

Approximating quantum many-body wave functions using artificial neural networks

Zi Cai*

*Key Laboratory of Artificial Structures and Quantum Control, Department of Physics and Astronomy,
Shanghai Jiao Tong University, Shanghai 200240, People's Republic of China*

Jinguo Liu

Beijing National Lab for Condensed Matter Physics and Institute of Physics, Chinese Academy of Sciences, Beijing 100190, China

(Received 25 April 2017; revised manuscript received 19 November 2017; published 9 January 2018)

In this paper, we demonstrate the expressibility of artificial neural networks (ANNs) in quantum many-body physics by showing that a feed-forward neural network with a small number of hidden layers can be trained to approximate with high precision the ground states of some notable quantum many-body systems. We consider the one-dimensional free bosons and fermions, spinless fermions on a square lattice away from half-filling, as well as frustrated quantum magnetism with a rapidly oscillating ground-state characteristic function. In the latter case, an ANN with a standard architecture fails, while that with a slightly modified one successfully learns the frustration-induced complex sign rule in the ground state and approximates the ground states with high precisions. As an example of practical use of our method, we also perform the variational method to explore the ground state of an antiferromagnetic J_1 - J_2 Heisenberg model.

DOI: [10.1103/PhysRevB.97.035116](https://doi.org/10.1103/PhysRevB.97.035116)**I. INTRODUCTION**

A central challenge in quantum many-body physics is developing efficient numerical tools for strongly correlated systems, whose Hilbert space dimensionality grows exponentially with the system size, so does the information required for characterizing a generic state of the system. However, for many physical systems of practical interest, the ground states may have a simplified structure, thus, it can be appropriately approximated using an exponentially smaller number of parameters than that required for characterizing generic states. Typical examples include low-dimensional strongly correlated systems, whose ground states can be represented in terms of matrix product states by taking advantage of limited entanglement entropy in the ground states [1–3]. For higher-dimensional systems, a different routine involving stochastic sampling applies for certain types of strongly correlated systems with positive-definite ground states, with quantum Monte Carlo (QMC) algorithm providing a well-controlled method to evaluate the physical quantities of interest based on an exponentially small fraction of all possible configurations [4–6]. In spite of the remarkable achievement of these numerical methods, developing a general strategy to represent typical many-body wave functions that bypasses the exponential complexity remains a formidable, if not impossible, task and is a question of principal interest in the condensed matter community.

In general, a many-body wave function can be expanded in terms of a set of orthogonal bases (e.g., the Fock basis) and can be fully characterized by a function where one feeds in a basis and gets the output the set of corresponding coefficients in the wave function. Consider a spin lattice system as an example; in

this case, an arbitrary wave function can be expanded as $\Psi = \sum_{\sigma} C[\sigma]|\sigma\rangle$, where $|\sigma\rangle = |\sigma_1 \dots \sigma_L\rangle$ are $\{S_i^z\}$ eigenstates spanning the Hilbert space of the spin configurations. Therefore, the task amounts to efficiently approximating the function $C[\sigma]$ using an exponentially smaller number of parameters than the Hilbert space dimensionality $\sim 2^L$. Among the existing modern techniques used for approximating functions, artificial neural networks (ANNs), as a powerful tool for data fitting and feature extraction, have not only achieved remarkable successes in machine learning and cognitive science fields in the past decades [7,8], but have also recently attracted considerable attention of researchers in the condensed matter community. Applying machine learning methods to problems in condensed matter physics is not only interesting in its own right [9–12], but it may also potentially provide new ideas and have practical applications for solving complex physics problems, such as identifying classical and quantum phases of matter and locating the phase transition points [13–21], categorizing and designing materials [22–24], improving existing numerical techniques [25–29], and even developing new methods in the quantum many-body physics [30,31]. Owing to its tremendous capability in function approximations, ANNs can also be considered as novel representations of many-body wave functions [30,32–35], e.g., in a seminal work, Carleo *et al.* proposed a new kind of variational wave functions using the restricted Boltzmann machine [30].

Even though it can be proven mathematically that ANNs can in principle approximate any smooth function to any accuracy [36–38], what matters in practice is the efficiency of the method: the amount of resources an ANN needs to approximate a given multivariable function. In this paper, we will demonstrate the expressibility of neural networks in approximating and characterizing quantum many-body wave functions using some notable examples of physical interest.

*zcai@sjtu.edu.cn

For a wave function with completely random coefficients, the information encoded in it cannot be compressed; thus, an exponentially large number of parameters are needed. The goal of this paper is to express the ground states of several notable many-body systems in terms of neural networks of feasible size and a small number of hidden layers, and, most importantly, networks learn in a reasonable time. The success of this approach relies on the specialty of the ground state compared to a generic eigenstate, where the underlying physical laws encoded in the ground state’s wave function can be extracted by the neural networks through a large amount of training.

In the following, we show which ground-state wave functions can be efficiently expressed by a simple neural network, and for those who cannot, how the ANN’s architecture should be modified to achieve this goal. The scaling of the computation resources with the system size is also investigated. Special attention is devoted to wave functions with the “sign problem,” where the function $C[\sigma]$ may alter its sign even for a local change in the input $[\sigma]$, and thus cannot be considered a “smooth” function. The sign problem not only hinders the application of QMC methods, but also make it more difficult for a simple ANN to learn a wave function. However, as we have shown, this problem can be circumvented by dividing the task into two subtasks and by designing ANNs with corresponding architectures. This is a typical example that illustrates the neural network’s ability of extracting the physics laws, even for those too complex to be written explicitly.

