

---

# Supplementing Recurrent Neural Network Wave Functions with Symmetry and Annealing to Improve Accuracy

---

Mohamed Hibat-Allah<sup>1,2</sup>, Roger G. Melko<sup>2,3</sup>, Juan Carrasquilla<sup>1,2</sup>

<sup>1</sup> Vector Institute, MaRS Centre, Toronto, Ontario, M5G 1M1, Canada

<sup>2</sup> Department of Physics and Astronomy, University of Waterloo, Ontario, N2L 3G1, Canada

<sup>3</sup> Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada

mohamed.hibat.allah@uwaterloo.ca

## Abstract

Recurrent neural networks (RNNs) are a class of neural networks that have emerged from the paradigm of artificial intelligence and has enabled lots of interesting advances in the field of natural language processing. Interestingly, these architectures were shown to be powerful ansatz to approximate the ground state of quantum systems [1]. Here, we build over the results of Ref. [1] and construct a more powerful RNN wave function ansatz in two dimensions. We use symmetry and annealing to obtain accurate estimates of ground state energies of the two-dimensional (2D) Heisenberg model, on the square lattice and on the triangular lattice. We show that our method is superior to Density Matrix Renormalisation Group (DMRG) for system sizes larger than or equal to  $14 \times 14$  on the triangular lattice.

## 1 Introduction

In recent years, deep learning has proven useful in quantum many-body physics [2, 3, 4, 5]. Namely in tasks such as quantum state tomography [6, 7] and finding ground states of quantum many-body Hamiltonians [8, 9, 10, 11, 12, 13, 1, 14]. Yet, up to now, the numerical study of frustrated quantum systems remains challenging due to the debilitating sign problem [15, 16].

Interestingly, recurrent neural networks (RNNs) [17, 18, 19] have shown a significant advantage in the study of quantum systems [7, 1, 14, 20]. In this paper, we aim to show the potential of RNNs in the study of frustrated quantum systems that suffer from the sign problem [15, 16]. To do this, we focus on the prototypical Heisenberg model and show that by applying symmetries and annealing leads to accurate representations of the ground state. In Sec. 2, we describe our improved RNN architecture and annealing scheme. In Sec. 3, we show results for the Heisenberg model on the square and the triangular lattice. We demonstrate that our method is superior compared to the Density Matrix Renormalization Group (DMRG) [21, 22] on system sizes larger than  $14 \times 14$  up to  $16 \times 16$ .

## 2 Methods

### 2.1 Recurrent neural network wave functions

To model a complex-valued wave function  $|\Psi\rangle$ , we use the following decomposition

$$\Psi_{\theta}(\sigma) = \sqrt{p_{\theta}(\sigma)} \exp(i\phi_{\theta}(\sigma)),$$

where  $\theta$  denotes the variational parameters of the ansatz wave function  $|\Psi_{\theta}\rangle$ , and  $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$  is a spin configuration. In our paper, we use a complex RNN (cRNN) wave

function to handle states that change sign in the computational basis [1]. A key feature of RNN wave functions is the ability to obtain uncorrelated samples to estimate observables through autoregressive sampling [1, 23]. More details about this sampling procedure in 1D and 2D can be found in Refs. [1, 24].

In order to efficiently encode 2D correlations, we use a two-dimensional recursion relation as follows [24]:

$$\mathbf{h}_{i,j} = F([\boldsymbol{\sigma}_{i-1,j}; \boldsymbol{\sigma}_{i,j-1}]T[\mathbf{h}_{i-1,j}; \mathbf{h}_{i,j-1}] + \mathbf{b}). \quad (1)$$

$\mathbf{h}_{i,j}$  is a hidden state with two indices for each spin in the 2D lattice. The symbol  $[\cdot; \cdot]$  denotes concatenation operation of two vectors, while  $\boldsymbol{\sigma}_{i,j}$  is the one-hot encoding of the spin  $\sigma_{i,j}$ . The tuneable arrays  $T$  and  $\mathbf{b}$  are a tensor and bias vector, respectively. Since  $\mathbf{h}_{i,j}$  can be considered a summary of the history of spins that were previously generated, it is used to compute the conditional probability and phase as follows:

$$p_{\boldsymbol{\theta}}(\sigma_{i,j} | \boldsymbol{\sigma}_{<i,j}) = \text{Softmax}(U\mathbf{h}_{i,j} + \mathbf{c}) \cdot \boldsymbol{\sigma}_{i,j}, \quad (2)$$

$$\phi_{\boldsymbol{\theta}}(\sigma_{i,j} | \boldsymbol{\sigma}_{<i,j}) = \pi \text{Softsign}(V\mathbf{h}_{i,j} + \mathbf{d}) \cdot \boldsymbol{\sigma}_{i,j}. \quad (3)$$

