



# CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines

S. Pilati, E. M. Inack, and P. Pieri

Phys. Rev. E **100**, 043301 — Published 2 October 2019

DOI: [10.1103/PhysRevE.100.043301](https://doi.org/10.1103/PhysRevE.100.043301)

# Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines

S. Pilati

*School of Science and Technology, Physics Division,  
Università di Camerino, 62032 Camerino (MC), Italy*

E. M. Inack

*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

P. Pieri

*School of Science and Technology, Physics Division,  
Università di Camerino, 62032 Camerino (MC), Italy and  
INFN, Sezione di Perugia, 06123 Perugia (PG), Italy*

The projective quantum Monte Carlo (PQMC) algorithms are among the most powerful computational techniques to simulate the ground state properties of quantum many-body systems. However, they are efficient only if a sufficiently accurate trial wave function is used to guide the simulation. In the standard approach, this guiding wave function is obtained in a separate simulation that performs a variational minimization. Here we show how to perform PQMC simulations guided by an adaptive wave function based on a restricted Boltzmann machine. This adaptive wave function is optimized along the PQMC simulation via unsupervised machine learning, avoiding the need of a separate variational optimization. As a byproduct, this technique provides an accurate ansatz for the ground state wave function, which is obtained by minimizing the Kullback-Leibler divergence with respect to the PQMC samples, rather than by minimizing the energy expectation value as in standard variational optimizations. The high accuracy of this self-learning PQMC technique is demonstrated for a paradigmatic sign-problem-free model, namely, the ferromagnetic quantum Ising chain, showing very precise agreement with the predictions of the Jordan-Wigner theory and of loop quantum Monte Carlo simulations performed in the low-temperature limit.

## I. INTRODUCTION

The similarity between the imaginary-time Schrödinger equation and the diffusion equation allows one to simulate quantum many-body systems by stochastically evolving a (typically large) population of random walkers [1]. This is the basis of so-called projective quantum Monte Carlo algorithms (PQMC). These algorithms are particularly efficient when they simulate the ground state of sign-problem-free Hamiltonians, since the corresponding wave function is real and nonnegative in a suitable basis. Using fixed-node and released node methods [2–5], PQMC simulations provide accurate predictions even when the sign problem occurs, in particular for many-fermion systems, albeit in general at a larger computational cost. PQMC algorithms have been used to simulate fundamental quantum systems, including the electron gas [6], molecular systems [2], liquid and solid Helium [7], electrons in solids [8, 9], ultracold gases [10], quantum spin and lattice models [11, 12], and nuclear matter [13]. However, these algorithms are efficient only if they are provided with a sufficiently accurate trial wave function, which is used to guide the random walkers towards the most relevant regions of the configuration space. If this guiding wave function is not included, or if it is not sufficiently accurate, the computational cost of a PQMC simulation scales exponentially with the system size [14–16]. For various relevant systems, the general form of a suitable guiding

wave function is provided by some physical theory, and the details can be further tuned via a variational minimization of the energy expectation value. However, an appropriate theory is not always available, and the variational optimization might turn out to be an extremely challenging computational task. Even more, an inaccurate ansatz for the guiding wave function might lead to biased predictions, in particular when different phases closely compete, as in many strongly-correlated and/or disordered many-body systems. This problem is especially relevant in the field of adiabatic quantum computing, where quantum Monte Carlo (QMC) algorithms are being used to simulate how quantum annealers solve complex optimization problems [17–21]. In fact, the Hamiltonians corresponding to typical instances of hard optimization problems can be written in the form of random Ising models. These models are characterized by glassy ground states, for which an accurate ansatz is hard to guess [22].

Wave functions based on a certain type of generative neural network, specifically, restricted Boltzmann machines (RBM) [see Fig. 1, upper panel], have been recently proposed as generic variational ansatzes for quantum spin models [23]. Their representational power, their entanglement properties, and more elaborate versions are currently the subject of very intense research activity [27–43]. They appear to be particularly useful when no other physics-inspired ansatz is available. In this paper, we show how to employ them as guid-




FIG. 1. (color online). Connectivity structure of the restricted Boltzmann machine (RBM) [23] employed in this work as PQMC guiding wave function (upper panel) and of the unrestricted Boltzmann machine employed in Ref. [24] (lower panel). The latter is analogous to the shadow wave function used for quantum fluids and solids [25, 26]. The visible spins are labelled by  $x_j$ , with  $j = 1, \dots, N$ , while the hidden spins by  $h_i$ , with  $i = 1, \dots, N_h$ . The segments indicate the allowed interactions.

ing wave functions in PQMC simulations. In a previous article [24], a neural network state that mimics a type of an unrestricted Boltzmann machine (see Fig. 1, lower panel) has been employed as guiding wave function. In the case of Ref. [24], the neural network parameters had to be obtained via a separate variational optimization, and the PQMC algorithm had to be overloaded with the sampling of additional hidden variables. Here we show that, using unsupervised machine learning algorithms, the RBM can be trained along the PQMC simulations directly from random-walker configurations. One obtains an adaptive guiding wave function that improves during the PQMC simulation, allowing one to completely eliminate the bias due to the finite random-walker population [16] and to drastically reduce the statistical fluctuations. As a testbed, we consider a paradigmatic sign-problem free model, namely, the ferromagnetic quantum Ising chain, making comparison against the exact predictions of the Jordan-Wigner theory and of loop quantum Monte Carlo (QMC) simulations performed with the ALPS library [44] in the low-temperature limit. Excellent agreement is obtained for ground-state energies and for the average magnetization, above, below, and at the

ferromagnetic quantum critical point. A (possibly useful) byproduct of the self-learning procedure described here is an optimized RBM ansatz. This ansatz is developed by minimizing a Kullback-Leibler divergence, rather than the common minimization of the variational energy [45]. In the model considered here, the self-learned RBM wave function turns out to be comparatively as accurate as the one obtained via variational minimization using a sophisticated optimization method such as the stochastic reconfiguration algorithm [46] as implemented in the NetKet library [47].

The paper is organized as follows. Section II introduces the model Hamiltonian we address and describes the PQMC algorithm, the RBM variational wave-function, as well as the adaptive unsupervised learning procedure. Numerical results and comparison with previous theories are presented in Section III. Section IV summarizes our findings and discusses possible future extensions.

## II. METHOD

As a test bed for the self-learning PQMC simulations we consider the one-dimensional ferromagnetic quantum Ising Hamiltonian, defined as:

$$\hat{H} = -J \sum_{j=1}^{N-1} \sigma_j^z \sigma_{j+1}^z - \Gamma \sum_{j=1}^N \sigma_j^x - h \sum_{j=1}^N \sigma_j^z, \quad (1)$$

where  $\sigma_j^x$  and  $\sigma_j^z$  indicate conventional Pauli matrices acting on the spin at the lattice site  $j = 1, \dots, N$ .  $N$  is the total number of spins, and we consider open boundary conditions. The parameter  $J > 0$  fixes the strength of the ferromagnetic interactions between nearest-neighbor spins. In the following, we use  $J = 1$  as unit of the energy scale.  $\Gamma$  and  $h$  are the intensities of the transverse and longitudinal magnetic field, respectively. Below, the eigenstates of the Pauli matrix  $\sigma_j^z$  are denoted as  $|x_j\rangle$ . The eigenvalue is  $x_j = 1$  when  $|x_j\rangle = |\uparrow\rangle$  and  $x_j = -1$  when  $|x_j\rangle = |\downarrow\rangle$ . The quantum states of  $N$  spins  $|\mathbf{x}\rangle = |x_1 x_2 \dots x_N\rangle$ , with  $\mathbf{x} = (x_1, \dots, x_N)$ , form the computational basis considered in the present work. We denote as  $|\psi\rangle$  the quantum state corresponding to the wave-function  $\langle x | \psi \rangle = \psi(\mathbf{x})$ .

### A. Projective Quantum Monte Carlo simulations

The PQMC algorithms allow one to simulate the ground-state of a generic Hamiltonian by stochastically evolving the imaginary-time Schrödinger equation. The algorithm's accuracy and efficiency greatly improve if one introduces a suitable ansatz for the ground-state wave-function — usually called guiding (or trial) wave function and indicated as  $\psi_G(\mathbf{x})$  — and let the product  $f(\mathbf{x}, \tau) = \psi(\mathbf{x}, \tau) \psi_G(\mathbf{x})$  evolve according to the modified imaginary-time Schrödinger equation. This equation

reads:

$$f(\mathbf{x}, \tau + \Delta\tau) = \sum_{\mathbf{x}'} \tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) f(\mathbf{x}', \tau). \quad (2)$$

Here,  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) = G(\mathbf{x}, \mathbf{x}', \Delta\tau) \frac{\psi_G(\mathbf{x})}{\psi_G(\mathbf{x}'')}$ , where  $G(\mathbf{x}, \mathbf{x}', \Delta\tau) = \langle \mathbf{x} | \exp[-\Delta\tau(\hat{H} - E_{\text{ref}})] | \mathbf{x}' \rangle$  is the imaginary-time Green's function for a (short) time step  $\Delta\tau$  (we set  $\hbar = 1$  throughout this paper).  $E_{\text{ref}}$  is a reference energy introduced to stabilize the numerics, as explained below.

