

CONTINUOUS-VARIABLE NEURAL-NETWORK QUANTUM STATES AND THE QUANTUM ROTOR MODEL

JAMES STOKES^{1*}, SAIBAL DE^{1*}, SHRAVAN VEERAPANENI^{1†}, AND GIUSEPPE CARLEO²

ABSTRACT. We initiate the study of neural-network quantum state algorithms for analyzing continuous-variable lattice quantum systems in first quantization. A simple family of continuous-variable trial wavefunctions is introduced which naturally generalizes the restricted Boltzmann machine (RBM) wavefunction introduced for analyzing quantum spin systems. By virtue of its simplicity, the same variational Monte Carlo training algorithms that have been developed for ground state determination and time evolution of spin systems have natural analogues in the continuum. We offer a proof of principle demonstration in the context of ground state determination of a stoquastic quantum rotor Hamiltonian. Results are compared against those obtained from partial differential equation (PDE) based scalable eigensolvers. This study serves as a benchmark against which future investigation of continuous-variable neural quantum states can be compared, and points to the need to consider deep network architectures and more sophisticated training algorithms.

Keywords. Neural Quantum States, Continuous-variable Quantum Information, Quantum Rotors

1. INTRODUCTION

Variational Monte Carlo (VMC) approaches to the quantum many-body problem [1] have witnessed a recent resurgence in activity fueled by the realization that when neural networks are exploited as systematically improvable trial wavefunctions, direct attack on otherwise intractable quantum many-body systems becomes possible. The success of so-called neural-network quantum states [2] is closely paralleled by the ability of deep neural networks to overcome a related curse-of-dimensionality in a variety of high-dimensional machine learning tasks. The key features shared by these learning tasks is that they involve iterative fitting of a high-dimensional function approximator using various forms of stochastic approximation to an objective function, for which the learner has incomplete knowledge in the form of samples. Applying this perspective to the VMC for ground state determination, the variational wavefunction acts as a data generating process from which IID samples can be drawn. These samples provide noisy estimates for a Rayleigh quotient to be optimized, whose exact dependence on the variational parameters is unknown.

In this paper, we argue that quantum many-body simulation can leverage the success of geometric deep learning [3, 4] from two different perspectives which are based on group invariances of the Hamiltonian and the space of states, respectively. The first perspective applies whenever the configuration space of the Hamiltonian has a non-Euclidean structure, as in the hyperbolic lattices [5, 6] currently under investigation using circuit quantum electrodynamics [7]. In the second perspective, one is concerned with efficiently describing states of the Hilbert space which are invariant or equivariant to group invariances of the Hamiltonian. In the theory of many-body Schrödinger operators, for example, it can be shown that the unrestricted ground state inherits any group invariances of the parent Hamiltonian.

Motivated by the above simulation prospects, in this paper we initiate the study of continuous-variable VMC on non-Euclidean spaces, proposing a Hamiltonian which can be summarized by a small amount of geometric data, whose invariances are respected by the quantum theory. We limit our investigation to the simplest non-trivial target space geometry corresponding to a system of quantum rotors. Ref. [8] undertook a variational ground state study of the quantum rotor model in one spatial dimension using basis truncation and matrix product states. The neural network approach advocated here, in contrast, does not rely on a choice of basis and could potentially be advantageous in the analysis of more complicated geometries such

¹DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109 USA

²INSTITUTE OF PHYSICS, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE (EPFL), CH-1015 LAUSANNE, SWITZERLAND

*THESE AUTHORS CONTRIBUTED EQUALLY TO THIS WORK.

†Corresponding author. *Email address:* shravan@umich.edu .

as hyperbolic manifolds, which are central to the study of quantum chaos [9]. Another possible advantage of the neural network approach compared to [8] is the lack of reliance on the density matrix renormalization group training algorithm [10, 11], which tends to struggle beyond one spatial dimension or in applications to disordered systems without a lattice.

Rather than pursuing the full machinery of geometric deep learning for variational simulation, we choose to focus on the introduction of baselines and illustration of general techniques using minimalist architectures which generalize those originally introduced for quantum spin systems [2]. In particular, experiments are conducted using a rotor variant of the restricted Boltzmann machine applicable to the planar quantum rotor model, which can be understood as a continuous-variable relaxation of the quantum Ising model. The results are compared against scalable PDE-based eigensolvers and the extension of the method to other geometries is discussed.

The paper is structured as follows: we first introduce a geometrically motivated Hamiltonian, highlighting some subtleties that arise in non-Euclidean space. The planar quantum rotor model is identified as the simplest candidate system for illustrating the applicability of geometric machine learning techniques. Baselines are introduced for ground state preparation of the rotor model, which we investigate using a combination of variational techniques inspired by shallow neural networks, as well as techniques based on scalable PDE eigensolvers. We conclude by summarizing future directions.

2. STATES AND HAMILTONIAN

The quantum systems under consideration describe finitely many particles moving on a Riemannian target space subject to two-body interactions. In particular, given a choice of Riemannian manifold (M, g) , a finite simple undirected graph $G = (\mathcal{V}, \mathcal{E})$ decorated by vertex and edge weights functions $h : \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ and $\beta : \mathcal{E} \rightarrow \mathbb{R}$ respectively, and a choice of potential function $V : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$, we define the generator of time evolution H in the infinite-dimensional Hilbert space of states $\mathcal{H} = \bigotimes_{i \in \mathcal{V}} L^2(M)$ as follows,

$$(1) \quad H\psi = -\frac{1}{2} \sum_{i \in \mathcal{V}} h_i \Delta_{g_i} \psi + \sum_{\{i, j\} \in \mathcal{E}} \beta_{ij} V(\text{dist}(x_i, x_j)) \psi ,$$

where Δ_{g_i} denotes the Laplace-Beltrami operator acting on M_i and $\text{dist} : M \times M \rightarrow \mathbb{R}_{\geq 0}$ denotes the Riemannian distance function. Since the Hamiltonian H depends only on the intrinsic geometry of the Riemannian manifold (via the distance and Laplace operator), the associated quantum theory inherits the invariances of the Riemannian space (M, g) and likewise for the invariances of the interaction graph. Under appropriate assumptions on the potential, these invariances are inherited by the unique positive ground state of H [12]. The geometry of the target space (M, g) has non-trivial implications for the smoothness of the potential energy. In particular, the distance function $\text{dist}(x, \cdot)$, viewed as a function of its second argument, will generically suffer from cusp singularities at x and at any points $y \in M$ admitting multiple minimizing geodesics to x (the so-called cut locus).

