

PAPER

Learning a compass spin model with neural network quantum states

To cite this article: Eric Zou *et al* 2022 *J. Phys.: Condens. Matter* **34** 125802

View the [article online](#) for updates and enhancements.

You may also like

- [Classical versus quantum models in machine learning: insights from a finance application](#)
Javier Alcazar, Vicente Leyton-Ortega and Alejandro Perdomo-Ortiz
- [Stabilization of Short Wavelength Resistive Ballooning Modes by Ion-to-Electron Temperature and Gradient Ratios in Tokamak Edge Plasmas](#)
Jian-Qiang Xu, , Xiao-Dong Peng et al.
- [Restricted Boltzmann machine: Recent advances and mean-field theory](#)
Aurélien Decelle and Cyril Furtlehner



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Learning a compass spin model with neural network quantum states

Eric Zou, Erik Long and Erhai Zhao* 

Department of Physics and Astronomy, George Mason University, Fairfax, Virginia 22030, United States of America

E-mail: ezhao2@gmu.edu

Received 12 October 2021, revised 2 December 2021

Accepted for publication 16 December 2021

Published 7 January 2022



Abstract

Neural network quantum states provide a novel representation of the many-body states of interacting quantum systems and open up a promising route to solve frustrated quantum spin models that evade other numerical approaches. Yet its capacity to describe complex magnetic orders with large unit cells has not been demonstrated, and its performance in a rugged energy landscape has been questioned. Here we apply restricted Boltzmann machines (RBMs) and stochastic gradient descent to seek the ground states of a compass spin model on the honeycomb lattice, which unifies the Kitaev model, Ising model and the quantum 120° model with a single tuning parameter. We report calculation results on the variational energy, order parameters and correlation functions. The phase diagram obtained is in good agreement with the predictions of tensor network ansatz, demonstrating the capacity of RBMs in learning the ground states of frustrated quantum spin Hamiltonians. The limitations of the calculation are discussed. A few strategies are outlined to address some of the challenges in machine learning frustrated quantum magnets.

Keywords: frustrated quantum spin models, neural network quantum states, machine learning, numerical many-body algorithms

(Some figures may appear in colour only in the online journal)

1. Introduction

Finding the ground-state wave functions of frustrated quantum spin models [1, 2] in two dimensions (2D) remains an outstanding theoretical challenge despite the great strides made in recent decades in the field of numerical many-body algorithms [3]. On the one hand, exact diagonalization and density matrix renormalization group (DMRG), when applied to 2D, are restricted by the finite system size. On the other hand, variational Monte Carlo approaches, while being unbiased and valid in the thermodynamic limit, depend on the quality of the trial wave functions and tend to lose its prediction power due to the lack of convergence associated with the prevalent ‘negative sign problem’ in frustrated quantum spin systems. Another efficient variational ansatz for quantum spin models is based on the tensor network (TN) representation of the

many-body wave functions [4, 5]. One type of TNs known as projected entangled pair states, which generalize the matrix product states from one dimension to 2D, have been successfully applied to quantum spin models [6]. The accuracy of TN ansatz however depends on the approximations employed when truncating and contracting the tensors. Given that each method has its advantages as well as drawbacks, a holistic approach will benefit from new numerical methods that can shed fresh light on this persistent problem.

Inspired by its tremendous success in machine learning and artificial intelligence [7], neural networks were recently proposed to solve quantum spin models [9, 10]. A generic wave function of N interacting spins, say $S = 1/2$, is a superposition of 2^N basis states with complex coefficients. Mathematically, the wave function $\psi(\mathbf{s})$ defines a mapping from vector $\mathbf{s} = (s_1, s_2, \dots, s_N)$, where $s_i = \uparrow$ or \downarrow , to a complex number. Therefore, it can be thought of as a machine that gobbles up \mathbf{s} and spits out a complex number. It is

* Author to whom any correspondence should be addressed.

conjectured that such mapping can be represented accurately by neural networks with sufficient number of nodes, layers, and connections [8]. Then, to find the ground state of a given interacting spin Hamiltonian, all one needs to do is to train the network by adjusting its parameters stochastically so that the energy expectation value is minimized. This approach was pioneered by Carleo and Troyer, who represented the many-spin wave function as restricted Boltzmann machines (RBMs) and successfully applied it to solve the Heisenberg model on the square lattice [9]. In addition to RBM, the wave function can also be expressed as feed-forward neural networks or other neural network architectures. For example, Choo *et al* employed convolutional feed-forward neural networks to solve the J_1 – J_2 model, a canonical example of frustrated magnets believed to host a quantum spin liquid, and obtained excellent energetics comparable to exact diagonalization and DMRG [11]. The J_1 – J_2 model has also been investigated in references [12–16] using the neural network ansatz. In a broader context, recent work has revealed a remarkable connection between the neural-network and the tensor-network representation of quantum many-body states despite their differences in appearance and origin [17–19]. It was shown that neural network quantum states can describe states with topological order, even with entanglement entropy beyond the area law [20–22]. Beyond quantum spin models, neural network quantum states have been also applied to strongly correlated fermions [23–25].