We employ the continuous-time algorithm of Refs. [48, 49], which allows one to stochastically simulate the exact (modified) imaginary-time Green's function  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau)$ , avoiding any finite time-step error which would occur with the use of Trotter approximations. This is achieved by using the formal zero-time limit of the Green's function and by appropriately sampling the time interval  $\delta\tau$  between consecutive single-spin flips. More flips might occur within a time-step  $\Delta\tau$ . Since the modified imaginary-time Green's function does not define a standard Markov process, i.e., one driven by a column normalized transition matrix, the simulation has to be performed by stochastically evolving a (large) population of random walkers; these walkers are subjected to updates in configuration space and to a branching process in which they are annihilated or replicated. The spin-configuration updates  $\mathbf{x}' \rightarrow \mathbf{x}$  (with  $\mathbf{x}' \neq \mathbf{x}$ ) are randomly selected from a probability distribution proportional to the modified imaginary-time Green's function. The guiding wave function favors updates towards relevant regions of the configuration space. In the branching process, the random walkers are annihilated or replicated according to the weight factor  $w_{\mathbf{x}'} = \exp[-\delta\tau(E_{\text{loc}}(\mathbf{x}') - E_{\text{ref}})]$ , where the local energy is  $E_{\text{loc}}(\mathbf{x}') = \sum_{\mathbf{x}} H_{\mathbf{x}', \mathbf{x}} \frac{\psi_G(\mathbf{x})}{\psi_G(\mathbf{x}'')}$  [50]. This weight accounts for the column normalization of the transition matrix. By dynamically tuning  $E_{\text{ref}}$ , the size of the random-walker population can be kept very close to a desired target value  $N_w$ . To implement the branching process and for the tuning of  $E_{\text{ref}}$  we use the textbook recipe of Ref. [51].

The continuous-time PQMC algorithm sketched here is more exhaustively described in Refs. [24, 48, 49], and we refer the interested readers to those references for additional implementation details. The projection can reach long imaginary times  $\tau = N_{\text{MC}}\Delta\tau$  by iterating a large number  $N_{\text{MC}}$  of Monte Carlo steps, each step corresponding to a short time step  $\Delta\tau$ . In the long imaginary-time limit, the walkers sample spin configurations with a probability distribution proportional to  $f(\mathbf{x}, \tau \rightarrow \infty) = \psi_0(\mathbf{x})\psi_G(\mathbf{x})$  (if  $N_w$  is large enough, as explained below), where  $\psi_0(\mathbf{x})$  is the ground-state wave function. Unbiased predictions of ground-state energies are obtained via Monte Carlo integration of the sum

$$E_{\tau \rightarrow \infty} \equiv \frac{\sum_{\mathbf{x}} f(\mathbf{x}, \tau \rightarrow \infty) E_{\text{loc}}(\mathbf{x})}{\sum_{\mathbf{x}} f(\mathbf{x}, \tau \rightarrow \infty)}. \quad (3)$$

Unbiased prediction of ground-state expectation values of other observables that commute with the Hamiltonian are obtained with analogous formulas. However, for operators  $\hat{O}$  that do not commute with the Hamiltonian, the analogous formulas would predict so-called mixed estimators, namely  $\langle \psi_0 | \hat{O} | \psi_G \rangle / \langle \psi_0 | \psi_G \rangle$ . In general, these are affected by a bias due to the guiding wave function, unless the latter coincides with the ground-state wave function. Nevertheless, if the operator  $\hat{O}$  is diagonal in the chosen computational basis, the pure estimator corresponding to  $\langle \psi_0 | \hat{O} | \psi_0 \rangle / \langle \psi_0 | \psi_0 \rangle$  can be determined via the standard forward-walking technique (see, e.g., Ref. [7]). An accurate guiding wave function reduces the computational cost of the forward walking technique.

In the large random-walker population limit,  $N_w \rightarrow \infty$ , the above Monte Carlo estimates are unbiased, being affected only by statistical fluctuations, which can be systematically reduced. For finite  $N_w$ , a systematic bias might arise due to the spurious correlations among replicated walkers [14–16, 52–54]. In Ref. [16], it was indeed shown that if one does not introduce a guiding wave function, which corresponds to setting  $\psi_G(\mathbf{x}) = 1$  in the equations above, the random-walker population required to keep this systematic bias below a chosen (small) threshold increases exponentially with the system size. This, in turn, implies a computational cost which scales exponentially. Instead, if the guiding wave function is exact, i.e. if  $\psi_G(\mathbf{x}) = \psi_0(\mathbf{x})$ , the local energy  $E_{\text{loc}}(\mathbf{x})$  is a constant function. This freezes the annihilations and replications of the random walkers in the branching process, eliminating any potential bias. If  $\psi_G(\mathbf{x})$  is a reasonable approximation for  $\psi_0(\mathbf{x})$ , the fluctuations of the random-walker number are anyway reduced compared to the case of the simple PQMC algorithm performed with  $\psi_G(\mathbf{x}) = 1$ , giving a much faster convergence to the exact  $N_w \rightarrow \infty$  limit.

## B. Boltzmann machines for PQMC algorithms

In Ref. [24], it has been shown that a guiding wave function based on an unrestricted Boltzmann machine is sufficiently accurate to drastically reduce the required random-walker population, leading to a polynomially scaling computational cost, at least for the ferromagnetic quantum Ising chain. Remarkably, this has been achieved by optimizing just three variational parameters. The small number of variational parameters to be optimized is the main benefit of the unrestricted architecture of Refs. [24, 41]. However, using the unrestricted Boltzmann machine requires to sample additional hidden spins, both during the variational optimization [55] and, chiefly, during the PQMC simulations. In fact, the PQMC algorithm has to be extended, combining visible-spin and hidden-spin sampling as explained in Refs. [24, 56]. An inefficient sampling of the addi-

tional hidden spins might introduce further correlations among replicated walkers, leading to a larger finite- $N_w$  bias. This deleterious effect has indeed been observed in PQMC simulation performed close to the ferromagnetic quantum critical point [24], where the statistical correlations among subsequent hidden-spin configurations are larger. This effect did not lead to an exponentially scaling computational cost in the ferromagnetic quantum Ising chain, but it might be more detrimental for more challenging Hamiltonians.

Below, we show that a guiding wave function based on an RBM allows one to eliminate the finite- $N_w$  bias. The additional benefit the RBM ansatz provides is that the optimization of the guiding function can be performed via unsupervised machine learning directly from the random-walker population, with no need for a separate variational optimization. Furthermore, as opposed to the unrestricted Boltzmann machines, the RBM ansatz does not require to sample additional classical hidden-spin variables. This allows one to avoid overloading the PQMC algorithm with the Monte Carlo updates for the hidden spins. We argue that these two benefits might be particularly relevant in two setups: in simulation of disordered models, where frustration effects might increase the correlations among successive hidden spin configurations, and in quantum annealing simulations, where the Hamiltonian varies in time and an automatically adaptive guiding wave function would allow one to perform more efficient simulations of the quantum annealers' dynamics.

