Approximately-invariant neural networks for quantum many-body physics

Dominik Kufel*, **Jack Kemp**, **Norman Y. Yao**
Department of Physics,
Harvard University

Abstract

We propose *approximately* group-invariant neural networks for quantum many-body physics problems. These tailored-made architectures are parameter-efficient, scalable, significantly outperform existing symmetry-unaware neural network architectures and are competitive with the state-of-the-art iPEPS methods as we demonstrate on a perturbed toric code toy model on a 10×10 lattice. This paves way towards studying traditionally challenging quantum spin liquid problems within interpretable neural network architectures.

1 Introduction

We focus on the problem of finding the lowest-eigenvalue eigenvector ("ground state") of a Hamiltonian matrix. This problem is paramount in the many-body physics community where the ground state governs the behaviour of quantum mechanical systems at low temperature. The challenge in finding the ground state eigenvector is that it lies in a vector (Hilbert) space which is exponentially large in the size of the system. This makes exact diagonalization of the Hamiltonian matrix infeasible even for small system sizes. To deal with this exponential complexity, approximate methods such as e.g., tensor networks [1] or Quantum Monte Carlo [2] are traditionally used.

More recently, methods using neural networks as variational ansatze [3] have been introduced. Ground states are found by minimizing an expectation value of the Hamiltonian matrix (energy) within a manifold of the Hilbert space parameterized by the neural network. These neural-network-based variational methods (known as neural quantum states or NQS) have sparked interest in the many-body physics community due to their expressibility; in particular, the expressibility of NQS is asymptotically ensured by a universal approximation theorem [4] and is strictly more powerful than efficiently contractible tensor networks [5]. Furthermore, NQS have recently become state-of-the-art for computing the ground state in certain archetypal models such as the 2D transverse field Ising model [6] or the square-lattice Rydberg blockade model [7].

However, there are still many physical systems, whose associated Hamiltonians, remain challenging to tackle using NQS. In the context of spin models, a particularly interesting subclass exhibiting long-range entanglement, is that of quantum spin liquid Hamiltonians [8]. Although some progress has been made in this direction (for example, on the J1-J2 Heisenberg model [9, 10]), in general, finding and characterizing spin-liquid ground states remains an extremely difficult problem owing to the inherently complicated optimization landscape [11]. A promising approach to simplify these optimization landscapes without sacrificing the expressibility of the network is to exploit our physical understanding of the group of *emergent* symmetries G associated with a particular spin-liquid. By imposing these symmetries on the neural network via group equivariant methods [12, 9], one can significantly reduce the number of optimization parameters. Indeed, for lattice translation or point group symmetries (e.g. rotation), imposing such symmetries on the NQS has been crucial for finding accurate ground states and avoiding local minima [13, 11].

*dkufel@g.harvard.edu

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

Symmetry is especially important for models expected to describe quantum spin liquids (such as $\mathbb{Z}_2$ lattice gauge theory), because such ground states exhibit an exponentially large "gauge" group of emergent symmetries. Group-equivariant neural networks [14] thus yield significant improvements over naive methods such as restricted Boltzmann machines or multi-layered perceptrons for lattice gauge theories. Unfortunately, for more general spin liquid problems, these ground state symmetries are only known at specific, exactly soluble points in the space of Hamiltonians. Away from these special regions, the exactly-known symmetry operations become only approximate. It is thus interesting to ask: *is it possible to improve NQS methods for quantum spin liquids by imposing approximate symmetries as an inductive bias?*

In this work, we answer this question in the affirmative. We show that approximately-invariant networks impose a soft inductive bias on the ground-state search while maintaining the flexibility to capture states that are not exactly symmetric. To impose approximate symmetries on NQS, we leverage results from the field of approximately group-equivariant networks [15, 16]. We modify these constructions for many-body physics problems, incorporating physical insight into the nature of the ground state (Sec. 2). In Sec. 3, we demonstrate the accuracy of our approach on a paradigmatic quantum spin liquid model: the $\mathbb{Z}_2$ toric code perturbed with longitudinal and transverse fields. In particular we demonstrate that the variational energies obtained: (i) outperform conventional NQS methods; (ii) converge to exact diagonalization results for small system sizes; and (iii) match state-of-the-art tensor network results for large system sizes. Finally, we discuss how the approximate-symmetries framework can pave the road towards NQS interpretability.