The primary goals of this paper are the development of scalable variational approaches to solve both the ground state eigenvalue problem and the time evolution problem corresponding to the Schrödinger operator (1). In this paper, however, we only consider the ground state eigenvalue problem, which can be rephrased as the following unconstrained optimization problem over $\mathcal{H}$,

$$(2) \quad \underset{\psi \in \mathcal{H}}{\text{minimize}} R_H(\psi), \quad R_H(\psi) = \frac{\langle \psi, H\psi \rangle}{\langle \psi, \psi \rangle} .$$

The simplest Riemannian manifold after Euclidean space is the d -dimensional unit sphere S^d . In the interests of simplicity we focus in this initial work on $d = 1$, corresponding to the unit circle S^1 . In terms of the implicit angular parametrization $\theta \in [0, 2\pi)$, the Riemannian distance function on the circle is given by

$$(3) \quad \text{dist}(\theta, \theta') = \min\{|\theta - \theta'|, 2\pi - |\theta - \theta'|\} ,$$

which exhibits the expected cusp singularities at antipodal points defined by $|\theta - \theta'| \in \{0, \pi\}$. In this example the cut locus is a single point. The singularities of the distance function can be smoothed out by a suitable choice of potential, which we take to be of cosinusoidal form. The resulting Hamiltonian is that of the

quantum rotor model formulated on an arbitrary graph,

$$(4) \quad H = -\frac{1}{2} \sum_{i \in \mathcal{V}} h_i \frac{\partial^2}{\partial \theta_i^2} + \sum_{\{i,j\} \in \mathcal{E}} \beta_{ij} [2 - 2 \cos(\theta_i - \theta_j)] .$$

The analysis of manifolds other than the sphere is also clearly of interest. In the Appendix A we explain how the Hamiltonian (1) provides a lattice regularization for quantum non-linear sigma models.

3. ROTOR RESTRICTED BOLTZMANN MACHINE

Inspired by the restricted Boltzmann neural network quantum state introduced in [2], we introduce a class of trial wavefunctions suitable for time evolution and ground state determination for the quantum rotor model (4). Since everything generalizes to rotors of arbitrary dimension, we first consider the S^d target space and subsequently specialize to the circle ($d = 1$). Denote the classical configuration of $n := |\mathcal{V}|$ visible rotors by $x := (\vec{x}_1, \dots, \vec{x}_n) \in (S^d)^n$. Following [2], we assign a probability amplitude to each configuration of rotors $x \in (S^d)^n$ by integrating a Boltzmann factor over a space consisting of m hidden rotors, whose collective coordinates are denoted $z := (\vec{z}_1, \dots, \vec{z}_m) \in (S^d)^m$. In order to construct a suitable Boltzmann factor, we consider the isometric embedding of the target space $S^d \subseteq \mathbb{R}^{d+1}$ into $(d+1)$ -dimensional Euclidean space and choose the exponent to be of restricted Boltzmann form,

$$(5) \quad \psi(x) = \int_{(S^d)^m} d\mu(z) \exp \left[\sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle \vec{z}_i, \vec{x}_j \rangle + \sum_{i=1}^m \langle \vec{b}_i, \vec{z}_i \rangle + \sum_{j=1}^n \langle \vec{c}_j, \vec{x}_j \rangle \right] ,$$

where $d\mu(z)$ denotes the surface measure on $(S^d)^m$ (the counting measure for $d = 0$). In the above expression $\langle \cdot, \cdot \rangle$ denotes the dot product for $\mathbb{R}^{d+1}$ and the variational parameters are given by $a_{ij} \in \mathbb{R}$, $\vec{b}_i \in \mathbb{R}^{d+1}$ and $\vec{c}_j \in \mathbb{R}^{d+1}$ for all $(i, j) \in [m] \times [n]$. If all the bias terms vanish, then the amplitude is invariant under global $O(d+1)$ transformations of the visible rotors, the proof of which follows by a change of integration variables combined with invariance of the integration measure. For $d = 0$ one has $S^0 = \{\pm 1\} \subseteq \mathbb{R}$ and $O(1) = \mathbb{Z}_2$, reproducing the proposal of [2]. The weights and biases can also be promoted to complex numbers, resulting in a holomorphic parametrization suitable for time evolution. In this paper, however, we only consider real parametrizations since the Hamiltonian (4) is known to have a non-negative ground state. The integration over the hidden rotors can be performed for any value of $d \in \mathbb{N}$. In the case of relevance to (4) we obtain the following logarithmic probability amplitude

$$(6) \quad \log \psi(x) = \sum_{j=1}^n \langle \vec{c}_j, \vec{x}_j \rangle + \sum_{i=1}^m \log [2\pi I_0(\|\vec{y}_i\|)] ,$$

where $y = ax + b$ is an affine transformation of the embedded rotor configurations $x \in \mathbb{R}^{n \times (d+1)}$, defined in terms of $a = (a_{ij})$ and $b \in \mathbb{R}^{n \times (d+1)}$, and I_0 denotes a modified Bessel function. The model can be trained using an efficient Markov Chain Monte Carlo method generalizing [2], which is summarized in Appendix B.

4. BENCHMARKS: JASTROW VARIATIONAL WAVEFUNCTION

It is instructive to compare the form of the continuous-variable RBM with that of a $O(d+1)$ -invariant Jastrow wavefunction, which is defined by local interactions dictated by the choice of interaction graph,

$$(7) \quad \log \psi_J(x) = \sum_{\{i,j\} \in \mathcal{E}} w_{ij} \langle \vec{x}_i, \vec{x}_j \rangle ,$$

where w_{ij} denote the variational parameters characterizing the trial wavefunction. Since the number of parameters is dictated by the choice of interaction graph, the Jastrow wavefunction, unlike the RBM, lacks the property of systematic improvability. The Jastrow wavefunction has the advantage that the Rayleigh quotient $R_H(\psi_J)$ can be computed analytically for certain interaction graphs. We carry out this calculation in the case of a linear (path) graph in Appendix C.

5. BENCHMARKS: HIGH-DIMENSIONAL PDE SOLVERS

As a second validation of the ground states obtained from VMC simulations, we compare them against those obtained using partial differential equation (PDE) based eigensolvers. Here, we give a brief overview of our implementation of these PDE solvers. We restrict ourselves to the Hamiltonian (4) associated with the unit circle S^1 with fixed positive weight $h_i = h$ on vertices $\mathcal{V} = \{1, \dots, n\}$. The associated eigenvalue problem is given by

$$(8) \quad H\psi(\boldsymbol{\theta}) = \lambda\psi(\boldsymbol{\theta}) \quad , \quad H = -\frac{h}{2} \sum_{i=1}^n \frac{\partial^2}{\partial \theta_i^2} + \sum_{\{i,j\} \in \mathcal{E}} \beta_{ij} [2 - 2\cos(\theta_i - \theta_j)] \quad ,$$

where we denote $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$, $\theta_i \in [0, 2\pi)$ for $1 \leq i \leq n$.