Here  $\boldsymbol{\sigma}_{<i,j}$  stands for the spins that were previously generated on the zigzag path as shown in Fig. 1(a).  $U, V$  and  $\mathbf{c}, \mathbf{d}$  are additional variational parameters. From the two previous relations, the probability  $p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})$  and the phase  $\phi_{\boldsymbol{\theta}}(\boldsymbol{\sigma})$  can be computed respectively by taking the product (resp. sum) of  $p_{\boldsymbol{\theta}}(\sigma_{i,j} | \boldsymbol{\sigma}_{<i,j})$  (resp.  $\phi_{\boldsymbol{\theta}}(\sigma_{i,j} | \boldsymbol{\sigma}_{<i,j})$ ).

We note that to compute  $\mathbf{h}_{i,j}$  in the recursion relation (1), we use the neighbouring hidden states and spins, as demonstrated in Fig. 1(a). We also remark that a tensorized operation of the hidden states and the spins is chosen instead of a concatenation operation (see Eq. (1)) to improve the expressivity of the 2D RNN as shown in Ref. [24]. In this paper, we also use a gated version of the 2D RNN to further improve the accuracy of our 2D cRNN wave function ansatz. More details are shown in Appendix. A.

Furthermore, symmetries in a system of interest can be incorporated in RNN wave functions, namely  $U(1)$  symmetry and discrete lattice symmetries as it was shown in Ref. [1]. We highlight that adding symmetries significantly improves the accuracy of a variational calculation, as motivated in the literature with the example of RBM wave functions [25]. We finally emphasize that the variational parameters  $\boldsymbol{\theta}$  are optimized with the variational Monte Carlo (VMC) scheme to minimize  $E_{\boldsymbol{\theta}} = \langle \Psi_{\boldsymbol{\theta}} | \hat{H} | \Psi_{\boldsymbol{\theta}} \rangle$  where  $\hat{H}$  is the Hamiltonian describing a quantum system of interest. More details about this scheme can be found in the Appendix of Ref. [1].

## 2.2 Variational Annealing

Finding the ground state of a non-stoquastic Hamiltonian is limited by the sign problem [15, 16], and has shown to induce rugged landscapes during the VMC optimization [26]. The latter can make our ansatz stuck in a local minima, similarly to what is commonly observed in spin glass models at low temperatures. We take inspiration from the concept of annealing, where we supplement VMC with an artificial temperature to overcome local minima during a VMC optimization. This procedure has been already explored in Ref. [24] for the purpose of finding ground states of spin-glass models. Here, we explore this idea in the context of finding ground states of frustrated quantum systems.

In order to supplement the VMC scheme with annealing, we add a pseudo-entropy [14] on top of variational energy  $E_{\boldsymbol{\theta}}$  to obtain a variational free energy:

$$F_{\boldsymbol{\theta}}(n) = E_{\boldsymbol{\theta}} - T(n)S_{\text{classical}}(p_{\boldsymbol{\theta}}), \quad (4)$$

where  $S_{\text{classical}}$  is the von Neumann entropy, given by

$$S_{\text{classical}}(p_{\boldsymbol{\theta}}) = - \sum_{\boldsymbol{\sigma}} p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) \log(p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})). \quad (5)$$

Here the sum goes over all possible configurations  $\{\boldsymbol{\sigma}\}$ . In this case  $T(n)$  is a pseudo-temperature. The latter is annealed starting from some initial temperature  $T_0$  to zero temperature as follows:  $T(n) = T_0(1 - n/N_{\text{annealing}})$  where  $n \in [0, N_{\text{annealing}}]$  and  $N_{\text{annealing}}$  is the total number of annealing steps.