2 Methods

2.1 Neural quantum states basics

We first present an overview of a general NQS framework. Consider a many-body system on N-qubits with associated Hamiltonian H. The many-body quantum state vector $|\psi\rangle$ may be decomposed into a complete basis of binary vectors $|s\rangle$ (e.g., $|s\rangle = |-1, 1, 1, -1, \cdots\rangle$ with 2^N elements: $|\psi\rangle = \sum_s \psi_s |s\rangle$ where $\psi_s \in \mathbb{C}$ is a complex amplitude. The key idea of the NQS is to represent ψ_s as a neural network with a complex scalar output for a particular bit string input s. We would like to find the lowest-eigenvalue eigenvector $|\psi\rangle$ of the Hamiltonian operator H (represented as a $2^N \times 2^N$ matrix in the basis $\{|s\rangle\}$). We approach the problem variationally by turning the original problem into an energy minimization problem with respect to the vectorized parameters W of the neural network, $\min_W \langle H \rangle = \min_W \frac{\langle \psi(W)|H|\psi(W)\rangle}{\langle \psi(W)|\psi(W)\rangle}$. The energy $\langle H \rangle$ is found by expanding the formula in the $\{|s\rangle\}$ basis, and evaluating the summation through Monte Carlo Markov chain sampling [3]. The network parameters W are optimized either via gradient descent or more complicated second order methods, such as stochastic reconfiguration [17] (closely connected to natural gradient [18, 19]).

2.2 Group equivariant framework

In this subsection we describe the group-equivariant neural network setup for many-body physics problems. Suppose a ground state $|\psi\rangle$ of the Hamiltonian exhibits a particular group of symmetries G i.e. $g|\psi\rangle = |\psi\rangle \; \forall g \in G$. We will assume that the basis $\{|s\rangle\}$ is chosen such that the group G acts as a permutation on bit-string basis elements. In such a basis, the action of G can be written in terms of the familiar condition of group-invariance for a neural network, $\psi_{gs} = \psi_s$; that is, two inputs of the network connected by a symmetry element should give the same scalar output.

The usual way of imposing group-invariance on a neural network is to ensure group-equivariance of each of its layers, $\Xi : V \to W$, and non-linearities. In the final layer, the scalar output of the network imposes group-invariance , i.e. $\Xi(\rho_{in}(g)x) = \rho_{out}(g)\Xi(x) \;\; \forall x \in V, g \in G$ where V is an N-dimensional input vector space, W is an n-dimensional output space, and $\rho_{in} : G \to GL(N, \mathbb{C})$ and $\rho_{out} : G \to GL(n, \mathbb{C})$ are the input and output representations of the symmetry group G.

Here we present a general approach for imposing group equivariance within quantum many-body physics problems. Following [20], we construct group-equivariant / group-invariant layers of the network by using equivariant multi-layered perceptrons (EMLP). In special cases, this approach reduces to other group-equivariant frameworks such as G-convolutional, G-steerable or deep set architectures. We consider linear layers with dimensions $O(N)$, which are constructed by appropriately restricting weights of the otherwise fully-connected layer. This is feasible since the cost of imposing

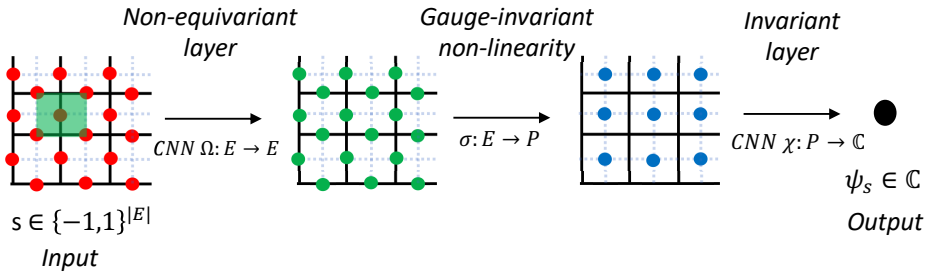

Figure 1: The approximately group-invariant NQS architecture for a perturbed toric code model. Output of the network is $\psi_s = \chi(\sigma(\Omega(s)))$. Gauge-invariant non-linearity is explicitly defined as $\sigma_p(x) = \prod_{j\in p} x_j$. Non-equivariant first layer breaks exact group-invariance.

equivariance couples only to the size of the generating set of the group [20]—at worst $O(poly(N))$ for the groups we consider. Due to the non-regular hidden-layer representations, one needs to ensure that the activation functions are also equivariant. Instead of traditionally used gated/norm [21] non-linearities, we utilize other gauge-equivariant non-linearities derived from a physical model in mind (an example for a $\mathbb{Z}_2$ lattice gauge theory will be described explicitly below). Finally, in the last layer one applies a gauge-invariant non-linearity constructed in a similar fashion.

