

SYMMETRY-ENHANCED NEURAL QUANTUM STATES FOR THE TRANSVERSE FIELD ISING MODEL

Samak Boonpan and Weerachai Sarakorn*

Received: July 03, 2025; Revised: September 02, 2025; Accepted: September 29, 2025

Abstract

The quantum many-body problem presents a significant computational challenge due to the exponential scaling of the Hilbert space. In this work, we address this problem by applying Neural Quantum States (NQS) within a Variational Monte Carlo (VMC) framework to determine the ground state of the Transverse Field Ising Model (TFIM) in one dimension. We employ neural network architectures tailored to the system's dimensionality: a Feed-Forward Network for the 1D chain. A key feature of our approach is the explicit enforcement of the global Z_2 spin-flip symmetry of the Hamiltonian onto the NQS wavefunction, which restricts the search to the correct physical subspace. We benchmark our results against exact diagonalization for systems of various sizes and model parameters. The calculated ground state energies show excellent agreement with the exact values, with relative errors falling in the range of 0.4% - 0.7%. This study demonstrates that tailored, symmetric NQS are a powerful, accurate, and flexible tool for investigating quantum many-body systems.

Keywords: Neural Quantum States; Variational Monte Carlo; Transverse Field Ising Model; Spin-flip Symmetry; Quantum Many-Body Systems

Introduction

The study of quantum many-body systems is a cornerstone of condensed matter physics (Altland and Simon, 2010). A primary challenge in this field is the curse of dimensionality, where the size of the Hilbert space grows exponentially with the number of particles, N . For a system of N spin-1/2 particles, the state vector requires 2^N complex numbers for a complete description. This exponential scaling makes the direct solution of the Schrödinger equation computationally intractable for all but the smallest systems (Feynman, 1982).

A canonical model for studying many-body phenomena is the Transverse Field Ising Model (TFIM) (Pfeuty, 1970; Stinchcombe, 1973). The TFIM is significant because it is one of the simplest models to exhibit a quantum phase transition (Sachdev, 2011; Suzuki *et al.*, 2013). The competition between the classical Ising term and the quantum transverse field term (proportional to external field) drives the system from an ordered ferromagnetic phase to a disordered paramagnetic phase at a critical value (Dutta *et al.*, 2015). This

Department of Mathematics, Faculty of Science, Khon Kaen University, Khon Kaen, Thailand. E-mail: wsarakorn@kku.ac.th; samak_b@kkumail.com

* Corresponding author

DOI: <https://doi.org/10.55766/sujst10600>

Suranaree J. Sci. Technol. 33(2):030377(1-8)

makes it an essential benchmark for testing new numerical methods.

To overcome the limitations of analytical solutions, various numerical techniques have been developed. Exact Diagonalization (ED) provides the complete eigenspectrum of the Hamiltonian but is severely limited by memory, typically to systems of around 40 spins (Sandvik, 2010). Other methods, such as the Density Matrix Renormalization Group (DMRG), are highly effective for one-dimensional systems but are less straightforward to apply in higher dimensions (White, 1992; Schollwöck, 2011). Therefore, the development of new, scalable numerical approaches remains a crucial area of research.

A promising new direction was introduced by Carleo and Troyer in 2017, who proposed the use of artificial neural networks as a variational ansatz for the many-body wavefunction, known as a Neural Quantum State (NQS) (Carleo and Troyer, 2017). The core idea is to leverage the high expressivity of neural networks, particularly Restricted Boltzmann Machines (RBMs) in the original work (Molensky, 1986), to efficiently represent the complete of the entanglement structure of quantum states (Deng *et al.*, 2017). The network's parameters are optimized variationally using the Variational Monte Carlo (VMC) algorithm to find the state that minimizes the energy expectation value.

In this work, we systematically investigate the performance of NQS for finding the ground state of the TFIM in one dimension. Our approach has two main features. First, we use network architectures that are well-suited to the problem of geometry: a simple Feed-Forward Network (FFN) for 1D systems to exploit spatial locality (Saito and Kato, 2018). Second, we explicitly incorporate the global Z_2 spin-flip symmetry of the TFIM Hamiltonian into the NQS wavefunction (Nomura *et al.*, 2017; Choo *et al.*, 2018). We validate our method by comparing the computed ground state energies against results from exact diagonalization.