Boltzmann machines are generative stochastic neural networks often used to approximate the probability distribution corresponding to a given population of stochastic samples. Beyond the visible spin variables  $\mathbf{x}$ , they include  $N_h$  additional hidden spin variables  $h_i = \pm 1$ , with  $i = 1, \dots, N_h$ , collectively denoted in the following as  $\mathbf{h} = (h_1, \dots, h_{N_h})$ . The probability distribution is written in the form of the Boltzmann weight  $P(\mathbf{x}, \mathbf{h})$  corresponding to a classical Ising Hamiltonian. In the case of the RBM, this classical Hamiltonian function reads:

$$H_{\text{RBM}}(\mathbf{x}, \mathbf{h}) = - \sum_{i,j} J_{ij} h_i x_j - \sum_j a_j x_j - \sum_i b_i h_i. \quad (4)$$

Notice that no visible-visible interactions nor hidden-hidden interactions are allowed. Instead, all visible spins interact with all hidden spins via the coupling parameters  $J_{ij}$ . One obtains two layers, respectively referred to as visible and hidden layer, with the bipartite all-to-all connectivity illustrated in the upper panel of Fig. 1. The parameters  $a_j$  and  $b_i$ , called biases, play the role of longitudinal local magnetic fields. The couplings  $J_{ij}$ , together with the bias terms  $a_j$  and  $b_i$ , define the RBM. In the following, all these parameters will be collectively denoted as  $\mathbf{W} = \{W_m\} = \{J_{ij}, a_j, b_i\}$ , for  $i = 1, \dots, N_h$  and  $j = 1, \dots, N$ . The index  $m$  labels all the RBM parameters. In practical applications the required number of hidden units is typically  $N_h \sim N$ . Thus, the number of parameters scales, to leading order, as  $N^2$ . This

quadratic scaling can be reduced to a linear scaling in translationally invariant models [23, 57]. However, in view of future studies on disordered systems, we avoid translational invariance adopting open boundary conditions. This allows us to test the unsupervised learning of RBM wave functions in the most generic setup. In the unrestricted Boltzmann machine considered in Ref. [24], intra-layer visible-visible and hidden-hidden direct interactions were included, reflecting the same (nearest-neighbor) connectivity of the quantum Hamiltonian in Eq. 1. Furthermore,  $N_h$  was taken equal to  $N$ , and inter-layer interactions were allowed only between visible and hidden spins corresponding to the same index, i.e., for  $i = j$ . The resulting connectivity structure is shown in the lower panel of Fig. 1. This structure is analogous to that of the shadow wave-function used to describe the liquid and solid phases of Helium-4 [25, 26].

In general, the probability to sample a visible-spin configuration  $\mathbf{x}$  is the marginal distribution over all possible hidden-spin configurations  $\mathbf{h}$ :

$$P(\mathbf{x}) \equiv \sum_{\mathbf{h}} P(\mathbf{x}, \mathbf{h}) = \frac{1}{Z} \sum_{\mathbf{h}} \exp[-H_{\text{RBM}}(\mathbf{x}, \mathbf{h})]. \quad (5)$$

Notice that the fictitious temperature and the Boltzmann constant  $k_B$  are here equal to unity. The normalization factor is the partition function  $Z = \sum_{\mathbf{x}, \mathbf{h}} \exp[-H_{\text{RBM}}(\mathbf{x}, \mathbf{h})]$ . Due to the absence of intra-layer interactions in the RBM, the hidden-spin configurations can be analytically traced out, resulting in a marginal distribution  $P(\mathbf{x}) \propto \exp\left(\sum_j a_j x_j\right) \prod_i F_i(\mathbf{x})$ , where  $F_i(\mathbf{x}) = 2 \cosh\left[b_i + \sum_j J_{ij} x_j\right]$ . In Ref. [23], it has been proposed to use the function  $P(\mathbf{x})$  to define an (unnormalized) ground-state wave-function. In order to describe both amplitude and phase, the RBM parameters  $\mathbf{W}$  should, in general, be complex valued. However, we consider here models whose ground-state wave function can be assumed to be real and nonnegative in a suitable basis; therefore, the RBM parameters can be restricted to have real parameters. Extensions to complex-valued ground-states for, e.g., fermionic systems, have recently been addressed [36, 58, 59]. In Ref. [23], the RBM parameters  $\mathbf{W}$  were determined via variational minimization of the expectation value  $\langle \psi_{\text{RBM}} | \hat{H} | \psi_{\text{RBM}} \rangle / \langle \psi_{\text{RBM}} | \psi_{\text{RBM}} \rangle$ , where  $|\psi_{\text{RBM}}\rangle$  indicates the quantum state corresponding to the (unnormalized) wave-function  $\langle \mathbf{x} | \psi_{\text{RBM}} \rangle \propto P(\mathbf{x})$ . The optimization of the variational parameters was performed using the stochastic reconfiguration algorithm [49]. In the following, the corresponding minimal variational energy, obtained using the Netket [47] library, will be indicated as  $E_{\text{var.min.}}$ . In Ref. [23], the variational-minimization approach has been referred to as reinforcement learning, due to the close similarity with the reinforcement-learning techniques used in the field of machine learning. However, in typical machine learning applications of RBMs the parameters  $\mathbf{W}$  are determined via unsuper-



vised machine-learning techniques [60], trying to learn the (unknown) distribution corresponding to a (typically large) dataset of stochastic samples. The unsupervised learning approach has already been used in Ref. [45] to perform quantum states tomography. Here, we adopt it to extract reasonably accurate variational ansatzes from the random-walker population generated by the PQMC algorithm. As discussed below, these ansatzes can then be used in an adaptive scheme as guiding wave functions to boost the efficiency of the PQMC simulation itself. On the one hand, this scheme allows one to eliminate the bias originating from the finite  $N_w$ . On the other hand, it allows one to eliminate the residual error in the RBM ansatz, providing unbiased predictions of ground-state properties.

### C. Unsupervised learning of adaptive RBM guiding wave functions

In the unsupervised learning approach, the RBM parameters are determined by maximizing the log-likelihood of a (typically large) training set:  $L(\mathbf{W}) = \sum_w \ln P(\mathbf{x}_w)$ , where  $w$  labels the instances in the training set. It can be shown that this corresponds to the minimization of the so-called Kullback-Leibler divergence (see, e.g., [61]). In general, the Kullback-Leibler divergence between two distributions  $p(\mathbf{x})$  and  $q(\mathbf{x})$  is defined as:

$$\text{KL}(q||p) = \sum_{\mathbf{x}} q(\mathbf{x}) \ln(q(\mathbf{x})/p(\mathbf{x})). \quad (6)$$

It represents a measure of the distance between two distributions (non-symmetric with respect to the exchange  $q \leftrightarrow p$ ). In the case discussed here,  $q(\mathbf{x})$  is identified with the distribution of the random walkers obtained via PQMC simulations at equilibrium, while  $p(\mathbf{x})$  with the RBM marginal distribution  $P(\mathbf{x})$ . The optimization of the RBM parameters  $\mathbf{W}$  can be performed using the gradient ascent algorithm. It consists in performing iterative updates starting from an initial (random) guess  $\mathbf{W}^0$ . At the step  $n = 0, 1, \dots, N_{\text{steps}}$ , one applies the rule:  $W_m^{n+1} = W_m^n + \eta \frac{\partial}{\partial W_m} L(\mathbf{W}^n)$ , where the coefficient  $\eta$  is the learning rate. This plain vanilla rule can be improved in various ways, e.g., by including a momentum term proportional to the update performed at the  $n - 1$  step, by annealing the learning rate, by adopting the adaptive gradient algorithm (AdaGrad) [62] or the adaptive moment estimation algorithm (Adam) [63]. For all results reported in this manuscript, the plain vanilla rule is augmented only by adding a momentum term corresponding to  $\nu \frac{\partial}{\partial W_m} L(\mathbf{W}^{n-1})$ , with  $\nu$  a rate tuned as discussed below. Also the learning-rate annealing is adopted, as explained in the following.

The main task in the implementation of the gradient ascent algorithm is the computation of the gradients of the log-likelihood. It is common to use stochastic estimates computed on mini-batches of  $N_b$  instances (typi-

cally,  $N_b \sim 10 - 100$ ) randomly sampled from the (much larger) training set. The formula for the gradients with respect to the couplings  $J_{ij}$  is usually stated in the following from:

$$\frac{\partial L(\mathbf{W})}{\partial J_{ij}} \propto \langle x_j h_i \rangle_{\text{data}} - \langle x_j h_i \rangle_{\text{model}}. \quad (7)$$

The first term on the right-hand side indicates the average obtained when the visible units are clamped to the data in the mini-batch. To determine its value, one uses the probability distribution of the hidden spins  $h_i$ . Due to the bipartite connectivity, it depends only on the visible-spin configurations. One has  $h_i = 1$  with probability  $p_{h_i=1}(\mathbf{x}) = 1 / [1 + \exp(-2 \sum_j x_j J_{ij} - 2b_i)]$ , and  $h_i = -1$  with probability  $1 - p_{h_i=1}(\mathbf{x})$ .

Determining the second term on the right-hand side of eq. (7) is more challenging. It represents the average obtained when both visible and hidden variables are sampled according to the probability distribution  $P(\mathbf{x}, \mathbf{h})$  defined by the RBM model. In machine learning jargon, this term is often referred to as the dreaming phase of the learning process, in contrast to the first term which would correspond to the awake phase. The average can be determined via Monte Carlo estimation, starting from the visible variables corresponding to the mini-batch instances, and then alternating Gibbs sampling of all hidden variables with fixed visible variables, followed by the sampling of all visible variables with hidden variables fixed at the previously sampled values. The probability to sample  $x_j = 1$  is analogous to the formula for the hidden spins:  $p_{x_j=1}(\mathbf{h}) = 1 / [1 + \exp(-2 \sum_i h_i J_{ij} - 2a_j)]$ .