The rest of the paper is organized as follows: First, we introduce the structure of the ANN, then use it to investigate some notable examples, including one- (1D) and two-dimensional (2D) free bosons and fermions, whose exact ground states are compared to those predicted by the ANN. To determine whether the ANN approach can work for large systems, we adopt the importance sampling algorithm to calculate physical quantities instead of wave functions, and compare them to the exact ones. Then, we attack the most difficult part: approximating the ground state of a frustrated quantum magnetism whose characteristic function can dramatically change its sign, thus being very different from a smooth function, which makes it extremely difficult for a regular ANN to approximate. In spite of this, we find that a slight modification of neurons in the ANN allows to capture the sign rule of the frustrated quantum magnetism, even at the phase transition point. Finally, we discuss the practical application of this method based on the variational method.

II. METHODS

Before we proceed further to discuss specific examples, let us describe the details of the ANNs and the optimization methods we will use. We consider a fully connected feed-forward neural network consisting of an interconnected group of nodes (neurons) with a stacked layered structure, and its expressibility is encoded in sets of adaptive weights of connections between neurons in adjacent layers. As shown in Fig. 1, we consider a four-layer ANN, with two hidden layers (each containing N_h neurons) that are sandwiched between the input layer [accepting the Fock basis ($|\mathbf{n}\rangle$ or $|\sigma\rangle$)] and the output layer that output the corresponding coefficient predicted by the ANN ($C_P[\sigma]$). A neuron can be considered to be an




FIG. 1. (Left) The structure of the feed-forward ANN we used to approximate quantum many-body ground-state wave functions; (right) three types of neurons with different activation functions.

elemental processor:

$$o_i^{[n+1]} = f \left(\sum_j W_{ij}^{[n]} o_j^{[n]} + b_i^{[n]} \right), \tag{1}$$

where $o_j^{[n]}$ is the output of the j th neuron of the n th hidden layer, $W_{ij}^{[n]}$ denotes the connection weight between the n th and $(n - 1)$ th hidden layers, and $b_i^{[n]}$ is the bias in this neuron. The activation function $f(x)$ can be any smooth nonlinear function. However, as we show below, choosing a proper nonlinear function may significantly increase the ANN’s efficiency for approximating certain target functions. We introduce a fidelity function

$$F = 1 - \left| \sum_{\sigma} (C_T^*[\sigma] C_P[\sigma]) \right| \tag{2}$$

to measure the difference between the target function $C_T[\sigma]$ and the one predicted by the ANN $C_P[\sigma]$. Here, we should point out that to calculate this quantity, the target functions are known in advance. We will discuss how to generalize the current method to explore new quantum many-body systems with previously unknown ground states in Sec. V. The problem now reduces to an optimization problem with the goal of finding the minimum of F in the landscape of ANN parameters (weights of the connection and bias, which are denoted as $\{W\}$ in the following). In the following, we adopt the ANN construction methods and the optimization techniques that are readily available in the machine learning libraries TensorFlow [39], with the training time measured in the units of T_0 , corresponding to the period of a single optimization iteration that depends on the details of the ANN; more details about the initialization and training can be found in the Appendixes.

III. FREE BOSONS/FERMIONS SYSTEMS

The first class of wave functions that we investigate are the ground states of the simplest many-body systems composed of free bosons or fermions. In spite of their extreme simplicity, these wave functions are perfect touchstones to test the expressibility of the ANN, for two reasons: first, the well-known analytic forms of these wave functions serve not only as a target function during the training process of a neural network,

but also a tester of the accuracy of its predictions; second, in some cases, e.g., the ground state of 2D free fermions away from half-filling, the wave function is simple but not trivial, in the sense that it is difficult to be characterized using the existing numerical methods such as the matrix product state (MPS) and the path-integral QMC method, because this wave function suffers from the entanglement area law and the sign problem simultaneously.

The examples we studied in this section include the ground state of free bosons in 1D lattice, and those of free fermions in 1D and 2D lattices, with the filling factor away from half-filling. For free bosons (FB) in a 1D lattice with length L and unit filling factor (with the overall number of bosons $N = L$), the ground-state wave function can be written as $\Psi_{\text{FB}} = \sum_{\mathbf{n}} C_{\text{FB}}[\mathbf{n}]|\mathbf{n}\rangle$ where $|\mathbf{n}\rangle = |n_1 \dots n_L\rangle$ is the occupation number basis spanning the Hilbert space under the constraint $N = L$. The analytic form of the multivariable characteristic function is $C_{\text{FB}}[\mathbf{n}] = \sqrt{L! / n_1! \dots n_L!} / L^{L/2}$. The wave function of a free-fermion (FF) lattice system has a similar form while n_i can only be 0 or 1. We assume there are N particles and $C_{\text{FF}}[\mathbf{n}] = \det[\mathbb{M}]$, where \mathbb{M} is an $N \times N$ matrix with the matrix elements $M_{ij} = f_i(x_j)$, with x_j denoting the position of the j th fermion, and $f_i(x)$ denoting the i th single-particle eigenstate, in a 1D chain with a periodic boundary condition (PBC), $f_i(x) = \frac{1}{\sqrt{L}} e^{ik_i x}$ with k_i being the i th momentum. In a 2D lattice, both x_j and k_i are replaced by a 2D vector \mathbf{x}_j and \mathbf{k}_i . Based on the above exact results, we implemented the training process, aiming to minimize the fidelity function F by adjusting the parameters of the ANN. To avoid the problem of overfitting, the number of the variational parameters in the ANN was chosen to be on the order of $\mathcal{O}(N_b^2) \sim 10^3$, significantly smaller than the typical Hilbert space dimensionality of the systems we studied here ($\sim 10^6$). As shown in Fig. 2(a), for cases of 1D free bosons or fermions, an ANN with only a few neurons $N_b \sim \mathcal{O}(10^1)$ can easily approximate the corresponding target function with an extraordinarily high precision. For 2D fermions, a direct approximation of $C_{\text{FF}}[\mathbf{n}]$ based on the ANN with the current structure seems to fail because the sign of $C_{\text{FF}}[\mathbf{n}]$ can dramatically change owing to the fermionic statistics in two dimensions. We will reconsider this point later on. To avoid this problem, here we chose