This paper is structured as follows. Section 2 details the materials and methods used in our study. Section 3 presents and discusses our numerical results, providing a detailed analysis of their accuracy and the implications of our findings. Finally, Section 4 provides concluding remarks and potential avenues for future research.

Materials and Methods

This section outlines the theoretical and computational methods used in this work. We first define the target quantum system, the Transverse Field Ising Model. We then describe the Variational

Monte Carlo framework, the specific Neural Quantum State architectures, and the implementation of symmetry constraints.

The Transverse Field Ising Model (TFIM)

The TFIM is a fundamental model in quantum statistical mechanics. We study it on one-dimensional (1D) chains of length N , where we assume periodic boundary conditions. The Hamiltonian of the system is given by:

$$H = -g \sum_{(i,j)} \sigma_z^{(i)} \sigma_z^{(j)} - \eta \sum_i \sigma_x^{(i)}. \quad (1)$$

The first term is the Ising interaction between nearest-neighbor spins (i, j) , with g as the coupling constant. Throughout this work, and therefore all energies, including the transverse field η and the ground state energy (E_0), are reported in units of g . This term favors the alignment of spins in the z -direction. The second term represents a uniform magnetic field of strength applied in the x -direction. The operator $\sigma_x^{(i)}$ flips the spin at site i in the z -basis, introducing quantum fluctuations into the system.

The Variational Monte Carlo (VMC) Framework

Our goal is to find the ground state wavefunction Ψ_0 for the Hamiltonian in Eq.(1). We use the VMC method (McMillan, 1965; Ceperley *et al.*, 1977), which is based on the variational principle of quantum mechanics. The principle states that the expectation value of the energy for any trial wavefunction $|\Psi_\theta(\sigma)\rangle$ is always greater than or equal to the true ground state energy (Griffiths and Schroeter, 2018)

$$E_\theta = \frac{\langle \Psi_\theta | H | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle} \geq E_0. \quad (2)$$

Here, $|\Psi_\theta\rangle$ is a variational ansatz parameterized by a set of parameters θ . The objective is to minimize E_θ by adjusting θ . In our case, the parameters θ are the weights and biases of a neural network.

For the NQS ansatz, we represent the wavefunction $\Psi_\theta(\sigma) = \langle \sigma | \Psi_\theta \rangle$ for a given spin configuration $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ using a neural network. The network takes a configuration σ as input and outputs a single complex number. For stability, we typically have the network output the logarithm of the amplitude, $\text{Log } \Psi_\theta(\sigma)$. The expectation value E_θ involves a sum over all 2^N spin configurations, which is intractable. This sum is estimated using Monte Carlo integration. Instead, we rewrite E_θ as the expectation value of the local energy $E_{\text{local}}(\sigma)$ over the probability of spin

distribution $p(\sigma) = |\Psi_\theta(\sigma)|^2 / \sum_{\sigma'} |\Psi_\theta(\sigma')|^2$ as the Metropolis - Hastings algorithm (Metropolis *et al.*, 1953; Hastings, 1970). So we can write $E_\theta = \sum_{\sigma} p(\sigma) E_{loc}(\sigma)$ where

$$E_{loc}(\sigma) = \frac{1}{\Psi_\theta(\sigma)} \sum_{\sigma'} \langle \sigma | H | \sigma' \rangle \Psi_\theta(\sigma'). \quad (3)$$

This sum can be estimated by sampling a set of M configurations $\{\sigma_1, \sigma_2, \dots, \sigma_M\}$ from $p(\sigma)$ using the Metropolis-Hastings algorithm. The energy is then approximated by the sample mean:

$$E_\theta \approx \frac{1}{M} \sum_{k=1}^M E_{loc}(\sigma_k). \quad (4)$$

The network parameters θ are updated iteratively using gradient descent (Sorella, 2007). The gradient of the energy ($\nabla_\theta E_\theta$) is computed as:

$$\nabla_\theta E_\theta = 2 \langle (E_{loc} - \langle E_{loc} \rangle) \log \{ \Psi_\theta(\sigma) \} \rangle_{p(\theta)}, \quad (5)$$

where $\langle E_{loc} \rangle$ is the mean local energy and $\langle \dots \rangle_{p(\theta)}$ is weight with a probability of spin distribution. We use the Adam optimizer to perform this minimization eigenspectrum (Kingma and Ba, 2014).

