

Convolutional restricted Boltzmann machine aided Monte Carlo: An application to Ising and Kitaev models

Daniel Alcalde Puente  and Ilya M. Eremin

Institut für Theoretische Physik III, Ruhr-Universität Bochum, D-44780 Bochum, Germany



(Received 20 April 2020; revised 1 October 2020; accepted 19 October 2020; published 30 November 2020)

Machine learning is becoming widely used in analyzing the thermodynamics of many-body condensed matter systems. Restricted Boltzmann machine (RBM) aided Monte Carlo simulations have sparked interest recently, as they manage to speed up classical Monte Carlo simulations. Here we employ the convolutional restricted Boltzmann machine (CRBM) method and show that its use helps to reduce the number of parameters to be learned drastically by taking advantage of translation invariance. Furthermore, we show that it is possible to train the CRBM at smaller lattice sizes, and apply it to larger lattice sizes. To demonstrate the efficiency of CRBM we apply it to the paradigmatic Ising and Kitaev models in two dimensions.

DOI: [10.1103/PhysRevB.102.195148](https://doi.org/10.1103/PhysRevB.102.195148)

I. INTRODUCTION

In recent years, machine learning and specifically restricted Boltzmann machines (RBMs) have been used in the field of many-body quantum systems to explicitly parametrize a probability distribution function of a quantum many-body state [1]. For example, RBMs have been found to be a powerful tool to obtain the ground state of quantum models through variational Monte Carlo (MC) [2] and to reconstruct quantum states from a set of qubit measurements [3,4]. RBMs have also been used to construct the exact ground state of quantum systems by reproducing the exact imaginary time evolution [5]. Neural networks have been shown to be useful in the classification of phases of matter in MC simulations [6], which has awakened interest in the classification of topological and nematic phases [7–9], as finding a suitable order parameter for phase transitions can be challenging. Increasing the resolution of already sampled states [10] has been another area of study. Another promising avenue in the application of the machine learning techniques is sizable speeding up of the quantum MC simulations obtained via the so-called self-learning Monte Carlo (SLMC) method [11–15] applied to the models of interacting fermionic systems where the Trotter decomposition can be employed. In this method, the effective energy is inexpensive to compute and is mostly composed of two-particle interactions, which enables cluster updates. The strength of these two-particle interactions is then learned by applying the linear regression method. For more complex systems neural networks are employed to model the effective energy [16,17]. This, however, makes cluster updates very hard to realize. At the same time, RBMs can be used as an alternative to SLMC for models where no Trotter decomposition is needed because the former can model more complex interactions than the original SLMC (learned by linear regression), and they are faster to sample from than SLMC with neural networks due to Gibbs sampling.

The RBM is a probabilistic model that has two main features. First, it is possible to sample states from the model's

probability distribution with global updates (Gibbs sampling). In addition, the nonnormalized probability distribution $\hat{P}_{\text{RBM}}(x; W)$ is well defined. Here, the matrix W determines the actual form of the probability distribution and the task is to choose a W such that $P_{\text{RBM}}(x; W) \approx P_{\text{target}}(x)$ for all x and then use its global updates for sampling. This has been done for example for the Ising and Falikov-Kimball model [18,19] for small lattice sizes, $L = 8$ and $L = 10$. However, the training appears to be slow and major difficulties arise for RBM training for larger lattices.

One of the drawbacks of using RBMs in its current form is that for their training a Metropolis aided Monte Carlo simulation, which is expensive, needs to be performed in advance. A well-behaving Metropolis simulation is needed to train an RBM, which poses a problem since the scenario where Metropolis is not behaving well is exactly the one where we want to apply the RBM. The second drawback is that for larger lattice size L , the number of parameters that need to be learned scale with L^4 in two-dimensions. So the bigger L is, the more training time is required.

