Hibat-Allah, Estelle M Inack, Roeland Wiersema, Roger G Melko, and Juan Carrasquilla. Variational neural annealing. *arXiv preprint arXiv:2101.10154*, 2021.

[9] Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

[10] William Lauchlin McMillan. Ground state of liquid he 4. *Physical Review*, 138(2A):A442, 1965.

[11] Alexander G Nikolaev and Sheldon H Jacobson. Simulated annealing. In *Handbook of metaheuristics*, pages 1–39. Springer, 2010.

[12] Ralf Schützhold and Sarah Mostame. Quantum simulator for the o (3) nonlinear sigma model. *Journal of Experimental and Theoretical Physics Letters*, 82(5):248–252, 2005.

[13] S. Steinerberger. Max-cut via kuramoto-type oscillators. *ArXiv*, abs/2102.04931, 2021.

[14] James Stokes, Saibal De, Shravan Veerapaneni, and Giuseppe Carleo. Continuous-variable neural-network quantum states and the quantum rotor model. *arXiv preprint arXiv:2107.07105*, 2021.

[15] Guillaume Verdon, Juan Miguel Arrazola, Kamil Brádler, and Nathan Killoran. A quantum approximate optimization algorithm for continuous problems. *arXiv preprint arXiv:1902.00409*, 2019.

[16] Tianchen Zhao, Giuseppe Carleo, James Stokes, and Shravan Veerapaneni. Natural evolution strategies and variational Monte Carlo. *Machine Learning: Science and Technology*, 2(2):02LT01, 2020.

Appendix A. NQS applied to simple graphs

In this section, we consider a collection of Max-Cut problems defined on simple and/or small graphs for which the optimal solution can be obtained easily via brute force search. The graphs, optimal cut value, and the cut value obtained by the NQS approach coupled with the Procedure-Cut algorithm for converting continuous variable solution to cut are given in Table 5. For graphs with 4 nodes, we use $N_{\text{iter}} = 300$, $N_{\text{samp}} = 10$, $N_{\text{warm}} = 0$ and $\lambda_{\text{reg}} = 10^{-9}$. For graphs with 6 nodes, we use $N_{\text{iter}} = 1000$, $N_{\text{samp}} = 40$, $N_{\text{warm}} = 0$ and $\lambda_{\text{reg}} = 10^{-9}$. For graphs with 12 nodes, we use $N_{\text{iter}} = 4000$, $N_{\text{samp}} = 40$, $N_{\text{warm}} = 0$ and $\lambda_{\text{reg}} = 10^{-9}$. Note that the observed standard deviation for the 10 runs (each with a different random seed) for all the tests cases is zero. This suggests that for simple graphs with appropriate parameter choices the NQS approach, although being Stochastic in nature, always gives the rotor configurations that leads to the optimal cut value for the graphs.
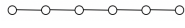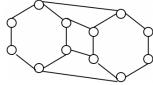
Graph		Cut value via NQS and Procedure-Cut		
(all edge weight=1)	Optimal cut value	Mean	Standard deviation	Maximum
	3	3	0	3
	4	4	0	4
	4	4	0	4
	5	5	0	5
	6	6	0	6
	3	3	0	3
	9	9	0	9
	16	16	0	16

TABLE 5. Results from applying the NQS and Procedure-Cut algorithm to the Max-Cut problems defined on simple graphs with uniform edge weight. For each problem, the NQS and Procedure-Cut algorithm is run 10 times with 10 different random seeds.