Imposition of Global Z_2 Spin-Flip Symmetry

To approximate the ground state of the quantum system, a variational ansatz in the form of an NQS is employed. The chosen architecture for this ansatz is a Feed-Forward Network (FFN). The network is designed as a universal function approximator, $\Psi_\theta(\sigma)$ which maps a spin configuration of the system, $\sigma \in \{+1, -1\}^N$, to the logarithm of its corresponding wavefunction amplitude.

The Hamiltonian of the TFIM with periodic boundary conditions possesses a global Z_2 spin-flip symmetry. This is characterized by a symmetry operator, S , which commutes with the Hamiltonian, $[H, S] = 0$. A consequence of this commutation is that the energy eigenstates of H are also eigenstates of S . It is well-established that the ground state of this system, $|\Psi_0\rangle$ belongs to the symmetric sector, meaning it is an eigenstate of S with an eigenvalue of $+1$: $S|\Psi_0\rangle = (+1)|\Psi_0\rangle$. To ensure our variational search is confined to this correct physical subspace, we construct a new ansatz, Ψ_{Z_2} , that is manifestly symmetric by definition. This is achieved by employing a product-based construction that combines evaluations of the basis wavefunction, ψ_θ :

$$\Psi_{\theta, Z_2}(\sigma) = \psi_\theta(\sigma) \psi_\theta(-\sigma), \quad (6)$$

where $-\sigma$ is the configuration obtained by flipping all spins of σ . This construction inherently satisfies the symmetry requirement, as $\Psi_{\theta, Z_2}(\mp\sigma) = \psi_\theta(-\sigma) \psi_\theta(\sigma) = \psi_\theta(\sigma) \psi_\theta(-\sigma)$. For numerical stability in the practical implementation, we operate in logarithmic space, where the product in equation (7) becomes a summation:

$$\log \Psi_{\theta, Z_2}(\sigma) = \log \psi_\theta(\sigma) + \log \psi_\theta(-\sigma). \quad (7)$$

In this procedure, the FFN serves as the foundational unit for computing Ψ_θ . The symmetry principle is then applied as described to construct the final, physically-constrained wavefunction, $\Psi_{\theta, Z_2}(\sigma)$. This resulting ansatz is the function that is ultimately optimized by the VMC algorithm.

Experimental Setup

The choice of a neural network architecture for the NQS ansatz is not arbitrary. It is a critical decision that embeds physical priors into the model. An effective architecture should reflect the expected correlation structure of the system's ground state, making the variational search more efficient and accurate. We designed two distinct architectures to match the differing physical nature of a one dimensional systems.

In this system quantum systems, correlations can be complex. While gapped systems typically exhibit short-range correlations that decay exponentially with distance, systems near a quantum critical point can develop long-range, algebraically decaying correlations. A suitable NQS must be flexible enough to capture either behavior. For this reason, we employ a Feed-Forward Network (FFN), a highly general ansatz capable of representing arbitrary patterns of long-range correlations due to its all-to-all connectivity (Goodfellow *et al.*, 2016).

Concretely, the FFN architecture used for an N -spin chain in this work consists of an input layer of size N , a single hidden layer, and a final output layer with one neuron. The size of the hidden layer, N_{hidden} , is set to be proportional to the number of spins, a common practice in NQS studies. This relationship is governed by a hyperparameter called the density parameter, α defined as:

$$N_{\text{hidden}} = \alpha N. \quad (8)$$

In this work, we choose $\alpha = \{4, 2\}$. This value is widely reported in the literature to provide a good balance between the network's expressive capacity and its computational cost for models. It ensures the network has sufficient parameters to

accurately represent the ground states of the small systems benchmarked here, without being excessively complex.

The FFN architecture used in this study shown in Figure 1: consists of an input layer, two hidden layers ($L=2$), and a linear output layer. The component of our structure are comprised of:




Figure 1. The Feed-Forward Network (FFN) architecture used to represent the base wavefunction, Ψ_{base} . The network maps an N -dimensional spin configuration, σ , through two hidden layers (sizes $4N$ and $2N$, with Tanh activation) to a single linear output representing the $\Psi_{\theta}(\sigma)$.

