

**Projective quantum Monte Carlo simulations guided by unrestricted neural network states**E. M. Inack,<sup>1,2,3</sup> G. E. Santoro,<sup>1,2,4</sup> L. Dell'Anna,<sup>5,6</sup> and S. Pilati<sup>7</sup><sup>1</sup>*The Abdus Salam International Centre for Theoretical Physics, 34151 Trieste, Italy*<sup>2</sup>*SISSA - International School for Advanced Studies, 34136 Trieste, Italy*<sup>3</sup>*INFN, Sezione di Trieste, 34136 Trieste, Italy*<sup>4</sup>*CNR-IOM Democritos National Simulation Center, Via Bonomea 265, 34136 Trieste, Italy*<sup>5</sup>*Dipartimento di Fisica e Astronomia "Galileo Galilei," Università di Padova, 35131 Padova, Italy*<sup>6</sup>*CNISM, Sezione di Padova, 35131 Padova, Italy*<sup>7</sup>*School of Science and Technology, Physics Division, Università di Camerino, 62032 Camerino (MC), Italy*

(Received 24 September 2018; published 21 December 2018)

We investigate the use of variational wave functions that mimic stochastic recurrent neural networks, specifically, unrestricted Boltzmann machines, as guiding functions in projective quantum Monte Carlo (PQMC) simulations of quantum spin models. As a preliminary step, we investigate the accuracy of such unrestricted neural network states as variational *Ansätze* for the ground state of the ferromagnetic quantum Ising chain. We find that by optimizing just three variational parameters, independently on the system size, accurate ground-state energies are obtained, comparable to those previously obtained using restricted Boltzmann machines with few variational parameters per spin. Chiefly, we show that if one uses optimized unrestricted neural network states as guiding functions for importance sampling, the efficiency of the PQMC algorithms is greatly enhanced, drastically reducing the most relevant systematic bias, namely, the one due to the finite random-walker population. The scaling of the computational cost with the system size changes from the exponential scaling characteristic of PQMC simulations performed without importance sampling, to a polynomial scaling, apparently even at the ferromagnetic quantum critical point. The important role of the protocol chosen to sample hidden-spin configurations, in particular at the critical point, is analyzed. We discuss the implications of these findings for what concerns the problem of simulating adiabatic quantum optimization using stochastic algorithms on classical computers.

DOI: [10.1103/PhysRevB.98.235145](https://doi.org/10.1103/PhysRevB.98.235145)**I. INTRODUCTION**

Quantum Monte Carlo (QMC) algorithms are generally believed to be capable of predicting equilibrium properties of quantum many-body systems at an affordable computational cost, even for relatively large system sizes, at least when the sign problem does not occur. However, it has recently been shown that the computational cost to simulate the ground state of a quantum Ising model with a simple projective QMC (PQMC) algorithm that does not exploit importance sampling techniques scales exponentially with the system size, making large-scale simulations unfeasible [1]. This happens in spite of the fact that the Hamiltonian is sign-problem free.

PQMC methods have found vast use in condensed matter physics, in chemistry, and beyond (see, e.g., Refs. [2–5]). Shedding light on their computational complexity, and possibly improving it by using importance sampling techniques based on novel variational wave functions, are therefore very important tasks. We address them in this paper.

PQMC algorithms have recently emerged as useful computational tools also to investigate the potential efficiency of adiabatic quantum computers in solving large-scale optimization problems via quantum annealing [6–10].

In particular, it has been shown that the (stochastic) dynamics of simple PQMC simulations allows to tunnel through

tall barriers of (effectively) double-well models even more efficiently than an adiabatic quantum computer, which exploits incoherent quantum tunneling [1,11–13]. This result seems to suggest that there might be no systematic quantum speed-up in using a quantum annealing device to solve an optimization problem, compared to a stochastic QMC simulation performed on a classical computer [11].

Remarkably, this computational advantage of the PQMC simulations with respect to the expected behavior of a quantum annealing device occurs also in more challenging models with frustrated couplings [1], as in the recently introduced Shamrock model, where QMC algorithms based on the (finite temperature) path-integral formalism display instead an exponential slowdown of the tunneling dynamics [14].

This result further stresses the importance of shedding light on the computational complexity of PQMC algorithms: if these computational techniques allowed one to simulate, with a polynomially scaling computational cost, both the ground-state properties of a model Hamiltonian, and also the tunneling dynamics of a quantum annealing device described by such Hamiltonian [1], then the quantum speedup mentioned above would be very unlikely to be achieved. We focus in this paper on the first of the two aspects, specifically, on analyzing and improving the scaling of the computational cost to simulate the ground-state properties of quantum Ising models.




FIG. 1. Structure of the unrestricted Boltzmann machine. The lower (yellow) nodes depict visible spins, the upper (magenta) nodes depict the hidden spins. The horizontal segments indicate intralayer visible-visible and hidden-hidden correlations. The vertical (blue) segments represent the interlayer correlations between the corresponding visible and hidden spins. The green lines allude to a possible extension to deep layers architectures.

It is well known that the efficiency of PQMC algorithms can be enhanced by implementing importance sampling techniques using as guiding functions accurate variational *Ansätze* [4]. However, building accurate variational wave functions for generic many-body systems is a highly nontrivial task. Recently, variational wave functions that mimic the structure of neural networks have been shown to accurately describe ground-state properties of quantum spin and lattice models [15–17]. The representational power and the entanglement content of such variational states, now referred to as neural network states, have been investigated [18–22], showing, among other properties, that they are capable of describing volume-law entanglement.

The authors of Ref. [15] considered neural network states that mimic restricted Boltzmann machines (RBM), i.e., such that no interaction among hidden spins is allowed. One very appealing feature of such restricted neural network states is that the role of the hidden spins can be accounted for analytically, without the need of Monte Carlo sampling over hidden variables. Furthermore, such states provide very accurate ground-state energy predictions, which can be systematically improved by increasing the number of hidden spins per visible spin (later on referred to as hidden-spin density). However, this high accuracy is obtained at the cost of optimizing a number of variational parameters that increases with the system size. This optimization task can be tackled using powerful optimization algorithms such as the stochastic reconfiguration method (see, e.g., Ref. [23]). Yet, having to optimize a large number of variational parameters is not desirable in the context of quantum annealing simulations, since one would be dealing with a variational optimization problem, potentially even more difficult than the original classical optimization problem.

In this paper, we consider instead neural network states that mimic unrestricted Boltzmann machines (uRBMs), allowing intralayer correlations among hidden spins, beyond the interlayer hidden-visible correlations and the intralayer visible-visible correlations (see Fig. 1). The structure of these states resembles the one of the shadow wave functions originally introduced to describe quantum fluids and solids [24,25]. We

test their representational power considering as a testbed the ferromagnetic quantum Ising chain. We find that by optimizing just three variational parameters, independently of the system size, very accurate ground-state energies are obtained, comparable to the case of restricted neural network states with one hidden spin per visible spin. Such a small number of variational parameters is a particularly appealing feature in the context of quantum annealing problems. However, it comes at the prize of having to perform Monte Carlo sampling over hidden-spin configurations.