One straightforward approach to solving (8) is to discretize the domain using a regular grid on $[0, 2\pi)^n$, applying the finite difference approximation to the Schrödinger operator and solving the resulting algebraic eigenvalue problem. This will lead to a solution scheme whose error decays only polynomially depending on the order of the finite difference scheme used to approximate the Hamiltonian [13]. Instead, we can take advantage of the periodicity of the domain and switch to the frequency domain using Fourier series expansion for the wavefunction:

$$(9) \quad \psi(\boldsymbol{\theta}) = \frac{1}{(2\pi)^n} \sum_{\boldsymbol{\omega} \in \mathbb{Z}^n} \hat{\psi}(\boldsymbol{\omega}) \exp(-i\boldsymbol{\omega} \cdot \boldsymbol{\theta}) \quad , \quad \boldsymbol{\omega} \cdot \boldsymbol{\theta} = \sum_{i=1}^n \omega_i \theta_i \quad .$$

As is well known, assuming ψ is smooth, its truncated Fourier approximation will yield spectral accuracy. A numerical solver for (8) can be constructed based on transforming the eigenvalue problem into the Fourier domain; its description is provided in Appendix D, along with details on the eigenvalue algorithms.

One major obstacle to implementing either the finite difference or the Fourier spectral schemes described above is the memory cost: assuming m degrees of freedom per dimension (number of grid points for finite difference or number of Fourier modes for spectral schemes), the discretized wavefunction will require storing m^n scalar variables for a full representation. As the dimensionality of the problem (number of rotors) increases, it quickly becomes impossible to store the eigenfunction in the memory of a single computing node. This is a direct consequence of the *curse of dimensionality*. To alleviate this issue and help scale our benchmark solvers, we adopt the distributed memory computing model, where we split the wavefunction across multiple computing nodes and use the message passing interface (MPI) library to implement necessary communication between the nodes. For our particular form of the Hamiltonian, the discretized operator ends up being sparse (both in the case of finite difference and Fourier spectral schemes), and this allows us to design fast matrix-vector products with minimal communication [14].


6. NUMERICAL RESULTS

We implemented the PDE eigensolvers and VMC algorithm in C++ and using state-of-the-art open source libraries; the PDE eigensolvers were built on top of Trilinos [15] to support distributed computing. The code is available publicly at <https://github.com/shravanvn/cnqs>. The numerical experiments described in this section were run using version 1.0.0 of the code. All simulations were run on the Great Lakes cluster at the University of Michigan. Each compute node of this cluster is equipped with two 18-core 3.0 GHz Intel Xeon Gold 6154 processors and 192 GB RAM.


6.1. Convergence of the Benchmark PDE Solver. We performed self-convergence analysis of our PDE eigensolver on two four-rotor networks depicted in Figure 1 with vertex weights $h_i = 5$ and edge weights $\beta_{ij} = 1$. We ran the Fourier PDE eigensolver for maximum frequencies $1 \leq \omega_{\max} \leq 32$, and Conjugate Gradient (CG) and inverse power iteration tolerances $\tau_{\text{cg}} = \tau_{\text{inv}} = 10^{-15}$.

We plot the state amplitude for both networks corresponding to the $\omega_{\max} = 32$ discretization in Figure 2; as we can see, the amplitudes vanish for approximately $\|\boldsymbol{\omega}\|_{\infty} \geq 10$. This leads to spectral convergence of our Fourier PDE eigensolver; we demonstrate this in Figure 3, where we plot the error in the ground state energy (using $\omega_{\max} = 32$ solution as reference) as we increase the maximum frequency. The error decays exponentially until it reaches machine precision.

In Figure 4 we plot the number of inverse power iterations necessary for convergence; we note that the iteration count is largely independent of the maximum frequency discretization parameter. Additionally,



(A) 2×2 lattice



(B) 4-node complete graph

FIGURE 1. Four-rotor networks used in convergence analysis of the PDE eigensolver.




FIGURE 2. Decay of the ground state amplitudes with increasing frequency corresponding to the four-rotor networks computed from $\omega_{\max} = 32$ discretization of the Fourier PDE eigensolver.




FIGURE 3. Spectral decay of error in the minimum eigenvalue estimate for the four-rotor networks with increasing maximum frequency. The error is estimated using the ground state computed from $\omega_{\max} = 32$ discretization.




FIGURE 4. Number of inverse power iterations required before convergence.




FIGURE 5. Average number of CG iterations per inverse power iteration; shaded region depict the range of the number of CG iterations.




FIGURE 6. Evolution of relative error in the energy estimate (taking the energy computed using the PDE eigensolver as reference) during the stochastic gradient descent in the VMC simulations for four-rotor networks. The solid line represents the relative error when using 250-point rolling average to compute the VMC energy estimate.

from Figure 5 we see that with our preconditioner for the Fourier problem, the number of CG iterations per inverse iteration plateaus very quickly.

6.2. VMC Simulations. We ran VMC simulations with the same four-rotor Hamiltonians. In these simulations, 20 hidden nodes were used to construct the restricted Boltzmann machine and 10000 stochastic gradient descent steps were performed with a learning rate of 10^{-2} . During each of these steps 24000 Metropolis-Hastings samples were generated, the first 4000 of which were discarded as burn-in and then every 20th sample were cherry-picked to compute the expected value quantities (e.g. energy and gradient). A stochastic reconfiguration parameter of 10^{-6} was selected for the simulations. For details on the stochastic reconfiguration algorithm we refer the reader to the appendices of [2].

In Figure 6, we demonstrate the convergence of the energy estimates for the ground state of the corresponding Hamiltonians over the optimization process.

6.3. Comparison of PDE and VMC Solvers. In Table 1, we compare the performances of the PDE and VMC solvers on linear chain networks. For each $2 \leq n \leq 7$, we ran inverse power iteration with tolerances $\tau_{\text{inv}} = 10^{-15}$ and $\tau_{\text{cg}} = 10^{-15}$ for the eigenproblem in Fourier domain with two discretizations: $\omega_{\text{max}} = 5$ and $\omega_{\text{max}} = 7$. This was done to ensure that the PDE eigensolver converged to the correct ground state energy. We then compared these reference eigenvalues against those obtained using the VMC simulations; as we can clearly see from the table, the eigenvalues match to one significant digit. We also note the mild increase in elapsed time in VMC simulations as the number of rotors is increased, compared to the exponential increase in the PDE eigensolver time. This implies that using VMC, we will be able to solve much larger problems than using the PDE eigensolver.