These promising developments have raised many open questions. So far the neural network ansatz was able to identify relatively simple states such as the Neel or stripe order. Is it capable of finding more exotic phases with complicated symmetry breaking patterns? How can it be applied to extract the entanglement signatures of quantum spin liquids? Does it provide an accurate, practical method to determine the phase transitions by computing the order parameters and correlation functions? Very recently, certain limitations to the expressive power of RBMs as well as the stochastic reconfiguration algorithm have been noted. For example, in some cases, the algorithm suffers from inherent numerical instabilities [14, 26]. For highly frustrated quantum spin models, e.g. near $J_2/J_1 = 0.5$ in the J_1 – J_2 model, the energy landscape is believed to be rugged, the approach to the global minimum may not be guaranteed in practice. These limitations led to ongoing efforts to represent the amplitude and the phase of the wave function separately using two real-valued networks, and to learn the sign structures of the wave function to facilitate the convergence [14].

In light of these open questions, in this work we apply the neural network ansatz to the tripod model [27], a frustrated quantum spin model in 2D. It contains the Kitaev model [28] as a special limit and has an extended spin liquid phase. At the same time, its phase diagram also includes the Neel order and a nontrivial bond order which, according to TN calculations [27], can be viewed as a periodic lattice of spin vortices. Thus, this model provides an ideal playground to test the performance and limitations of the neural network ansatz. We note that previously, there have been several works that applied neural networks to study the ground state and excitations of the

Kitaev model or its generalizations, e.g. with external magnetic field or Heisenberg terms [29–32]. The model here is rather different: it overlaps with the Kitaev model only at one special point. Moreover, our primary focus is on the phase diagram and phase transitions between the spin liquid and the long-range ordered states.

This paper is organized as follows. In section 2, we introduce the tripod model and summarize existing numerical results from TN ansatz. Then we outline the RBM ansatz in section 3. Section 4 gives a detailed discussion of our main numerical results, including the energy, the order parameters, and the resulting phase diagram. In section 5, we discuss the limitations of the neural network ansatz as implemented in our work, and directions for future improvement. We hope our results, including the strategies employed to facilitate the learning process, can be useful for applying the neural network ansatz to other quantum spin models, and more generally, to quantum many-body systems.

2. The tripod model

The tripod model is a quantum spin model defined on the two dimensional honeycomb lattice. It belongs to compass spin models, a broad class of Hamiltonians in which the exchange interaction between two neighboring spins depends on the spatial direction of the bond. The study of compass models has a long history, for review see reference [33]. Perhaps the best known example is the Kitaev model [28]: along the three bond directions of the honeycomb lattice, the spin exchange interaction is given by $S_x S_x$, $S_y S_y$, and $S_z S_z$ respectively (in this shorthand notation, the first spin operator is for one lattice site and the second for a neighboring site). Another interesting example of compass models is the quantum 120° model discovered by Zhao and Liu [34], and independently by Wu [35], in the study of strongly interacting p -orbital fermions. In this model, the spin exchanges along the three bonds of the honeycomb lattice are given by $S_1 S_1$, $S_2 S_2$, and $S_3 S_3$ respectively. In spin space, each spin operator is represented by a vector, and here the three spin vectors $S_{1,2,3}$ lie within a plane forming 120° angle with each other. It is apparent that the 120° model is a cousin of the Kitaev model where the three corresponding spin vectors $S_{x,y,z}$ form an orthogonal triad in spin space (i.e. 90° angles with each other). This intimate connection between the two models motivated the authors of reference [27] to unify the 120° model, the Kitaev model, and the Ising model into a single compass model parameterized by an angle θ . Its Hamiltonian is given by

$$H = J \sum_{\mathbf{r}, \gamma} S_\gamma(\mathbf{r}) S_\gamma(\mathbf{r} + \mathbf{e}_\gamma). \quad (1)$$

Here $J > 0$ is the antiferromagnetic coupling, \mathbf{r} labels the lattice site, and \mathbf{e}_γ with $\gamma = 1, 2, 3$ denotes the three bond vectors of the honeycomb lattice, i.e. $\mathbf{r} + \mathbf{e}_\gamma$ is a neighboring site of \mathbf{r} (we have set the lattice spacing to one). The spin $1/2$ operator S_γ is defined as

$$S_\gamma = \frac{1}{2}(\tau_z \cos \phi_\gamma + \tau_x \sin \phi_\gamma) \cos \theta + \frac{1}{2} \tau_y \sin \theta, \quad (2)$$

where τ_x, τ_y, τ_z are the Pauli matrices, and $\phi_\gamma = 0, 2\pi/3, 4\pi/3$ are the azimuthal angle of the corresponding bond direction \mathbf{e}_γ . For brevity, we have suppressed the θ dependence of S_γ and H , and the notation $S_\gamma(\mathbf{r})$ means the spin operator S_γ is localized at site \mathbf{r} .