Input Layer: Comprises N neurons, where N is the number of sites in the one-dimensional spin chain. Each neuron receives the value $\sigma_i \in \{+1, -1\}$ representing the spin state at site i .

Hidden Layers: The network employs a deep architecture whose dimensions are scaled relative to the system size N to maintain sufficient expressive capacity. The first hidden layer contains $4N$ neurons, and the second hidden layer contains $2N$ neurons. The hyperbolic tangent ($\tanh$) is utilized as the non-linear activation function for all hidden neurons.

Output Layer: Consists of a single linear neuron that performs a final affine transformation to produce a single scalar value, which represents the log-amplitude, $\Psi_{\theta}(\sigma)$.

The network's computation can be expressed mathematically. Defining $h_i^{(0)} = \sigma_i$ as the input, the output of a neuron k in a hidden layer l is given by:

$$h_k^{(l)} = \tanh\left(\sum_{j=1}^{N_{l-1}} W_{kj}^{(l)} h_j^{(l-1)} + b_k^{(l)}\right), \quad (9)$$

and the output layer

$$\Psi_{\theta}(\sigma) = \sum_{k=1}^{N_L} W_k^{(\text{out})} h_k^{(L)} + b_k^{(\text{out})}. \quad (10)$$

The set of all weights and biases, $\theta = \{W^{(l)}, b^{(l)}, W^{(\text{out})}, b^{(\text{out})}\}$, constitutes the trainable variational parameters. This function serves as the fundamental building block for the final, symmetrized ansatz.

Throughout all simulations, the coupling constants were fixed at $g=1.0$ for the ferromagnetic interaction term and $\eta=0.5$ for the transverse field. This choice of parameters places the system within the ferromagnetic phase while including significant quantum fluctuations. To investigate the scaling behavior of the method, we conducted experiments on systems of varying sizes, specifically $N \in \{8, 10, 14, 16\}$.

The hyperparameters used for training are specified as follows:

Optimizer: Parameter optimization was performed using the Adam optimizer.

Learning Rate: A fixed learning rate of $r = 2 \times 10^{-3}$ was used for all experiments.

Batch Size: The stochastic gradient was estimated using a batch size of $M = 128$ spin configurations for each optimization step.

VMC Sampling: The Metropolis-Hastings algorithm was employed for sampling configurations. To ensure decorrelated samples, the Markov chain was advanced by $2N$ single-spin-flip steps between each parameter update.

Training Epoch: Each training run was configured for a maximum of 300 epochs. An early stopping criterion was implemented to halt training if the energy expectation value did not improve by a tolerance of 10^{-8} over a patience window of 40 epochs.

The primary investigation involved a direct comparison between two model configurations for each system size N : a Non-Symmetrized Model and Symmetrized Model (Z_2): The ansatz constructed with the imposed Z_2 spin-flip symmetry

To ensure the statistical reliability of our findings, each experiment (for a given N and symmetry choice) was repeated 5 times, with each run initiated with a different pseudo-random number generator seed. The results presented in this paper represent the mean and standard deviation computed over these 5 independent runs. The accuracy of the VMC results was benchmarked against the exact ground state energy, E_0 . For the system sizes studied, E_0 was computed using the Lanczos algorithm via sparse exact diagonalization (ED), providing a high-precision reference value for comparison.

Results and Discussion

Following the experimental protocol detailed in this section, simulations were performed for both the non-symmetrized and Z_2 -symmetrized ansatzes across system sizes of $N = 8, 10, 14$, and 16 . The resulting data from these experiments systematically reveal the profound impact of imposing a physical symmetry on the variational search process. The analysis herein begins with a qualitative examination of the training dynamics, presented as convergence plots that depict the energy per site (E/N) as a function of the training epoch for each independent run. These plots provide a visual representation of the stability and efficiency of the optimization process. Subsequently, a rigorous quantitative comparison is conducted using key performance metrics, namely the final mean energy, the relative error with respect to the exact value, and the standard deviation, which have been collated for all experimental conditions. The primary finding that emerges from both the qualitative and quantitative data is a dramatic divergence in performance between the symmetrized and non-symmetrized models as the system size, N , increases. The convergence plots for each condition and the summary of quantitative results, which form the basis for the subsequent discussion, are presented accordingly.