6.4. VMC State Parametrization. The main advantage of the VMC approach over the PDE approach is the ability to parametrize the infinite-dimensional quantum rotor state using a finite-dimensional manifold of parameters. The dimensionality of the manifold is dependent on the number of hidden nodes; as the number of hidden nodes is increased, we are able to capture increasingly complicated quantum states. Note however, this also increases the number of parameters to learn (on top of increasing the computational complexity), and this may lead to poor performance if the model is not trained for long enough. In Table 2, we record the eigenvalue standard deviation and gradient norm of a ten-rotor chain quantum rotor system with various number of hidden RBM nodes during various stages of the training process. We note that after 100 stochastic gradient descent steps, the RBM model with 60 hidden nodes has the smallest standard deviation in the eigenvalue; however as the training continues, all of the models end up with comparable converged eigenvalue.

# Rotors (n)	PDE ($\omega_{\max} = 5$)		PDE ($\omega_{\max} = 7$)		VMC			Jastrow
	$\lambda_{\min}$	T_{elapsed}	$\lambda_{\min}$	T_{elapsed}	avg	std	T_{elapsed}	$R_H(\psi_J)$
2	1.625	0.04	1.625	0.04	1.625	1.2×10^{-3}	364	1.627
3	3.235	0.2	3.235	0.2	3.235	2.1×10^{-2}	721	3.254
4	4.844	1	4.844	4	4.844	4.4×10^{-2}	2029	4.882
5	6.452	20	6.452	116	6.454	6.6×10^{-2}	2667	6.509
6	8.059	392	8.059	3768	8.058	8.9×10^{-2}	3934	8.136
7	9.667	894	9.667	9579	9.669	1.0×10^{-1}	5573	9.763

TABLE 1. Comparison of the PDE and VMC based solvers for the quantum rotor model on a linear graph. The PDE eigensolvers were run with CG and inverse power iteration tolerances $\tau_{\text{cg}} = \tau_{\text{inv}} = 10^{-15}$ and two discretization settings, $\omega_{\max} = 5$ and $\omega_{\max} = 7$, were used to ensure convergence to the smallest eigenvalue. 16 cores were utilized for the $n = 7$ rotor PDE simulations, as such the reported wall-time should be multiplied by a factor of 16 to obtain the actual CPU compute-time. The VMC eigensolvers were run for 10000 stochastic gradient descent steps with learning rate 10^{-2} , $5n$ hidden RBM nodes, $6000n$ Metropolis-Hastings samples per gradient descent step ($1000n$ initial samples were discarded as burn-in and then every $5n$ samples were cherry-picked). All elapsed times are in seconds. For the VMC, avg and std refer to the average and standard deviation of the energy computed in the optimized RBM state. The last column shows the optimal energy $1.62718(n-1)$ of the Jastrow wavefunction (7) assuming uniform weights w_{ij} .

# Hidden nodes	Metric	SGD Step				Average Energy
		100	500	5000	9999	
20	std. dev.	4.4×10^{-1}	1.9×10^{-1}	1.7×10^{-1}	1.7×10^{-1}	1.4489×10^1
	grad. norm	1.1×10^{-2}	7.2×10^{-4}	1.3×10^{-3}	4.4×10^{-3}	
40	std. dev.	4.1×10^{-1}	1.8×10^{-1}	1.7×10^{-1}	1.5×10^{-1}	1.4488×10^1
	grad. norm	1.3×10^{-2}	3.7×10^{-3}	6.5×10^{-3}	3.5×10^{-3}	
60	std. dev.	3.6×10^{-1}	2.0×10^{-1}	1.6×10^{-1}	1.6×10^{-1}	1.4495×10^1
	grad. norm	1.7×10^{-2}	3.9×10^{-3}	5.0×10^{-3}	2.7×10^{-3}	
80	std. dev.	4.1×10^{-1}	1.9×10^{-1}	1.5×10^{-1}	1.5×10^{-1}	1.4493×10^1
	grad. norm	5.9×10^{-3}	2.2×10^{-2}	1.7×10^{-3}	2.4×10^{-3}	
100	std. dev.	4.4×10^{-1}	1.9×10^{-1}	1.5×10^{-1}	1.5×10^{-1}	1.4496×10^1
	grad. norm	5.0×10^{-4}	1.8×10^{-4}	6.6×10^{-4}	1.0×10^{-5}	

TABLE 2. Evolution of the energy standard deviation and gradient norm in the VMC simulations with ten-rotor chain network as we vary the number of hidden units. The VMC eigensolvers were run for 10000 stochastic gradient descent steps with learning rate of 10^{-2} and 120000 Metropolis-Hastings samples per gradient descent step (20000 samples were discarded as burn-in and then every 100th sample were cherry-picked). The last column reports the converged energy for each of these models.

7. CONCLUSION AND FUTURE DIRECTIONS

We introduced continuous-variable neural quantum states as a variational ansatz for finding the ground-states of quantum Hamiltonian operators on continuous manifolds. We demonstrated the ability of these

neural states to converge to the minimal eigenpair of the rotor Hamiltonian by comparing the obtained eigenvalue against those obtained using a baseline PDE based eigensolver. We observed that the scalability of our variational solver increases far slower when compared to the PDE eigensolver as the number of dimensions increase.

The baseline PDE eigensolver introduced in this paper leverages simple techniques from scalable scientific computing algorithms. While the implementation supports distributed computing, allowing us to scale beyond the memory limits of a single node, it does not address the issues related to the curse of dimensionality. Tensor factorization techniques can be used to compress the quantum state of high-dimensional systems to enable memory efficient computing. In [16, 17], the authors use the canonical polyadic (CP) decomposition to find the ground state of a Hamiltonian similar to the one we consider and in [18], the author uses the tensor-train (TT)/matrix-product state (MPS) decomposition on the same problem. [19, 20] use the TT decomposition to find the vibrational spectra and ground states of molecules.

A line of investigation promising improved scalability of the VMC is the exploitation of symmetries of the interaction graph $G = (\mathcal{V}, \mathcal{E})$. Although simple convolutional architectures are likely sufficient for square grid graphs with discrete translational symmetry, a detailed investigation of the interplay between the automorphism group of $G = (\mathcal{V}, \mathcal{E})$ and the isometry group of the target space (M, g) would be desirable.