Model equation (1) is dubbed the tripod model, because geometrically the three S_γ form a tripod in the spin space as shown in figure 1 of reference [27]. They are tilted out of the xz plane by angle θ and, when projected onto the xz plane, are 120° from each other. In addition to the tilting angle θ that defines S_γ , it is convenient to follow reference [27] to introduce θ' , the angle between S_1 and S_2 , i.e. the angle subtended by the two adjacent legs of the tripod. The two angles are related by $\cos \theta' = 1 - (3/2)\cos^2 \theta$. Three limits can be identified as we change θ from 0 to 90° . At $\theta = 0$ (and correspondingly $\theta' = 120^\circ$), the three legs of the tripod $S_{1,2,3}$ are fully open and lie within the xz plane. In this limit, H reduces to the quantum 120° model. As θ is increased, the three legs are increasingly tilted out of the xz plane, corresponding to a partially open tripod. At $\theta = \theta_K = \arccos(\sqrt{2/3}) \simeq 35.26^\circ$, θ' becomes exactly 90° , then the tripod model becomes the Kitaev model: now that the three operators S_γ are orthogonal to each other, we can carry out a spin rotation and redefine them as $S_{x,y,z}$. Finally, when θ is increased all the way to 90° , we have $\theta' = 0$ and all three S_γ collapse to the y axis. The tripod is now fully closed. In this limit, H reduces to the Ising model, $H = (J/4) \sum_{\mathbf{r}, \gamma} \tau_y(\mathbf{r})\tau_y(\mathbf{r} + \mathbf{e}_\gamma)$. Note that usually the Ising interaction is written in the form of $\tau_z\tau_z$. Here to make it easier to compare with previous literature, we follow the convention to choose τ_y as the vertical axis in spin space [27]. This choice of the axes is particularly convenient to recover the 120° model defined in earlier work reference [34]. In passing, we note that the tripod model is not only of theoretical interest due to its synthesis of three important models in quantum magnetism. Recent experiments on honeycomb antiferromagnet $\text{NaNi}_2\text{BiO}_{6-\delta}$ suggest that its dominant exchange interactions resemble those in the tripod model with additional terms, such as the Heisenberg exchange, also playing a role [37].

Some limits of the tripod model are easy to understand. For example, in the Ising limit, the ground state has Neel order, and the order parameter is the staggered magnetization along y . At the Kitaev point, the model is analytically solvable, and its ground state is a spin liquid and has no long-range magnetic order [28]. Aside from these two limits, for general θ , the tripod model must be solved numerically. This is challenging because the model is frustrated and hosts highly nontrivial quantum phases. In particular, the ground state of the 120° model has been somewhat controversial [34–36]. It was conjectured to be long-range ordered despite the geometric frustration [35]. Reference [27] for the first time solved the tripod model for general θ' and obtained its ground state phase diagram using TN ansatz. The main conclusion is that there are three phases separated by two phase transitions, see figure 1 of reference [27]. In particular it predicted that the ground state of the 120° model has valence bond order. In this state, all the spin are confined within the xz plane to form a periodic pattern which can be viewed as a triangular lattice of hexagons. Around each hexagon, the spin winds successively at a 60°

interval, forming a discrete spin vortex, see figure 5 of reference [27]. Note that this phase was referred to as ‘dimer phase’ in reference [27], because along the bonds connecting these hexagons, two neighboring spins point in opposite directions. Such a terminology is unconventional, because in the literature ‘dimer’ is usually synonymous to spin singlet. To avoid potential confusion, we prefer to call this phase having *valence bond order*, because it features spatially periodic modulations of the bond energy. TN ansatz also predicted that the quantum spin liquid is stabilized within the finite window $\theta' \in [87^\circ, 94^\circ]$ enclosing the Kitaev point $\theta' = 90^\circ$. Judging from the variation of the order parameters with θ' , the valence bond to spin liquid transition seems continuous, while the spin liquid to Neel transition seems to be first order [27].

The main goal of the present work is to investigate the ground state phase diagram of the tripod model using an independent method. This serves two purposes. On the one hand, the variational calculation with neural network ansatz here provides a check for the TN results, especially regarding the ground state in the 120° limit as well as the location and nature of the phase transitions. On the other hand, the calculation tests the capacity of the neural network ansatz by applying it to solve a frustrated quantum spin model which has not only spin liquid but also nontrivial long-range order with an intricate spatial pattern. *A priori*, it is unclear whether these ground states, the order parameters, or phase transitions can be captured by the neural network ansatz. Overall, our calculation benchmarks the efficiency, stability, and accuracy of the neural network algorithm by comparing to the state-of-the-art TN results.