In our setup, we follow the same scheme in Ref. [24], where we first do  $N_{\text{warmup}}$  training steps to prepare our RNN wave function close to the minimum of  $F_{\boldsymbol{\theta}}(n = 0)$ . In the second step, we keep

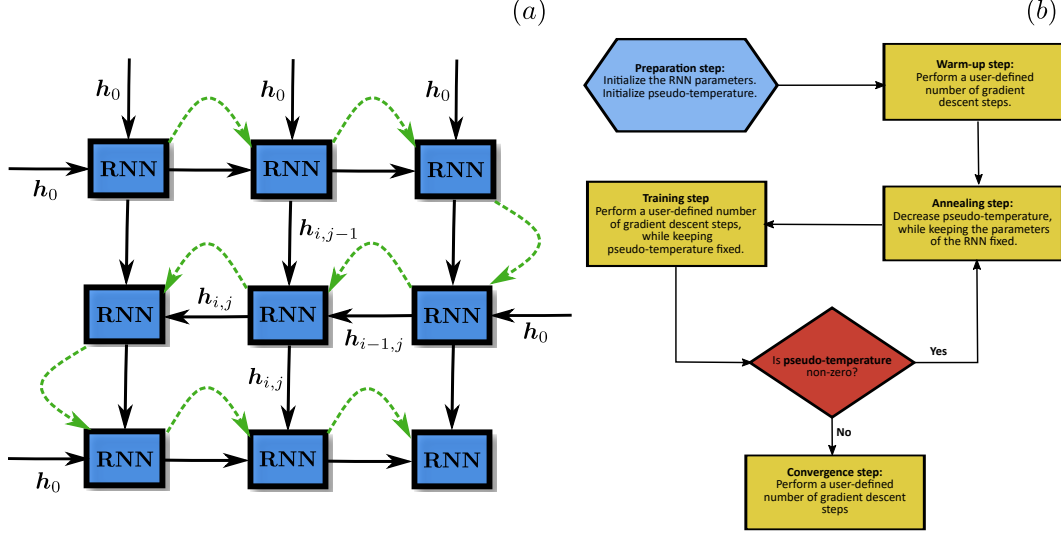


Figure 1: (a) A graphical illustration of a 2D RNN. Each RNN cell receives two hidden states  $h_{i,j-1}$  and  $h_{i-1,j}$ , as well as two input vectors  $\sigma_{i,j-1}$  and  $\sigma_{i-1,j}$  (not shown) as illustrated by the black arrows. The green arrows correspond to the zigzag path we use for 2D autoregressive sampling. The initial memory state  $h_0$  of the RNN and the initial inputs  $\sigma_0$  (not shown) are null vectors. (b) A flowchart describing our annealing implementation.

the parameters of our ansatz fixed and we change  $n$  to  $n + 1$ . In step 3, we perform  $N_{\text{train}}$  gradient steps to bring the RNN ansatz closer to the minimum of  $F_{\theta}(n + 1)$ . After iterating over steps 2 and 3, we arrive at  $n = N_{\text{annealing}}$ . At this point, we perform a user-defined number of gradient steps  $N_{\text{convergence}}$  to further improve our estimate of the ground state energy of  $\hat{H}$ . A flowchart detailing this procedure is provided in Fig. 1(b).

### 3 Results

We focus our attention on the task of finding the ground state of the two dimensional anti-ferromagnetic Heisenberg model on the square lattice and on the triangular lattice with open boundary conditions (OBC). The Hamiltonian is given as follows:

$$\hat{H} = \frac{1}{4} \sum_{\langle i,j \rangle} (\hat{\sigma}_i^x \hat{\sigma}_j^x + \hat{\sigma}_i^y \hat{\sigma}_j^y + \hat{\sigma}_i^z \hat{\sigma}_j^z),$$

where the sum is over nearest neighbours and  $\hat{\sigma}_i^{x,y,z}$  are Pauli matrices. In the square lattice case,  $\hat{H}$  has a  $C_{4v}$  symmetry<sup>1</sup>, whereas for the case of the triangular lattice,  $\hat{H}$  has a  $C_{2d}$  symmetry<sup>2</sup>. We also remark that in both cases, the ground state has zero magnetization [27], due to the  $U(1)$  symmetry of the Hamiltonian  $\hat{H}$ . We show that enforcing these symmetries in our ansatz allows to obtain a better accuracy without adding more parameters. More details about how to apply symmetries to the 2D cRNN ansatz can be found in the Appendix of Ref. [1].

#### 3.1 Heisenberg model on the square lattice

Since the Hamiltonian  $\hat{H}$  can be made stoquastic on the square lattice [28, 29], we set  $T_0 = 0$ , i.e., we do not use annealing in this case. The latter choice is motivated by the observations shown in Appendix. B.

First, we train our 2D cRNN wave function to find the ground state of the  $6 \times 6$  square lattice with and without applying the symmetries of the Hamiltonian  $\hat{H}$ . We find in Fig. 2(a) that increasing the

<sup>1</sup>A point group with four rotations and four mirror reflections (vertical, horizontal and diagonal).