The main goal of this paper is to show that the above-mentioned unrestricted neural network states can be used as a guide for importance sampling in PQMC simulations. This also implies that the development of neural network states can be limited to obtaining reasonably accurate, but not necessarily exact, variational *Ansätze*, since the residual error can be eliminated within the PQMC simulation. In particular, we provide numerical evidence that the major source of systematic bias of the PQMC algorithms, namely the bias originating from the finite size of the random-walker population which has to be stochastically evolved in any PQMC simulation, can be drastically reduced using optimized unrestricted neural network states, even at the point of changing the scaling of the required population size from exponential (corresponding to the case without importance sampling) to polynomial in the system size. This also implies a change of computational complexity from exponential to polynomial. For comparison, we show that a conventional variational wave function of the Boltzmann type (with no hidden spins), instead, does not determine a comparable efficiency improvement.

The rest of the paper is organized as follows. In Sec. II, we define the conventional Boltzmann-type variational wave functions and the unrestricted neural network states, and we then analyze how accurately they predict the ground-state energy of the quantum Ising chain via optimization of, respectively one and three, variational parameters. Section III deals with the continuous-time PQMC algorithm and with the implementation of importance sampling using both Boltzmann-type wave functions and, chiefly, unrestricted neural network states, showing how the systematic bias due to the finite random-walker population is affected, both at and away from the quantum critical point. The important effect of choosing different sampling protocols for the hidden spins is also analyzed. Our conclusions and the outlook are reported in Sec. IV.

## II. UNRESTRICTED NEURAL NETWORK STATES FOR QUANTUM ISING MODELS

In this paper, we consider as a test bed the one-dimensional ferromagnetic quantum Ising Hamiltonian:

$$\hat{H} = \hat{H}_{\text{cl}} + \hat{H}_{\text{kin}}, \quad (1)$$

where  $\hat{H}_{\text{cl}} = -J \sum_{i=1}^N \sigma_i^z \sigma_{i+1}^z$  and  $\hat{H}_{\text{kin}} = -\Gamma \sum_{i=1}^N \sigma_i^x \sigma_{i+1}^x$ ,  $\sigma_i^x$ ,  $\sigma_i^y$ , and  $\sigma_i^z$  indicate Pauli matrices acting on spins at the lattice site  $i$ .  $N$  is the total number of spins, and we adopt periodic boundary conditions, i.e.,  $\sigma_{N+1}^\alpha = \sigma_1^\alpha$ , with  $\alpha = x, y, z$ . The parameter  $J > 0$  fixes the strength of the ferromagnetic interactions among nearest-neighbor spins. In the following, we set  $J = 1$ . All energy scales are henceforth

expressed in units of  $J$ . The parameter  $\Gamma$  fixes the intensity of a transverse magnetic field. Given  $|x_i\rangle$  an eigenstate of the Pauli matrix  $\sigma_i^z$  with eigenvalue  $x_i = 1$  when  $|x\rangle = |\uparrow\rangle$  and  $x_i = -1$  when  $|x\rangle = |\downarrow\rangle$ , the quantum state of  $N$  spins is indicated by  $|\mathbf{x}\rangle = |x_1 x_2 \dots x_N\rangle$ . Notice that the function  $E_{\text{cl}}(\mathbf{x}) = \langle \mathbf{x} | \hat{H}_{\text{cl}} | \mathbf{x} \rangle$  [with  $\mathbf{x} = (x_1, x_2, \dots, x_N)$ ] corresponds to the Hamiltonian function of a classical Ising model, while the operator  $\hat{H}_{\text{kin}}$  introduces quantum (kinetic) fluctuations.

Our first goal is to develop trial wave functions that closely approximate the ground-state wave function  $\Psi_0(\mathbf{x}) = \langle \mathbf{x} | \Psi_0 \rangle$  of the Hamiltonian (1). A simple *Ansatz* can be defined as

$$\Psi_{\kappa}(\mathbf{x}) = e^{-\beta E_{\text{cl}}(\mathbf{x})} = e^{-K_1 \sum_{i=1}^N x_i x_{i+1}}. \quad (2)$$

$\kappa$  is here a set of real variational parameters to be optimized. Their values are obtained by minimizing the average of the energy, as in standard variational quantum Monte Carlo approaches. In this case, only one parameter  $K_1 = \beta$  is present,  $\kappa = \{K_1\}$ . This choice is inspired by the classical Boltzmann distribution where  $\beta$  would play the role of a fictitious inverse temperature. The above *Ansatz* will be referred to as Boltzmann-type wave function.

A more sophisticated *Ansatz* can be constructed by using a generative stochastic artificial neural network, namely an uRBM (see Fig. 1). Beyond the visible spin variables  $\mathbf{x} = (x_1, x_2, \dots, x_N)$ , one introduces  $N$  hidden spin variables  $\mathbf{h} = (h_1, h_2, \dots, h_N)$ , taking values  $h_i = \pm 1$  (with  $i = 1, \dots, N$ ). Periodic boundary conditions within the layers are also incorporated, i.e.,  $x_{N+1} = x_1$  and  $h_{N+1} = h_1$ . The trial wave function is thus written in the following integral form:

$$\Psi_{\kappa}(\mathbf{x}) = \sum_{\mathbf{h}} \phi_{\kappa}(\mathbf{x}, \mathbf{h}), \quad (3)$$

where

$$\phi_{\kappa}(\mathbf{x}, \mathbf{h}) = e^{-\sum_{i=1}^N (K_1 x_i x_{i+1} + K_2 h_i h_{i+1} + K_3 x_i h_i)}. \quad (4)$$

Notice that the architecture of this uRBM includes correlations between nearest-neighbor visible spins, between nearest-neighbor hidden spins, as well as between pairs of visible and hidden spins with the same index  $i$ . These three correlations are parametrized by the three constants  $K_1$ ,  $K_2$ , and  $K_3$ , respectively. With this uRBM trial *Ansatz*, the set of variational parameters is  $\kappa = \{K_1, K_2, K_3\}$ . It is straightforward to generalize the uRBM *Ansatz* including more layers of hidden spins. Every additional hidden-spin layer adds two more variational parameters, and it effectively represents the application of an imaginary-time Suzuki-Trotter step  $e^{-\Delta\tau \hat{H}}$  for a certain time step  $\Delta\tau$ . Thus a deep neural network state with many hidden layers can represent a long imaginary-time dynamics, which projects out the ground state provided that the initial state is not orthogonal to it. In fact, the mapping between deep neural networks and the imaginary time projection has been exploited in Refs. [22,26] to construct more complex neural network states. In this article, we consider only the single hidden-spin layer uRBM, since this *Ansatz* turns out to be adequate for the ferromagnetic quantum Ising chain. The multi hidden-spin layer *Ansatz* might be useful to address more complex models as, e.g., frustrated Ising spin glasses. Extensions along these lines are left as future work.

In a recent work [15], Carleo and Troyer considered a restricted Boltzmann machine (RBM), where direct correlations among hidden spins were not allowed. Their *Ansatz* included a larger number of hidden spins, as well as more connections between visible and hidden spins, leading to an extensive number of variational parameter proportional to  $\alpha N$ , where  $\alpha = 1, 2, \dots$ . One advantage of the RBM, due to the absence of hidden-hidden correlations, is that the role of hidden spins can be analytically traced out. The uRBM we employ, which is analogous to the shadow wave functions used to describe quantum fluids and solids, includes only three variational parameters, independently of the system size. However, their effect has to be addressed by performing sampling of hidden spins configurations, as described below [27].