$|C_{\text{FF}}[\mathbf{n}]|$ instead of $C_{\text{FF}}[\mathbf{n}]$ as our target function, which still allowed us to calculate the average values of diagonal operators in the Fock basis [e.g., the nearest-neighbor (NN) density correlation $O_{nn} = \frac{1}{L} \sum_i \langle n_i n_{i+1} \rangle$ and the local density operator $O_{n1} = \langle n_1 \rangle$]. As shown in Fig. 2(c), an ANN with $N_b \sim \mathcal{O}(10^2)$ can give rise to values of O_{nn} and O_{n1} with precisions $\sim \mathcal{O}(10^{-3})$. We also notice that the ground state of interacting quantum models (e.g., the quantum spin model [30] and Bose-Hubbard model [40,41]) have been studied by other machine learning methods, where the ANNs are trained to minimize the ground-state energy, instead of the fidelity of the wave function.

In all of the above-studied cases, the system size is relatively small since we want to compare our results with the exact ones. In what follows, we show that, in principle, there is no intrinsic difficulty to use ANNs to simulate larger systems, whose Hilbert space dimensionality is much larger than the memory of any computer, thus making it impossible to store the predicted wave functions and calculate its overlap with the exact ones. As a consequence, we focus on physical quantities in the ground states instead of focusing on the wave functions themselves, to test the accuracy of the ANN. The key point is that we approximate the characteristic functions only training the ANN only over a tiny fraction of the entire configuration space and we assume that the predicted functions are also valid for other configurations once the ANN finds the correct form of the target function. Obviously, the efficiency of the function approximating strongly depends on the choice of training configurations, which need to be “representative” in the Hilbert space. In other words, the probability of choosing a certain configuration $[\mathbf{n}]$ should be proportional to its weight $|C[\mathbf{n}]|^2$ in a given wave function, which can be achieved by the importance sampling in the Monte Carlo algorithm. Consider the ground state of a 1D $L = 64$, $N = 31$ free-fermion system as an example. For this system, we first implemented importance sampling to generate millions of “representative” configurations based on the exact value of $|C_T[\mathbf{n}]|^2$, and use these generated configurations to train the ANN and approximate the target function. After completing the training, a new set of “representative” configurations were generated according to their weights predicted by the ANN




FIG. 2. (a) Fidelity $F = 1 - |\langle \Psi_{ED} | \Psi_{\text{ANN}} \rangle|$ as a function of the number of neurons N_b , for 1D free bosons with $L = 12$, $N = 12$, and the corresponding Hilbert space dimensionality $\mathbb{D} = 1\,352\,078$ and for 1D free fermions with $L = 24$, $N = 11$, and $\mathbb{D} = 2\,496\,144$. (b) The precision function $\delta = |O_{\text{ANN}} - O_{ED}| / O_{ED}$ for two different quantities $O_{nn} = \frac{1}{L} \sum_i \langle n_i n_{i+1} \rangle$ and $O_{n1} = \langle n_1 \rangle$ for 2D free fermions with $L = 24(4 \times 6)$, $N = 13$, and $\mathbb{D} = 2\,496\,144$, the inset shows a typical Fock configuration. (c) The precision function δ for $O_{nn} = \frac{1}{L} \sum_i \langle n_i n_{i+1} \rangle$ as a function of N_b obtained by the importance sampling algorithm for a 1D free fermion, where both the amount of the training and the sampling data are chosen as $\mathbb{H} = 10^6$ in the simulations of various system sizes; the inset shows δ as a function of L with a fixed $\mathbb{H} = 10^6$ and $N_b = 200$.

$|C_P[\mathbf{r}]|^2$, and we calculated the value of the physical quantities based on these new configurations and compared them with their exact values. The result is shown in Fig. 2(c), where we observe that the precision of O_{nn} evaluated based on the above-described sampling scheme can reach $\mathcal{O}(10^{-3})$. For a fixed amount of training and sampling data, the scaling relation between the precision and system size is shown in the inset of Fig. 2(c).

IV. NON-POSITIVE-DEFINITE WAVE FUNCTIONS

In general, neural networks perform much better on approximating smooth functions rather than rough ones, which may restrict their applicability in quantum many-body physics, because for certain types of important wave functions, their characteristic functions may drastically change and even alter their sign in response to a slight change in the input configuration. The sign problem, in its different forms, imposes challenges on both existing methods (e.g., QMC) and the current neural network function approximating. For example, the ANN with a simple structure illustrated in Fig. 1 fails to approximate the ground state of one of the simplest models: a 1D antiferromagnetic Heisenberg model. In QMC, the sign problem in such a bipartite lattice can be avoided by performing a basis rotation. In the current method, we attack the problem by separately approximating the amplitude and the sign of the target function by two different ANNs. The ANN that approximates the amplitude is similar to those studied above, with the only difference that the target function is replaced by its absolute value. This ‘‘amplitude ANN’’ performs as well as that in the previously described ones in Sec. III.