In principle, this alternated Gibbs sampling should be iterated till the Markov chain equilibrates. This would provide an unbiased estimate of the average. In practice, repeating a sufficient number of iterations to guarantee that equilibrium has been achieved is often computationally overwhelming. The algorithm corresponding to a finite number  $k$  of iterations is referred to as  $k$ -step contrastive divergence, since it can be derived by minimizing the difference of two Kullback-Leibler divergences [64]. In the  $k \rightarrow \infty$  limit, this algorithm is provably unbiased, meaning that it corresponds to maximizing the log-likelihood of the dataset [65]. However, it often turns out to be very accurate also for small  $k$ . Formulas for the log-likelihood derivatives with respect to the other RBM parameters, namely the biases  $a_j$  and  $b_i$ , can also be derived, analogously to eq. (7). Rather than reporting them here, we provide in *Algorithm 1* the detailed procedure, which is adapted from Ref. [61] to the case of interest to us with binary values  $x_j, h_i = \pm 1$ , rather than the values 1 and 0, more common in the machine-learning literature. The input of the algorithm is a mini-batch of  $N_b$  randomly sampled instances, while the output provides the partial derivatives of the log-likelihood  $L_{W_m} \equiv \frac{\partial L(\mathbf{W})}{\partial W_m}$ . It is convenient to normalize the output, dividing by  $N_b$ . For a recent review on the topic, the interested reader is referred to Ref. [66].

**Algorithm 1**  $k$ -step contrastive divergence.

---

Input: mini-batch  $\mathbf{x}_l$ , with  $l = 1, \dots, N_b$   
Output: partial derivatives  $L_{J_{ij}}, L_{b_i}, L_{a_j}$ 


---

Initialization:  $L_{J_{ij}} = L_{b_i} = L_{a_j} = 0$ 


---

**for**  $l = 1, \dots, N_b$  **do**  
   $\mathbf{x}^0 \leftarrow \mathbf{x}_l$   
  **for**  $t = 0, \dots, k - 1$  **do**  
    set  $h_i^t = 1$  w. prob.  $p_{h_i=1}(\mathbf{x}^t)$ , else  $h_i^t = -1$   
    set  $x_j^{t+1} = 1$  w. prob.  $p_{x_j=1}(\mathbf{h}^t)$ , else  $x_j^{t+1} = -1$   
  **end for**  
  **for**  $i = 1, \dots, N_h$  and  $j = 1, \dots, N$  **do**  
     $L_{J_{ij}} \leftarrow L_{J_{ij}} + (2p_{h_i=1}(\mathbf{x}^0) - 1)x_j^0 - (2p_{h_i=1}(\mathbf{x}^k) - 1)x_j^k$   
     $L_{a_j} \leftarrow L_{a_j} + x_j^0 - x_j^k$   
     $L_{b_i} \leftarrow L_{b_i} + (2p_{h_i=1}(\mathbf{x}^0) - 1) - (2p_{h_i=1}(\mathbf{x}^k) - 1)$   
  **end for**  
**end for**


---

In the adaptive scheme presented in this paper, the above unsupervised learning algorithm is used to train the RBM to describe the (unnormalized) probability distribution  $f(\mathbf{x}, \tau \rightarrow \infty)$  corresponding to a large random-walker population produced by a PQMC simulation at equilibrium. It is in fact known that, in principle, for a sufficiently large  $N_h$  an RBM can approximate any discrete distribution (in the sense of the Kullback-Leibler divergence), and many researchers have found that the training algorithms described above are capable of finding optimal, or close to optimal, RBM parameters.

The scheme we propose here involves many consecutive stints, labelled in the following by the index  $s = 0, 1, 2, \dots, N_{\text{stints}} - 1$ , each including a PQMC simulation with a guiding function  $\psi_{G_s}(\mathbf{x})$  for a long imaginary time  $\tau_s$ , followed by a learning stage which is used to construct the new guiding function  $\psi_{G_{s+1}}(\mathbf{x})$ . The imaginary time  $\tau$  runs up to  $\tau_s N_{\text{stints}}$ . In each learning stage, occurring at the imaginary times  $(s + 1)\tau_s$ , the RBM is trained to describe the equilibrium random-walker distribution produced during the last stint. For the initial PQMC stint  $s = 0$ , a quite crude guiding wave function  $\psi_{G_{s=0}}(\mathbf{x})$  can be chosen, e.g., a constant function or a wave function based on (the square root of) an RBM with random parameters,  $\psi_{G_{s=0}}(\mathbf{x}) \propto \sqrt{P(\mathbf{x})}$ . After the  $s$  stint, the learned RBM distribution is  $P_s(\mathbf{x}) \propto \psi_{G_s}(\mathbf{x})\psi_0(\mathbf{x})$  (assuming  $N_w$  and  $\tau_s$  large enough). Notice that the learning process converges much faster if in each learning stage the RBM parameters are initialized at the values found in the previous learning stage.

The new guiding function for the stint  $s + 1$  is obtained by setting  $\psi_{G_{s+1}}(\mathbf{x}) = \sqrt{P_s(\mathbf{x})}$ . With this choice the accuracy of the guiding function increases stint after stint. Indeed, under the idealized conditions of  $N_h$  sufficiently large, successful unsupervised optimization, and  $N_w$  and  $\tau_s$  large enough, at the  $s$  stint the random walkers sample the (unnormalized) distribution  $f(\mathbf{x}, \tau \rightarrow \infty) = \psi_{G_{s=0}}^{1/2^s}(\mathbf{x})\psi_0^{2^{-1/2^s}}(\mathbf{x})$ . This indicates an extremely rapid convergence  $\psi_{G_{s \rightarrow \infty}}(\mathbf{x}) \rightarrow \psi_0(\mathbf{x})$ . As a consequence, in less idealized conditions, the (possible)

bias due to the finite  $N_w$  is expected to disappear after a relatively small number of stints/training stages. In the next section, we demonstrate that this is indeed the case. This is achieved with an affordable number of hidden spins  $N_h$ , using a simple implementation of the stochastic gradient ascent algorithm described above.

#### D. Simulation details

For all results reported in this article, the simulation details are the following. The PQMC simulations are performed with a target number of walkers  $N_w = 10^4$  and time step  $\Delta\tau = 0.04$ . In each stint, the PQMC simulation runs for an imaginary time  $\tau_s = 20$ , which turns out to be sufficiently large to represent the infinite imaginary-time limit  $\tau \rightarrow \infty$ . A small (randomly selected) fraction of the walker population, namely  $N_w/20$ , is stored at each PQMC step for the learning stage, excluding the initial time segment of each stint corresponding to  $\tau \in [s\tau_s, s\tau_s + 8]$ . This avoids correlations and non-equilibrium effects. This protocol provides  $N_{\text{train}} \simeq 15 \times 10^4$  instances for each unsupervised learning stage. The number of stints ranges from  $N_{\text{stints}} = 20$  to  $N_{\text{stints}} = 50$ . This appears to be sufficient to approach the  $s \rightarrow \infty$  limit. Unsupervised learning is performed after each PQMC stint using the  $k$ -step contrastive divergence algorithm, performing a number  $N_{\text{steps}} = 5 \times 10^4$  to  $10^5$  of stochastic gradient ascent steps, computed on mini-batches of size  $N_b = 10$  to  $50$ . The coefficient of the momentum term is  $\nu = \eta/10$  (see description above). The learning rate  $\eta$  is kept fixed within each learning stage, but is reduced stage after stage following the simple empirical annealing protocol  $\eta(s) = \eta_0 0.75^s$ , where  $\eta_0 = 0.01$  is the learning rate at the first learning stage  $s = 0$ . Remarkably,  $k = 1$  is found to suffice. Tests with  $k \approx 30$  provide comparable results. This strongly suggests that in this problem the  $k$ -step contrastive divergence algorithm is correctly minimizing the Kullback-Leibler divergence. The initial guiding function  $\psi_{G_{s=0}}(\mathbf{x})$  is the square root of an RBM distribution with uniform random couplings  $J_{ij} \in [-0.025 : 0.025]$ . In the absence of longitudinal magnetic field, i. e. with  $h = 0$ , the bias terms are initialized to zero, while for finite  $h$  they are set to uniform random values  $a_j, b_i \in [0 : 0.05]$ . It is worth mentioning that, while the setup described here turns out to be adequate for the problems addressed in the present work, it is plausible that even more efficient simulations could be implemented, e.g., by performing more frequent learning stages, storing larger training sets, or changing the optimization algorithm.