It is worth pointing out that correlations beyond nearest-neighbor spins could also be included in the uRBM *Ansatz*, with straightforward modifications in the sampling algorithms described below. We mention here also that, as shown in Ref. [21], neural network states with intralayer correlations can be mapped to deep neural networks with more hidden layers, but no intralayer correlations.

For what concerns the previously mentioned Boltzmann-type trial wave function, one can show that it can be mapped to an RBM [28,29] where each hidden spin is connected only to two visible spins; thus, it provides less representational power compared with standard RBMs, and the results shown below indeed confirm its limited accuracy.

In the case of an uRBM variational wave function, the average value of the energy  $E = \langle \hat{H} \rangle$  is computed as follows:

$$\begin{aligned} \langle \hat{H} \rangle &= \frac{\langle \Psi_{\kappa} | \hat{H} | \Psi_{\kappa} \rangle}{\langle \Psi_{\kappa} | \Psi_{\kappa} \rangle} = \frac{\sum_{\mathbf{x}, \mathbf{x}'} \Psi_{\kappa}(\mathbf{x}) H_{\mathbf{x}, \mathbf{x}'} \Psi_{\kappa}(\mathbf{x}')}{\sum_{\mathbf{x}} \Psi_{\kappa}(\mathbf{x}) \Psi_{\kappa}(\mathbf{x})} \\ &= \frac{\sum_{\mathbf{x}, \mathbf{x}'} [\sum_{\mathbf{h}_a} \phi_{\kappa}(\mathbf{x}, \mathbf{h}_a)] H_{\mathbf{x}, \mathbf{x}'} [\sum_{\mathbf{h}_b} \phi_{\kappa}(\mathbf{x}', \mathbf{h}_b)]}{\sum_{\mathbf{x}} [\sum_{\mathbf{h}_a} \phi_{\kappa}(\mathbf{x}, \mathbf{h}_a)] [\sum_{\mathbf{h}_b} \phi_{\kappa}(\mathbf{x}, \mathbf{h}_b)]} \\ &= \langle \langle E_{\text{loc}}(\mathbf{x}, \mathbf{h}_b) \rangle \rangle, \end{aligned} \quad (5)$$

where the local energy  $E_{\text{loc}}(\mathbf{x}, \mathbf{h})$  is defined as

$$E_{\text{loc}}(\mathbf{x}, \mathbf{h}) = \frac{\sum_{\mathbf{x}'} H_{\mathbf{x}, \mathbf{x}'} \phi_{\kappa}(\mathbf{x}', \mathbf{h})}{\phi_{\kappa}(\mathbf{x}, \mathbf{h})}, \quad (6)$$

with  $H_{\mathbf{x}, \mathbf{x}'} = \langle \mathbf{x} | \hat{H} | \mathbf{x}' \rangle$ .  $\mathbf{h}_a$  and  $\mathbf{h}_b$  indicate two hidden spin configurations. Notice that the formula for the local energy can be symmetrized with respect to the two sets of hidden spins  $\mathbf{h}_a$  and  $\mathbf{h}_b$ , providing results with slightly reduced statistical fluctuations. The double brackets  $\langle \langle \dots \rangle \rangle$  indicate the expectation value over the visible-spin configurations  $\mathbf{x}$  and two sets of hidden spins configurations  $\mathbf{h}_a$  and  $\mathbf{h}_b$ , sampled from the following normalized probability distribution:

$$p(\mathbf{x}, \mathbf{h}_a, \mathbf{h}_b) = \frac{\phi_{\kappa}(\mathbf{x}, \mathbf{h}_a) \phi_{\kappa}(\mathbf{x}, \mathbf{h}_b)}{\sum_{\mathbf{x}, \mathbf{h}_a, \mathbf{h}_b} \phi_{\kappa}(\mathbf{x}, \mathbf{h}_a) \phi_{\kappa}(\mathbf{x}, \mathbf{h}_b)}. \quad (7)$$

As in standard Monte Carlo approaches, this expectation value is estimated as the average of  $E_{\text{loc}}(\mathbf{x}, \mathbf{h})$  over a (large) set of uncorrelated configurations, sampled according to  $p(\mathbf{x}, \mathbf{h}_a, \mathbf{h}_b)$ . The statistical uncertainty can be reduced at will by increasing the number of sampled configurations. The

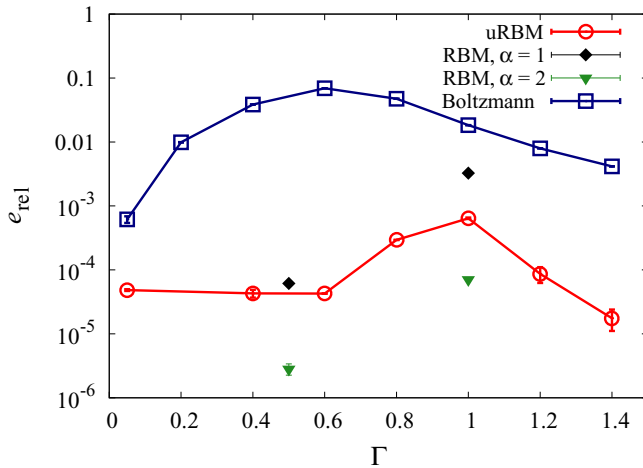


FIG. 2. Relative error  $e_{\text{rel}}$  in the variational estimates of the ground-state energy, see Eq. (8), as a function of the transverse field  $\Gamma$ , obtained using the simple Boltzmann wave function and for the unrestricted Boltzmann machine (uRBM) *Ansatz*. The system size is  $N = 80$ . For comparison, we also show the data corresponding to the restricted Boltzmann machine (RBM) obtained using the code from Ref. [15], where  $\alpha$  indicates the hidden-spin density. The thin lines are guides to the eyes.

optimal variational parameters  $\kappa_{\text{opt}}$  that minimize the energy expectation value can be found using a stochastic optimization method. We adopt a relatively simple yet quite efficient one, namely the stochastic gradient descent algorithm (see, e.g., Ref. [30]). While more sophisticated algorithms exist as, e.g., the stochastic reconfiguration method [23], such methods are not necessary here since the *Ansätze* that we consider include a very small number of variational parameters, one or three. In fact, in these cases, the optimal variational parameters can be obtained also by performing a scan on a fine grid. By doing so, we obtain essentially the same results provided by the stochastic gradient descent algorithm.

We assess the accuracy of the optimized variational wave functions by calculating the relative error

$$e_{\text{rel}} = \frac{|E - E_{\text{JW}}|}{|E_{\text{JW}}|}, \quad (8)$$

in the obtained variational estimate  $E$  of the ground-state energy of the Hamiltonian in Eq. (1).  $E_{\text{JW}}$  is the exact finite size ground-state energy of the quantum Ising chain. It is obtained by performing the Jordan-Wigner transformation, followed by a Fourier and the Bogoliubov transformations.