<sup>2</sup>A point group with two rotations and two diagonal mirror reflections.




Figure 2: (a) A plot of the relative error  $\epsilon$  after applying different symmetries of the Heisenberg model on the square lattice with size  $6 \times 6$ . The relative error is computed with respect to the DMRG energy [14].  $C_4$  is the point group of four rotations. (b) A comparison of the energy per site obtained with our 2DRNN ansatz and PEPS [30], PixelCNN [13] and QMC [30] on the Heisenberg model on the square lattice with size  $10 \times 10$ .

number of symmetries encoded in the cRNN leads to a more accurate estimation of the ground state energy.

We also use our 2D cRNN wave function to estimate the ground state of the  $10 \times 10$  square lattice after applying  $U(1)$  and  $C_{4v}$  symmetries. The results are shown in Fig. 2(b), where we compare our estimates with projected-pair entangled states (PEPS) [30], Pixel CNN wave functions [13], as well as Quantum Monte Carlo (QMC) [30]. These results show that our ansatz is more accurate compared to PEPS and PixelCNN. We also find that QMC is the best among all the three methods.

### 3.2 Heisenberg model on the triangular lattice

Due to the frustrated, non-bipartite nature of the triangular lattice, the Hamiltonian  $\hat{H}$  can no longer be made stoquastic with a simple unitary transformation. Such a Hamiltonian can make the VMC optimization landscape rough and filled with local minima [26]. Here, we use annealing to overcome local minima and to obtain accurate estimates of ground state energies.

To demonstrate the idea of annealing, we target the ground state of the  $4 \times 4$  triangular lattice at different number of annealing steps  $N_{\text{annealing}}$ . The results in Fig. 3(a) show that starting from a non-zero  $T_0$  allows to obtain a lower value of the relative error as opposed to traditional VMC where the initial pseudo-temperature  $T_0 = 0$ . We also remark that a higher accuracy is obtained by increasing the number of annealing steps, which underlines the importance of adiabaticity in our scheme.

We now focus our attention on larger system sizes. Here we train our 2D cRNN wave function using annealing for a system size  $6 \times 6$ . We then optimize our ansatz at larger system sizes after increasing the lattice side length by 2 until we arrive at size  $16 \times 16$ . This step is done at zero pseudo-temperature and without reinitializing the parameters of the RNN. This idea was already proposed in the literature in Ref. [14]. The latter takes advantage of the translation invariance property encoded in the weight-sharing of the parameters of the RNN (see Eqs. 1, 2 and 3). The results in Fig. 3(b) show that for system sizes larger than  $14 \times 14$ , the variational energy obtained by our RNN ansatz is more accurate compared to the energy obtained by DMRG. Interestingly, our RNN ansatz uses less than 0.1% of the parameters of DMRG for system sizes larger than  $14 \times 14$ .

## Conclusion

In this paper, we have applied a modified variant of the two dimensional RNN wave function [1, 24] for the task of finding the ground state of the 2D Heisenberg model on the square and the triangular lattice. We have shown that applying symmetry and annealing helps to obtain highly accurate ground states. We have also demonstrated the power of our ansatz compared to DMRG for large system size on the triangular lattice, where frustration makes the problem more challenging. We believe that




Figure 3: (a) A scaling of the relative error  $\epsilon$  against the number of annealing steps  $N_{\text{annealing}}$  for the triangular Heisenberg model with size  $4 \times 4$ , ‘VMC’ corresponds to an initial pseudo-temperature  $T_0 = 0$  whereas for ‘VMC with annealing’, we start with  $T_0 = 1$ . (b) A plot of the energy difference per site between the 2DRNN and the DMRG. Negative values show that our ansatz is superior compared to DMRG for system sizes larger than  $14 \times 14$ .

more advanced versions of RNN wave functions along with the idea of applying symmetries and annealing can make our ansatz a highly competitive tool for studying quantum many-body systems.

**Note:** the hyperparameters used to produce the results of this paper are detailed in Appendix. C.

## Acknowledgments and Disclosure of Funding

Our RNN implementation is based on Tensorflow [31] and NumPy [32]. Computer simulations were made possible thanks to the Vector Institute computing cluster and the Compute Canada cluster. J.C. acknowledges support from Natural Sciences and Engineering Research Council of Canada (NSERC), the Shared Hierarchical Academic Research Computing Network (SHARCNET), Compute Canada, Google Quantum Research Award, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade.