The difficult part is to approximate the sign function: $S[\sigma] = (C[\sigma]/|C[\sigma]| + 1)/2$, which takes the values of 1(0) if $C[\sigma]$ is positive (negative). For a 1D lattice (or more generally, a bipartite lattice) model, it is well known that the sign of $C[\sigma]$ obeys the Marshall sign rule: $S[\sigma] = 1/0$ if in σ the total number of down spins in the odd sites is even or odd. This mathematical theorem enables us to perform a basis rotation in the even or odd sites to eliminate the sign problem in QMC simulations. For the current method, the question is without the prior knowledge of the Marshall sign rule, whether a ANN can automatically extract it from the training set data? Based on our numerical tests, we found that a standard ANN, as the one shown in Fig. 1, fails to extract the Marshall sign rule, but a modified ANN with the activation function of neurons in the first hidden layer replaced by a cosine function (hereafter denoted as ‘‘sign ANN’’) succeeds. This modification can significantly increase the efficiency of the ANN with the accuracy of 100% because the cosine function is more capable of capturing the even/odd features in the input data. Another important wave function is the ground state of the Majumdar-Ghosh model: the dimerized state $\Psi = \bigotimes_{i=1}^{\frac{L}{2}} \frac{1}{\sqrt{2}} [|\uparrow\rangle_{2i-1} |\downarrow\rangle_{2i} - |\downarrow\rangle_{2i-1} |\uparrow\rangle_{2i}]$, with the sign function $S[\sigma] = [\prod_{i=1}^{\frac{L}{2}} (S_{2i-1}^z - S_{2i}^z) + 1]/2$. We found that this sign rule can also be satisfactorily learned by the sign ANN with the accuracy of 100%.

In the two examples above, the sign rules are rather simple in the sense that they can be written explicitly in a simple form. However, for a generic wave function with the sign problem,

this is not the case. The sign rule may be too complex to be captured by programming or designing explicit algorithms, which on the other hand is exactly what ANNs are good at. Approximating the sign function becomes a classification problem, which reminds us of one of the most successful applications of ANN: recognizing handwritten digits. In this classical problem, the ANN was shown to automatically and successfully infer the rules of classification using the training set examples. Here, we adopt a similar strategy to extract the elusive sign rule for the ground state of a frustrated quantum magnetism model: the 1D J_1 - J_2 antiferromagnetic (AF) Heisenberg model with the Hamiltonian $H = \sum_i [J_1 S_i S_{i+1} + J_2 S_i S_{i+2}]$ (where both J_1 and J_2 are non-negative). The two cases studied above are exactly the ground state of this model in two limits: $\alpha = 0$ and 0.5 with $\alpha = J_2/J_1$.

For a general α , there is no exact solution of the ground state; thus, we used the Lanczos method to calculate the sign function for a finite-size system and compared it to the one predicted by the sign ANN. The accuracy versus α is plotted Fig. 3(a) for various N_b , and we observe that in the entire region $0 \leq \alpha \leq 0.5$, the accuracy predicted by the sign ANN is high, reaching 99%, while the minimum of the accuracy corresponds to the phase transition point. Aside from the accuracy, the efficiency of an ANN also depends on the typical training time and how this time scales with the system size. The ‘‘time’’ evolution of the accuracy during a training process is shown in Fig. 3(b), where we observe that the ANN first experiences a period of ‘‘confusion’’ with the prediction accuracy $P \simeq 0.5$ until a certain time T_c , after which the machine finally learns the correct approximation of the exact sign function and the accuracy will increase rapidly before saturating. This time scale T_c , together with the number of the neurons N_b , can be understood as the computation resources one needs to capture the sign rule using the ANN. The scaling relation between T_c and the system size L is plotted in the inset of Fig. 3(b). The system size we studied is relatively small, thus, it is difficult to tell whether T_c scales with L in a polynomial or exponential manner, which remains an open question.

Now, we discuss more details of the ANN. First, even though throughout this paper we choose a ANN with a two-hidden-layer structure, one may wonder whether further increasing the number of the layers can improve the prediction accuracy or not. To address this issue, we calculate P for different ANN with different hidden layer (up to three) with fixed iteration steps. As shown in Fig. 2(c), we found that, at least for this example, even though an ANN with two hidden layers indeed performs much better than that with only one hidden layer, further increasing the number of layers does not significantly improve the performances. We also check the correlation between the erroneously assigned signs for a certain coefficient and the coefficient’s absolute value. To do that, we define a parameter ε to measure the ratio between the average absolute value of the coefficient with the erroneously assigned sign and that of the whole training set:

$$\varepsilon = \frac{\frac{1}{D'} \sum_{\sigma'} |C[\sigma']|}{\frac{1}{D} \sum_{\sigma} |C[\sigma]|}, \quad (3)$$

where $[\sigma']$ denotes the set of input basis with the erroneous assigned sign, with the dimensionality D' , while $[\sigma]$ denotes