Figure 2 displays the relative error  $e_{\text{rel}}$  in Eq. (8) corresponding to the variational wave functions introduced above, as a function of the transverse field  $\Gamma$ . The system size is  $N = 80$ , which is here representative of the thermodynamic limit. The Boltzmann-type *Ansatz* does not provide particularly accurate predictions. In the ferromagnetic phase  $\Gamma < 1$ , the relative error is up to 10%. The uRBM, instead, provides very accurate predictions. The relative error is always below 0.1%. The largest discrepancy occurs at the quantum critical point  $\Gamma = 1$ . Such high accuracy is remarkable, considering that the uRBM *Ansatz* involves only 3 variational parameters. It is also worth mentioning that very similar accuracies

are obtained also for different system sizes. Therefore the uRBM *Ansatz* represents a promising guiding function for simulations of quantum annealing optimization of disordered models. As a term of comparison, we show in Fig. 2 the results obtained in Ref. [15] using the RBM *Ansatz*. The relative errors corresponding to the RBM with hidden-unit density  $\alpha = 1$  are larger than those corresponding to the uRBM, despite the fact that the RBM *Ansatz* involves a larger number of variational parameters. However, it is worth stressing that the RBM results can be systematically improved by increasing  $\alpha$ . For example, with  $\alpha = 2$ , the RBM relative errors are approximately an order of magnitude smaller than those corresponding to the uRBM *Ansatz*.

### III. IMPORTANCE SAMPLING GUIDED BY UNRESTRICTED NEURAL NETWORK STATES

In this section, we discuss how optimized variational wave functions can be utilized to boost the performance of PQMC simulations. First, we consider the implementation of the PQMC algorithm without guiding functions. PQMC methods allow one to extract ground-state properties of quantum many-body systems [31,32] by stochastically simulating the Schrödinger equation in imaginary time  $\tau = it$ . In the Dirac notation, this equation is written as

$$-\frac{\partial}{\partial \tau} |\Psi(\tau)\rangle = (\hat{H} - E_{\text{ref}}) |\Psi(\tau)\rangle. \quad (9)$$

The reduced Planck constant is set to  $\hbar = 1$  throughout this paper.  $E_{\text{ref}}$  is a reference energy introduced to stabilize the simulation, as discussed later. Equation (9) is simulated by iteratively applying the equation  $\Psi(\mathbf{x}, \tau + \Delta\tau) = \sum_{\mathbf{x}'} G(\mathbf{x}, \mathbf{x}', \Delta\tau) \Psi(\mathbf{x}', \tau)$ .  $\Delta\tau$  is a (short) time step and  $G(\mathbf{x}, \mathbf{x}', \Delta\tau) = \langle \mathbf{x} | e^{-\Delta\tau(\hat{H} - E_{\text{ref}})} | \mathbf{x}' \rangle$  is the Green's function of Eq. (9). Below, it is discussed how one can write a suitable explicit expression. Long propagation times  $\tau$  are achieved by iterating many (small) time steps  $\Delta\tau$ , allowing one to sample, in the  $\tau \rightarrow \infty$  limit, spin configurations with a probability density proportional to the ground-state wave function  $\Psi_0(\mathbf{x})$  (assumed to be real and non negative). One should notice that the Green's function  $G(\mathbf{x}, \mathbf{x}', \Delta\tau)$  does not define a stochastic matrix; while its elements are nonnegative, one has  $\sum_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}', \Delta\tau) \neq 1$ , in general. Therefore it cannot be utilized to define the transition matrix of a conventional Markov chain Monte Carlo simulation. This problem can be circumvented by rewriting the Green's function as  $G(\mathbf{x}, \mathbf{x}', \Delta\tau) = G_{\text{T}}(\mathbf{x}, \mathbf{x}', \Delta\tau) b_{\mathbf{x}'}$ , where  $G_{\text{T}}(\mathbf{x}, \mathbf{x}', \Delta\tau)$  is by definition stochastic, and the normalization factor is  $b_{\mathbf{x}'} = \sum_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}', \Delta\tau)$ . A stochastic process can then be implemented, where a large population of equivalent copies of the system, in jargon called *walkers*, is evolved. Each walker represents one possible spin configuration  $\mathbf{x}'_n$  (the index  $n$  labels different walkers), and is gradually modified by performing spin-configuration updates according to  $G_{\text{T}}(\mathbf{x}_n, \mathbf{x}'_n, \Delta\tau)$ . Thereafter, their (relative) weights  $w_n$  are accumulated according to the rule  $w_n \rightarrow w_n b_{\mathbf{x}'_n}$ , starting with equal initial weights  $w_n = 1$  for all the walkers in the initial population. While this implementation is, in principle, correct, it is known to lead to an exponentially fast signal loss as the number of Monte Carlo steps increases. This is due to the fact

that the relative weight of few walkers quickly becomes dominant, while most other walkers give a negligible contribution to the signal. An effective remedy consists in introducing a branching process, where each walker is replicated (or annihilated) a number of times corresponding, on average, to the weight  $w_n$ . The simplest correct rule consists in generating, for each walker in the population at a certain imaginary time  $\tau$ , a number of descendants  $n_d$  in the population at imaginary time  $\tau + \Delta\tau$ .  $n_d$  is defined as  $\text{int}[w_n + \eta]$ , where  $\eta \in [0, 1]$  is a uniform random number, and the function  $\text{int}[\cdot]$  gives the integer part of the argument [33]. Clearly, after branching has been performed, all walkers have the same weight  $w_n = 1$ . Therefore the number of walkers in the population fluctuates at each PQMC iteration and can be kept close to a target value by adjusting the reference energy  $E_{\text{ref}}$ . Introducing the branching process provides one with a feasible, possibly efficient algorithm. However, such a process might actually introduce a systematic bias if the average population size  $N_w$  is not large enough. The bias originates from the spurious correlations among walkers generated from the same ancestor [30]. This effect becomes negligible in the  $N_w \rightarrow \infty$  limit, but might be sizable for finite  $N_w$ . It is known to be the most relevant and subtle possible source of systematic errors in PQMC algorithms [34–36]. In fact, it was shown in Ref. [1] that in order to determine with a fixed target relative error, the ground-state energy of the ferromagnetic quantum Ising chain with the (simple) diffusion Monte Carlo algorithm (which belongs to the category of PQMC methods), the walker-population size  $N_w$  has to exponentially increase with the system size  $N$ . This implies an exponentially scaling computational cost.

A promising strategy to circumvent the aforementioned problem is to introduce the so-called importance sampling technique. This is, indeed, a well established approach to boost the efficiency of PQMC simulations (see, e.g., Ref. [4]) because it has the potential to reduce the number of walkers needed to attain a given accuracy [30]. It consists in evolving a function  $f(\mathbf{x}, \tau) = \Psi(\mathbf{x}, \tau)\psi_T(\mathbf{x})$  via a modified imaginary-time Schrödinger equation.  $\psi_T(\mathbf{x})$  is a guiding function designed to accurately approximate the ground-state wave function. Its role is to favor the sampling of configurations with high probability amplitude. The obtained modified imaginary-time Schrödinger equation is solved via a Markov process defined by the following equation:

$$f(\mathbf{x}, \tau + \Delta\tau) = \sum_{\mathbf{x}'} \tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) f(\mathbf{x}', \tau), \quad (10)$$

where the modified Green's function is given by  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) = G(\mathbf{x}, \mathbf{x}', \Delta\tau) \frac{\psi_T(\mathbf{x})}{\psi_T(\mathbf{x}'')}$ . A suitable approximation for the modified Green's function can be obtained by dividing the time step  $\Delta\tau$  into  $M$  shorter time steps  $\delta\tau = \Delta\tau/M$ . If  $\delta\tau$  is sufficiently short, one can employ a Taylor expansion truncated at the linear term,  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau) \cong [\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau)]^M$ , where

$$\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau) = \left[ \delta_{\mathbf{x}, \mathbf{x}'} - \delta\tau (H_{\mathbf{x}, \mathbf{x}'} - E_{\text{ref}} \delta_{\mathbf{x}, \mathbf{x}'}) \frac{\psi_T(\mathbf{x})}{\psi_T(\mathbf{x}')} \right]. \quad (11)$$