## Broader Impact

The scope of this paper is within quantum many-body physics, which aims to understand the behaviour of matter based on simple rules at the microscopic scale. In the best case, our method could enhance the state-of-the-art algorithms and enable a better understanding of unknown quantum systems.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
  - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) Appendix. A provides more details about the 2DRNN cell we use in our study. We also provide more details in Appendix. C about the training process to facilitate the reproducibility of our results. Making our code publicly available is still a work in progress.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) see Appendix. C.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) the error bars correspond to a one standard deviation. They are computed after obtaining a large number of samples from the RNN using autoregressive sampling.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) We provide the type of GPUs we used to produce our results in Appendix. C. However, the total amount of compute was not included due to the lack of data.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
  - (b) Did you mention the license of the assets? [\[N/A\]](#)
  - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[N/A\]](#)
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)




Figure 4: (a) A plot of the energy variance per spin  $\sigma^2$  against the number of gradient descent steps, for both a 2D non-gated RNN and a 2D gated RNN. Here we choose the Heisenberg model on a square lattice with size  $6 \times 6$  as a test bed. (b) A scaling of the relative error against the number of annealing steps  $N_{\text{annealing}}$  for the square Heisenberg model with size  $4 \times 4$ , ‘VMC’ corresponds to an initial pseudo-temperature  $T_0 = 0$ , whereas for ‘VMC with annealing’ we start with  $T_0 = 1$ .

## A Two dimensional gated recurrent neural networks

In this section, we show how we incorporate the gating mechanism in our 2D cRNN wave function ansatz. The 2D RNN cell we use is based on the following recursion relations:

$$\tilde{\mathbf{h}}_{i,j} = \tanh([\boldsymbol{\sigma}_{i-1,j}; \boldsymbol{\sigma}_{i,j-1}]T[\mathbf{h}_{i-1,j}; \mathbf{h}_{i,j-1}] + \mathbf{b}), \quad (6)$$

$$\mathbf{u}_{i,j} = \text{sigmoid}([\boldsymbol{\sigma}_{i-1,j}; \boldsymbol{\sigma}_{i,j-1}]T_g[\mathbf{h}_{i-1,j}; \mathbf{h}_{i,j-1}] + \mathbf{b}_g), \quad (7)$$

$$\mathbf{h}_{i,j} = \mathbf{u}_{i,j} \odot \tilde{\mathbf{h}}_{i,j} + (1 - \mathbf{u}_{i,j}) \odot (W[\mathbf{h}_{i-1,j}; \mathbf{h}_{i,j-1}]). \quad (8)$$

Here we obtain the state  $\mathbf{h}_{i,j}$  through a combination of the neighbouring hidden states  $\mathbf{h}_{i-1,j}, \mathbf{h}_{i,j-1}$  and the candidate hidden state  $\tilde{\mathbf{h}}_{i,j}$ . The update gate  $\mathbf{u}_{i,j}$  decides how much the candidate hidden state  $\tilde{\mathbf{h}}_{i,j}$  will be modified. The latter combination allows to circumvent some of the limitations related to the vanishing gradient problem [33, 34]. Note that the size of the hidden state  $\mathbf{h}_{i,j}$  that we denote as  $d_h$  is a hyperparameter that we choose before optimizing the parameters of our ansatz. The tensors  $T, T_g$ , the weight matrix  $W$  and the biases  $\mathbf{b}, \mathbf{b}_g$  are variational parameters of our RNN ansatz in addition to the parameters in Eqs. (2) and (3).

To show the advantage of a 2D gated RNN over the 2D non-gated RNN, we train them using the VMC scheme to find the ground state of the Heisenberg model on the square lattice with size  $6 \times 6$ . We find that the gated RNN allows to obtain a more accurate energy. More specifically, we plot the energy variance per spin  $\sigma^2 = (\langle \hat{H}^2 \rangle - \langle \hat{H} \rangle)^2 / N$  in Fig. 4(a). This quantity is a good indicator for the proximity of the ansatz to the ground state provided that local minima are avoided [35, 36, 37]. We observe that the gated RNN can get about an order of magnitude lower  $\sigma^2$  compared to the non-gated RNN. The latter results demonstrate the advantage of adding the gating mechanism to our wave function ansatz.