### III. RESULTS

This section focuses on verifying the accuracy of the self-learning PQMC scheme described in the previous




FIG. 2. (color online). Self-learning PQMC simulation of a ferromagnetic quantum Ising chain with  $N = 80$  spins and open boundary conditions. The transverse-field intensity is  $\Gamma = 0.85$ , the longitudinal-field intensity is  $h = 0$ . The energy per spin  $E/N$  is plotted as a function of the imaginary time  $\tau$ . The (violet) dashed curve indicates the running average in a simple PQMC simulation performed without a guiding wave function. The (dark-green) dot-dashed curve corresponds to a PQMC simulation guided by a fixed RBM wave function  $\psi_{G_{s=0}}(\mathbf{x})$  with  $N_h = 20$  hidden spins and random parameters. The (red) continuous curve corresponds to the PQMC simulation guided by the adaptive RBM wave function  $\psi_{G_s}(\mathbf{x})$  trained with unsupervised learning. The (blue) squares indicate the average energy  $E_{\psi_{G_s}}$  corresponding to the adaptive RBM wave function  $\psi_{G_s}(\mathbf{x})$  obtained after each stint  $s$ . The (brown) horizontal bar indicates the minimal variational energy  $E_{\text{var.min.}}$  for  $N_h = 20$  obtained with the NetKet library [47] using the stochastic reconfiguration algorithm. The (black) horizontal line corresponds to the exact ground-state energy computed via Jordan-Wigner transformation [67].

Section. The ferromagnetic quantum Ising chain (1) is used as a testbed. The PQMC predictions for the ground-state energy are compared against the exact value computed via the Jordan-Wigner transformation [67]. Notice that the comparison is made with the Jordan-Wigner results corresponding to the same system size and the same (open) boundary condition employed in the PQMC simulations. Let's first focus on a chain with  $N = 80$  spins, in the absence of a longitudinal field (i.e., with  $h = 0$ ), with a transverse field of intensity  $\Gamma = 0.85$ . This value is in the ferromagnetic phase, close to the quantum critical point, which corresponds to  $\Gamma = 1$ . This is the regime where the systematic bias due to the finite random-walker population is more sizeable [16, 24]. [For moderately large  $N_w$ , the bias is larger around  $\Gamma \approx 0.85$ . However, in the  $N_w \rightarrow \infty$  limit, the slowest vanishing of this bias occurs at the quantum critical point  $\Gamma = 1$ , in the thermodynamic limit.]

Indeed, as shown in Fig. 2, a simple PQMC simulation performed without guiding wave function, i.e., with




FIG. 3. (color online). Relative error  $(E - E_{\text{JW}}) / |E_{\text{JW}}|$  of the estimated ground-state energy  $E$  with respect to the (exact) Jordan-Wigner result  $E_{\text{JW}}$ , as a function of the hidden-spin number  $N_h$ . The system parameters are:  $N = 80$ ,  $\Gamma = 1$ , and  $h = 0$ . The (red) squares indicate the energy expectation value of the optimal adaptive RBM wave-function  $\psi_{G_{s \rightarrow \infty}}(\mathbf{x})$ , optimized via unsupervised learning along the PQMC simulation. This optimization corresponds to the minimization of the Kullback-Leibler divergence. The (blue) circles correspond to the minimal variational energy for an RBM wave-function obtained with the NetKet library using the stochastic reconfiguration algorithm. As a reference, the minimal variational energy corresponding to the unrestricted Boltzmann machine (uRBM) wave function [24] is indicated by the horizontal (green) bar. The width represents (twice) the statistical error. The uRBM wave function is optimized using the stochastic gradient descent algorithm.

$\psi_G(\mathbf{x}) = 1$ , provides energies with a considerable upward bias of  $\sim 1 - 2\%$ . This bias occurs despite the relatively large population of random walkers  $N_w = 10^4$ . To understand the occurrence of this large bias, one should consider that the constant wave function  $\psi_G(\mathbf{x}) = 1$  describes the ground state in the  $\Gamma \rightarrow \infty$  limit, where all spins align along the transverse axis and all elements of the computational basis have equal probability amplitude. Thus, for  $\Gamma = 0.85$ , the constant function represents a rather poor guiding wave-function, since here the ground state is drastically different, displaying long range ferromagnetic correlations along the longitudinal axis. Fig. 2 also displays a PQMC simulation guided by a fixed RBM wave function  $\psi_{G_{s=0}}(\mathbf{x})$  with  $N_h = 20$  hidden units and randomly chosen parameters (see previous section for more details). This simulation provides a comparable, but slightly smaller, bias compared to the  $\psi_G(\mathbf{x}) = 1$  case. Instead, the self-learning PQMC simulations guided by the adaptive wave function  $\psi_{G_s}(\mathbf{x})$  display a radically different behavior. In the first stint ( $s = 0$ ), the guiding wave function  $\psi_{G_{s=0}}(\mathbf{x})$  is the RBM with random parameters described in the previous section; the systematic bias is, again, noticeable. However, this bias completely disappears after a few



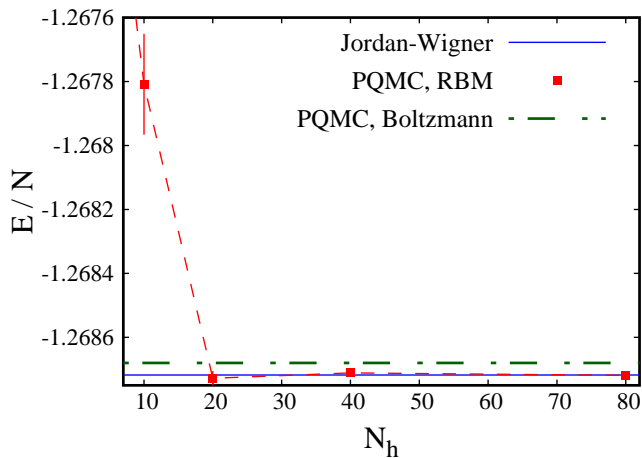


FIG. 4. (color online). Energy per spin  $E/N$  obtained from a PQMC simulation guided by the optimized adaptive RBM wave functions  $\psi_{G_{s \rightarrow \infty}}(\mathbf{x})$  with different numbers of hidden spins  $N_h$ . The system parameters are:  $N = 80$ ,  $\Gamma = 1$ , and  $h = 0$ . The (blue) horizontal line indicates the exact ground-state energy computed via the Jordan-Wigner (JW) transformation. For comparison, the (green) dot-dashed line indicates the results of a PQMC simulation guided by an optimized Boltzmann-type ansatz.

stints/learning-stages  $s \gtrsim 3$ . Not only the bias disappears, but also the statistical fluctuations are drastically suppressed, leading to an extremely accurate and precise agreement with the Jordan-Wigner prediction. This indicates that the RBM is capable of accurately learning the random-walker distribution, leading to an adaptive guiding wave function  $\psi_{G_{s \rightarrow \infty}}(\mathbf{x})$  sufficiently similar to the ground-state  $\psi_0(\mathbf{x})$  to eliminate the systematic bias in the PQMC simulation. To quantify the accuracy of the learned adaptive wave functions  $\psi_{G_s}(\mathbf{x})$ , we compute the corresponding energy expectation values

$$E_{\psi_{G_s}} \equiv \frac{\langle \psi_{G_s} | \hat{H} | \psi_{G_s} \rangle}{\langle \psi_{G_s} | \psi_{G_s} \rangle}. \quad (8)$$

These computations are performed via Monte Carlo integration using a simple Metropolis sampling algorithm. The corresponding results are displayed in Fig. 2, placed at the final imaginary-time of the corresponding PQMC stint  $\tau = \tau_s(s+1)$ . For small  $s$ , due to the memory of the initial guiding wave function  $\psi_{G_{s=0}}(\mathbf{x})$  and, to a lesser extent, to the bias affecting the PQMC algorithm, this expectation value largely deviates from the exact ground-state energy. This deviation might also originate from limitations in the learning process due to the finite number of learning updates  $N_{\text{steps}}$  and/or of the instances in the training set  $N_{\text{train}}$ . After a few stints/learning-stages the deviation from the ground-state energy significantly reduces, indicating that the adaptive RBM wave function better approximates the ground state. Remarkably, after a number  $s \gg 1$  of stints and learning stages,  $E_{\psi_{G_s}}$  approaches the optimal variational estimate  $E_{\text{var. min.}}$  (de-




FIG. 5. (color online). Main panel: energy per spin  $E/N$  as a function of the system size  $N$ , at the quantum critical point  $\Gamma = 1$  with  $h = 0$ . The (red) dots and interpolating line indicate the exact Jordan-Wigner result. The (blue) circles indicate the results of PQMC simulation guided by the optimized adaptive RBM wave function  $\psi_{G_{s \rightarrow \infty}}(\mathbf{x})$  with  $N_h = N/2$  hidden spins. Inset: ratio of the PQMC result  $E_{\text{PQMC}}$  and the exact (Jordan-Wigner) ground-state energy  $E_{\text{JW}}$ .