The convergence plots, presented in Figure 2(a) - Figure 2 (h), provide a qualitative visualization of the training dynamics and stability of the VMC process under different conditions. At the smallest system size of $N=8$ (Figure 2(a) and Figure 2(b)), the difference in convergence stability between the non-symmetrized model (a) and the Z_2 -symmetrized model (b) is not pronounced. Both models exhibit relatively well-clustered convergence paths and are capable of reaching the vicinity of the exact energy. However, this trend changes dramatically as the system size increases to $N=10$ (c, d), $N=14$ (e, f), and $N=16$ (g, h).

For the non-symmetrized models (Figures c, e, g), the convergence behavior becomes increasingly erratic and unstable. The energy curves from the independent runs are highly dispersed, reflecting an optimization process that frequently becomes trapped in suboptimal local minima at significantly different energy levels. Some training runs fail to approach the true ground state energy, instead plateauing at much higher values.

In stark contrast, the Z_2 -symmetrized models (Figure 2(d) -Figure 2(h)) demonstrate remarkable convergence stability across all larger system sizes. The energy curves from all independent runs are tightly bundled, consistently converging to the same low-energy manifold near the exact ground state value. This indicates that the imposition of the physical symmetry acts as a powerful regularizer, guiding the search process within the correct physical sector and making the optimization robust to different initial conditions.

Therefore, the qualitative analysis of the convergence plots provides compelling visual evidence that corroborates the quantitative data, showing that the symmetrized models exhibit remarkably stable and tightly-clustered convergence paths. This confirms that the application of symmetry is a critical factor that guides the search process, ensuring that the training of Neural Quantum State models is both stable and reliable for larger quantum systems.

A quantitative analysis of the results, summarized in Table 1, provides unambiguous evidence of a stark performance divergence between the non-symmetrized and Z_2 symmetrized models as the system size, N , increases. In terms of accuracy, the symmetrized model performs consistently well across all studied system sizes.

Table 1. Summary of quantitative results comparing the non - symmetrized (-) and Z_2 symmetrized models for system sizes $N=8, 10, 14$, and 16 . All metrics are computed from 5 independent runs. STD denotes the standard deviation, and Rel. Err(%) is the relative error against the exact ground state energy

N sites	Symm.	STD (E/N)	Mean (E/N)	Max (E/N)	Mean Epoch	Exact (E/N)	Rel. Err(%)
8	-	1.57e-03	-1.0692	-1.0666	185.0	-1.0636	0.5241
8	Z_2	2.59e-03	-1.0709	-1.0682	157.6	-1.0636	0.6915
10	-	1.20e-01	-0.9713	-0.8216	204.4	-1.0635	8.6708
10	Z_2	1.81e-03	-1.0694	-1.0677	133.6	-1.0635	0.5551
14	-	8.53e-02	-1.0014	-0.8926	212.6	-1.0635	5.8357
14	Z_2	1.08e-03	-1.0681	-1.0662	172.4	-1.0635	0.4339
16	-	7.46e-02	-1.0063	-0.9076	172.2	-1.0635	5.3790
16	Z_2	1.79e-03	-1.0693	-1.0671	181.4	-1.06354	0.5429




Figure 2. VMC training convergence for the non-symmetrized (left column) and Z_2 -symmetrized (right column) FFN ansatzes at system sizes $N=8$ (a, b), $N=10$ (c, d), $N=14$ (e, f), and $N=16$ (g, h). Each plot shows the energy per site (E/N) versus epoch for five independent runs. The dashed line indicates the exact ground state energy. The stability of the symmetrized model is evident from the tight clustered convergence paths, in contrast to the high variance of the non-symmetrized model at larger N .

In stark contrast, the non-symmetrized model exhibits a pronounced degradation in accuracy, with its relative error escalating to over 5% for systems with $N \geq 10$, indicating its failure to approximate the true ground state in larger Hilbert spaces.