If the first quantized approach of this paper can ultimately be made to scale, then it should become possible to analyze the quantum phase transitions corresponding to the Berezinskii–Kosterlitz–Thouless (BKT) transition via the quantum classical mapping. In two (Euclidean) dimensions the BKT topological phase transition is a well-known consequence of the proliferation of vortices. Can systematically improvable variational wavefunctions cast light on the nature of the corresponding quantum phase transition? Similar techniques adapted to variational real-time dynamics could potentially offer a window into many-body quantum chaos via disordered rotor models [21] or the models with negative curvature [22].

In closing, the approximation properties of continuous-variable neural-network quantum states is poorly understood. Considerable effort has been expended in the search for exact representations of many-body quantum states using restricted Boltzmann machines [23, 24, 25, 26, 27, 28] and it would be very interesting if similar techniques can be adapted to continuous-variable lattice systems.

8. ACKNOWLEDGEMENTS

J.S. thanks Di Luo and Tobias Osborne for discussion and encouragement. This research was supported in part through computational resources and services provided by the Advanced Research Computing (ARC) at the University of Michigan.

APPENDIX A. NONLINEAR SIGMA MODEL REGULARIZATION

A nonlinear sigma model is a theory of maps from a spacetime manifold to the target Riemannian manifold (M, g) . In this appendix we focus on Minkowski space $\mathbb{R}^{n,1}$ which is topologically the product $\mathbb{R}^n \times \mathbb{R}$. The following discussion can be easily adapted to replace $\mathbb{R}^n$ with a manifold of finite spatial volume such as the n -torus. The degrees of freedom of the theory consist of target-space-valued functions of spacetime, which we denote by $x^a = x^a(\sigma, \tau)$, where $a = 1, \dots, \dim M$ indexes the implicit coordinates for the target manifold M . Classical trajectories of the sigma model satisfy a nonlinear wave equation obtained by extremizing the following functional,

$$(10) \quad S[x] = \frac{1}{2f} \int_{\mathbb{R}} d\tau \int_{\mathbb{R}^n} d^n\sigma g_{ab}(x) \left[\dot{x}^a \dot{x}^b - (\vec{\nabla} x^a) \cdot (\vec{\nabla} x^b) \right] ,$$

where $f > 0$ is an overall normalization. Define the momentum density as the following functional derivative,

$$(11) \quad \pi_a := \frac{\delta S}{\delta \dot{x}^a} = \frac{1}{f} g_{ab}(x) \dot{x}^b ,$$

in terms of which the energy density $\mathcal{H} := \pi_a \dot{x}^a - \mathcal{L}$ is given by

$$(12) \quad \mathcal{H} = \frac{f}{2} g^{ab} \pi_a \pi_b + \frac{1}{2f} g_{ab} (\vec{\nabla} x^a) \cdot (\vec{\nabla} x^b) ,$$

where $\mathcal{L}$ is the Lagrangian density. In the simplest Euler discretization, the total energy $E(\tau) = \int d^n\sigma \mathcal{H}$ at time τ of a map $x^a = x^a(\sigma, \tau)$ can be obtained from the limiting procedure $E(\tau) = \lim_{a \rightarrow 0} E_a(\tau)$,

$$(13) \quad E_a(\tau) = \sum_{\sigma \in a\mathbb{Z}^n} \left[\frac{f}{2a^n} g^{ab}(x(\sigma, \tau)) p_a(\sigma, \tau) p_b(\sigma, \tau) + \frac{a^{n-2}}{2f} \sum_{\mu=1}^n g_{ab}(x(\sigma, \tau)) \delta_\mu x^a(\sigma, \tau) \delta_\mu x^b(\sigma, \tau) \right] ,$$

where $\delta_\mu x^a(\sigma, \tau) := x^a(\sigma + a \hat{e}_\mu, \tau) - x^a(\sigma, \tau)$ is the finite difference operator in the direction of the unit basis vector $\hat{e}_\mu \in \mathbb{R}^n$, and $p_a(\sigma, \tau) := a^d \pi_a(\sigma, \tau)$ is the momentum at vertex $\sigma \in a\mathbb{Z}^n$. The double summation runs over the edges of the cubic lattice graph. Observe that the potential $g_{ab} \delta_\mu x^a \delta_\mu x^b$ corresponding to each edge of the graph is the quadratic approximation to the squared Riemannian distance $\text{dist}^2(x(\sigma), x(\sigma + a \hat{e}_\mu))$, which provides an accurate approximation when the involved distances are much less than the radius of curvature of the target space (M, g) . The form of the energy (13) motivates a quantum Hamiltonian in which the quadratic potential is replaced by its nonlinear completion. It follows from the requirement of self-adjointness of the conjugate momentum operator, that the kinetic term of the quantum Hamiltonian is obtained by replacing $g^{ab} p_a p_b$ with minus the Laplace-Beltrami operator defined as follows

$$(14) \quad \Delta = \frac{1}{\sqrt{g}} \frac{\partial}{\partial x^a} \left(\sqrt{g} g^{ab} \frac{\partial}{\partial x^b} \right) ,$$

where $g = \det g_{ab}$. The resulting Hamiltonian is of the form (1) with $V(r) = r^2$ and with uniform vertex and edge weights given by $h_i = f/a^n$ and $\beta_{ij} = a^{n-2}/(2f)$, respectively.

APPENDIX B. TRAINING THE VMC

The components of the variational derivatives are given by

$$(15) \quad \frac{\partial \log \psi(x)}{\partial \vec{c}_j} = \vec{x}_j ,$$

$$(16) \quad \frac{\partial \log \psi(x)}{\partial \vec{b}_i} = g(r_i) \frac{\vec{y}_i}{r_i} ,$$

$$(17) \quad \frac{\partial \log \psi(x)}{\partial w_{ij}} = g(r_i) \frac{\langle \vec{y}_i, \vec{x}_j \rangle}{r_i} ,$$

where $g(x) := I_1(x)/I_0(x)$ and $r_i := \|y_i\|$. Each time a Metropolis update of the j th visible rotor occurs of the form $\vec{x}_j \leftarrow \vec{x}'_j$, the variable r_i is updated for all $i \in [m]$ according to the following rule:

$$(18) \quad \vec{y}_i \leftarrow y_i + w_{ij}(\vec{x}'_j - \vec{x}_j) ,$$

$$(19) \quad r_i \leftarrow \|\vec{y}_i\| .$$

In our quantum rotor setup, the vectors $\vec{x}_i \in S^1$ are parametrized by $\theta_i \in [-\pi, \pi)$ as $\vec{x}_i = (\cos \theta_i, \sin \theta_i)$. The Metropolis update is then simply updating this angle parameter

$$(20) \quad \theta'_i = \theta_i + \delta_i, \quad \delta_i \sim \text{Uniform}(-a, a)$$

while shifting θ'_i by multiples of 2π to keep it in the range $[-\pi, \pi)$. Here $a > 0$ is another hyperparameter for the training.