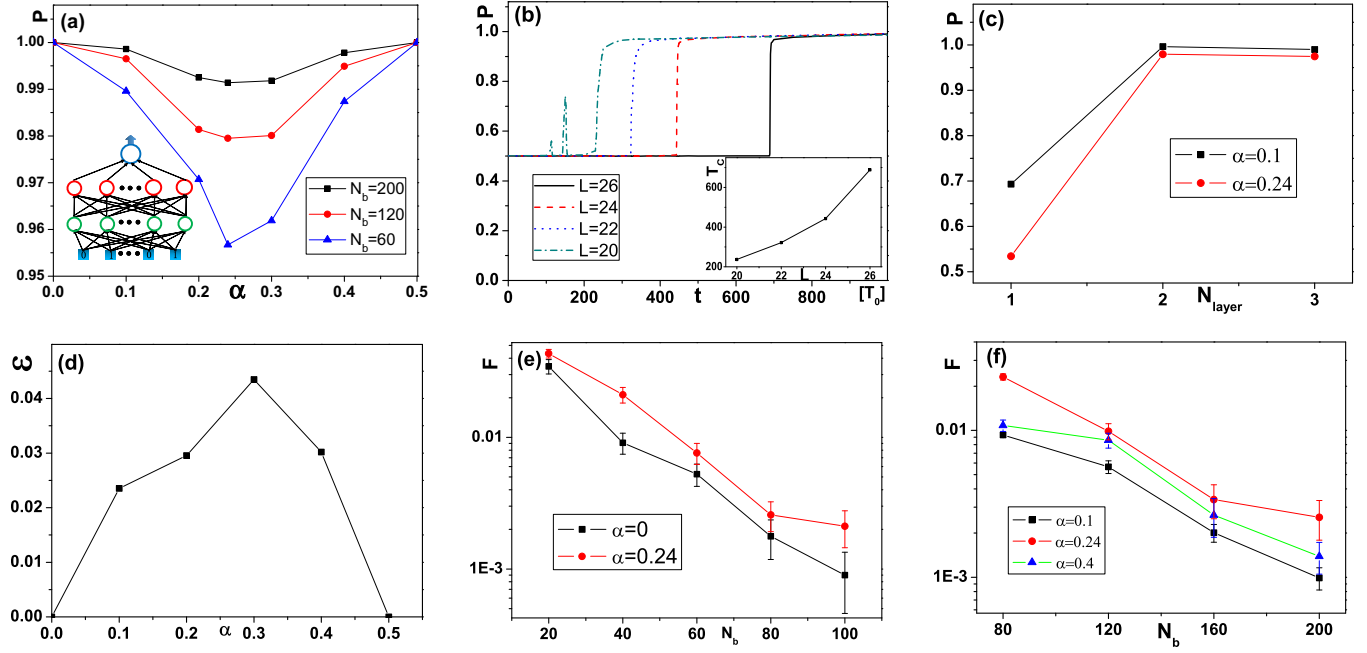


FIG. 3. (a) The accuracy function $P(\alpha)$ of the sign rule in the ground state of 1D J_1 - J_2 model with $L = 24$ and various N_b , with the inset showing the modified ANN used to capture the sign rule with neurons in the first hidden layer replaced by those with the cosine activation functions $f(x) = \cos(\pi x)$. (b) The evolution of P during the training with fixed $N_b = 200$ and various systems sizes, with the inset showing a typical “confusion” time T_c in the training as a function of system size with $N_b = 200$. (c) The accuracy P predicted by the sign ANN as a function of the layer of the ANN with a fixed $N_b = 120$. (d) The ratio between the average absolute value of the coefficient with the erroneously assigned sign and that of the whole training set with $N_b = 120$. (e) The fidelity as a function of N_b for the amplitude ANN. (f) The fidelity in the J_1 - J_2 model and various α as a function of N_b predicted by the combination of the sign (with $N_b^s = N_b$) and amplitude ANNs (with $N_b^a = N_b/2$). For (c)–(f), the system size $L = 24$.

all the training bases with the dimensionality \mathbb{D} . As shown in Fig. 3(d), we can find that $\varepsilon \sim \mathcal{O}(10^{-2}) \ll 1$, indicating that the erroneous predictions of the sign ANN tend to occur for those input bases whose coefficients have small absolute value.

By combining the results of the amplitude [as shown in Fig. 3(e) for examples] and sign ANNs, we can calculate the ground state of this frustrated quantum magnetic model, which agrees very well with the exact results, as shown in Fig. 3(f). In summary, by dividing the ANN into two parts with different architecture, we can approximate the ground state of a frustrated quantum magnetism with high precisions. This strategy may shed light on using ANNs to solve the complex quantum many-body systems with sign problems.

V. VARIATIONAL RESULTS

In all the cases studied above, because the target functions (the ground-state wave functions) are given, either analytically or numerically, one may expect that the current method would not be useful for exploring new quantum many-body systems with previously unknown ground states. However, our previous results have established that the ground states of some quantum many-body systems can be efficiently represented by ANNs, which enables us to consider an ANN as a variational wave function for the true ground state of a new system. The connection weights and the bias in an ANN are the variational parameters with respect to which we seek to minimize the expectation value of the Hamiltonian (energy), instead of the fidelity function as described above. For a given ANN with a

set of parameters $\{W\}$, the corresponding variational energy $E(\{W\})$ can be estimated using Monte Carlo simulations, with a procedure similar to those used for previously studied free-fermion cases in Sec. III. The minimum of $E(\{W\})$ in the space of parameters can be found using the stochastic reconfiguration optimization method. Similar strategy has been used in a different type to ANN: the restricted Boltzmann machine [30] to solve the nonfrustrated quantum magnetic models, e.g., the Heisenberg model and transverse Ising model. For the complex problems like the frustrated quantum magnetisms, we shall show that the strategy we proposed by dividing the ANNs into amplitude and sign parts can also improve the efficiency of this variational method.

In the following, we use the ANN combined with variational Monte Carlo (VMC) methods (the details of the method can be found in the Appendix) to solve the ground state of a 1D $L = 30$ J_1 - J_2 AF Heisenberg model. The structure of the ANN we used in the variational method is shown in Fig. 4(a). We adopt the strategy proposed in Sec. IV that the ANN has been divided into the amplitude and sign parts, and the final output is the product of them. However, one of the crucial differences is that the output of the sign ANN is a continuous number instead of a discrete one. The reason is that a discrete function is usually nondifferentiable, which may decrease the efficiency of the variational method. We also notice that the structure of the ANN in Fig. 4(a) is much simpler than the one we used in Sec. IV: there is only one hidden layer instead of two in both amplitude and sign ANN. We will turn back to this point later.




FIG. 4. (a) The structure of the ANN we used in the variational method calculations. (b) The “time” evolution of the excess energy per site $\Delta e = (E_v - E_0)/L$ during the training process for a 1D J_1 - J_2 AF Heisenberg model with $L = 30$ using the trial wave function as shown in (a) with the neuron number in the hidden layer of the amplitude ANN $N_b = 30$.