The results obtained in this study are consistent and comparable with foundational research in the field of Neural Quantum States (NQS). Specifically, when compared to the pioneering work of Carleo and Troyer (2017), which also utilized the 1D Transverse Field Ising Model (TFIM) as a primary testbed system, the relative error for the ground state that we calculated, ranging from 0.4% to 0.7%, is found to be highly accurate and consistent with previous work, which reported relative errors for similar systems in the range of 10^{-2} to 10^{-4} . This demonstrates that our chosen Feed-Forward Network (FFN) architecture is a highly effective alternative to the original Restricted Boltzmann Machine (RBM) architecture. Furthermore, our approach of enforcing the global Z_2 spin-flip symmetry directly aligns with the principles presented by Choo *et al.* (2018), who developed a general framework for constructing NQS wavefunctions with specific symmetries to study various states of quantum systems. Although their research was conducted on different models, such as the one-dimensional spin-1/2 Heisenberg model, their work provides an important validation of the principle that directly incorporating physical symmetry information into the network is a powerful method for constraining the search space and enhancing computational stability, as has been clearly demonstrated by our experimental results.

This disparity is even more pronounced when considering the stability and reliability of the training process. The standard deviation of the final energy $\text{STD}(E/N)$ for the non-symmetrized model is one to two orders of magnitude larger than that of its symmetrized counterpart for $N \geq 10$. This high variance, further evidenced by the wide range between the minimum and maximum energies obtained across independent runs, signifies an unreliable optimization process heavily dependent on initial conditions. Conversely, the extremely low standard deviation and tight energy range of the Z_2 - symmetrized model demonstrate its robustness and ability to reproducibly converge to a consistent, high-quality solution. Collectively, the quantitative data presented in Table 1 fully corroborates the qualitative trends observed in the convergence plots. The combination of high accuracy, low variance, and stable training dynamics decisively demonstrates the superiority of the symmetrized ansatz. This evidence underscores that imposing known physical symmetries is the key factor in achieving reliable and precise ground state

solutions with the NQS method for the systems studied.

Conclusion and Future Work

In this study, we have systematically investigated the impact of imposing global Z_2 spin-flip symmetry on a Feed-Forward Network-based Neural Quantum State ansatz for determining the ground state of the one-dimensional Transverse Field Ising Model. Our comparative analysis across a range of system sizes ($N=8$ to $N=16$) provides conclusive evidence that as system size increases, a non-symmetrized ansatz fails catastrophically in terms of both accuracy and training stability, highlighting how critical the role of symmetry becomes in ensuring a robust and accurate solution. We therefore conclude that the incorporation of known physical principles is not merely an incremental enhancement but an indispensable requirement for transforming NQS from a theoretical construct into a powerful and practical tool for computational physics.

The principles demonstrated in this one-dimensional study pave the way for several promising avenues for future research. A primary extension is the application of these concepts to two-dimensional quantum systems. For such problems, the FFN architecture would be replaced by a more suitable Convolutional Neural Network (CNN), whose inductive bias is well-suited for capturing local spatial correlations on a 2D lattice.

The investigation would involve studying the 2D TFIM on a square lattice and incorporating a richer set of symmetries beyond the Z_2 spin-flip. This includes the spatial symmetries of the lattice, such as translational symmetry, degree rotational (C_4) symmetry, and reflectional symmetries. By constructing a highly symmetric NQS ansatz that respects the full point group of the Hamiltonian, we hypothesize that the performance gains and necessity of symmetry will be even more significant than in the 1D case, given the more rapid growth in Hilbert space complexity. Further research could also explore more advanced network architectures, such as Graph Neural Networks or Transformers, and apply these symmetrized VMC techniques to other fundamental models in condensed matter physics, including the Heisenberg and Hubbard models.

Acknowledgement

The authors would like to express their sincere gratitude to the reviewers for their insightful comments and constructive suggestions. We are also grateful to the editor for the professional

handling and guidance throughout the review process.