With this approximation, Eq. (10) defines a stochastic implementation of the power method of linear algebra. Convergence to the exact ground state is guaranteed as long as  $\delta\tau$  is smaller than a finite value, sufficiently small to ensure that all matrix elements of  $\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau)$  are not negative [37]. As the system size increases, shorter and shorter time steps  $\delta\tau$  are required. This leads to pathologically inefficient simulations, since in this regime the identity operator dominates, resulting in extremely long autocorrelation times. This problem can be solved by adopting the continuous-time Green's function Monte Carlo (CTGFMC) algorithm. The derivation and the details of this algorithm are given in Refs. [30,38], and so we only sketch it here. The idea is to formally take the  $M \rightarrow \infty$  limit, and determine the (stochastic) time interval  $\delta\tau'$  that passes before the next configuration update occurs. It is convenient to bookkeep the remaining time  $\delta\tau_t$  left to complete a total interval of time  $\Delta\tau$ . This is to ensure that each iteration of the PQMC simulation corresponds to a time step of duration  $\Delta\tau$ . The time interval  $\delta\tau'$  is sampled using the formula  $\delta\tau' = \text{Min}(\delta\tau_t, \frac{\ln(1-\xi)}{E_{\text{loc}}(\mathbf{x}') - E_{\text{ref}}(\mathbf{x}')} )$  with  $\xi \in (0, 1)$  being a uniform random number. The spin-configuration update  $\mathbf{x}' \rightarrow \mathbf{x}$  (with  $\mathbf{x}' \neq \mathbf{x}$ ) is randomly selected from the probability distribution

$$t_{\mathbf{x}, \mathbf{x}'} = \frac{p_{\mathbf{x}, \mathbf{x}'}}{\sum_{\mathbf{x} \neq \mathbf{x}'} p_{\mathbf{x}, \mathbf{x}'}} \quad (12)$$

$$p_{\mathbf{x}, \mathbf{x}'} = \frac{\tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau')}{\sum_{\mathbf{x}} \tilde{g}(\mathbf{x}, \mathbf{x}', \delta\tau')}.$$

Notice that, with the Hamiltonian (1),  $\mathbf{x}$  differs from  $\mathbf{x}'$  only for one spin flip. The weight-update factor for the branching process takes the exponential form  $b_{\mathbf{x}'} = e^{-\delta\tau' [E_{\text{loc}}(\mathbf{x}') - E_{\text{ref}}]}$ , where the local energy is now  $E_{\text{loc}}(\mathbf{x}') = \sum_{\mathbf{x}} H_{\mathbf{x}, \mathbf{x}'} \frac{\psi_T(\mathbf{x})}{\psi_T(\mathbf{x}'')}$ .

In summary, the CTGFMC algorithm requires to perform, for each walker  $n$  in the population, the following steps: (i) initialize the time interval  $\delta\tau_t = \Delta\tau$ , and the weight factor  $w_n = 1$ ; (ii) sample the time  $\delta\tau'$  at which the configuration update  $\mathbf{x}' \rightarrow \mathbf{x}$  might occur; (iii) if  $\delta\tau' < \delta\tau_t$ , update  $\mathbf{x}'$  with a transition probability  $t_{\mathbf{x}, \mathbf{x}'}$  in Eq. (12), else set  $\delta\tau' = \delta\tau_t$ ; (iv) accumulate the weight factor according to the rule  $w_n \rightarrow w_n b_{\mathbf{x}'}$  and set  $\delta\tau_t \rightarrow \delta\tau_t - \delta\tau'$ ; (v) Go back to step (ii) until  $\delta\tau_t = 0$ ; and (vi) finally, perform branching according to the total accumulated weight factor  $w_n$ . This continuous-time algorithm implicitly implements the exact imaginary-time modified Green's function  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau)$ .

In the long imaginary-time limit, the walkers sample spin configurations with a probability distribution proportional to  $f(\mathbf{x}, \tau \rightarrow \infty) = \Psi_0(\mathbf{x})\psi_T(\mathbf{x})$ . If  $\psi_T(\mathbf{x})$  is a good approximation of the ground-state wave function, this distribution closely approximates the quantum-mechanical probability of finding the system in the spin configuration  $\mathbf{x}$ . It is important to notice that if our guiding wave function was exact, i.e., if  $\psi_T(\mathbf{x}) = \Psi_0(\mathbf{x})$ , then the local energy  $E_{\text{loc}}(\mathbf{x})$  would be a constant function. This would completely suppress the fluctuations of the number of walkers, therefore eliminating the bias due to the finite walkers population  $N_w$ . If  $\psi_T(\mathbf{x})$  is, albeit not exact, a good approximation of  $\Psi_0(\mathbf{x})$ , the fluctuations of the number of walkers are still reduced compared to the case of the simple CTGFMC algorithm [which corresponds to setting  $\psi_T(\mathbf{x}) = 1$ ] giving a faster convergence to the exact

$N_w \rightarrow \infty$  limit. Below we consider the use of the variational wave functions  $\Psi_\kappa(\mathbf{x})$  described in Sec. II as guiding wave functions for the PQC algorithm, setting the variational parameters  $\kappa$  at their optimal values.

In order to employ the unrestricted neural-network states as guiding functions, the PQC algorithm has to be modified. One has to implement a combined dynamics of the visible-spin configurations  $\mathbf{x}$  and of the hidden-spin configurations  $\mathbf{h}$ . We will indicate the global configuration as  $\mathbf{y} = (\mathbf{x}, \mathbf{h})$ . The goal is to sample global configurations with the (normalized) probability distribution

$$p(\mathbf{y}) = \frac{\Psi_0(\mathbf{x})\phi_\kappa(\mathbf{x}, \mathbf{h})}{\sum_{\mathbf{x}, \mathbf{h}} \Psi_0(\mathbf{x})\phi_\kappa(\mathbf{x}, \mathbf{h})}. \quad (13)$$