3. Restricted Boltzmann machines

We will represent the many-spin wave function using one of the simplest neural networks, the RBMs. The implementation follows the original work of reference [9]. To avoid repetition, here we only outline the main ideas. More details can be found in reference [9, 10]. A RBM is a shallow neural network with two layers, the visible layer consisting of N nodes characterized by spin variables s_i ($i = 1, 2, \dots, N$) and a hidden layer of M nodes described by variables h_j ($j = 1, 2, \dots, M$). The coupling between node s_i and node h_j is described by a connection weight w_{ij} . The quantum mechanical wave function takes the form of joint Boltzmann weight [9],

$$\begin{aligned} |\Psi\rangle &= \sum_{\{s_i\}} \psi(\mathbf{s}) |\{s_i\}\rangle \\ &= \sum_{\{s_i\}} \sum_{\{h_j\}} e^{a_i s_i + b_j h_j + w_{ij} s_i h_j} |\{s_i\}\rangle. \end{aligned} \quad (3)$$

Here, repeated indices in the exponent are summed over, $\{s_i\} = \{s_1, s_2, \dots, s_N\}$ are all possible spin configurations (similarly for $\{h_j\}$), the biases a_i and b_j as well as the connection weights w_{ij} are all complex variational parameters. Even though there is no direct intra-layer connection in a ‘restricted’ Boltzmann machine, the hidden nodes induce correlations among the visible nodes. For real biases and connections, it is known that RBMs can represent any classical

distribution to desired accuracy with sufficient numbers of hidden units [8]. The expressive power of complex RBMs is less known. It has been argued that a fully connected RBM can capture entanglement bounded by volume law and hence efficiently describe the ground states of many Hamiltonians [17]. In our calculations, we consider a finite honeycomb lattice with $L \times L$ unit cells with periodic boundary conditions. Then the number of sites $N = 2L^2$ and the number of bonds $N_b = 3L^2$. The layer density ratio $\alpha = M/N$ is a tuning parameter, we find $\alpha = 2$ gives satisfactory performance for $L = 4$.

Starting from some initial guess, e.g. random values, the variational parameters are adjusted iteratively to minimize the variation energy $E = \langle \Psi | H | \Psi \rangle$, the expectation value of the Hamiltonian equation (1) for the current wave function equation (3), computed approximately by Monte Carlo sampling [9]. This is done by making stochastic moves in a large parameter space based on estimating the energy gradient. This stochastic optimization procedure is often called *learning*, or training the RBM. Here many mature algorithms from the machine learning literature can be applied [7, 38]. For example, we have tested and compared several algorithms including stochastic gradient descent, adagrad, and adamax [38, 39, 41]. The actual computation is carried out using the powerful Netket library [41], aided by custom-made routines to manipulate the variational wave functions directly.

We emphasize that while the model and learning algorithm are relatively straightforward to set up, the actual training of the RBM with a vast parameter space is by no means a trivial task. This is analogous to many other complex machine learning tasks: efficient training a neural network hinges on understanding the particularities of the model, the parameter space, and the quantity and quality of the data etc [38, 39]. For example, starting from a random configuration of the RBM, the algorithm may lead to a quick convergence to a local minimum and stall there. This becomes especially problematic in regions where a few orders compete: for example, a blind stochastic search often yields wildly fluctuating results for two neighboring parameters that belong to the same phase. In this case, one may find the best energetics by trying different optimization algorithms or starting from different initial guesses. Even when the true ground state is being approached, the accuracy of the converged energy depends critically on the proper choice of the parameters such as learning rate and sampling batch size. A more serious problem is the sporadic occurrence of numerical instability, presumably due to the parameters being complex, which may manifest as a fast runaway of the energy toward divergence. These numerical complexities complicate the task of finding the ground state phase diagram. (The performance of RBM is discussed in reference [40] for classical spin models). Some of the strategies we employ to alleviate these problems are discussed below in section 4.

4. Ground-state phase diagram

The procedure to learn the phase diagram of the tripod model is as follows. For a given value of the tilting angle θ , the RBM is started from random parameter values, then stochastic moves are made to lower the variational energy until

convergence is achieved. A crucial parameter here is the learning rate r , or step size of the stochastic moves [38, 39]. For stochastic gradient descent, choosing a r that is too large can easily end up with numerical instability, while having r too small may slow the learning to a crawl and trap it inside a local minimum. The optimal value of r depends on the model and the optimizer (many popular optimizers use adaptive learning rates determined from gradient and/or momentum). Its order of magnitude is determined by trial and error, and its value is adjusted on the fly, for example, when entering a flat energy landscape. When the algorithm fails to reach the anticipated energy, different optimizers or learning parameters are tried to shake things up. If no further progress can be made, the machine is restarted. Some states, for example the Neel state, are rather easy to reach with fast and robust convergence, e.g. after hundreds of iterations. Other states, such as the spin liquid or bond order, require many more steps for the energy to relax, e.g. at least thousands of iterations even with reasonable learning parameters and initial guess. This is expected because of the frustration and the presence of many competing orders. To facilitate the search for ground states in highly frustrated regimes, it is useful to start from wave functions that were learned previously for parameters nearby and have competitive energies. In this case, random fluctuations are introduced to the wave function before the run, and the result must be compared to those obtained from random, blind guesses. During the run, the statistical errors (variance or standard deviation) of the observables are also monitored.