## B Annealing on the square lattice

In the results section, we have used  $T_0 = 0$  on the square lattice which is equivalent to not doing annealing. This is motivated by the observation that the Heisenberg model on the square lattice is not frustrated and can be made stoquastic using the Marshall sign rule [28, 29]. To further corroborate this observation, we compare VMC with and without annealing on the square lattice for a system size  $4 \times 4$  and report our results in Fig. 4(b). The latter illustrates that annealing does not show an advantage for large  $N_{\text{annealing}}$  as opposed to the case of the triangular lattice demonstrated in Fig. 3(a). This result suggests that in the absence of frustration, there is no need to supplement VMC with annealing.

## C Hyperparameters

In this section, we summarize the hyperparameters used to produce the main results of this paper in Tab. 1. We note that the training of our ansatz was performed using P100 GPUs.

In order to produce the results of Fig. 3(b), we use  $d_h = 300$  and  $N_s = 100$  samples for training. We have first performed annealing on the system size  $6 \times 6$  with  $N_{\text{warmup}} = 1000$  and  $N_{\text{annealing}} = 10000$  and an initial pseudo-temperature  $T_0 = 0.25$ , while using a fixed learning rate  $\eta = 5 \times 10^{-4}$ . In this initial phase, we only apply the U(1) symmetry. In the next phase, we perform an additional 25000 gradients steps at zero pseudo-temperature and add an additional 25000 gradients steps after applying  $C_{2d}$  symmetry. During this convergence phase, the learning rate is decayed as  $\eta = 5 \times 10^{-5}/(1 + t/2000)$ , where  $t$  corresponds to the current number of convergence steps. We then increase the system size to  $8 \times 8$ , while keeping our ansatz parameters fixed, applying  $C_{2d}$  symmetry and while using zero pseudo-temperature. In this phase, the learning rate is fixed to  $\eta = 10^{-5}$ . After this step, we train our RNN ansatz until convergence. We repeat the same procedure for system sizes  $10 \times 10$ ,  $12 \times 12$ ,  $14 \times 14$  and  $16 \times 16$ . We note that for  $8 \times 8$ , we add 40000 gradient steps. For  $10 \times 10$ , we continue training with 20000 gradient steps. For  $12 \times 12$ , we converge using 10000 training steps. For  $14 \times 14$ , we continue training with 5000 gradient steps. Finally, for  $16 \times 16$ , we add 2000 convergence steps.

To produce the DMRG results in Fig. 3(b), we used a bond dimension  $D = 4000$  for sizes  $6 \times 6$ ,  $8 \times 8$  and  $10 \times 10$ . For  $12 \times 12$ , we used  $D = 3000$  and for the sizes  $14 \times 14$  and  $16 \times 16$ , we used  $D = 2000$  since we were not able to obtain the DMRG energy at  $D = 4000$  with a limit of 100 GB memory allocation. The DMRG calculations were run using ITensor [39].

We finally note that the Marshall sign [28, 29] is applied on top of our cRNN wave function on the square lattice and on the triangular lattice during all our numerical experiments to speed up convergence. For DMRG, we observed that applying the Marshall sign does not affect the accuracy.

## References

- [1] Mohamed Hibat-Allah, Martin Ganahl, Lauren E. Hayward, Roger G. Melko, and Juan Carrasquilla. Recurrent neural network wave functions. *Physical Review Research*, 2(2), Jun 2020. ISSN 2643-1564. doi: 10.1103/physrevresearch.2.023358. URL <http://dx.doi.org/10.1103/PhysRevResearch.2.023358>.
- [2] J. Androsiuk, L. Kulak, and K. Sienicki. Neural network solution of the Schrödinger equation for a two-dimensional harmonic oscillator. *Chemical Physics*, 173(3):377–383, 1993. ISSN 0301-0104. doi: 10.1016/0301-0104(93)80153-Z.
- [3] Artificial neural network methods in quantum mechanics. *Computer Physics Communications*, 104(1):1 – 14, 1997. ISSN 0010-4655. doi: [https://doi.org/10.1016/S0010-4655\(97\)00054-4](https://doi.org/10.1016/S0010-4655(97)00054-4). URL <http://www.sciencedirect.com/science/article/pii/S0010465597000544>.
- [4] Juan Carrasquilla and Roger Melko. Machine learning phases of matter. *Nature Physics*, 14: 431–434, 2017. doi: <https://doi.org/10.1038/nphys4035>.
- [5] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, Feb 2017. ISSN 1095-9203. doi: 10.1126/science.aag2302. URL <http://dx.doi.org/10.1126/science.aag2302>.
- [6] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. Neural-network quantum state tomography. *Nature Physics*, 14:447–450, 2018. doi: <https://doi.org/10.1038/s41567-018-0048-5>.
- [7] Juan Carrasquilla, Giacomo Torlai, Roger G. Melko, and Leandro Aolita. Reconstructing quantum states with generative models. *Nature Machine Intelligence*, 1(3):155–161, Mar 2019. ISSN 2522-5839. doi: 10.1038/s42256-019-0028-1. URL <http://dx.doi.org/10.1038/s42256-019-0028-1>.
- [8] Zi Cai and Jinguo Liu. Approximating quantum many-body wave functions using artificial neural networks. *Phys. Rev. B*, 97:035116, Jan 2018. doi: 10.1103/PhysRevB.97.035116. URL <https://link.aps.org/doi/10.1103/PhysRevB.97.035116>.