This allows one to compute the ground-state energy as  $E = \lim_{N_c \rightarrow \infty} \sum_{i=1}^{N_c} E_{\text{loc}}(\mathbf{x}_i, \mathbf{h}_i)/N_c$ , where  $N_c$  is a number of uncorrelated configurations  $\{\mathbf{y}_i\}$  sampled from  $p(\mathbf{y})$ . The local energy  $E_{\text{loc}}(\mathbf{x}, \mathbf{h})$  is defined as in Eq. (6). A suitable algorithm was implemented in Ref. [39] in the case of the continuous-space Green's function Monte Carlo algorithm, where importance sampling was implemented using shadow wave functions. Here we modify the approach of Ref. [39] to address quantum spin models. The visible-spin configurations  $\mathbf{x}$  are evolved according to the CTGFMC described above, keeping the hidden-spin configuration  $\mathbf{h}'$  fixed. The modified imaginary-time Green's function is now  $\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau|\mathbf{h}') = G(\mathbf{x}, \mathbf{x}', \Delta\tau) \frac{\phi_\kappa(\mathbf{x}, \mathbf{h}')}{\phi_\kappa(\mathbf{x}', \mathbf{h})}$ . As discussed above, this has to be rewritten as the product of a stochastic matrix, which defines how the visible-spin configurations updates are selected, and a weight term, which is taken into account with the branching process. The weight-update factor is  $b_{\mathbf{y}'} = \sum_{\mathbf{x}} \tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau|\mathbf{h}')$ . The dynamics of the hidden-spins configurations is dictated by a (classical) Markov chain Monte Carlo algorithm. Considering  $\phi_\kappa(\mathbf{x}, \mathbf{h})$  as an unnormalized probability distribution allows one to write—for any fixed visible-spin configuration  $\mathbf{x}$ —the Master equation:

$$\phi_\kappa(\mathbf{x}, \mathbf{h}) = \sum_{\mathbf{h}'} T(\mathbf{h}, \mathbf{h}'|\mathbf{x})\phi_\kappa(\mathbf{x}, \mathbf{h}'), \quad (14)$$

where  $T(\mathbf{h}, \mathbf{h}'|\mathbf{x})$  is the transition matrix that defines the Markov process. Clearly, the following condition must be fulfilled  $\sum_{\mathbf{h}} T(\mathbf{h}, \mathbf{h}'|\mathbf{x}) = 1$ , for any  $\mathbf{x}$ .

Our choice is a single spin flip METROPOLIS algorithm, where the flip of a randomly selected spin is proposed, and accepted with the probability

$$A(\mathbf{h}' \rightarrow \mathbf{h}|\mathbf{x}) = \text{Min} \left\{ 1, \frac{\phi_\kappa(\mathbf{x}, \mathbf{h})}{\phi_\kappa(\mathbf{x}, \mathbf{h}')} \right\}. \quad (15)$$

Here,  $\mathbf{h}$  differs from  $\mathbf{h}'$  only for the (randomly selected) flipped spin. One could perform a certain number, call it  $k$ , of METROPOLIS updates, without modifying the formalism. In fact, this turns out to be useful, as discussed below. The combined dynamics of the visible and the hidden spins is driven by the following equation:

$$f(\mathbf{y}, \tau + \Delta\tau) = \sum_{\mathbf{y}'} G(\mathbf{y}, \mathbf{y}', \Delta\tau)f(\mathbf{y}', \tau), \quad (16)$$

with  $G(\mathbf{y}, \mathbf{y}', \Delta\tau) = T(\mathbf{h}, \mathbf{h}'|\mathbf{x})\tilde{G}(\mathbf{x}, \mathbf{x}', \Delta\tau|\mathbf{h}')$ . It can be shown [39] that the equilibrium probability distribution of this equation is the desired joint probability distribution  $p(\mathbf{y})$  in Eq. (13). The stochastic process corresponding to this equation can be implemented with the following steps: (i) perform the visible-spin configuration update  $\mathbf{x}' \rightarrow \mathbf{x}$ , keeping  $\mathbf{h}'$  fixed, according to the CTGFMC algorithm described above (including accumulation of the weight factor); (ii) perform  $k$  single-spin METROPOLIS updates of the hidden-spin configuration  $\mathbf{h}'$ , keeping  $\mathbf{x}$  fixed; and (iii) perform branching of the global configuration. It is easily shown that the hidden-spin dynamics does not directly affect the weight factor since the normalization of the Green function of the combined dynamics is set by  $b_{\mathbf{y}'}$ .

Since the optimized uRBM describes the ground-state wave function with high accuracy, one expects that its use as a guiding function leads to a drastic reduction of the systematic errors due to the finite random walker population. However, one should take into account that there might be statistical correlations among subsequent hidden-spin configurations along the Markov chain. This might in turn affect the systematic error. Clearly, increasing the number of Metropolis steps  $k$  per CTGFMC visible-spin configuration update allows one to suppress such correlations, possibly reducing the systematic error. This will indeed turn out to be important, in particular at the quantum critical point where statistical correlations along the Markov chain are more significant.

Following Ref. [1], we analyze the computational complexity of the PQC algorithm by determining the number of walkers  $N_w$  needed to determine the ground-state energy of the Hamiltonian (1) with a prescribed accuracy. All data described below have been obtained with a time step  $\Delta\tau = 0.1$ , and all simulations have been run for a long enough total imaginary time to ensure equilibration.

First, we consider the simple PQC algorithm, i.e., performed without importance sampling. Figure 3 displays the




FIG. 3. Number of random walkers  $N_w$  required to determine, using the PQC algorithm without importance sampling, the ground-state energy with a relative error  $e_{\text{rel}}$ , see Eq. (8), as a function of the system size  $N$ . Different data sets correspond to different transverse field intensities  $\Gamma$  and different relative errors. The lines represent exponential fitting functions.




FIG. 4. Number of random walkers  $N_w$  required to determine, using the optimized Boltzmann-type wave function to guide importance sampling in the PQMC simulation, the ground-state energy with a relative error  $e_{\text{rel}}$ , see Eq. (8), as a function of the system size  $N$ . Different data sets correspond to different transverse field intensities  $\Gamma$ . The (red) dotted and (blue) dot-dashed lines represent exponential fitting functions, while the (green) dashed line represents a power-law fit with power  $b = 0.54(5)$ .

scaling with the system size  $N$  of the number of walkers  $N_w$  required to keep the relative error  $e_{\text{rel}}$ , defined in Eq. (8), at the chosen threshold. This scaling is evidently exponential, below, above, and also at the quantum critical point. The most severe scaling comes from the ordered phase and could be attributed to the fact that the simple PQMC is formally equivalent to PQMC with a constant  $\psi_T(\mathbf{x})$  for importance sampling. This turns out to be a very poor choice of the guiding function in the ordered regime given that it treats all configurations on an equal footing. Analogous results have been obtained in Ref. [1] using the diffusion Monte Carlo algorithm. This is another PQMC method—in fact very similar to the CTGFMC algorithm employed here—whose transition matrix is defined from the imaginary-time Green’s function derived within the symmetrized Trotter decomposition. Introducing importance sampling using the optimized Boltzmann-type *Ansatz* as a guiding function significantly reduces the systematic error due to the finite random walker population, allowing one to reach quite small relative errors. In particular, in the paramagnetic phase at  $\Gamma = 1.4$ , the scaling of  $N_w$  versus  $N$  is quite flat (see Fig. 4); it appears to be well described by the power-law  $N_w \sim N^b$  with the small power  $b = 0.54(5)$ , rather than by an exponential. However, in the ferromagnetic phase at  $\Gamma = 0.6$  and at the quantum critical point  $\Gamma = 1$  the scaling is still clearly exponential. This means that the simple Boltzmann-type *Ansatz* is, in general, insufficient to ameliorate the exponentially scaling computational cost of the PQMC algorithm. Figure 5 shows the scaling of  $N_w$  obtained using the optimized uRBM *Ansatz* as the guiding function. The number of hidden-spin Metropolis steps per visible-spin update is set to a (small) fraction of the system size  $N$ , namely to  $k = 0.1N$ . At  $\Gamma = 0.6$ , the required walker population size  $N_w$  turns out to be essentially independent on the system size  $N$ . It is worth noticing that the prescribed relative error is here as small as  $e_{\text{rel}} = 10^{-6}$ , and that this high accuracy is