Figure 1 shows the energy per bond in units of J , $\epsilon = E/N_b J$, as a function of the tilting angle θ . The energy is the lowest in the Ising limit $\theta = 90^\circ$, and the neural network ansatz accurately reproduces the analytical result $\epsilon = -1/4$, corresponding to the antiparallel alignment of neighboring spins. The convergence to ground state in this limit is rather fast, perhaps due to the classical nature of the Ising model. (In comparison, for $\theta < 50^\circ$, reaching the ground state is not as straightforward and requires some of the strategies outlined above.) As θ is reduced, ϵ rises quickly; and after going through the Kitaev point $\theta_K \simeq 35.26^\circ$, it reaches its peak value of $\epsilon \simeq -0.122$ at $\theta = 34^\circ$. The elevation in energy is in accordance with the fact that in this region around θ_K the system is most frustrated. Upon further reduction of θ , the energy starts to decrease. The noticeable cusp in energy located at $\theta = 34^\circ$ marks the transition to the bond ordered phase, inside which the frustration is relieved to some degree but not entirely. This trend continues until the 120° model limit is reached at $\theta = 0$. Note the spin liquid to Neel transition is not obvious by inspecting the energy alone. But plotting $d\epsilon/d\theta$ reveals a sudden change at $\theta \sim 40^\circ$. In what follows, we present a better way to reveal the phase boundaries by computing the order parameters and spin correlation functions.

The first marker of phase transition is provided by the spin-spin correlation function. It measures the antiferromagnetic long-range order and is defined by

$$C_y = \frac{4}{N-1} \sum_{\mathbf{r} \neq \mathbf{r}'} \eta_{\mathbf{r}, \mathbf{r}'} [\langle S_y(\mathbf{r}) S_y(\mathbf{r}') \rangle - \langle S_y(\mathbf{r}) \rangle \langle S_y(\mathbf{r}') \rangle]. \quad (4)$$




Figure 1. The converged variational ground state energy per bond in units of J for the tripod model equation (1) obtained from RBMs. Here θ is the tilting angle out of the xz plane. $\theta = 90^\circ$ is the Ising limit, $\theta \simeq 35.26^\circ$ is the Kitaev point, and $\theta = 0^\circ$ realizes the quantum 120° model. The vertical dotted lines at $\theta = 34^\circ$ and 39° are guide to the eye. The system has 32 sites with periodic boundary conditions, $L = 4$, $\alpha = 2$. No symmetry is enforced. The statistical error is smaller than the symbol size.

Here \mathbf{r}, \mathbf{r}' label the sites, and $\eta_{\mathbf{r},\mathbf{r}'} = 1 (-1)$ if \mathbf{r} and \mathbf{r}' belong to the same (different) sublattice. As shown in figure 2, C_y decays from 1 in the Ising limit and drops sharply as the spin liquid phase is approached. The drop experiences a glitch at $\theta = 39^\circ$, and C_y vanishes for $\theta \leq 34^\circ$. This plot clearly demonstrates the existence of Neel order at $\theta > 39^\circ$, as well as the lack of out-of-plane antiferromagnetic correlation for $\theta < 34^\circ$.

The second marker for phase transition is the expectation value of the in-plane spin

$$S_\perp = \frac{1}{N} \sum_{\mathbf{r}} \sqrt{\langle S_x(\mathbf{r}) \rangle^2 + \langle S_z(\mathbf{r}) \rangle^2}, \quad (5)$$

where $\langle S_x(\mathbf{r}) \rangle$ is the expectation value of the local spin operator $S_x(\mathbf{r})$, and the sum $\sum_{\mathbf{r}}$ is over all the sites. Figure 3 shows the variation of S_\perp with θ , which exhibits a trend opposite to C_y : it assumes large values within the bond ordered phase, drops sharply around $\theta_1 \sim 34^\circ$, followed by a small glitch at $\theta_2 \sim 39^\circ$. Afterwards, it remains suppressed and vanishes in the Ising limit where all the spins align parallel to the y -axis. This result confirms that the spins are predominantly in-plane within the bond ordered phase, in good agreement with TN results (see plot of O_2 in figure 2 of reference [27]). Combining figures 2 and 3 together, it is clear that an intermediate (the spin liquid) region is bounded by the lower critical point $\theta_B = 34^\circ$ and the upper critical point $\theta_N = 39^\circ$. These critical values are close to, but not identical with, the TN results $\theta_B = 33^\circ$ and $\theta_N = 38^\circ$ (the phase boundaries were given in reference [27] in terms of θ' , which can be easily converted into θ).