2.3 Approximately-invariant networks

Having established how group-equivariance may be imposed in the many-body physics, we now to turn to constructing *approximately* group-equivariant NQS. In particular, we ensure that $|\psi_{gs} - \psi_s| < \epsilon \; \forall g \in G$. Starting from a fully invariant neural network, $\epsilon = 0$, one can lift the strict constraints by adding an extra non-equivariant layer to the network or by adding a non-equivariant skip-connection to the network. These methods are in the spirit of the previously described "combo" architecture of [16] and "residual pathway priors" of [15], respectively. In principle, for sufficiently large invariance breaking, ϵ, a neural network constructed in such a fashion can target any vector in the Hilbert space. In practice, ϵ can be controlled by the size of the non-equivariant layers and how they are initialized.

3 Results

We demonstrate the usefulness of this approximately symmetric approach for a paradigmatic quantum spin liquid problem: the $\mathbb{Z}_2$ toric code model under a magnetic field in the x-z plane.

We define a $\mathbb{Z}_2$ toric code model with $N = (2L^2 - 2L)$ qubits on the set of edges E of an $L \times L$ square lattice (Fig. 1) with open-boundary conditions. The Hamiltonian of the model is $H = -\sum_v A_v - \sum_p B_p - h_x \sum_i X_i - h_z \sum_i Z_i$ where vertex operators $A_v = \prod_{j\in v} X_j$ act on qubits surrounding a particular lattice vertex $v \in V$, plaquette operators $B_p = \prod_{j\in p} Z_j$ act on qubits around each fundamental plaquette $p \in P$ (square) of the lattice, h_x and h_z are strengths of the magnetic fields in x and z direction respectively, and $X_j = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $Z_j = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ are Pauli matrices. To understand the symmetry of the quantum spin liquid ground state we hope to exploit, notice that $[B_p, A_v] = 0 \; \forall v \in V, p \in P$. It is well-understood [22, 23] that the phase diagram of the perturbed toric code hosts a quantum spin liquid up to finite values of h_x and h_z. Along the $h_z = 0$ line, the generators of the emergent $\mathbb{Z}_2$ gauge symmetry are precisely given by the A_v operators, i.e. $|\psi\rangle$ fulfills $A_v|\psi\rangle = |\psi\rangle \; \forall v \in V$. These operators form a group, $G = \mathbb{Z}_2^{\times N/2}$, which physically corresponds to creating so-called "Wilson loops". For any $|h_z| > 0$, the precise form of the emergent gauge symmetry is unknown. However, since the system is "topologically ordered", many of its interesting properties (such as long-range entanglement and topological entanglement entropy) should extend to a finite portion of the phase diagram, even for $|h_z| > 0$ [24]. This means that by applying a finite-depth unitary operation, one can transform the ground state back to the $h_z = 0$ toric code with exact gauge symmetries.

As discussed in Sec. 2 we construct approximately-invariant neural networks based on the combo architecture [16]. In particular, we impose $A_v|\psi\rangle = |\psi\rangle$ gauge symmetries on the neural network and weakly break them by first transforming the input with a non-equivariant layer. The proposed