Since the target wave function is unknown, the cost function to be minimized during the training process is the variational energy of the system, which can be evaluated by the VMC method for a given trial wave function (ANN). In Fig. 4(b), we plot the “time” evolution of the excess energy per site $\Delta e = (E_v - E_0)/L$ during the training process, where E_v is the variational energy evaluated by VMC and E_0 is the exact ground-state energy calculated by the density matrix renormalization group method. As shown in Fig. 4(b), an ANN with a simple structure shown in Fig. 4(a) can give rise to a value of the ground-state energy with precisions $\sim \mathcal{O}(10^{-3})$.

It is interesting to notice that compared to the ANN we used in Sec. IV, the structure of the ANN we used in this variational method is much simpler (only one hidden layer), but its performance is even better. There are several reasons for this counterintuitive fact: (i) In the variational method, the sign and amplitude ANN are trained as a whole, while in Sec. IV they are trained separately. However, for the sign ANN, an extremely high accuracy may not be very necessary for the final results, since as shown in Fig. 3(d) the erroneous predictions of the sign ANN tend to occur for those input bases whose coefficients have very small absolute value, which give very little contribution to the final results, especially to the energy. (ii) In all the sections except Sec. V, during the optimization process the training data sets are divided into N batches and the optimization is performed batch by batch, while in the variational method, all the entries in wave function are optimized simultaneously. (iii) In Sec. IV, we choose the cost function as the overlap of the wave functions (fidelity), which is a global quantity, thus is usually difficult to be minimized, while in the variational method, we only need to minimize a local quantity: the variational energy.

VI. CONCLUSION AND OUTLOOK

In this paper, we demonstrated the powerful applicability of a simple ANN in approximating the ground-state wave functions of some notable quantum many-body systems. Even though an ANN with a simple structure can already approximate some of them with a high precision, there is still a long way to go before this method can learn to solve problems that remain inaccessible by any other well-established numerical method, and the efficiency improvement plays a key role in this process.

Some avenues along this line suggest the directions for future studies. First, as with many existing numerical methods, imposing the Hamiltonian symmetries on the ANN will significantly improve its efficiency. Consider a 2D system for example. An ANN with translational symmetry does not only reduce the number of the variational parameters and training data, but also learns the lattice geometry from the beginning. Inspired by the impressive success of deep learning techniques, we expect that an ANN can be more powerful when networks are made deeper. In numerical simulations, we found that even though an ANN with two hidden layers indeed performs much better than that with only one hidden layer, further increasing the number of layers does not significantly improve the performances. One of the possible reasons is that for an ANN with more hidden layers, even though its expressibility may be more powerful, the computational cost to find the optimal representation is significantly increased since the landscape in the parameter space is more complex. Incorporation of “deep learning” into our simulations remains an open issue and deserves further studies.

A fully connected ANN used in our simulations may contain many redundant connections that complicate the optimization process. Adopting ANN with more advanced architectures (e.g., convolutional neuron networks) may help to avoid this redundancy, thus likely significantly improve its efficiency. Recently, the relation of the quantum entanglement and the range of the connections in the ANN have been built, and we expect the convolutional neuron networks may work well for those ground states with short-range entanglement [16]. Last but not least, it is known that the artificial neural network, in general, is a heuristic algorithm whose efficiency largely depends on the designer’s experience and intuition; however, to make this approach valuable in the physics, a systematic understanding of its validity and limitation, at least for this concrete problem, is still needed, which may be also beneficial to the artificial intelligence community.

ACKNOWLEDGMENTS

We wish to thank L. Wang and J. Carrasquilla for fruitful discussions. Z.C. is supported by the National Key Research and Development Program of China (Grant No. 2016YFA0302001), the National Natural Science Foundation of China under Grant No. 11674221, and the Shanghai Rising-Star Program. J.G.L. is supported by the National Natural Science Foundation of China under Grant No. 11774398.

APPENDIX A: DETAILS OF THE OPTIMIZATION ALGORITHM

During the training process of the artificial neural network, one needs to adjust the parameters in ANN to minimize the distance function, which turns to a multivariable optimization problem. In our simulations, we use two optimization techniques depending on the concrete problems: the stochastic gradient descent (SGD) [8] and an adaptive learning rate optimization algorithm: adaptive moments (Adam). Here, we only explain the details of the SGD method, and the Adam optimization algorithm has been explicitly illustrated in Ref. [42].

In our simulation, each input data is composed of the basis and the corresponding coefficient (label) calculated by exact diagonalization (ED) or other methods (e.g., $\{|\sigma\rangle, C[\sigma]\}$). For a state with large Hilbert space dimension \mathbb{D} , the typical coefficient is pretty small $C[\sigma] \sim \mathcal{O}(1/\sqrt{\mathbb{D}})$. In general, the nonlinear activation function is not sensitive in the regime where the input/output is too small, therefore, to make the ANN more efficient, we multiply the target function by a factor of $\sqrt{\mathbb{D}}$, thus $C[\sigma] \sim \mathcal{O}(1)$. This renormalization does not change the results of any physical observable, but can significantly improve the efficiency of the fitting. The input data set are randomly divided into two groups: 80% of them are used for training and the rest are testing data sets, and we choose a random set of ANN parameters \mathbf{W}_0 as the initial parameters. The training set of data are randomly reshuffled and divided into N batches, each of which contain M data. In each batch, the training data are labeled as $\{\sigma^{(1)}, \dots, \sigma^{(M)}\}$ with corresponding target $C[\sigma^{(i)}]$.