To further elucidate the nature of the bond ordered phase, we compute the bond energies

$$B_\gamma(\mathbf{r}) = 4 \langle S_\gamma(\mathbf{r}) S_\gamma(\mathbf{r} + \mathbf{e}_\gamma) \rangle \quad (6)$$

for all three bonds connected to a given site at \mathbf{r} . One then notices that for $\theta < \theta_B$, one of the bond is stronger than the other two, and a bond modulation pattern develops in space which breaks the underlying lattice symmetry. There are three ways to break the symmetry of the three bonds locally. For example, in symmetry-breaking pattern p_1 , B_1 is stronger (negative with larger magnitude) while B_2 and B_3 are roughly (up to some small fluctuations) the same but weaker. The other two patterns $p_{2,3}$ are obtained by permuting $\gamma = 1, 2, 3$, e.g. bond B_2 is stronger in pattern p_2 . Obviously, these three patterns are related to each other by C_3 rotations in real space. Let us define the bond modulation ΔB as the difference between the stronger bond and the average of the two weaker bonds, for example,

$$\Delta B = \frac{1}{N} \sum_{\mathbf{r}} \frac{1}{2} [B_2(\mathbf{r}) + B_3(\mathbf{r})] - B_1(\mathbf{r}), \quad (7)$$

where the average over all sites is taken. A finite ΔB is expected if the bond modulation pattern p_1 is repeated throughout the lattice. As shown in figure 4, the bond ordered phase is characterized by a finite bond modulation, whereas in both the Neel and spin liquid phase, bond energies are approximately uniform in space. Thus, the bond ordered phase found here has a solid order of periodically modulated bonds, i.e. a valence bond solid. It breaks the C_3 symmetry of the underlying honeycomb lattice, but differs from the spin vortex state discussed in reference [27]. And its energy per bond $\epsilon = -0.143$ is higher than the best TN result -0.148 [27]. The reason behind this difference is addressed in the next section.

5. Limitations and outlook

To summarize, neural network quantum states based on RBM have performed very well to identify the main phases (Neel,




Figure 2. The spin–spin correlation function C_y defined in equation (4) as function of θ . It measures the out-of-plane antiferromagnetic order, and approaches 1 in the Ising limit. Its rapid decay to almost zero at $\theta = 34^\circ$ (indicated by the vertical dotted line on left) and the visible kink at $\theta = 39^\circ$ (vertical dotted line, right) suggest two quantum phase transitions.




Figure 3. The in-plane spin S_\perp defined in equation (5). Its maximum is achieved at $\theta = 0$ and remains finite within the bond ordered phase. The sudden drop at $\theta = 34^\circ$ provides a clear marker for the bond order to spin liquid transition. Transition to Neel phase is accompanied by a small bump around $\theta \sim 40^\circ$. Within the Neel phase, S_\perp is very small and vanishes in the Ising limit.

bond order, and spin liquid) and phase transitions of the frustrated tripod model. We find it remarkable that with some judicious control over the learning parameters and learning strategy, the algorithm can efficiently navigate the 2^{32} -dimensional Hilbert space stochastically to find variational ground states that have competitive energies. In particular the two phase boundaries are close to the state-of-the-art TN ansatz. And for smaller systems, e.g. $L = 3$, the energetics is also in excellent agreement with exact diagonalization.

Our study also exposes some limitations of the unconstrained RBM ansatz as implemented here. In our calculation,

we did not impose any symmetry constraints on the RBM wave functions. While this has the advantage of being completely unbiased, it also makes it exceedingly hard, if not at all impossible, to reach intricate states such as the spin vortex lattice proposed in reference [27] for the 120° model. As shown in figure 4 above, our calculation reaches one of the valence bond states, where the bond modulation pattern p_1 is repeated periodically in space. There are two other states with degenerate energies, where pattern p_2 or p_3 is repeated instead. In fact, all possible coverings of the lattice by a suitable combination of local patterns $p_{1,2,3}$ have the same classical energy,




Figure 4. Bond modulation ΔB defined by equations (6) and (7). The bond ordered phase features bond modulations that grow with decreasing θ . The Neel phase only has small fluctuations in bond energy.

giving rise to a large residual entropy similar to those found in spin ice [42]. The quantum Hamiltonian equation (1) induces transitions between different coverings and lifts the classical degeneracy. Then a particular covering, or superposition of coverings, acquires lower energy to become the quantum mechanical ground state. For example, the vortex state of reference [27] represents a particular periodic covering of the whole lattice with local bond pattern p_1 , p_2 , and p_3 . In principle, this state may be eventually reached by RBM with further refinement in energy. In practice, this turns out to be hard, due to the flatness of the energy landscape (since different covering patterns have close energies) and the diminishing probability of settling into a highly symmetric configuration with a large unit cell by pure stochastic moves in a huge parameter space. Our attempts to further improve the energy frequently encounter numerical instabilities. We conjecture that this barrier can be overcome by applying symmetry constraints [11, 43, 44] to the RBM states, e.g. by enforcing C_6 symmetry and fixing the unit cell shape and size. This should also improve the convergence and numerical stability. The downside is that one must compare the energies of all candidate states with different symmetries. Testing this proposal is left for future work.