To overcome these problems, we propose to use convolutional restricted Boltzmann machines (CRBMs), which utilizes translation invariance and was originally developed in the context of feature extraction in images [20,21]. In this paper, we apply CRBMs to the paradigmatic Ising and Kitaev models in two dimensions. Using the fact that CRBMs are translationally invariant we demonstrate that using them reduces the number of parameters that need to be learned, which leads to faster training. In addition, the same CRBM can be applied to different lattice sizes, which means that after the CRBM has learned a probability distribution with a certain lattice size L_{small} , it can then be scaled to larger lattice sizes without extra computational costs. Furthermore, we show that if the CRBMs does not fit the probability distribution exactly, it can be corrected using a version of parallel tempering. Convolutions have been employed before in the context of Neural Network SLMC [16,17] but after the convolution a fully connected layer is applied, which breaks translations




FIG. 1. One-dimensional representation of (a) fully connected RBMs and (b) CRBMs. Left panels show a connection between visible and hidden units. On the right panel the effective connections between visible units after summing over h , representing the free energy [see Eq. (3)] are shown. Observe that only nearest neighbors are connected through h in the CRBM and so only nearest neighbors interact.

invariance. Note that CRBM can only be applied to statistical mechanics models where translation invariance is preserved. This, for example, rules out its application to random spin systems.

The paper is organized as follows. In the next section, we present the details of the CRBM and in Sec. III the auto-correlation times are compared between the Metropolis MC and CRBM for the Ising model. We extend the results to the two-dimensional Kitaev model in Sec. IV where the specific heat at different temperatures is computed for both periodic and open boundary conditions. In this section, the CRBM is also compared to the fully connected RBM. The main results are summarized in the Conclusion.

II. FROM FULLY CONNECTED RBM TO CONVOLUTIONAL RBM

A restricted Boltzmann machine (RBM) is a probabilistic generative neural network model, which can be used to learn an approximate probability distribution and then sample from it using the Gibbs sampling. The model has two distinct groups of statistical variables, the visible variables v and the hidden variables h as shown in Fig. 1(a). Here, v and h are the vectors of length N where each component takes either the value 0 or 1, i.e., $v \in \{0, 1\}^{N_p}$ and $h \in \{0, 1\}^{N_h}$, where N_p is the size of the vector v and N_h is the size of the vector h . Note that for a square lattice $N_p = L^2$. The

probability distribution over both the visible and the hidden variables is $P_{\text{RBM}}(v, h; W)$. Summing over the hidden variables, the probability distribution over the visible units is given as $P_{\text{RBM}}(v; W) = \sum_h P_{\text{RBM}}(v, h; W)$, which after training will approximate the target distribution. The hidden units are required as instruments for Gibbs sampling and to mediate the interaction between visible units. Visible and hidden variables form a bipartite system, connected through a matrix W to each other but not to themselves. The probability distribution over both visible and hidden variables can be expressed as

$$P_{\text{RBM}}(v, h) = \frac{e^{-E(v, h)}}{Z}, \quad (1)$$

$$E(v, h) = - \sum_i h_{\text{bias}}^i h^i - \sum_j v_{\text{bias}}^j v^j - \sum_{ij} h^i W_{ij} v^j, \quad (2)$$

where Z is the normalization constant and W , v_{bias} , and h_{bias} are the model parameters. Summing over h , the visible units v_i are no longer independent of each other [see Fig. 1(a)] in the sense that each h_j now represents an interaction of the v_i that are connected via W_{ij} . The probability over the visible units is then given by

$$P_{\text{RBM}}(v) = \sum_{h \in \{0, 1\}^N} P(v, h) = \frac{e^{-F(v)}}{Z}, \quad (3)$$

$$F(v) = - \sum_j v_{\text{bias}}^j v^j - \sum_i \log 1 + e^{h_{\text{bias}}^i + \sum_j W_{ij} v^j}. \quad (4)$$

The physical model with energy $E_{\text{phys}}(x)\beta$ will then be approximated by the energy $F_{\text{RBM}}(x; W)$ of the RBM. Training is then done supervised similarly to Ref. [19] by minimizing the loss

$$\text{loss}(W) = \frac{1}{M} \sum_{i=1}^M [E_{\text{phys}}(x_i)\beta - F_{\text{RBM}}(x_i; W) - C(W)]^2 \quad (5)$$

where β is the inverse temperature and C is a constant that can be chosen freely since the probability function is invariant under the addition of a constant to the energy. Note that if $\text{loss}(W) = 0$ then $F_{\text{RBM}}(x_i; W) = E_{\text{phys}}(x_i)\beta + C(W)$ and so $P_{\text{RBM}}(x; W) = P_{\text{phys}}(x)$. C is chosen such that the loss is minimal

$$C(W) = \frac{1}{M} \sum_{i=1}^M E_{\text{phys}}(x_i)\beta - F_{\text{RBM}}(x_i; W). \quad (6)$$