fined in the previous section) for an RBM wave function obtained using the NetKet library [47] for the same value of  $N_h$ . This library implements various sophisticated optimization algorithms. The results reported here are obtained using the stochastic reconfiguration method, which appears to be the most effective algorithm for the model addressed here. The agreement with the NetKet variational prediction indicates that the unsupervised learning algorithm is capable of identifying an optimal, or close to optimal, ground-state wave function from the random walkers sampled by the PQMC algorithm. It is worth mentioning that the convergence to the optimal RBM wave function could be accelerated by a more clever (i.e., not random) choice of the parameters of the initial guiding wave function  $\psi_{G_{s=0}}(\mathbf{x})$ . This would reduce the required number of stints. An effective initialization of the RBM parameters could be obtained by exploiting the mapping between the RBM wave functions and other ansatzes, such as Jastrow wave functions or correlator product states [45, 68]

The adaptive RBM wave function  $\psi_{G_s}(\mathbf{x})$  with  $N_h = 20$  is sufficiently accurate to eliminate the bias and to boost the efficiency of the PQMC simulation. However, the corresponding energy expectation value  $E_{\psi_{G_s}}$  does not precisely agree with the exact ground-state energy, even after a number of stints  $s \rightarrow \infty$ . As shown in Fig. 3, this residual deviation can be systematically reduced by increasing the number of hidden neurons  $N_h$  in the RBM. This analysis is shown for the quantum critical point  $\Gamma = 1$ , where it is more challenging to accurately approximate the ground-state with neural network ansatzes [23, 24, 39, 41]. The system size is, again,

$N = 80$ . One also observes that here the NetKet variational estimates  $E_{\text{var. min.}}$  for small  $N_h$  are slightly lower than  $E_{\psi_{G_s \rightarrow \infty}}$ . This might indicate that the stochastic reconfiguration algorithm is able to better optimize the RBM parameters. However, one should also consider that when a variational ansatz is not very accurate, and the exact ground state is not contained in the manifold defined by the variational parameters, minimizing the energy expectation value does not necessarily provide a better overall description of the ground state. One cannot exclude that the adaptive RBM wave function  $\psi_{G_s \rightarrow \infty}(\mathbf{x})$  produced by the unsupervised learning protocol, which aims at minimizing the Kullback-Leibler divergence with respect to the PQMC samples rather than at minimizing the energy expectation value, would provide more accurate predictions of physical properties other than the energy. As a reference, Fig. 3 displays also the minimal variational energy corresponding to the unrestricted Boltzmann machine ansatz of Ref. [24]. This is obtained by minimizing the energy expectation value with an implementation of the stochastic gradient descent algorithm which samples the additional hidden spin variables. This variational estimate is comparable to the one corresponding to an RBM wave-function with  $N_h = 80 - 100$  hidden spins. More importantly, if the number of hidden units is too small, the adaptive RBM wave function might not be sufficiently accurate to eliminate the systematic bias of the PQMC simulation, even after many stints/learning-stages. This effect is visualized in Fig. 4, again for  $\Gamma = 1$ . Indeed, one notices that, when the adaptive guiding function has  $N_h \leq 10$  hidden spins, the PQMC results are affected by a small bias  $\sim 0.1\%$ . This bias completely disappears for  $N_h \geq 20$ , within statistical uncertainty. For comparison, Fig. 4 displays also the PQMC result obtained by using a conventional Boltzmann-type guiding wave function, defined as  $\psi_G(\mathbf{x}) = \exp(-\beta H_{\text{cl}})$ , where  $H_{\text{cl}} = -\sum_j^{N-1} x_j x_{j+1}$  is the Hamiltonian function of a classical Ising chain, and the fictitious temperature  $\beta$  is determined via variational energy minimization. At  $\Gamma = 1$ , this optimization provides  $\beta = 0.274(1)$ . One notices that the Boltzmann guiding wave function is not sufficiently accurate to completely eliminate the bias in the PQMC result due to the finite  $N_w$ . Indeed, it was shown in Ref. [24] that, in the vicinity of the ferromagnetic critical point, the computational cost of PQMC simulations guided by Boltzmann ansatzes increases exponentially with the system size if the systematic bias is kept below a chosen small threshold. These findings further highlight the importance of implementing a systematically improvable ansatz, such as the adaptive RBM guiding wave functions adopted here.

By exploiting the adaptive RBM wave function, unbiased PQMC results can be obtained also for larger system sizes. Figure 5 shows data for different system sizes (up to  $N = 160$ ). With a number of hidden spins  $N_h = N/2$ , one sees from the inset that the relative error with respect to the exact result is as small as  $\sim 10^{-3}\%$  even for the largest system size.

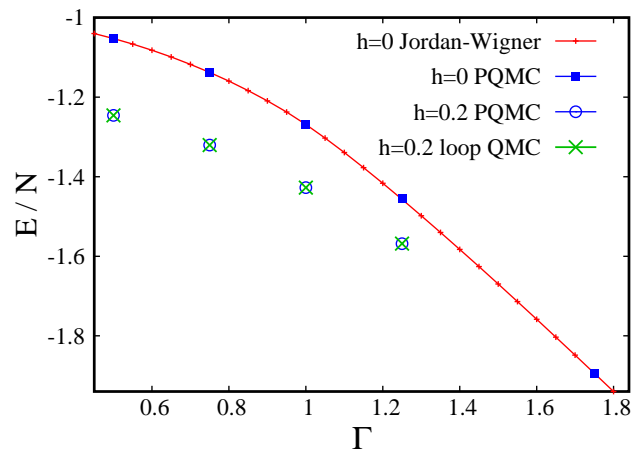


FIG. 6. (color online). Energy per spin  $E/N$  as a function of the transverse magnetic field  $\Gamma$ , for vanishing longitudinal magnetic field, i. e.  $h = 0$ , and for  $h = 0.2$ . The system size is  $N = 80$ . At  $h = 0$ , the results of PQMC simulations guided by the optimized adaptive RBM wave function  $\psi_{G_s \rightarrow \infty}(\mathbf{x})$  (blue squares) are compared with the Jordan-Wigner theory (red dots and interpolating line). At  $h = 0.2$ , comparison is made with loop QMC simulations performed at low temperature  $T = 0.05$  using the ALPS library [44].

The PQMC ground-state energy predictions are shown in Fig. 6 as a function of the transverse field  $\Gamma$ . This plot explores a moderately broad region around the critical point. The system size is  $N = 80$ . Perfect agreement with the Jordan-Wigner prediction is found over the whole region. Results obtained in the presence of a longitudinal field of intensity  $h = 0.2$  are also shown. In this case, the comparison is made against the predictions of loop QMC simulations [69] performed with the ALPS library [44]. Notice that the loop QMC algorithm provides finite-temperature results, but we consider here a sufficiently low temperature  $T = 0.05$  (we set  $k_B = 1$ ) such that thermal effects are negligible. In the PQMC simulations, the guiding function is, for both  $h = 0$  and  $h = 0.2$ , an optimized adaptive RBM wave function  $\psi_{G_s \rightarrow \infty}(\mathbf{x})$  with  $N_h = 40$  hidden spins. For  $h = 0$ , the bias terms of the RBM are set to zero; in fact, we find that optimizing them does not lead to a more accurate guiding wave function. This is expected, since nonzero bias terms would break the  $Z_2$  symmetry of the (finite system) ground state. Instead, for  $h = 0.2$  the unsupervised learning algorithm leads to sizable bias terms.

The PQMC algorithm provides predictions also for observables other than the energy. In Fig. 7, the estimates of the ground-state magnetization  $M \equiv \langle \psi_0 | \sum_j \sigma_j^z | \psi_0 \rangle$ , obtained using the forward-walking technique, are shown as a function of the transverse field. A finite longitudinal field  $h = 0.2$  is included, so that the  $Z_2$  symmetry, which would lead to a vanishing average magnetization in finite systems, is broken. The adaptive RBM wave function has

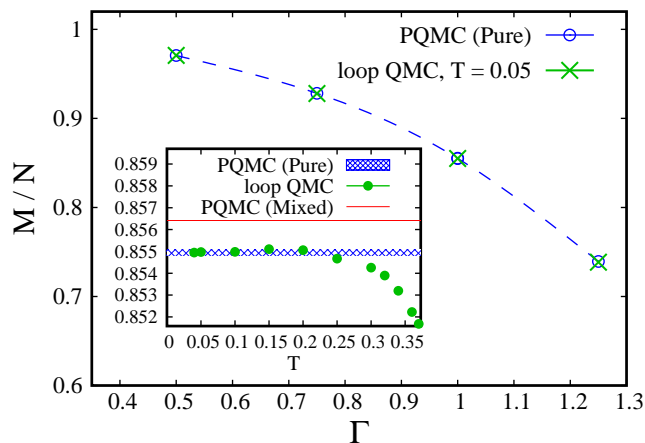


FIG. 7. (color online). Main panel: magnetization per spin  $M/N$  as a function of the transverse magnetic field  $\Gamma$ . The system size is  $N = 80$  and the longitudinal-field intensity is  $h = 0.2$ . The PQMC results obtained with the pure estimator (blue empty circles) are compared with the results of loop QMC simulations (green crosses) performed with the ALPS library at the temperature  $T = 0.05$ . Inset: loop QMC results (green bullets) for  $M/N$  as a function of the temperature  $T$ , for  $\Gamma = 1$  and  $h = 0.2$ . The (red) horizontal thin bar corresponds to the zero-temperature PQMC prediction obtained using the mixed estimator. The horizontal bar with (blue) diagonal pattern corresponds to the pure estimator.