References

- Altland, A., & Simons, B. D. (2010). Condensed matter field theory (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511789984>
- Carleo, G., & Troyer, M. (2017). Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325), 602-606. <https://doi.org/10.1126/science.aag2302>
- Ceperley, D., Chester, G. V., & Kalos, M. H. (1977). Monte Carlo simulation of a many-fermion system. *Physical Review B*, 16(8), 3081-3099. <https://doi.org/10.1103/PhysRevB.16.3081>
- Choo, K., Carleo, G., Regnault, N., & Neupert, T. (2018). Symmetries and many-body localization with neural-network quantum states. *Physical Review Letters*, 121(16), 167204. <https://doi.org/10.1103/PhysRevLett.121.167204>
- Cohen, T., & Welling, M. (2016). Group equivariant convolutional networks. In Proceedings of the 33rd International Conference on Machine Learning (ICML 2016) (pp. 2990-2999). PMLR.
- Deng, D. L., Li, X., & Das Sarma, S. (2017). Quantum entanglement in neural network states. *Physical Review X*, 7(2), 021021. <https://doi.org/10.1103/PhysRevX.7.021021>
- Dutta, A., Aeppli, G., Chakrabarti, B. K., Divakaran, U., Rosenbaum, T. F., & Sen, D. (2015). Quantum phase transitions in transverse field spin models: From statistical physics to quantum information. Cambridge University Press. <https://doi.org/10.1017/CBO9781107706057>
- Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21, 467-488. <https://doi.org/10.1007/BF02650179>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
- Griffiths, D. J., & Schroeter, D. F. (2018). Introduction to quantum mechanics (3rd ed.). Cambridge University Press.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv. <https://doi.org/10.48550/arXiv.1412.6980>
- Lanczos, C. (1950). An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45(4), 255-282. <https://doi.org/10.6028/jres.045.026>
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), 2278-2324. <https://doi.org/10.1109/5.726791>
- McMillan, W. L. (1965). Ground state of liquid He4. *Physical Review Journals Archive*, 138(2A), A442-A451. <https://doi.org/10.1103/PhysRev.138.A442>
- Mehta, P., Bukov, M., Wang, C. H., Day, A. G. R., Richardson, C., Fisher, C. K., & Schwab, D. J. (2019). A high-bias, low-variance introduction to machine learning for physicists. *Physics Reports*, 810, 1-124. <https://doi.org/10.1016/j.physrep.2019.03.001>
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6), 1087-1092. <https://doi.org/10.1063/1.1699114>
- Nomura, Y., Darmawan, A. S., Yamaji, Y., & Imada, M. (2017). Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Physical Review B*, 96(20), 205152. <https://doi.org/10.1103/PhysRevB.96.205152>
- Pfeuty, P. (1970). The one-dimensional Ising model with a transverse field. *Annals of Physics*, 57(1), 79-90. [https://doi.org/10.1016/0003-4916\(70\)90270-8](https://doi.org/10.1016/0003-4916(70)90270-8)
- Sachdev, S. (2011). Quantum phase transitions (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511973765>
- Saito, H., & Kato, M. (2018). Machine learning technique to find quantum many-body ground states of bosons on a lattice. *Journal of the Physical Society of Japan*, 87(1), 014001. <https://doi.org/10.7566/JPSJ.87.014001>
- Sandvik, A. W. (2010). Computational studies of quantum spin systems. *AIP Conference Proceedings*, 1297(1), 135-338. <https://doi.org/10.1063/1.3518900>
- Schollwöck, U. (2011). The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1), 96-192. <https://doi.org/10.1016/j.aop.2010.09.012>
- Sharir, O., Levine, Y., Wies, N., Carleo, G., & Shashua, A. (2020). Deep autoregressive models for the efficient variational simulation of many-body quantum systems. *Physical Review Letters*, 124(2), 020503. <https://doi.org/10.1103/PhysRevLett.124.020503>
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 194-281). MIT Press.
- Sorella, S. (2007). Weak-binding electrons in a trapped Fermi gas. *Physical Review B*, 75(18), 180519(R).
- Stinchcombe, R. B. (1973). Ising model in a transverse field. *Journal of Physics C: Solid State Physics*, 6(15), 2459. <https://doi.org/10.1088/0022-3719/6/15/009>
- Suzuki, S., Inoue, J., & Chakrabarti, B. K. (2013). Quantum Ising phases and transitions in transverse Ising models (Lecture Notes in Physics, Vol. 862). Springer. <https://doi.org/10.1007/978-3-642-33039-1>
- White, S. R. (1992). Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69(19), 2863-2866. <https://doi.org/10.1103/PhysRevLett.69.2863>