Observe that this minimization of the loss is equivalent to minimizing the Kullback-Leibler divergence

$$D_{\text{KL}}(P_{\text{RBM}} || P_{\text{phys}}) = \sum_{x \in X} P_{\text{RBM}}(x; W) \log \frac{P_{\text{RBM}}(x; W)}{P_{\text{phys}}(x)}. \quad (7)$$

In particular, It can be rewritten as

$$D_{\text{KL}}(P_{\text{RBM}} || P_{\text{phys}}) = \sum_{x \in X} P_{\text{RBM}}(x; W) [E_{\text{phys}}(x)\beta - F_{\text{RBM}}(x)] + \log Z_{\text{phys}} - \log Z_{\text{RBM}}, \quad (8)$$

and its gradient is given by

$$\frac{\partial D_{\text{KL}}(P_{\text{RBM}} || P_{\text{phys}})}{\partial W} = \left\langle \left[\left\langle \frac{\partial F(x; W)}{\partial W} \right\rangle_{x \in P_{\text{RBM}}(x)} - \frac{\partial F(x; W)}{\partial W} \right] [E_{\text{phys}}(x)\beta - F_{\text{RBM}}(x; W)] \right\rangle_{x \in P_{\text{RBM}}(x)}, \quad (9)$$

which is equivalent to performing the derivative on our loss

$$\left\langle \frac{\partial \frac{1}{2} (E_{\text{phys}}(x)\beta - F_{\text{RBM}}(x; W) - C(W))^2}{\partial W} \right\rangle_{x \in P_{\text{RBM}}(x)} \quad (10)$$

Note that, in practice, we use a mixture of states sampled from both the physical probability distributions and the RBM's probability distributions. Also, note that this version of the Kullback-Leibler divergence will converge faster than the one commonly used in machine learning for unsupervised learning $D_{\text{KL}}(P_{\text{phys}}||P_{\text{RBM}})$. In our approach, we can use it since we have prior knowledge of the energy of the target distribution, which is not the case for unsupervised machine learning tasks. Note that ADAM batch-gradient descent [22] is used to minimize this loss. The only remaining question is which states x_i do we train on? The authors of Ref. [19] proposed to train with the states sampled from the physical distribution. The problem with this is that if a state x' has a tiny realization probability $P_{\text{phys}}(x') \ll 1$, then the RBM cannot learn the state x' due to its rare realization. This leaves low probability regions of P_{phys} undefined. This is especially a problem at low temperatures. Therefore, the solution to the problem should be not only to train RBM with samples from P_{phys} but also to include samples from the RBM itself. This will suppress the development of high probability regions of P_{RBM} in areas where no physical states are sampled. In practice, two steps have to be performed. First, states from the physical distribution are sampled through the METROPOLIS algorithm. Second, in each training step, we sample from the RBM using Gibbs sampling (see Appendix B 1), combine it with samples obtained previously with METROPOLIS MC, and then perform one training step. In addition, a pretraining step is also performed. In this step, the CRBM is only trained using random samples as discussed in Appendix A. The biggest drawback of using the RBM in Monte Carlo simulation is that a METROPOLIS simulation (numerically expensive) needs to be performed before training. This usually makes the use of RBMs redundant since they can only be trained if the METROPOLIS algorithm performs well, in which case expectation values could have been computed directly. The second drawback as mentioned in the Introduction is that for larger lattice size L , the number of parameters that need to be learned scales with L^4 in two dimensions. So the larger L is, the more training time is required. To overcome these difficulties, we propose to employ convolutional RBM (CRBM), a model used originally in image feature extraction [20,21], which exploits the fact that models in question are translationally invariant and the interactions are local. In particular, the matrix W can be chosen such that the probability distribution $P_{\text{RBM}}(v; W)$ is translationally invariant. Consequently, instead of all lattice points being connected to each other like in the fully connected RBM, in the CRBM only the neighboring lattice sites are connected. This is illustrated in Fig. 1(b) for the one-dimensional case. Translation symmetry also requires that v_1 should be connected to v_2 in the same way as v_3 is connected to v_4 . With nearest-neighbor interactions, applying this condition, only two weights remain. After summing over h an Ising-like interaction is obtained.