$N_h = N/2 = 40$  hidden units. The agreement with the loop QMC predictions is extremely precise. This can be appreciated in the inset of Fig. 7, where the loop QMC results are plotted as a function of the temperature  $T$ , showing that thermal effects become indeed irrelevant at the lowest temperatures we consider. The inset displays also the prediction provided by the PQMC mixed estimator (see the definition in the previous section). This is affected by a remarkably small bias  $\sim 0.1\%$ , indicating that the optimized adaptive RBM guiding function is a good approximation for the ground state. As expected, the pure estimator removes even this small bias.

#### IV. CONCLUSIONS

Machine learning techniques have been employed in the recent past to implement smart updates in classical and in path-integral Monte Carlo simulations [70–74], leading to a substantial efficiency improvement. In this paper, we used them to boost the performance of QMC simulations based on projective algorithms. Specifically, we have introduced a self-learning PQMC algorithm in which the guiding wave function is adaptively developed along the PQMC simulation, avoiding the need of a separate variational optimization. This adaptive guiding wave function is based on an RBM, and it is optimized via unsupervised machine learning using the  $k$ -step contrastive divergence

algorithm. If the number of hidden units is sufficiently large, this adaptive guiding wave function allows one to eliminate any systematic bias in the PQMC simulation and to drastically boost its performance by suppressing the statistical fluctuations.

The self-learning PQMC algorithm appears to be particularly suitable when it is otherwise challenging to guess an accurate ansatz for the ground-state wave function. This is the case, e.g., of the disordered Ising Hamiltonians that describe the combinatorial optimization problems commonly addressed in the field of adiabatic quantum optimization. As a side product of the self-learning PQMC simulations, one obtains an accurate approximation for the ground-state wave function written in the form of an RBM. By construction, this function minimizes the Kullback-Leibler divergence with respect to the random-walker distribution. We argue that this unsupervised learning approach is complementary to other common techniques for wave function optimization based on the variational minimization of energy expectation values. Following Ref. [23], one might refer to these latter techniques as reinforcement learning. It is indeed plausible that, when the available ansatz is not sufficiently flexible to exactly describe the ground state, the minimization of the Kullback-Leibler divergence might lead to a fairer description of ground state properties other than the energy. Our unsupervised learning approach is also complementary to the optimization techniques for neural network states based on supervised machine learning introduced very recently in Ref. [37].

The self-learning PQMC algorithm has been tested in the ferromagnetic quantum Ising chain, obtaining excellent agreement with the Jordan-Wigner theory and with the loop QMC simulations performed in the low temperature limit. The model considered here is not affected by a sign problem. As a perspective, one can envision the extension of the self-learning PQMC technique to models where the sign problem occurs. Examples are many-fermion models and other non-stoquastic Hamiltonians relevant in adiabatic quantum computing. Fermionic PQMC simulations can be implemented using random walkers that carry a sign. They have been performed in the past adopting fixed-node and released node approaches, eventually performing annihilation of walkers with opposite sign [5] to alleviate the suppression of the signal-to-noise ratio associated to the sign problem. In such fermionic simulations, the adaptive RBM guiding function would remain nonnegative, since its aim is to learn the probability distribution of the random walkers. It might allow one to further alleviate the sign problem by learning the nodal regions, and so reducing the random-walker crossings that lead to the suppression of the signal-to-noise ratio. We leave these endeavors to future investigations.

## ACKNOWLEDGEMENTS

We acknowledge useful discussions with Ferran Mazzanti, Juan Carrasquilla, Giacomo Torlai, and Giuseppe Santoro. E. M. I. thanks Dustin Lang for his help in performing the simulations on the Perimeter Institute HPC architectures.

S. P. and P. P. acknowledge financial support from the FAR2018 project of the University of Camerino and from the Italian MIUR under the project PRIN2017 CEn-

TraL 20172H2SC4. S. P. also acknowledges the CINECA award under the ISCRA initiative, for the availability of high performance computing resources and support. 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. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958.

- 
- [1] J. B. Anderson, “A random-walk simulation of the Schrödinger equation:  $H+3$ ,” *J. Chem. Phys.* **63**, 1499–1503 (1975).
- [2] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester Jr, “Fixed-node quantum Monte Carlo for molecules,” *J. Chem. Phys.* **77**, 5593–5603 (1982).
- [3] D. M. Ceperley and B. Alder, “Quantum Monte Carlo for molecules: Greens function and nodal release,” *J. Chem. Phys.* **81**, 5833–5844 (1984).
- [4] D. M. Ceperley and B. Alder, “Quantum Monte Carlo,” *Science* **231**, 555–560 (1986).
- [5] G. H. Booth, A. J. Thom, and A. Alavi, “Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space,” *J. Chem. Phys.* **131**, 054106 (2009).
- [6] D. M. Ceperley and B. J. Alder, “Ground state of the electron gas by a stochastic method,” *Phys. Rev. Lett.* **45**, 566 (1980).
- [7] J. Boronat, “Monte Carlo simulations at zero temperature: helium in one, two, and three dimensions,” in *Microscopic Approaches to Quantum Liquids in Confined Geometries*, edited by E. Krotscheck and J. Navarro (World Scientific, Singapore, 2002, 2002) Chap. 2, pp. 21–90.
- [8] W. Foulkes, L. Mitas, R. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Rev. Mod. Phys.* **73**, 33 (2001).
- [9] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, “Towards an exact description of electronic wavefunctions in real solids,” *Nature* **493**, 365 (2013).
- [10] S. Giorgini, J. Boronat, and J. Casulleras, “Ground state of a homogeneous Bose gas: A diffusion Monte Carlo calculation,” *Phys. Rev. A* **60**, 5129 (1999).
- [11] N. Trivedi and D. M. Ceperley, “Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study,” *Phys. Rev. B* **41**, 4552 (1990).
- [12] M. Calandra Buonauro and S. Sorella, “Numerical study of the two-dimensional heisenberg model using a green function monte carlo technique with a fixed number of walkers,” *Phys. Rev. B* **57**, 11446–11456 (1998).
- [13] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa, “Quantum Monte Carlo methods for nuclear physics,” *Rev. Mod. Phys.* **87**, 1067 (2015).
- [14] N. Nemeč, “Diffusion Monte Carlo: Exponential scaling of computational cost for large systems,” *Phys. Rev. B* **81**, 035119 (2010).
- [15] M. Boninsegni and S. Moroni, “Population size bias in diffusion Monte Carlo,” *Phys. Rev. E* **86**, 056712 (2012).
- [16] E. M. Inack, G. Giudici, T. Parolini, G. Santoro, and S. Pilati, “Understanding quantum tunneling using diffusion Monte Carlo simulations,” *Phys. Rev. A* **97**, 032307 (2018).
- [17] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, “Theory of quantum annealing of an Ising spin glass,” *Science* **295**, 2427–2430 (2002).
- [18] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, “Evidence for quantum annealing with more than one hundred qubits,” *Nat. Phys.* **10**, 218–224 (2014).
- [19] E. M. Inack and S. Pilati, “Simulated quantum annealing of double-well and multiwell potentials,” *Phys. Rev. E* **92**, 053304 (2015).
- [20] S. Boixo, V. N. Smelyanskiy, A. Shabani, S. V. Isakov, M. Dykman, V. S. Denchev, M. H. Amin, A. Y. Smirnov, M. Mohseni, and H. Neven, “Computational multiqubit tunnelling in programmable quantum annealers,” *Nat. Comm.* **7**, 10327 (2016).
- [21] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, “Quantum versus classical annealing of Ising spin glasses,” *Science* **348**, 215–217 (2015).
- [22] L. Stella and G. E. Santoro, “Quantum annealing of an Ising spin-glass by Greens function Monte Carlo,” *Phys. Rev. E* **75**, 036703 (2007).
- [23] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science* **355**, 602–606 (2017).
- [24] E. M. Inack, G. E. Santoro, L. Dell’Anna, and S. Pilati, “Projective quantum Monte Carlo simulations guided by unrestricted neural network states,” *Phys. Rev. B* **98**, 235145 (2018).
- [25] L. Reatto and G. L. Masserini, “Shadow wave function for many-boson systems,” *Phys. Rev. B* **38**, 4516–4522 (1988).
- [26] S. Vitiello, K. Runge, and M. H. Kalos, “Variational calculations for solid and liquid He 4 with a shadow wave function,” *Phys. Rev. Lett.* **60**, 1970 (1988).
- [27] D.-L. Deng, X. Li, and S. Das Sarma, “Quantum entanglement in neural network states,” *Phys. Rev. X* **7**, 021021 (2017).
- [28] X. Gao and L.-M. Duan, “Efficient representation of quantum many-body states with deep neural networks,” *Nat. Commun.* **8**, 662 (2017).
- [29] H. Saito, “Solving the Bose-Hubbard Model with Machine Learning,” *J. Phys. Soc. Jpn.* **86**, 093001 (2017).