- [9] Di Luo and Bryan K. Clark. Backflow transformations via neural networks for quantum many-body wave functions. *Phys. Rev. Lett.*, 122:226401, Jun 2019. doi: 10.1103/PhysRevLett.122.226401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.122.226401>.
- [10] David Pfau, James S. Spencer, Alexander G. D. G. Matthews, and W. M. C. Foulkes. Ab initio solution of the many-electron schrödinger equation with deep neural networks. *Physical Review Research*, 2(3), Sep 2020. ISSN 2643-1564. doi: 10.1103/physrevresearch.2.033429. URL <http://dx.doi.org/10.1103/PhysRevResearch.2.033429>.
- [11] Jan Hermann, Zeno Schätzle, and Frank Noé. Deep-neural-network solution of the electronic schrödinger equation. *Nature Chemistry*, 12(10):891–897, Sep 2020. ISSN 1755-4349. doi: 10.1038/s41557-020-0544-y. URL <http://dx.doi.org/10.1038/s41557-020-0544-y>.
- [12] Kenny Choo, Antonio Mezzacapo, and Giuseppe Carleo. Fermionic neural-network states for ab-initio electronic structure. *Nature Communications*, 11:2368, 2020. doi: <https://doi.org/10.1038/s41467-020-15724-9>.
- [13] 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), Jan 2020. ISSN 1079-7114. doi: 10.1103/physrevlett.124.020503. URL <http://dx.doi.org/10.1103/PhysRevLett.124.020503>.
- [14] Christopher Roth. Iterative retraining of quantum spin models using recurrent neural networks, 2020.
- [15] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Physical Review Letters*, 94(17), May 2005. ISSN 1079-7114. doi: 10.1103/physrevlett.94.170201. URL <http://dx.doi.org/10.1103/PhysRevLett.94.170201>.
- [16] Tom Westerhout, Nikita Astrakhantsev, Konstantin S. Tikhonov, Mikhail I. Katsnelson, and Andrey A. Bagrov. Generalization properties of neural network approximations to frustrated magnet ground states. *Nature Communications*, 11(1):1593, March 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-15402-w.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735.
- [18] Alex Graves. Supervised sequence labelling. In *Supervised sequence labelling with recurrent neural networks*, pages 5–13. Springer, 2012. URL <https://www.cs.toronto.edu/~graves/preprint.pdf>.
- [19] Zachary C. Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning, 2015.
- [20] Corneel Casert, Tom Vieijra, Stephen Whitelam, and Isaac Tamblyn. Dynamical large deviations of two-dimensional kinetically constrained models using a neural-network state ansatz, 2020.
- [21] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, Nov 1992. doi: 10.1103/PhysRevLett.69.2863. URL <https://link.aps.org/doi/10.1103/PhysRevLett.69.2863>.
- [22] U. Schollwöck. The density-matrix renormalization group. *Reviews of Modern Physics*, 77(1):259–315, Apr 2005. ISSN 1539-0756. doi: 10.1103/revmodphys.77.259. URL <http://dx.doi.org/10.1103/RevModPhys.77.259>.
- [23] S. Bengio and Y. Bengio. Taking on the curse of dimensionality in joint distributions using neural networks. *IEEE Transactions on Neural Networks*, 11(3):550–557, May 2000. ISSN 1045-9227. doi: 10.1109/72.846725.
- [24] Mohamed Hibat-Allah, Estelle M. Inack, Roeland Wiersema, Roger G. Melko, and Juan Carrasquilla. Variational neural annealing. *Nature Machine Intelligence*, October 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00401-3. URL <https://doi.org/10.1038/s42256-021-00401-3>.
- [25] Yusuke Nomura. Helping restricted boltzmann machines with quantum-state representation by restoring symmetry. *Journal of Physics: Condensed Matter*, 33(17):174003, Apr 2021. ISSN 1361-648X. doi: 10.1088/1361-648x/abe268. URL <http://dx.doi.org/10.1088/1361-648x/abe268>.