The interaction part of the “free” energy for Fig. 1(b) is

$$\begin{aligned} F_{\text{CRBM}}(v) &= f(v_1, v_2) + f(v_2, v_3) \\ &\quad + f(v_3, v_4) + f(v_4, v_1), \\ f(v_i, v_k; W) &= \log 1 + e^{h+w^1 v_i + w^2 v_j}. \end{aligned} \quad (11)$$

If W is chosen such that $f(v_i, v_j; W) = \beta J s_i s_j + C$ then the RBM is equivalent to Ising model ($s_i = 2v_i - 1$). In Appendix C, an analytical mapping between RBM and the Ising model is calculated as in Ref. [23]. This is a good point to emphasize that the RBM does not directly learn statistical correlations but learns how to reproduce the energy of the target distribution. Even though the statistics differ greatly between small lattices and large lattice sizes, the structure of the energy remains unchanged and so the CRBM can be trained at a small lattice size and then applied to large lattice sizes.

For the two-dimensional case, the energy is

$$\begin{aligned} E(v, h) &= - \sum_{k,i,j} h_{ij}^k (W^k * v)_{ij} \\ &\quad - \sum_k h_{\text{bias}}^k \sum_{i,j} h_{ij}^k - v_{\text{bias}} \sum_{i,j} v_{ij}, \end{aligned} \quad (12)$$

where the symbol $*$ represents a wraparound convolution between the lattice and the kernel. Note that the visible units have dimension $L \times L$ and the hidden units have dimension $K \times L \times L$. They are connected through K convolutional kernels W^k . After summing over the hidden units, this gives the free energy

$$F(v) = -v_{\text{bias}} \sum_{i,j} v_{ij} - \sum_{k,i,j} \log 1 + e^{(v * W^k)_{ij} + h_{\text{bias}}^k}. \quad (13)$$

Note that the boundary conditions of the CRBM will be adapted to the ones of the physical problem. In the following sections, we will compare METROPOLIS MC and CRBM results in the application to the Ising and the Kitaev model in two dimensions.

III. A PRIMER: 2D ISING MODEL

The 2D Ising model with nearest-neighbor ferromagnetic interaction, $J < 0$, among the spins s_i

$$E(s) = -J \sum_{\langle i,j \rangle} s_i s_j \quad (14)$$

is one of the simplest models in numerical statistical physics and RBMs have been successfully applied [18]. Here, the local update METROPOLIS is compared with the CRBM as a proof of concept. We employ THEANO [24], a python library, for optimization and sampling. Differentiation of the loss is performed automatically by the library. Calculations for the METROPOLIS-HASTING algorithm were performed using a CPU and the convolutions needed for the CRBM were computed using a GPU.

The CRBM consists of two kernels with size 2×2 and is first trained for the temperature $T = 2.2$. The model has $2 \times 2 \times 2$ weights, one visible bias, and two hidden biases. These are 11 free parameters for the free energy $F_{\text{RBM}}(x)$. The two trained kernels as expected only show interactions between nearest neighbors (Fig. 2). The CRBM can be trained




FIG. 2. Convolution kernels W_k [see Eq. (13)] for the CRBM with two convolutional kernels, two hidden bias, and one visible bias. The CRBM was trained at $T = 2.2$.

with $L = 3$ states since the interaction is only nearest neighbors. Furthermore, we stress that to train the CRBM for the Ising model no METROPOLIS MC simulation is necessary as the structure of the energy distribution is simple. At $L = 3$ there are only $2^3 = 8$ possible states so instead of first sampling with METROPOLIS we can directly train the CRBM with those states or alternatively directly use the analytical mapping between the Ising model and the CRBM found in Appendix C. At the same time, once the CRBM has been trained, the MC simulation to compute expectation values can be performed. The emergent statistical behavior of the simple structure of the energy leads to finite-size effects.