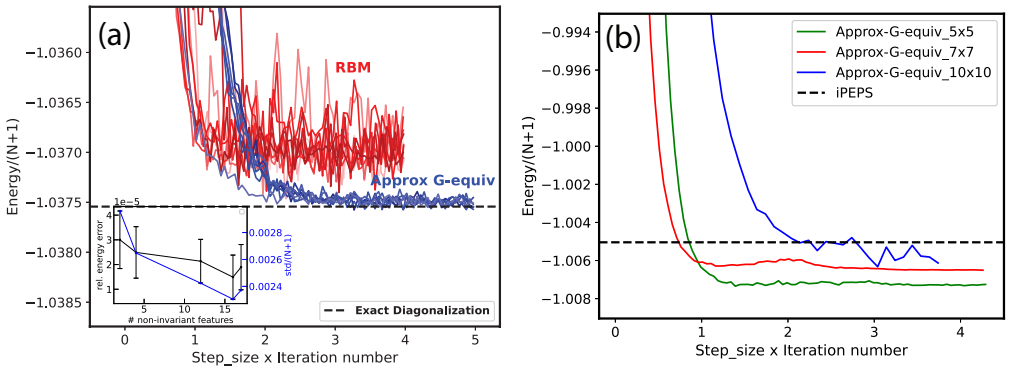

Figure 2: (a) Energy density convergence curve for 4×4 ($N = 24$) system in $h_x = h_z = 0.2$ field for approximately-invariant networks Approx-G-equiv NQS and naive NQS methods as compared with exact diagonalization. The relative error $\delta E = |E_{NQS} - E_{ED}|/E_{ED}$ is of order 10^{-5} and its scaling with the number of non-invariant features is shown in the inset. The runtime of Approx-G-equiv NQS on an NVIDIA V100 GPU for this system size is approximately a few minutes. (b) Energy density convergence curve for 5×5, 7×7, and 10×10 systems for Approx-G-equiv NQS and iPEPS at $h_x = h_z = 0.1$ field.

architecture, shown in Fig. 1, is defined by $\psi_s = \chi(\sigma(\Omega(s)))$ where s is a bit string input, $\Omega : E \to E$ is a non-equivariant convolutional layer, $\sigma : E \to P$ is a gauge-invariant non-linearity and $\chi : P \to \mathbb{C}$ is a convolutional layer. The non-equivariant convolutional layer Ω has a kernel centered at each link of the lattice and explicitly breaks $G = \mathbb{Z}_2^{\times N/2}$ symmetry. The non-linear layer, σ, is constructed using gauge-invariant operators in the model given by products over the $B_p = \prod_{j \in p} Z_j$ operators [14]. This layer ensures the gauge invariance of any further layers within the representation, since it maps the input to the trivial representation of the symmetry group. Finally, the convolutional layer χ consists of square-shaped kernels. We do not use any linear equivariant layers since, in an N-dimensional vector space, they would simply be proportional to the identity.

In Fig. 2a we demonstrate the accuracy of the approximately-symmetric NQS architecture by showing the energy convergence curves for the perturbed toric code Hamiltonian on a 4×4 lattice. At these small system sizes, we find excellent agreement with the result from exact diagonalization, in contrast to conventional baseline NQS approaches (i.e. Restricted Boltzmann Machines incorporating lattice symmetries [3, 25]). Conventional NQS are limited by the intractability of the training landscape, and thus increasing the number of parameters does not improve the performance. On the other hand, we do observe such an improvement for our approximately-invariant architecture (Fig. 2a inset). Finally, in Fig. 2b, we demonstrate the scalability of our approach by plotting the energy density for $L = 5, 7, 10$ (at $h_x = h_z = 0.1$) and comparing with state-of-the-art iPEPS results obtained for $L \to \infty$ [26]. As expected, we see that the NQS energy approaches the iPEPS result strictly from below as the effect of the boundary diminishes.

4 Conclusions

In this work, we have detailed the construction of approximately group-invariant neural networks for addressing problems in quantum many-body physics with approximate symmetry. We demonstrated the accuracy and scalability of our approach on the perturbed toric code, paving the way for tackling challenging problems in quantum many-body physics such as $U(1)$ spin liquids. We anticipate that incorporating transfer learning techniques and enhancing the NQS architecture could further improve our results. Furthermore, the choice of the NQS architecture described above allows us to make first steps towards neural network interpretability in NQS: intuitively the non-equivariant layer transforms the input from a a space with approximate symmetries to one with exact symmetries. This transformation should be equivalent to a finite-depth unitary that reverts the perturbed toric code model to its exactly-symmetric state on the phase diagram. We hope to rigorously demonstrate the equivalence of these statements in the forthcoming work.

Acknowledgments and Disclosure of Funding

We thank DinhDuy Vu, Andrea Pizzi, Johannes Feldmeier, Quynh Nguyen, Rui Wang, Marc Machaczek, Roger Melko, Ruben Verresen, Phil Crowley, Leo Lo, Marcello Dalmonte, Eun-Ah Kim, Joaquin Rodriguez-Nieva, Lode Pollet, Arthur Pesah and Yi Tan for interesting discussions. We are indebted to Wen-Tao Xu for providing us iPEPS data. We acknowledge support from the NSF via the STAQ II program and the QLCI program (grant no. OMA-2016245). D.K. acknowledges support from a Generation-Q AWS fellowship. N.Y.Y. acknowledges support from a Simons Investigator Award.