After the initializations are finished, we start the optimization process. In each step of the SGD update, we choose one batch and calculate the gradient estimate in the parameter landscape \mathbf{W} :

$$\mathbf{g} = \nabla_{\mathbf{W}} \frac{1}{M} \sum_{i=1}^M L(C_p[\sigma^{(i)}], C[\sigma^{(i)}]), \quad (\text{A1})$$

where $C_p[\sigma]$ is the coefficient predicted by the ANN and $L(C_p[\sigma], C[\sigma]) = (C[\sigma] - C_p[\sigma])^2$ is a function of \mathbf{W} , denoted as loss function. Once we obtain the gradient, the parameter \mathbf{W} is updated as

$$\mathbf{W} \leftarrow \mathbf{W} - \epsilon \mathbf{g}, \quad (\text{A2})$$

where ϵ is the parameter controlling the learning rate, which gradually decreases over time. The above optimization processes continue until all the batches are chosen, then one iteration of the training is finished. In our simulations, the training time is measured in the unit of the time of a single iteration T_0 . The typical training time ranges from $10^2 \sim 10^3 T_0$ depending on the convergency of the problems. The SGD method is the most used optimization technique for machine learning. In our simulations, we use the GPU to speed up the computational efficiency of the training. A set of raw data generated during the training process can be found in the Supplemental Material [43].

APPENDIX B: DETAILS OF THE VARIATIONAL METHOD

In this Appendix, we will show the variational analysis of using the artificial neural network to explore new ground states. In this case, the target function is previously unknown, therefore, during the training process, what we need to minimize is not the distance function F , but the variational energy (the expectation value of the Hamiltonian over the ANN trial wave function), which can be calculated using the variational Monte Carlo method, as we will show in the following. In general,

a Hamiltonian can be split into the diagonal and nondiagonal parts: $\hat{H} = \hat{T} + \hat{V}$ in the basis $|\sigma\rangle$ (or Fock basis). For a given trial wave function in terms of the ANN with a set of variational parameter $\{W\}$, the expectation value of the diagonal part

$$V(\{W\}) = \langle \hat{V} \rangle = \sum_{[\sigma]} p_{\sigma}\{W\} \langle \sigma | \hat{V} | \sigma \rangle, \quad (\text{B1})$$

where $p_{\sigma}\{W\} = |C_{\text{ANN}}[\sigma]|^2$ is the probability predicted by the ANN. The summation $\sum_{[\sigma]}$ is over the whole Hilbert space, whose dimensionality exponentially grows with the system size. Following in the spirit of Monte Carlo, the summation over the whole Hilbert space can be replaced by the summation over those important configurations:

$$V(\{W\}) \simeq \sum_{[\bar{\sigma}]} \frac{p_{\bar{\sigma}}\{W\}}{Z\{W\}} \langle \bar{\sigma} | \hat{V} | \bar{\sigma} \rangle \quad (\text{B2})$$

with $Z\{W\} = \sum_{[\bar{\sigma}]} p_{\bar{\sigma}}\{W\}$, and $[\bar{\sigma}]$ is an exponentially small fraction of the whole configuration, denoting the set of representative configurations chosen by the importance sampling using Metropolis algorithm according to its probability $p_{\bar{\sigma}}$. The expectation value of the off-diagonal term can be evaluated in a similar way, without loss of generality, we assume $T = |\sigma'\rangle \langle \sigma|$ with $|\sigma\rangle \neq |\sigma'\rangle$, thus,

$$T(\{W\}) \simeq \sum_{[\bar{\sigma}]} \frac{p_{\bar{\sigma}}\{W\}}{Z\{W\}} \times \frac{C_{\text{ANN}}[\bar{\sigma}']}{C_{\text{ANN}}[\bar{\sigma}]} \langle \bar{\sigma}' | \hat{T} | \bar{\sigma} \rangle, \quad (\text{B3})$$

where the observable we need to calculate during the sampling is not only $\langle \bar{\sigma}' | \hat{T} | \bar{\sigma} \rangle$, but $\langle \bar{\sigma}' | \hat{T} | \bar{\sigma} \rangle C_{\text{ANN}}[\bar{\sigma}'] / C_{\text{ANN}}[\bar{\sigma}]$. By combining Eqs. (B2) and (B3), we can obtain the variational energy for a given trail ANN wave function. To minimize the variational energy, one needs to calculate the derivative of $H(\{W\})$ with respect to $\{W\}$, and the variational parameters can be determined by solving the equation

$$\frac{\partial H(\{W\})}{\partial w_i} = \frac{\partial V(\{W\})}{\partial w_i} + \frac{\partial T(\{W\})}{\partial w_i} = 0. \quad (\text{B4})$$

Notice that an ANN is a combination of a set of nonlinear functions whose explicit forms are already known; as a consequence, one can easily obtain its derivative $\frac{\partial C_{\text{ANN}}^{(W)}[\sigma]}{\partial w_i}$, thus,

$$\begin{aligned} \frac{\partial V(\{W\})}{\partial w_i} &= \sum_{[\sigma]} p_{\sigma}\{W\} \times \frac{\partial C_{\text{ANN}}^{(W)}[\sigma]}{\partial w_i} \frac{2\langle \sigma | \hat{V} | \sigma \rangle}{C_{\text{ANN}}^{(W)}[\sigma]}, \\ \frac{\partial T(\{W\})}{\partial w_i} &= \sum_{[\sigma]} p_{\sigma}\{W\} \times \frac{\langle \sigma' | \hat{T} | \sigma \rangle}{C_{\text{ANN}}^{(W)}[\sigma]} \\ &\quad \times \left[\frac{\partial C_{\text{ANN}}^{(W)}[\sigma']}{\partial w_i} + \frac{\partial C_{\text{ANN}}^{(W)}[\sigma]}{\partial w_i} \frac{C_{\text{ANN}}^{(W)}[\sigma']}{C_{\text{ANN}}^{(W)}[\sigma]} \right]. \end{aligned}$$

By performing the importance sampling, one can calculate the $\frac{\partial V(\{W\})}{\partial w_i}$ and $\frac{\partial T(\{W\})}{\partial w_i}$, then substitute them into Eq. (B4) to solve the variational parameters.