Figure 4 illustrates another caveat of unconstrained learning in large systems: there are small but visible fluctuations in the bond energy even in the Neel phase. For an ideal Neel state, one expects $\Delta B = 0$. While the algorithm successfully approaches the antiferromagnetic ground state with excellent energy, the RBM rarely settles into a completely frozen state with uniform bond energy. The stochastic nature of the algorithm unavoidably introduces low lying excitations, which for larger systems are increasingly harder to eliminate. A similar situation is observed in figure 2, where the order parameter C_y drops to almost zero within the bond ordered phase, but small fluctuations are still visible. This presents a dilemma: on the one hand

we need large clusters to accommodate orders with long modulation periods, on the other hand for large systems it becomes more challenging to relax to pristine long-range ordered states.

Given these considerations, we advocate the following strategy to make the best out of the neural network ansatz. First, the unconstrained network is trained to find the rough phase diagram and symmetry breaking patterns. It has the virtue of being unbiased. Then, other methods, such as analytical variational wave functions or neural network with symmetry, are used to further improve the energetics and elucidate the long-range order. We envision such a hybrid approach will be especially useful in understanding complex spin systems, for example models inspired by a large class of Kitaev materials [45, 46]. Our results suggest that, with further refinements and complemented by other approaches, variational ansatz based on neural network quantum states can serve as a powerful tool to understand frustrated quantum spin models and more generally strongly interacting many-body systems.

Acknowledgments

This work is supported by NSF Grant No. PHY-2011386 (EZ). EZ would like to thank Christian McGuirk for preliminary work implementing and benchmarking RBM and Mahmoud Lababidi, Ahmet Keles and Haiyuan Zou for illuminating discussions. The numerical simulation is based on the Netket library version 2.1.1 [41]. Part of the calculation was carried out on the ARGO clusters provided by the Office of Research Computing at George Mason University.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