Samples are generated with both the METROPOLIS MC algorithm and the previously trained CRBM. For the METROPOLIS algorithm, k steps are performed between each recorded step, while only one Gibbs step is performed between each

recorded step by the CRBM. k is chosen such that both simulations take the same amount of time (see also Appendix D). For example, at $L = 500$, $k = 10^4$ METROPOLIS steps take the same time as one CRBM step. The behavior of the thermodynamic observables at different temperatures matches between METROPOLIS and CRBM [see Figs. 3(c) and 3(d)]. The autocorrelation of the energy is compared between different lattice sizes at $T = 2.2$ [Fig. 3(a)] and $T = T_c = \frac{2}{\ln(1+\sqrt{2})}$. At small lattice sizes, the METROPOLIS MC algorithm performs better than the CRBM, which reverses at $100 < L$. At $L = 500$ and $T = 2.2$ ($T = T_c$) the CRBM is seven times (four times) faster than METROPOLIS MC. It is important to realize that for the Ising model there are cluster update algorithms that would outperform the CRBM. The only remarkable thing is that for interactions where these techniques are not available, the CRBM can be used. The use of the CRBM will be shown to be especially useful in the application in systems where it is computationally expensive to compute the energy as will be shown in the next section. A JUPYTER notebook with a showcase for the Ising model can be found in [25].

IV. KITAEV MODEL ON HONEYCOMB LATTICE

KitaeV's honeycomb lattice model [26] is the actual model to which we would like to apply the CRBM. It acquired significant attention recently due to the nontrivial spin liquid ground state with Majorana excitations [27], which can be computed analytically. The Hamiltonian on the honeycomb