- [30] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, “Equivalence of restricted Boltzmann machines and tensor network states,” *Phys. Rev. B* **97**, 085104 (2018).
- [31] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, “Neural-network quantum states, string-bond states, and chiral topological states,” *Phys. Rev. X* **8**, 011006 (2018).
- [32] N. Freitas, G. Morigi, and V. Dunjko, “Neural Network Operations and Susuki-Trotter evolution of Neural Network States,” arXiv preprint arXiv:1803.02118 (2018).
- [33] Z. Cai and J. Liu, “Approximating quantum many-body wave functions using artificial neural networks,” *Phys. Rev. B* **97**, 035116 (2018).
- [34] 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).
- [35] G. Carleo, Y. Nomura, and M. Imada, “Constructing exact representations of quantum many-body systems with deep neural networks,” *Nature Communications* **9**, 5322 (2018).
- [36] D. Luo and B. K. Clark, “Backflow transformations via neural networks for quantum many-body wavefunctions,” arXiv preprint arXiv:1807.10770 (2018).
- [37] D. Kochkov and B. K. Clark, “Variational optimization in the AI era: Computational graph states and supervised wave-function optimization,” arXiv preprint arXiv:1811.12423 (2018).
- [38] M. Ruggeri, S. Moroni, and M. Holzmann, “Nonlinear Network Description for Many-Body Quantum Systems in Continuous Space,” *Phys. Rev. Lett.* **120**, 205302 (2018).
- [39] M. J. Beach, R. G. Melko, T. Grover, and T. H. Hsieh, “Making Trotters sprint: A variational imaginary time ansatz for quantum many-body systems,” arXiv preprint arXiv:1904.00019 (2019).
- [40] K. McBrien, G. Carleo, and E. Khatami, “Ground state phase diagram of the one-dimensional Bose-Hubbard model from restricted Boltzmann machines,” arXiv preprint arXiv:1903.03076 (2019).
- [41] M. Collura, L. Dell’Anna, T. Felser, and S. Montangero, “On the descriptive power of Neural-Networks as constrained Tensor Networks with exponentially large bond dimension,” arXiv preprint arXiv:1905.11351 (2019).
- [42] J. Kessler, F. Calcavecchia, and T. D. Kühne, “Artificial neural networks as trial wave functions for quantum monte carlo,” arXiv preprint arXiv:1904.10251 (2019).
- [43] A. Nagy and V. Savona, “Variational quantum monte carlo with neural network ansatz for open quantum systems,” arXiv preprint arXiv:1902.09483 (2019).
- [44] B. Bauer, L. Carr, H. G. Evertz, A. Feiguin, J. Freire, S. Fuchs, L. Gamper, J. Gukelberger, E. Gull, S. Guertler, *et al.*, “The ALPS project release 2.0: open source software for strongly correlated systems,” *J. Stat. Mech.: Theory Exp.* **2011**, P05001 (2011).
- [45] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, “Neural-network quantum state tomography,” *Nat. Phys.* **14**, 447–450 (2018).
- [46] S. Sorella, “Green function Monte Carlo with stochastic reconfiguration,” *Phys. Rev. Lett.* **80**, 4558–4561 (1998).
- [47] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, *et al.*, “Netket: A machine learning toolkit for many-body quantum systems,” arXiv preprint arXiv:1904.00031 (2019).
- [48] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
- [49] S. Sorella and L. Capriotti, “Green function Monte Carlo with stochastic reconfiguration: An effective remedy for the sign problem,” *Phys. Rev. B* **61**, 2599–2612 (2000).
- [50] One has to accumulate the weight factors of the updates occurring within the time-step  $\Delta\tau$ . See Refs. [24, 48, 49] for more details.
- [51] J. Thijssen, *Computational physics* (Cambridge University Press, 2007).
- [52] J. H. Hetherington, “Observations on the statistical iteration of matrices,” *Phys. Rev. A* **30**, 2713 (1984).
- [53] O. Golinelli, T. Jolicoeur, and R. Lacaze, “Haldane gaps in a spin-1 heisenberg chain with easy-plane single-ion anisotropy,” *Phys. Rev. B* **45**, 9798 (1992).
- [54] L. Pollet, N. V. Prokof’ev, and B. V. Svistunov, “Stochastic lists: Sampling multivariable functions with population methods,” *Phys. Rev. B* **98**, 085102 (2018).
- [55] In one and in two dimensional models with short-range interactions, the unrestricted Boltzmann machine ansatzes can be mapped to constrained matrix product states and to constrained tensor network states, respectively [41]. In one dimension this mapping allows one to implement alternative variational minimization techniques, which avoid the sampling over hidden spins.
- [56] S. A. Vitiello and P. A. Whitlock, “Green’s-function Monte Carlo algorithm for the solution of the Schrödinger equation with the shadow wave function,” *Phys. Rev. B* **44**, 7373–7377 (1991).
- [57] K. Sohn and H. Lee, “Learning invariant representations with local transformations,” in *Proceedings of the 29th International Conference on Computational Conference on Machine Learning* (Omnipress, 2012) pp. 1339–1346.
- [58] K. Choo, T. Neupert, and G. Carleo, “Study of the two-dimensional frustrated J1-J2 model with neural network quantum states,” arXiv preprint arXiv:1903.06713 (2019).
- [59] F. Ferrari, F. Becca, and J. Carrasquilla, “Neural Gutzwiller-projected variational wave functions,” arXiv preprint arXiv:1906.00463 (2019).
- [60] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, “A learning algorithm for Boltzmann machines,” *Cogn. Sci.* **9**, 147–169 (1985).
- [61] A. Fischer and C. Igel, “An introduction to restricted Boltzmann machines,” in *Iberoamerican congress on pattern recognition* (Springer, 2012) pp. 14–36.
- [62] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *J. Mach. Learn. Res.* **12**, 2121–2159 (2011).
- [63] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv preprint arXiv:1412.6980 (2014).
- [64] G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural Computation* **14**, 1771–1800 (2002).
- [65] Y. Bengio and O. Delalleau, “Justifying and generalizing contrastive divergence,” *Neural Comput.* **21**, 1601–1621 (2009).
- [66] G. Torlai and R. G. Melko, “Machine learning quantum states in the NISQ era,” arXiv e-prints , arXiv:1905.04312 (2019), arXiv:1905.04312 [quant-ph].
- [67] G. G. Cabrera and R. Jullien, “Role of boundary conditions in the finite-size Ising model,” *Phys. Rev. B* **35**,

- 7062 (1987).
- [68] S. R. Clark, “Unifying neural-network quantum states and correlator product states via tensor networks,” *J. Phys. A* **51**, 135301 (2018).
- [69] S. Todo and K. Kato, “Cluster algorithms for general-S quantum spin systems,” *Phys. Rev. Lett.* **87**, 047203 (2001).
- [70] L. Huang and L. Wang, “Accelerated Monte Carlo simulations with restricted Boltzmann machines,” *Phys. Rev. B* **95**, 035105 (2017).
- [71] 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(R) (2017).
- [72] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, “Self-learning Monte Carlo method,” *Phys. Rev. B* **95**, 041101(R) (2017).
- [73] H. Shen, J. Liu, and L. Fu, “Self-learning Monte Carlo with deep neural networks,” *Phys. Rev. B* **97**, 205140 (2018).
- [74] S. Li, P. M. Dee, E. Khatami, and S. Johnston, “Accelerating lattice quantum Monte Carlo simulation using artificial neural networks: an application to the Holstein model,” arXiv preprint arXiv:1905.07440 (2019).