APPENDIX C. EXACT ENERGY FOR JASTROW WAVEFUNCTION

In this appendix we analytically compute the Rayleigh quotient $R_H(\psi_J)$ for the quantum rotor Hamiltonian on a linear graph,

$$(21) \quad H = -\frac{\hbar}{2} \sum_{i=1}^n \frac{\partial^2}{\partial \theta_i^2} + \sum_{i=1}^{n-1} \beta_i [2 - 2 \cos(\theta_i - \theta_{i+1})] ,$$

using a generalized Jastrow trial wavefunction of the form

$$(22) \quad \psi_J(\theta_1, \dots, \theta_n) = \frac{1}{\sqrt{2\pi}} \prod_{i=1}^{n-1} \varphi_i(\theta_i - \theta_{i+1}) ,$$

where for each $1 \leq i \leq n-1$, the function $\varphi_i : \mathbb{R} \rightarrow \mathbb{R}$ is even, 2π -periodic and satisfies the following normalization condition,

$$(23) \quad \int_0^{2\pi} d\theta \varphi_i(\theta)^2 = 1 ,$$

which ensures overall normalization

$$(24) \quad \int_{[0, 2\pi]^n} d\theta_1 \cdots d\theta_n \psi_J(\theta_1, \dots, \theta_n)^2 = 1 .$$

If we define constant functions $\varphi_0(\theta) = \varphi_n(\theta) = 0$, then

$$(25) \quad \left\langle \psi_J, \frac{\partial^2 \psi_J}{\partial \theta_i^2} \right\rangle := \int_{[0, 2\pi]^n} d\theta_1 \cdots d\theta_n \psi_J(\theta_1, \dots, \theta_n) \frac{\partial^2}{\partial \theta_i^2} \psi_J(\theta_1, \dots, \theta_n) ,$$

$$(26) \quad = \int_0^{2\pi} d\theta [\varphi_{i-1}(\theta) \varphi_{i-1}''(\theta) + \varphi_i(\theta) \varphi_i''(\theta)] - 2 \left[\int_0^{2\pi} d\theta \varphi_{i-1}(\theta) \varphi_{i-1}'(\theta) \right] \left[\int_0^{2\pi} d\bar{\theta} \varphi_i(\bar{\theta}) \varphi_i'(\bar{\theta}) \right] ,$$

$$(27) \quad = \int_0^{2\pi} d\theta [\varphi_{i-1}(\theta) \varphi_{i-1}''(\theta) + \varphi_i(\theta) \varphi_i''(\theta)] ,$$

where in the last equality we used the fact that φ_i' is odd since φ_i is even. Similarly,

$$(28) \quad \langle \psi_J, \cos(\theta_i - \theta_{i+1}) \psi_J \rangle = \int_0^{2\pi} \varphi_i(\theta) \cos \theta ,$$

and thus the total energy in the state ψ_J is

$$(29) \quad \langle \psi_J, H \psi_J \rangle = \sum_{i=1}^{n-1} \left[2\beta_i - \int_0^{2\pi} d\theta \varphi_i(\theta) (\hbar \varphi_i''(\theta) + 2\beta_i \cos \theta) \right] .$$

In the particular case of Eq. (7) of the main text we have,

$$(30) \quad \varphi_i(\theta) = \frac{\exp(w_i \cos \theta)}{\sqrt{2\pi I_0(2w_i)}} ,$$

and the energy is found to be

$$(31) \quad \langle \psi_J, H \psi_J \rangle = \sum_{i=1}^{n-1} \left[2\beta_i + \frac{I_1(2w_i)}{I_0(2w_i)} \left(\frac{\hbar}{2} w_i - 2\beta_i \right) \right] .$$

APPENDIX D. PRECONDITIONED FOURIER EIGENSOLVER

Here, we provide details on the Fourier-based spectral method for solving (8). Substituting the Fourier expansion (9) in the eigenvalue problem (8), multiplying both sides by $\exp(i\boldsymbol{\omega} \cdot \boldsymbol{\theta})$ and integrating over the domain $[0, 2\pi)^n$, we obtain

$$(32) \quad \frac{h}{2} \|\boldsymbol{\omega}\|^2 \hat{\psi}(\boldsymbol{\omega}) + \sum_{\{i,j\} \in \mathcal{E}} \beta_{ij} [2\hat{\psi}(\boldsymbol{\omega}) - \hat{\psi}(\boldsymbol{\omega} + \mathbf{e}_i - \mathbf{e}_j) - \hat{\psi}(\boldsymbol{\omega} - \mathbf{e}_i + \mathbf{e}_j)] = \lambda \hat{\psi}(\boldsymbol{\omega}) , \quad \boldsymbol{\omega} \in \mathbb{Z}^n ,$$

where $\mathbf{e}_i \in \mathbb{R}^n$ is the vector of all zeros except at the i -th entry, which is 1. Clearly, (32) is an eigenvalue equation for the Fourier coefficients

$$(33) \quad \hat{H} \hat{\psi}(\boldsymbol{\omega}) = \lambda \hat{\psi}(\boldsymbol{\omega}) ,$$

where the operator $\hat{H}$ defined as

$$(34) \quad \hat{H} \hat{\psi}(\boldsymbol{\omega}) := \frac{h}{2} \|\boldsymbol{\omega}\|^2 \hat{\psi}(\boldsymbol{\omega}) + \sum_{\{i,j\} \in \mathcal{E}} \beta_{ij} [2\hat{\psi}(\boldsymbol{\omega}) - \hat{\psi}(\boldsymbol{\omega} + \mathbf{e}_i - \mathbf{e}_j) - \hat{\psi}(\boldsymbol{\omega} - \mathbf{e}_i + \mathbf{e}_j)] .$$

Note that given $\psi \in L^2([0, 2\pi)^n)$ with periodic weak derivatives, the Fourier coefficients $\hat{\psi} \in L^2(\mathbb{Z}^d)$ satisfy

$$(35) \quad \sum_{\boldsymbol{\omega} \in \mathbb{Z}^d} \|\boldsymbol{\omega}\|^2 |\hat{\psi}(\boldsymbol{\omega})| < \infty .$$

This condition (35) suggests that the Fourier coefficients decay rapidly as infinity norm of the frequency vector $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)$ increases in size. Thus, setting the Fourier coefficients to zero outside a hypercube in the frequency space

$$(36) \quad \hat{\psi}(\boldsymbol{\omega}) = 0 \quad \text{if} \quad \|\boldsymbol{\omega}\|_\infty := \max\{|\omega_1|, \dots, |\omega_n|\} \geq \omega_{\max} ,$$

provides a sufficiently accurate approximation to the full wavefunction as long as the cut-off frequency $\omega_{\max}$ is not too small. This assumption also allows us to represent the wavefunction as a n -dimensional $(2\omega_{\max} + 1) \times \dots \times (2\omega_{\max} + 1)$ tensor $\hat{\psi}$ for computational purposes. Next, restriction $\hat{H}$ of the Hamiltonian operator $\hat{H}$ to the frequency range $[-\omega_{\max}, \omega_{\max}]^n$ is sparse, with at most $2|\mathcal{E}| + 1$ non-zero entries per row. This structure leads to efficient matrix-vector operations and minimal inter-node communications in distributed computing setups.