- [26] Marin Bukov, Markus Schmitt, and Maxime Dupont. Learning the ground state of a non-stoquastic quantum hamiltonian in a rugged neural network landscape. *SciPost Physics*, 10(6), Jun 2021. ISSN 2542-4653. doi: 10.21468/scipostphys.10.6.147. URL <http://dx.doi.org/10.21468/SciPostPhys.10.6.147>.
- [27] Elliott Lieb and Daniel Mattis. Ordering energy levels of interacting spin systems. *Journal of Mathematical Physics*, 3(4):749–751, 1962. doi: 10.1063/1.1724276.
- [28] W Marshall. Antiferromagnetism. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 232(1188):48–68, oct 1955. ISSN 2053-9169. doi: 10.1098/rspa.1955.0200. URL <http://www.royalsocietypublishing.org/doi/10.1098/rspa.1955.0200>.
- [29] Luca Capriotti. Quantum effects and broken symmetries in frustrated antiferromagnets. *International Journal of Modern Physics B*, 15(12):1799–1842, May 2001. ISSN 1793-6578. doi: 10.1142/S0217979201004605. URL <http://dx.doi.org/10.1142/S0217979201004605>.
- [30] Wen-Yuan Liu, Shao-Jun Dong, Yong-Jian Han, Guang-Can Guo, and Lixin He. Gradient optimization of finite projected entangled pair states. *Physical Review B*, 95(19), May 2017. ISSN 2469-9969. doi: 10.1103/physrevb.95.195154. URL <http://dx.doi.org/10.1103/PhysRevB.95.195154>.
- [31] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from [tensorflow.org](https://www.tensorflow.org/).
- [32] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with numpy. *Nature*, 585(7825):357–362, Sep 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [33] Guo-Bing Zhou, Jianxin Wu, Chen-Lin Zhang, and Zhi-Hua Zhou. Minimal gated unit for recurrent neural networks, 2016.
- [34] Huitao Shen. Mutual information scaling and expressive power of sequence models, 2019.
- [35] Claudius Gros. Criterion for a good variational wave function. *Phys. Rev. B*, 42:6835–6838, Oct 1990. doi: 10.1103/PhysRevB.42.6835. URL <https://link.aps.org/doi/10.1103/PhysRevB.42.6835>.
- [36] Roland Assaraf and Michel Caffarel. Zero-variance zero-bias principle for observables in quantum monte carlo: Application to forces. *The Journal of Chemical Physics*, 119(20):10536–10552, 2003. doi: 10.1063/1.1621615. URL <https://doi.org/10.1063/1.1621615>.
- [37] F. Becca and S. Sorella. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press, 2017. doi: 10.1017/9781316417041.
- [38] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [39] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. The ITensor software library for tensor network calculations, 2020.

Figures	Parameter	Value
Fig. 2(a)	Architecture Number of memory units Number of samples Learning rate Number of training steps	2D Tensorized Gated cRNN wave function $d_h = 300$ $N_s = 100$ $\eta = 5 \times 10^{-4} \times (1 + (t/5000))^{-1}$ 100000
Fig. 2(b)	Architecture Number of memory units Number of samples Learning rate Number of training steps Applied symmetries	2D Tensorized Gated cRNN wave function $d_h = 200$ $N_s = 100$ $\eta = 5 \times 10^{-4} \times (1 + (t/5000))^{-1}$ 150000 $U(1)$ and $C_{4v}$
Figs. 3(a) and 4(b)	Architecture Number of memory units Number of samples Learning rate Number of warmup steps Applied symmetries	2D Tensorized Gated cRNN wave function $d_h = 100$ $N_s = 100$ $\eta = 10^{-4}$ $N_{\text{warmup}} = 1000$ $U(1), C_{4v}$ and spin parity
Fig. 4(a)	Number of memory units Number of samples Learning rate Applied symmetries	$d_h = 300$ $N_s = 100$ $\eta = 5 \times 10^{-4} \times (1 + (t/5000))^{-1}$ $U(1)$ and $C_{4v}$

Table 1: Hyperparameters used to obtain the results reported in this paper. We note that we use the Adam optimizer to perform gradient descent steps [38].