References

[1] Jacob C Bridgeman and Christopher T Chubb. Hand-waving and interpretive dance: an introductory course on tensor networks. *Journal of physics A: Mathematical and theoretical*, 50 (22):223001, 2017.

[2] Stefan Wessel. 7 monte carlo simulations of quantum spin models. *Emergent Phenomena in Correlated Matter*, 2013.

[3] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.

[4] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[5] Or Sharir, Amnon Shashua, and Giuseppe Carleo. Neural tensor contractions and the expressive power of deep neural quantum states. *Physical Review B*, 106(20):205136, 2022.

[6] Or Sharir, Yoav Levine, Noam Wies, Giuseppe Carleo, and Amnon Shashua. Deep autoregressive models for the efficient variational simulation of many-body quantum systems. *Physical review letters*, 124(2):020503, 2020.

[7] Kyle Sprague and Stefanie Czischek. Variational monte carlo with large patched transformers. *arXiv preprint arXiv:2306.03921*, 2023.

[8] Lucile Savary and Leon Balents. Quantum spin liquids: a review. *Reports on Progress in Physics*, 80(1):016502, 2016.

[9] Christopher Roth, Attila Szabó, and Allan H MacDonald. High-accuracy variational monte carlo for frustrated magnets with deep neural networks. *Physical Review B*, 108(5):054410, 2023.

[10] Zhuo Chen, Laker Newhouse, Eddie Chen, Di Luo, and Marin Soljačić. Autoregressive neural tensornet: Bridging neural networks and tensor networks for quantum many-body simulation. *arXiv preprint arXiv:2304.01996*, 2023.

[11] Kevin Zhang, Samuel Lederer, Kenny Choo, Titus Neupert, Giuseppe Carleo, and Eun-Ah Kim. Hamiltonian reconstruction as metric for variational studies. *SciPost Physics*, 13(3):063, 2022.

[12] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.

[13] Moritz Reh, Markus Schmitt, and Martin Gärttner. Optimizing design choices for neural quantum states. *Physical Review B*, 107(19):195115, 2023.

[14] Di Luo, Giuseppe Carleo, Bryan K Clark, and James Stokes. Gauge equivariant neural networks for quantum lattice gauge theories. *Physical review letters*, 127(27):276402, 2021.

[15] Marc Finzi, Gregory Benton, and Andrew G Wilson. Residual pathway priors for soft equivariance constraints. *Advances in Neural Information Processing Systems*, 34:30037–30049, 2021.

[16] Rui Wang, Robin Walters, and Rose Yu. Approximately equivariant networks for imperfectly symmetric dynamics. In *International Conference on Machine Learning*, pages 23078–23091. PMLR, 2022.

[17] Sandro Sorella. Green function monte carlo with stochastic reconfiguration. *Physical review letters*, 80(20):4558, 1998.

[18] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2): 251–276, 1998.

[19] James Stokes, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020.

[20] Marc Finzi, Max Welling, and Andrew Gordon Wilson. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In *International conference on machine learning*, pages 3318–3328. PMLR, 2021.

[21] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *Advances in Neural Information Processing Systems*, 31, 2018.

[22] Fengcheng Wu, Youjin Deng, and Nikolay Prokof'ev. Phase diagram of the toric code model in a parallel magnetic field. *Physical Review B*, 85(19):195104, 2012.

[23] Simon Trebst, Philipp Werner, Matthias Troyer, Kirill Shtengel, and Chetan Nayak. Breakdown of a topological phase: Quantum phase transition in a loop gas model with tension. *Physical review letters*, 98(7):070602, 2007.

[24] Bei Zeng, Xie Chen, Duan-Lu Zhou, and Xiao-Gang Wen. Quantum information meets quantum matter–from quantum entanglement to topological phase in many-body systems. *arXiv preprint arXiv:1508.02595*, 2015.

[25] Filippo Vicentini, Damian Hofmann, Attila Szabó, Dian Wu, Christopher Roth, Clemens Giuliani, Gabriel Pescia, Jannes Nys, Vladimir Vargas-Calderón, Nikita Astrakhantsev, et al. Netket 3: machine learning toolbox for many-body quantum systems. *SciPost Physics Codebases*, page 007, 2022.

[26] SPG Crone and P Corboz. Detecting a z 2 topologically ordered phase from unbiased infinite projected entangled-pair state simulations. *Physical Review B*, 101(11):115143, 2020.