Following this restriction of the eigenstate and the Hamiltonian in the Fourier domain, we need to find the minimal eigenvalue-eigenvector pair corresponding to the linear system

$$(37) \quad \hat{H} \hat{\psi} = \lambda \hat{\psi} .$$

This is equivalent to finding the maximal eigenvalue-eigenvector pair of the system

$$(38) \quad \hat{H}' \hat{\psi} = \lambda' \hat{\psi} , \quad \hat{H}' = (\hat{H} - \mu I)^{-1} , \quad \lambda' = (\lambda - \mu)^{-1} ,$$

where μ is a lower bound on the eigenvalues of $\hat{H}$, implying $\hat{H} - \mu I$ is invertible. Such a lower bound is derived in Appendix E.

Power iteration provides a simple method for computing the maximal eigenpair of (38): starting from an arbitrary initial state $\hat{\psi}_0$, we iterate

$$(39) \quad \hat{\psi}_{k+1} = \frac{\hat{H}' \hat{\psi}_k}{\|\hat{H}' \hat{\psi}_k\|} , \quad \lambda_{k+1} = \frac{\hat{\psi}_{k+1} \cdot \hat{H}' \hat{\psi}_{k+1}}{\hat{\psi}_{k+1} \cdot \hat{\psi}_{k+1}} , \quad k \geq 0 .$$

As long as the minimal eigenpair of $\hat{H}$ is *non*-degenerate, and the initial state $\hat{\psi}_0$ is not orthogonal to the minimal eigenstate, this (inverse) power iteration is guaranteed to converge to the correct answer [29].

Finally, note that computing $\hat{\phi}_{k+1} = \hat{H}' \hat{\psi}_k$ in the iteration requires solving a linear system

$$(40) \quad (\hat{H} - \mu I) \hat{\phi}_{k+1} = \hat{\psi}_k .$$

This can be done by any iterative solvers such as conjugate gradient (CG) or generalized minimal residual (GMRES) methods. They only require a routine for computing $(\hat{H} - \mu I) \hat{\psi}$ for arbitrary vector $\hat{\psi}$ [29]. We have already pointed out how the sparsity structure of $\hat{H}$ can be used to compute the matrix-vector product efficiently.

In our implementation, we choose CG as the iterative linear solver. To accelerate the convergence, we utilize the diagonal preconditioner

$$(41) \quad \hat{M}\hat{\psi}(\omega) = \left(\frac{h}{2}\|\omega\|^2 + \sum_{\{i,j\} \in \mathcal{E}} 2\beta_{ij} - \mu \right)^{\frac{1}{2}} \hat{\psi}(\omega) ,$$

and solve the following linear system

$$(42) \quad \left[\hat{M}^{-1}(\hat{H} - \mu I) \hat{M}^{-1} \right] \left[\hat{M}\hat{\phi}_{k+1} \right] = \left[\hat{M}^{-1}\hat{\psi}_k \right] .$$

Algorithm 1 describes the full inverse power iteration process for determining the ground state of the quantum rotor Hamiltonian.

Algorithm 1 Inverse power iteration for ground state determination

Require: Laplacian prefactor h , edge weights $\{\beta_{ij} : (i, j) \in \mathcal{E}\}$, frequency cutoff $\omega_{\max}$, tolerances $\tau_{\text{cg}}, \tau_{\text{inv}}$

Ensure: Ground state $\hat{\psi}$ and energy λ

- 1: Construct Hamiltonian $\hat{H}$ using (34) and $\omega_{\max}$
 - 2: Construct preconditioner $\hat{M}$ using (41) and $\omega_{\max}$
 - 3: Determine eigenvalue lower bound μ from (43)
 - 4: Initialize $\psi_0 \in \mathbb{R}^{(2\omega_{\max}+1) \times \dots \times (2\omega_{\max}+1)}$ randomly and normalize
 - 5: Initialize $\lambda_0 \leftarrow \hat{\psi}_0 \cdot \hat{H}\hat{\psi}_0$
 - 6: Initialize $k \leftarrow 0$
 - 7: **repeat**
 - 8: Solve $(\hat{H} - \mu I)\hat{\phi}_{k+1} = \hat{\psi}_k$ using CG iteration with preconditioner $\hat{M}$ and tolerance τ_{cg}
 - 9: Design new state $\hat{\psi}_{k+1} \leftarrow \hat{\phi}_{k+1} / \|\hat{\phi}_{k+1}\|_2$
 - 10: Determine eigenvalue $\lambda_{k+1} \leftarrow \hat{\psi}_{k+1} \cdot \hat{\phi}_{k+1}$
 - 11: $k \leftarrow k + 1$
 - 12: **until** $|\lambda_k - \lambda_{k-1}| < \tau_{\text{inv}}$
 - 13: $\hat{\psi} \leftarrow \hat{\psi}_k$
 - 14: $\lambda \leftarrow \lambda_k$
-

APPENDIX E. LOWER BOUND ON MINIMAL EIGENVALUE

Here, we derive a simple lower bound on the magnitude of the eigenvalues of the rotor Hamiltonian H , and equivalently, the operator $\hat{H}$ in the Fourier basis:

Theorem E.1. *All eigenvalues λ of the Hamiltonian (8) satisfy $\lambda \geq \mu$ with*

$$(43) \quad \mu = -4 \sum_{\{i,j\} \in \mathcal{E}} |\beta_{ij}| .$$

Proof. Let us denote

$$(44) \quad h_{ij}(\theta) = \beta_{ij}(2 - 2\cos(\theta_i - \theta_j)) \implies |h_{ij}(\theta)| \leq 4|\beta_{ij}|$$

so that we can write

$$(45) \quad H = -\frac{h}{2}\Delta + \sum_{\{i,j\} \in \mathcal{E}} h_{ij}(\theta)$$