FIG. 5. Number of random walkers  $N_w$  required to determine, using the optimized uRBM *Ansatz* to guide importance sampling in the PQMC simulation, the ground-state energy with a relative error  $e_{\text{rel}}$ , see Eq. (8), as a function of the system size  $N$ . The number of single-spin Metropolis updates of the hidden spins per CTGFMC hidden-spin update is  $k = 0.1N$ . The (red) dotted line represents an exponential fit, while the (blue) dot-dashed line represents a linear fit.

achieved with a rather small walkers population  $N_w \lesssim 1000$ . However, at the quantum critical point,  $N_w$  still displays an exponential scaling with system size. This effect can be traced back to the diverging statistical correlations among subsequent hidden-spin configurations along the Markov chain, due to quantum criticality. As anticipated above, these statistical correlations can be suppressed by increasing the number of hidden-spin updates  $k$ . Fig. 6 displays the scaling of  $N_w$ , at

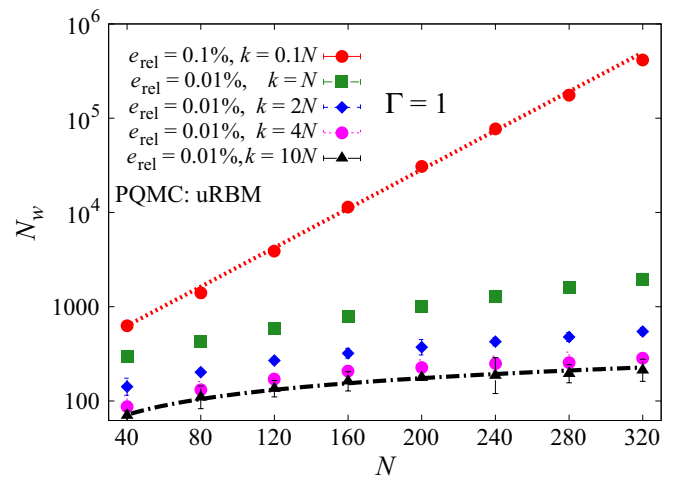


FIG. 6. Number of random walkers  $N_w$  required to determine, using the optimized uRBM to guide importance sampling in the PQMC simulation, the ground-state energy with a relative error  $e_{\text{rel}}$ , see Eq. (8), as a function of the system size  $N$ . The transverse field intensity is set at the ferromagnetic quantum critical point  $\Gamma = 1$ . Different data sets correspond to different values of the number of single-spin Metropolis updates  $k$ . The (red) dotted line represents an exponential fit, while the (black) dot-dashed line represents a power-law fit, with power  $b = 0.55(1)$ .

the quantum critical point, for different  $k$  values. One observes that the scaling substantially improves already for moderately larger  $k$  values, leading to a crossover from the exponential scaling obtained with  $k = 0.1N$ , to a scaling behavior that is well described by a square-root function  $N_w \sim N^{0.55(1)}$  when  $k = 10N$ . While it is, in principle, possible that, for system sizes  $N \gg 300$ ,  $N_w$  still approaches an exponential scaling with a very small base, we argue that the data we provide indicate that the scaling of  $N_w$  is not exponential, even at the ferromagnetic critical point, provided that  $k$  is large enough. It is important to point out that increasing  $k$  implies a correspondingly increasing contribution to the global computational cost of the QPMC algorithm. However, since  $k$  is here linear in the system size, this contribution does not modify, to leading order, the scaling of the global computational cost. Therefore one can conclude that the uRBM *Ansatz* is sufficient to change the scaling of the computational cost of the QPMC algorithm from exponential in the system size, to an amenable polynomial scaling. In the simulations presented here, single-spin flip METROPOLIS updates are employed for the hidden variables. It is possible that cluster spin updates would lead to an even faster convergence to the exact  $N_w \rightarrow \infty$  limit, due to the more efficient sampling of the hidden-spin configurations. However, such cluster updates cannot always be implemented, in particular for frustrated disordered Hamiltonians relevant for optimization problems; therefore, we do not consider them here.

#### IV. CONCLUSIONS

The accuracy of variational wave functions that mimic unrestricted Boltzmann machines, which we refer to as unrestricted neural network states, has been analyzed using the one-dimensional ferromagnetic Ising model as a testbed. By optimizing just three variational parameters, ground-state energies with a relative error smaller than  $10^{-3}$  have been obtained. The ferromagnetic quantum phase transition turns out to be the point where the relative error is the largest. This accuracy is comparable to the one previously obtained using restricted neural network states with few hidden variables per visible spin [15]. These restricted neural network states involve a number of variational parameters proportional to the system size, as opposed to the unrestricted neural network states considered here, where the (small) number of variational parameters is fixed. This feature of the unrestricted states makes them very suitable in the context of quantum annealing simulations for Ising-type models (which are sign-problem free). Indeed, it was shown in Ref. [40] that in order to efficiently simulate quantum annealing via QPMC it is essential to employ accurate variational *Ansätze* as guiding functions. However, since one has to integrate over hidden-spin configurations via Monte Carlo sampling, as opposed to the case of the restricted neural network states [15]—for which the hidden-spin configurations can be integrated out—they represent a less promising approach to model ground states of Hamiltonian where the negative sign problem occurs. Indeed, in such case, an accurate variational *Ansatz* might have to include also hidden-spin configurations with negative wave-function amplitude, making Monte Carlo integration via random sampling inapplicable.

The variational study summarized here represented a necessary preliminary step to investigate the use of optimized unrestricted neural network states as guiding functions for importance sampling in QPMC simulations. We have found that unrestricted neural network states allow one to drastically reduce the systematic bias of the QPMC algorithm originating from the finite size of the random-walker population. Specifically, the scaling of the population size required to keep a fixed relative error as the system size increases changes from the exponential scaling characteristic of simple QPMC simulations performed without guiding functions, to a polynomial scaling. This also implies a corresponding change in the scaling of the computational cost. This qualitative scaling change occurs above, below, and apparently also at the ferromagnetic quantum phase transition, at least for the system sizes  $N \lesssim 300$  considered in this article. While we cannot strictly rule out a weak exponential scaling for  $N \gg 300$ , the numerical data we provide suggest that this is an unlikely scenario. Instead, a conventional variational *Ansatz* of the Boltzmann type was found to provide a significant improvement of the computational cost only above the critical point (in the paramagnetic phase), but to provide only a marginal improvement at and below the transition. It is worth emphasizing that the use of unrestricted neural network states as guiding functions in QPMC simulations requires the sampling of both the visible and the hidden spins, using the combined algorithm described in Sec. III (more efficient variants might be possible). The role of the statistical correlations among hidden-spin configurations shows up in particular at the ferromagnetic quantum critical point. We found that these correlations can be eliminated by performing several single-spin updates, still without affecting, to leading order, the global computational complexity of the simulation.