ORCID iDs

Erhai Zhao  <https://orcid.org/0000-0001-8954-1601>

References

- [1] Lacroix C, Mendels P and Mila F (ed) 2011 *Introduction to Frustrated Magnetism: Materials, Experiments, Theory* (Berlin: Springer)
- [2] Diep H T (ed) 2013 *Frustrated Spin Systems* (Singapore: World Scientific)
- [3] Savary L and Balents L 2016 Quantum spin liquids: a review *Rep. Prog. Phys.* **80** 016502
- [4] Verstraete F and Cirac J I 2004 Renormalization algorithms for quantum-many body systems in two and higher dimensions (arXiv:cond-mat/0407066)
- [5] Orús R 2014 A practical introduction to tensor networks: matrix product states and projected entangled pair states *Ann. Phys., NY* **349** 117
- [6] Verstraete F, Murg V and Cirac J I 2008 Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems *Adv. Phys.* **57** 143
- [7] Mehta P, Bukov M, Wang C-H, Day A G R, Richardson C, Fisher C K and Schwab D J 2019 A high-bias, low-variance introduction to machine learning for physicists *Phys. Rep.* **810** 1
- [8] Le Roux N and Bengio Y 2008 Representational power of restricted Boltzmann machines and deep belief networks *Neural Comput.* **20** 1631
- [9] Carleo G and Troyer M 2017 Solving the quantum many-body problem with artificial neural networks *Science* **355** 602
- [10] Melko R G, Carleo G, Carrasquilla J and Cirac J I 2019 Restricted Boltzmann machines in quantum physics *Nat. Phys.* **15** 887
- [11] Choo K, Neupert T and Carleo G 2019 Two-dimensional frustrated J_1 – J_2 model studied with neural network quantum states *Phys. Rev. B* **100** 125124
- [12] Liang X, Liu W-Y, Lin P-Z, Guo G-C, Zhang Y-S and He L 2018 Solving frustrated quantum many-particle models with convolutional neural networks *Phys. Rev. B* **98** 104426
- [13] Cai Z and Liu J 2018 Approximating quantum many-body wave functions using artificial neural networks *Phys. Rev. B* **97** 035116
- [14] Bukov M, Schmitt M and Dupont M 2021 Learning the ground state of a non-stoquastic quantum Hamiltonian in a rugged neural network landscape *SciPost Phys.* **10** 147
- [15] Nomura Y and Imada M 2021 Dirac-type nodal spin liquid revealed by refined quantum many-body solver using neural-network wave function, correlation ratio, and level spectroscopy *Phys. Rev. X* **11** 031034
- [16] Kochkov D, Pfaff T, Sanchez-Gonzalez A, Battaglia P and Clark B K 2021 Learning ground states of quantum Hamiltonians with graph networks (arXiv:2110.06390)
- [17] Chen J, Cheng S, Xie H, Wang L and Xiang T 2018 Equivalence of restricted Boltzmann machines and tensor network states *Phys. Rev. B* **97** 085104
- [18] Glasser I, Pancotti N, August M, Rodriguez I D and Cirac J I 2018 Neural-network quantum states, string-bond states, and chiral topological states *Phys. Rev. X* **8** 011006
- [19] Huang Y and Moore J E 2021 Neural network representation of tensor network and chiral states *Phys. Rev. Lett.* **127** 170601
- [20] Deng D-L, Li X and Das Sarma S 2017 Quantum entanglement in neural network states *Phys. Rev. X* **7** 021021
- [21] Deng D-L, Li X and Das Sarma S 2017 Machine learning topological states *Phys. Rev. B* **96** 195145
- [22] Kaubruegger R, Pastori L and Budich J C 2018 Chiral topological phases from artificial neural networks *Phys. Rev. B* **97** 195136
- [23] Nomura Y, Darmawan A S, Yamaji Y and Imada M 2017 Restricted Boltzmann machine learning for solving strongly correlated quantum systems *Phys. Rev. B* **96** 205152
- [24] Broecker P, Carrasquilla J, Melko R G and Trebst S 2017 Machine learning quantum phases of matter beyond the fermion sign problem *Sci. Rep.* **7** 8823
- [25] Ch'ng K, Carrasquilla J, Melko R G and Khatami E 2017 Machine learning phases of strongly correlated fermions *Phys. Rev. X* **7** 031038
- [26] Park C Y and Kastoryano M J 2020 Are neural quantum states good at solving non-stoquastic spin Hamiltonians? (arXiv:2012.08889)
- [27] Zou H, Liu B, Zhao E and Liu W V 2016 A continuum of compass spin models on the honeycomb lattice *New J. Phys.* **18** 053040
- [28] Kitaev A 2006 Anyons in an exactly solved model and beyond *Ann. Phys., NY* **321** 2
- [29] Noormandipour M, Sun Y and Haghighat B 2021 Restricted Boltzmann machine representation for the ground state and excited states of Kitaev honeycomb model *Machine Learning: Science and Technology* **3** 015010
- [30] Alcalde Puente D and Eremín I M 2020 Convolutional restricted Boltzmann machine aided Monte Carlo: an application to Ising and Kitaev models *Phys. Rev. B* **102** 195148
- [31] Li C-X, Yang S and Xu J-B 2021 Learning spin liquids on a honeycomb lattice with artificial neural networks *Sci. Rep.* **11** 16667
- [32] Rao N, Liu K, Machaczek M and Pollet L 2021 Machine-learned phase diagrams of generalized Kitaev honeycomb magnets *Phys. Rev. Res.* **3** 033223
- [33] Nussinov Z and van den Brink J 2015 Compass models: theory and physical motivations *Rev. Mod. Phys.* **87** 1
- [34] Zhao E and Liu W V 2008 Orbital order in Mott insulators of spinless p-band fermions *Phys. Rev. Lett.* **100** 160403
- [35] Wu C 2008 Orbital ordering and frustration of p-band Mott insulators *Phys. Rev. Lett.* **100** 200406
- [36] Nasu J, Nagano A, Naka M and Ishihara S 2008 Doubly degenerate orbital system in honeycomb lattice: implication of orbital state in layered iron oxide *Phys. Rev. B* **78** 024416
- [37] Scheie A *et al* 2019 Counterrotating magnetic order in the honeycomb layers of $\text{NaNi}_2\text{BiO}_{6-\delta}$ *Phys. Rev. B* **100** 214421
- [38] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press)
- [39] Geron A 2019 *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* 2nd edn (Sebastopol, CA: O'Reilly & Associates)
- [40] Azizi A and Pleimling M 2021 A cautionary tale for machine learning generated configurations in presence of a conserved quantity *Sci. Rep.* **11** 1
- [41] Carleo G *et al* 2019 NetKet: a machine learning toolkit for many-body quantum systems *SoftwareX* **10** 100311
- [42] Ramirez A P, Hayashi A, Cava R J, Siddharthan R and Shastri B S 1999 Zero-point entropy in spin ice *Nature* **399** 333
- [43] Nomura Y 2021 Helping restricted Boltzmann machines with quantum-state representation by restoring symmetry *J. Phys.: Condens. Matter* **33** 174003
- [44] Roth C and MacDonald A H 2021 Group convolutional neural networks improve quantum state accuracy (arXiv:2104.05085)
- [45] Winter S M, Li Y, Jeschke H O and Valentí R 2016 Challenges in design of Kitaev materials: magnetic interactions from competing energy scales *Phys. Rev. B* **93** 214431
- [46] Trebst S 2017 Kitaev materials (arXiv:1701.07056)