It follows that

$$(46) \quad \langle \psi, H\psi \rangle = \frac{h}{2} \langle \psi, -\Delta\psi \rangle + \sum_{\{i,j\} \in \mathcal{E}} \langle \psi, h_{ij}\psi \rangle$$

and we can compute

$$(47) \quad \langle \psi, -\Delta\psi \rangle = - \int d\theta \psi^*(\theta) \Delta\psi(\theta) = \int d\theta \nabla\psi^*(\theta) \cdot \nabla\psi(\theta) = \int d\theta \|\nabla\psi(\theta)\|^2 \geq 0$$

where the integration is performed over the domain $[0, 2\pi)^n$ and the second equality follows from integration by parts (boundary terms vanish due to the periodic boundary conditions). Further

$$(48) \quad |\langle \psi, h_{ij} \psi \rangle| = \left| \int d\theta \psi^*(\theta) h_{ij}(\theta) \psi(\theta) \right| \leq \int d\theta |h_{ij}(\theta)| |\psi(\theta)|^2 \leq 4|\beta_{ij}| \int d\theta |\psi(\theta)|^2 = 4|\beta_{ij}| \langle \psi, \psi \rangle$$

It follows that

$$(49) \quad \langle \psi, H\psi \rangle \geq \frac{h}{2} \cdot 0 - 4 \sum_{\{i,j\} \in \mathcal{E}} |\beta_{ij}| \langle \psi, \psi \rangle \implies R_H(\psi) = \frac{\langle \psi, H\psi \rangle}{\langle \psi, \psi \rangle} \geq -4 \sum_{\{i,j\} \in \mathcal{E}} |\beta_{ij}|$$

Clearly,

$$(50) \quad \mu = -4 \sum_{\{i,j\} \in \mathcal{E}} |\beta_{ij}|$$

serves as a lower bound for the eigenvalues of the Hamiltonian H . $\square$

DECLARATIONS

Authors' contributions

S.D. and J.S. implemented the algorithms described here. J.S. wrote the formulation and S.D. compiled the results. All authors helped write and review the manuscript.

Funding

We acknowledge support from NSF under grant DMS-2038030.

Availability of data and materials

Not applicable.

REFERENCES

- [1] William Lauchlin McMillan. Ground state of liquid he 4. *Physical Review*, 138(2A):A442, 1965.
- [2] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- [3] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [4] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Velićković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- [5] Igor Boettcher, Przemyslaw Bienias, Ron Belyansky, Alicia J Kollár, and Alexey V Gorshkov. Quantum simulation of hyperbolic space with circuit quantum electrodynamics: From graphs to geometry. *Physical Review A*, 102(3):032208, 2020.
- [6] Xingchuan Zhu, Jiaojiao Guo, Nikolas P Breuckmann, Huaiming Guo, and Shiping Feng. Quantum phase transitions of interacting bosons on hyperbolic lattices. *arXiv preprint arXiv:2103.15274*, 2021.
- [7] Alicia J Kollár, Mattias Fitzpatrick, and Andrew A Houck. Hyperbolic lattices in circuit quantum electrodynamics. *Nature*, 571(7763):45–50, 2019.
- [8] S Iblisdir, R Orus, and JI Latorre. Matrix product states algorithms and continuous systems. *Physical Review B*, 75(10):104305, 2007.
- [9] Martin C Gutzwiller. The geometry of quantum chaos. *Physica Scripta*, 1985(T9):184, 1985.
- [10] Steven R White. Density matrix formulation for quantum renormalization groups. *Physical review letters*, 69(19):2863, 1992.
- [11] Steven R White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48(14):10345, 1993.
- [12] James Glimm and Arthur Jaffe. *Quantum physics: a functional integral point of view*. Springer Science & Business Media, 2012.
- [13] Randall J LeVeque. *Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems*. SIAM, 2007.
- [14] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [15] Michael A Heroux, Roscoe A Bartlett, Vicki E Howle, Robert J Hoekstra, Jonathan J Hu, Tamara G Kolda, Richard B Lehoucq, Kevin R Long, Roger P Pawlowski, Eric T Phipps, et al. An overview of the trilinos project. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):397–423, 2005.
- [16] Gregory Beylkin and Martin J Mohlenkamp. Numerical operator calculus in higher dimensions. *Proceedings of the National Academy of Sciences*, 99(16):10246–10251, 2002.
- [17] Gregory Beylkin and Martin J Mohlenkamp. Algorithms for numerical analysis in high dimensions. *SIAM Journal on Scientific Computing*, 26(6):2133–2159, 2005.
- [18] Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

- [19] Maxim Rakhuba and Ivan Oseledets. Calculating vibrational spectra of molecules using tensor train decomposition. *The Journal of chemical physics*, 145(12):124101, 2016.
- [20] Alexander Veit and L Ridgway Scott. Using the tensor-train approach to solve the ground-state eigenproblem for hydrogen molecules. *SIAM Journal on Scientific Computing*, 39(1):B190–B220, 2017.
- [21] Gong Cheng and Brian Swingle. Chaos in a quantum rotor model. *arXiv preprint arXiv:1901.10446*, 2019.
- [22] Steven Gubser, Zain H Saleem, Samuel S Schoenholz, Bogdan Stoica, and James Stokes. Nonlinear sigma models with compact hyperbolic target spaces. *Journal of High Energy Physics*, 2016(6):1–15, 2016.
- [23] Yichen Huang and Joel E Moore. Neural network representation of tensor network and chiral states. *arXiv preprint arXiv:1701.06246*, 2017.
- [24] Dong-Ling Deng, Xiaopeng Li, and S Das Sarma. Machine learning topological states. *Physical Review B*, 96(19):195145, 2017.
- [25] Giuseppe Carleo, Yusuke Nomura, and Masatoshi Imada. Constructing exact representations of quantum many-body systems with deep neural networks. *Nature communications*, 9(1):1–11, 2018.
- [26] Jing Chen, Song Cheng, Haidong Xie, Lei Wang, and Tao Xiang. Equivalence of restricted boltzmann machines and tensor network states. *Physical Review B*, 97(8):085104, 2018.
- [27] Ermal Rrapaj and Alessandro Roggero. Exact representations of many-body interactions with restricted-boltzmann-machine neural networks. *Physical Review E*, 103(1):013302, 2021.
- [28] Michael Y Pei and Stephen R Clark. Compact neural-network quantum state representations of jastrow and stabilizer states. *arXiv preprint arXiv:2103.09146*, 2021.
- [29] Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. SIAM, 1997.