[1] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).

[2] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization

- group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
- [3] U. Schollwoeck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [4] A. W. Sandvik and J. Kurkijärvi, Quantum monte carlo simulation method for spin systems, *Phys. Rev. B* **43**, 5950 (1991).
- [5] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, *Phys. Lett. A* **238**, 253 (1998).
- [6] J. Gubernatis, N. Kawashima, and P. Werner, *Quantum Monte Carlo Methods: Algorithms for Lattice Models* (Cambridge University Press, Cambridge, 2016).
- [7] M. A. Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).
- [8] I. Goodfellow, Y. Bengio, and A. Courville (unpublished).
- [9] P. Mehta and D. J. Schwab, An exact mapping between the Variational Renormalization Group and Deep Learning, [arXiv:1410.3831](https://arxiv.org/abs/1410.3831).
- [10] H. W. Lin and M. Tegmark, Why does deep and cheap learning work so well? *J. Stat. Phys.* **168**, 1223 (2017).
- [11] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, On the Equivalence of Restricted Boltzmann Machines and Tensor Network States, [arXiv:1701.04831](https://arxiv.org/abs/1701.04831).
- [12] E. M. Stoudenmire and D. J. Schwab, Supervised Learning with Quantum-Inspired Tensor Networks, *Adv. Neural Inf. Process. Syst.* **29**, 4799 (2016).
- [13] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
- [14] L. Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
- [15] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nat. Phys.* **13**, 435 (2017).
- [16] D.-L. Deng, X. Li, and S. Das Sarma, Quantum Entanglement in Neural Network States, *Phys. Rev. X* **7**, 021021 (2017).
- [17] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Many-body quantum state tomography with neural networks, [arXiv:1703.05334](https://arxiv.org/abs/1703.05334).
- [18] T. Ohtsuki and T. Ohtsuki, Deep Learning the Quantum Phase Transitions in Random Two-Dimensional Electron Systems, *J. Phys. Soc. Jpn* **85**, 123706 (2016).
- [19] Y. Zhang and E.-A. Kim, Quantum Loop Topography for Machine Learning, *Phys. Rev. Lett.* **118**, 216401 (2017).
- [20] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine Learning Phases of Strongly Correlated Fermions, *Phys. Rev. X* **7**, 031038 (2017).
- [21] P. Ponte and R. G. Melko, Kernel methods for interpretable machine learning of order parameters, *Phys. Rev. B* **96**, 205146 (2017).
- [22] S. Curtarolo, D. Morgan, K. Persson, J. Rodgers, and G. Ceder, Predicting Crystal Structures with Data Mining of Quantum Calculations, *Phys. Rev. Lett.* **91**, 135503 (2003).
- [23] S. V. Kalinin, B. G. Sumpter, and R. K. Archibald, Big-cdeepsmart data in imaging for guiding materials design, *Nat. Mater.* **14**, 973 (2015).
- [24] L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler, Big Data of Materials Science: Critical Role of the Descriptor, *Phys. Rev. Lett.* **114**, 105503 (2015).
- [25] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Machine learning quantum phases of matter beyond the fermion sign problem, *Sci. Rep.* **7**, 8823 (2017).
- [26] L. Huang, Y.-F. Yang, and L. Wang, Recommender engine for continuous-time quantum Monte Carlo methods, *Phys. Rev. E* **95**, 031301 (2017).
- [27] L. Huang and L. Wang, Accelerated monte carlo simulations with restricted boltzmann machines, *Phys. Rev. B* **95**, 035105 (2017).
- [28] J. Liu, H. Shen, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning monte carlo method and cumulative update in fermion systems, *Phys. Rev. B* **95**, 241104 (2017).
- [29] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning monte carlo method, *Phys. Rev. B* **95**, 041101 (2017).
- [30] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [31] L. Wang, Can Boltzmann Machines Discover Cluster Updates? *Phys. Rev. E* **96**, 051301 (2017).
- [32] X. Gao and L.-M. Duan, Efficient Representation of Quantum Many-body States with Deep Neural Networks, *Nat. Commun.* **8**, 662 (2017).
- [33] D.-L. Deng, X. Li, and S. Das Sarma, Exact Machine Learning Topological States, *Phys. Rev. B* **96**, 195145 (2017).
- [34] Y. Huang and J. E. Moore, Neural network representation of tensor network and chiral states, [arXiv:1701.06246](https://arxiv.org/abs/1701.06246).
- [35] Y. Levine, D. Yakira, N. Cohen, and A. Shashua, Deep Learning and Quantum Entanglement: Fundamental Connections with Implications to Network Design, [arXiv:1704.01552](https://arxiv.org/abs/1704.01552).
- [36] A. N. Kolmogorov, On the representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables, *Dokl. Akad. Nauk SSSR* **108**, 179 (1961).
- [37] K. Hornik, M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* **2**, 359 (1989).
- [38] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Math. Control, Signals Syst.* **2**, 303 (1989).
- [39] M. Abadi, Large-scale machine learning on heterogeneous systems, software available from [tensorflow.org](https://www.tensorflow.org).
- [40] H. Saito, Solving the bose-hubbard model with machine learning, *J. Phys. Soc. Jpn.* **86**, 093001 (2017).
- [41] H. Saito and M. Kato, Machine learning technique to find quantum many-body ground states of bosons on a lattice, *J. Phys. Soc. Jpn.* **87**, 014001 (2018).
- [42] D. P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [43] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.97.035116>, which includes a set of raw data generated during the training process of a 1D spinless fermion with $L = 24$, $N = 11$, $N_b = 60$.