In Ref. [41], it was proven that it is possible to devise polynomially-scaling numerical algorithms to determine the ground-state energy, with a small additive error, of various ferromagnetic spin models, including the ferromagnetic Ising chain considered here. However, practical implementations have not been provided. The numerical data we have reported in this manuscript indicate that the QPMC algorithm guided by an optimized unrestricted neural network state represents a practical algorithm with polynomial computational complexity for the ferromagnetic quantum Ising chain. More in general, it was shown in Ref. [42] that the problem of estimating the ground-state energy of a generic sign-problem free Hamiltonian with a small additive error is at least NP-hard. Indeed, this task encompasses hard optimization problems such as  $k$ -SAT and MAX-CUT. This suggests that there might be relevant models where the unrestricted neural network states discussed here are not sufficient to make the computational cost of the QPMC simulations affordable. Relevant candidates are Ising spin-glass models with frustrated couplings. Such systems might require more sophisticated guiding functions possibly obtained by using the same couplings of the classical Ising-glass Hamiltonian in the intralayer coupling of the uRBM with a global multiplicative variational parameter (one for each layer). More hidden layers may be added in a similar fashion in order to represent a path integral as previously discussed in Sec. II. This would result in a deep neural network with a number of variational parameters that is still



size independent. However, it is possible that in the disordered case a more sophisticated *Ansatz* with an extensive number of variational parameters is needed to reach an accuracy comparable to the one obtained here for the ferromagnetic system. In future work, we plan to search for models that make PQMC simulations problematic. We argue that this will help us in understanding if and for which models a systematic quantum speed-up in solving optimization problems using quantum annealing devices, instead of PQMC simulations performed on classical computer, could be achieved.

## ACKNOWLEDGMENTS

We acknowledge insightful discussions with Giuseppe Carleo, Rosario Fazio, Ivan Glasser, Mohammad Amin, Guglielmo Mazzola, Nicola Pancotti, Francesco Pederiva, Sandro Sorella, and Matteo Wauters. S.P. and L.D. acknowledge financial support from the BIRD2016 project “Superfluid properties of Fermi gases in optical potentials” of the University of Padova. G.E.S. acknowledges support by the EU FP7 under ERC-MODPHYSFRICT, Grant Agreement No. 320796.

- 
- [1] E. M. Inack, G. Giudici, T. Parolini, G. Santoro, and S. Pilati, *Phys. Rev. A* **97**, 032307 (2018).
- [2] D. Ceperley and B. Alder, *Science* **231**, 555 (1986).
- [3] B. L. Hammond, W. A. Lester, and P. J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry* (World Scientific, Singapore, 1994), Vol. 1.
- [4] W. Foulkes, L. Mitas, R. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- [5] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. Schmidt, and R. B. Wiringa, *Rev. Mod. Phys.* **87**, 1067 (2015).
- [6] A. B. Finnila, M. A. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, *Chem. Phys. Lett.* **219**, 343 (1994).
- [7] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, *Science* **295**, 2427 (2002).
- [8] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Nat. Phys.* **10**, 218 (2014).
- [9] E. M. Inack and S. Pilati, *Phys. Rev. E* **92**, 053304 (2015).
- [10] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, *Science* **348**, 215 (2015).
- [11] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, *Phys. Rev. Lett.* **117**, 180402 (2016).
- [12] Z. Jiang, V. N. Smelyanskiy, S. V. Isakov, S. Boixo, G. Mazzola, M. Troyer, and H. Neven, *Phys. Rev. A* **95**, 012322 (2017).
- [13] G. Mazzola, V. N. Smelyanskiy, and M. Troyer, *Phys. Rev. B* **96**, 134305 (2017).
- [14] E. Andriyash and M. H. Amin, [arXiv:1703.09277](https://arxiv.org/abs/1703.09277).
- [15] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [16] H. Saito, *J. Phys. Soc. Jpn.* **86**, 093001 (2017).
- [17] H. Saito and M. Kato, *J. Phys. Soc. Jpn.* **87**, 014001 (2018).
- [18] D.-L. Deng, X. Li, and S. Das Sarma, *Phys. Rev. X* **7**, 021021 (2017).
- [19] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, *Phys. Rev. B* **97**, 085104 (2018).
- [20] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, *Phys. Rev. X* **8**, 011006 (2018).
- [21] X. Gao and L.-M. Duan, *Nat. Commun.* **8**, 662 (2017).
- [22] N. Freitas, G. Morigi, and V. Dunjko, [arXiv:1803.02118](https://arxiv.org/abs/1803.02118).
- [23] S. Sorella, M. Casula, and D. Rocca, *J. Chem. Phys.* **127**, 014105 (2007).
- [24] S. Vitiello, K. Runge, and M. H. Kalos, *Phys. Rev. Lett.* **60**, 1970 (1988).
- [25] L. Reatto and G. L. Masserini, *Phys. Rev. B* **38**, 4516 (1988).
- [26] G. Carleo, Y. Nomura, and M. Imada, [arXiv:1802.09558](https://arxiv.org/abs/1802.09558).
- [27] One could map this specific uRBM in one-dimension to a matrix product state with small bond dimension [43,44], allowing to integrate out the hidden-spin configurations. However, we are interested in a flexible QMC algorithms that straightforwardly extends to general setups, in particular to higher-dimensional geometries and to models with disorder and/or long-range interactions, since these are relevant for optimization problems. In these cases, Monte Carlo sampling becomes essential [45].
- [28] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [29] Y. Huang and J. E. Moore, [arXiv:1701.06246](https://arxiv.org/abs/1701.06246).
- [30] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, 2017).
- [31] J. B. Anderson, *J. Chem. Phys.* **63**, 1499 (1975).
- [32] M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods* (Wiley, Weinheim, 2008).
- [33] J. Thijssen, *Computational Physics* (Cambridge University Press, Cambridge, 2007).
- [34] N. Nemeč, *Phys. Rev. B* **81**, 035119 (2010).
- [35] M. Boninsegni and S. Moroni, *Phys. Rev. E* **86**, 056712 (2012).
- [36] L. Pollet, N. V. Prokof'ev, and B. V. Svistunov, *Phys. Rev. B* **98**, 085102 (2018).
- [37] K. E. Schmidt, P. Niyaz, A. Vaught, and M. A. Lee, *Phys. Rev. E* **71**, 016707 (2005).
- [38] S. Sorella and L. Capriotti, *Phys. Rev. B* **61**, 2599 (2000).
- [39] S. A. Vitiello and P. A. Whitlock, *Phys. Rev. B* **44**, 7373 (1991).
- [40] L. Stella and G. E. Santoro, *Phys. Rev. E* **75**, 036703 (2007).
- [41] S. Bravyi and D. Gosset, *Phys. Rev. Lett.* **119**, 100503 (2017).
- [42] S. Bravyi, *Quantum Inf. Comput.* **15**, 1122 (2015).
- [43] E. Robeva and A. Seigal, *Information and Inference: A Journal of the IMA*, [iay009](https://arxiv.org/abs/1709.00099) (2017), doi: 10.1093/imaiai/iay009.
- [44] I. Glasser, N. Pancotti, and J. I. Cirac, [arXiv:1806.05964](https://arxiv.org/abs/1806.05964).
- [45] A. W. Sandvik and G. Vidal, *Phys. Rev. Lett.* **99**, 220602 (2007).