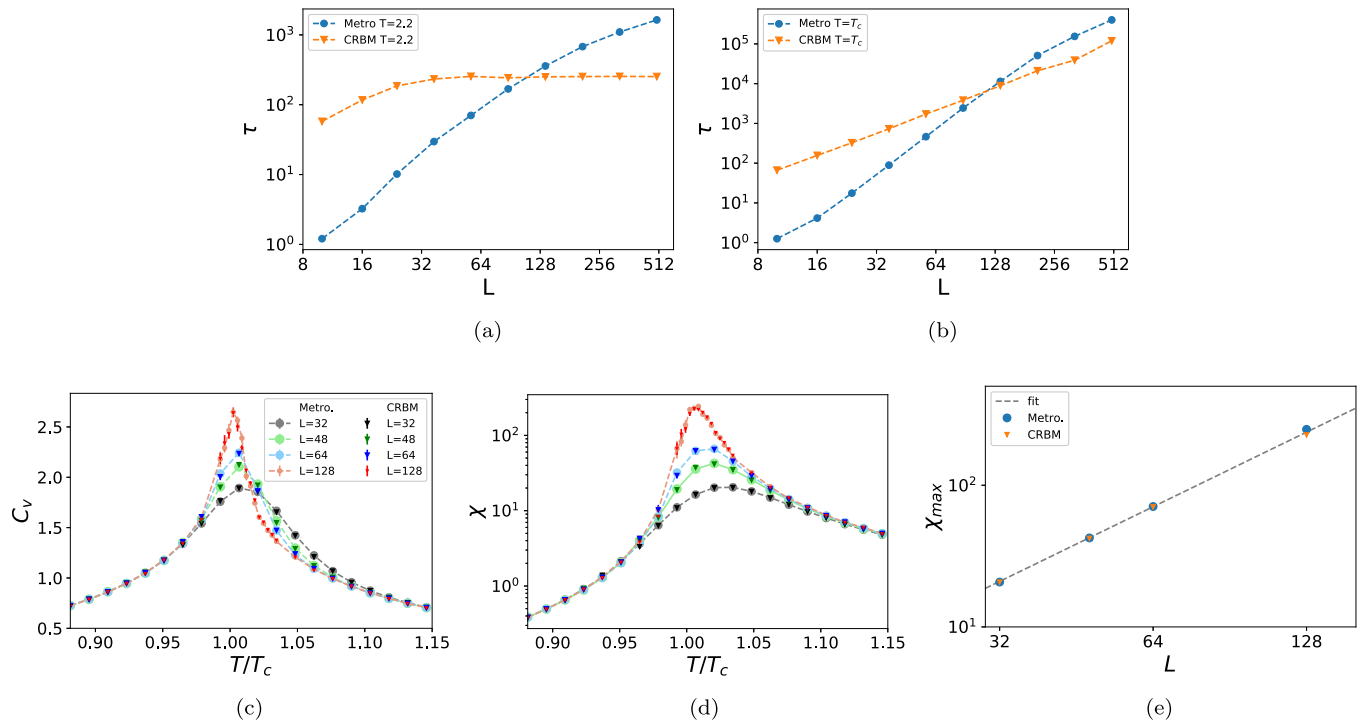


FIG. 3. Autocorrelation time at different lattice sizes at (a) $T = 2.2$ and (b) $T = T_c$ for simulations with 10^8 samples. (c) Specific heat and (d) magnetic susceptibility at $L = \{32, 48, 64, 128\}$ for an MC with 10^7 samples. (e) Finite-size scaling of the maximum of the susceptibility at different lattice sizes yielding the expected critical exponent of $\gamma_{\text{CRBM}} = 1.7$ and $\gamma_{\text{metro.}} = 1.8$ vs the theoretical value of $\gamma_{\text{theo.}} = 1.75$. For a higher accuracy in the critical exponent more samples at $L = 128$ and a more rigorous finite-size scaling would be necessary.




FIG. 4. Illustration of the honeycomb lattice. The unit cell consists of two atoms (black and white dots).

lattice, as shown in Fig. 4, has the following form:

$$H = -J_x \sum_{\langle ij \rangle_x} \sigma_i^x \sigma_j^x - J_y \sum_{\langle ij \rangle_y} \sigma_i^y \sigma_j^y - J_z \sum_{\langle ij \rangle_z} \sigma_i^z \sigma_j^z \quad (15)$$

with anisotropic exchange interaction, J_x , J_y , J_z , and Pauli matrices σ_i^α . Applying the Jordan-Wigner and the Majorana transformation [27] the Hamiltonian transforms to

$$H = -iJ_x \sum_{x \text{ bonds}} c_b c_w - iJ_y \sum_{y \text{ bonds}} c_b c_w - iJ_z \sum_{z \text{ bonds}} \eta_r c_b c_w, \quad (16)$$

where c_w (c_b) are Majorana operators defined at white (black) lattice points and $\eta_r = i\tilde{c}_b \tilde{c}_w = \pm 1$ are classical \mathbb{Z}_2 variables defined on each z bond. Through eigenvalue decomposition of the Hamiltonian, one obtains a free energy $F(\eta_r) = -T \sum_\lambda \log 2 \cosh(\beta \epsilon_\lambda / 2)$. The η configuration will be sampled from $P(\eta_r) = \frac{e^{-\beta F(\eta_r)}}{Z}$. The temperature is in units of the interaction strengths J_α and $\hbar = k_B = 1$. The usual METROPOLIS MC simulation is performed following Ref. [28] using periodic boundary conditions in the z -direction. As proof of concept, we first employ CRBM using periodic boundary conditions for both z - and for x - y -directions. Next, we will explore the open boundary condition only along the z -direction as in Ref. [28]. Note that for all our calculations the interaction strength is chosen to be $J_\alpha = \frac{1}{3}$.

It is important to note that the free energy of the Kitaev model is expensive to compute, which is in contrast to the Ising model where computing the energy is cheaper than one CRBM Gibbs step. For the Kitaev model, the computation of the free energy is as expensive as $k = 4$ Gibbs updates at lattice size $L = 8$ and as expensive as $k = 350$ Gibbs updates at $L = 30$. For more details see Appendix D.

Monte Carlo simulations of the Kitaev model have two major difficulties upon increasing the lattice size L . First, the computational complexity of the free energy scales with $\mathcal{O}(L^6)$. Second, the autocorrelation time $\tau_{\text{metro}}(L)$ of the METROPOLIS algorithm increases with larger L . In total, the complexity increase is $\mathcal{O}[\tau_{\text{metro}}(L)L^6]$.

The CRBM tackles the slowness of the METROPOLIS algorithm. Instead of having to do expensive METROPOLIS steps,

we employ cheap CRBM steps (convolution) $\mathcal{O}(L^2)$ until the original state is uncorrelated to the new state. This new state is then corrected through the parallel tempering exchange correction. The correction step involves computing the expensive free energy $\mathcal{O}(L^6)$. This means that if the CRBM is close enough to the physical distribution, this approach has a complexity $\mathcal{O}[L^6 + \tau_{\text{CRBM}}(L)L^2]$. So we conclude that a well-trained CRBM will perform Monte Carlo with a complexity of $\mathcal{O}(L^6)$ instead of $\mathcal{O}[\tau_{\text{metro}}(L)L^6]$. In the following, we apply CRBM to the Kitaev model with open and periodic boundary conditions.

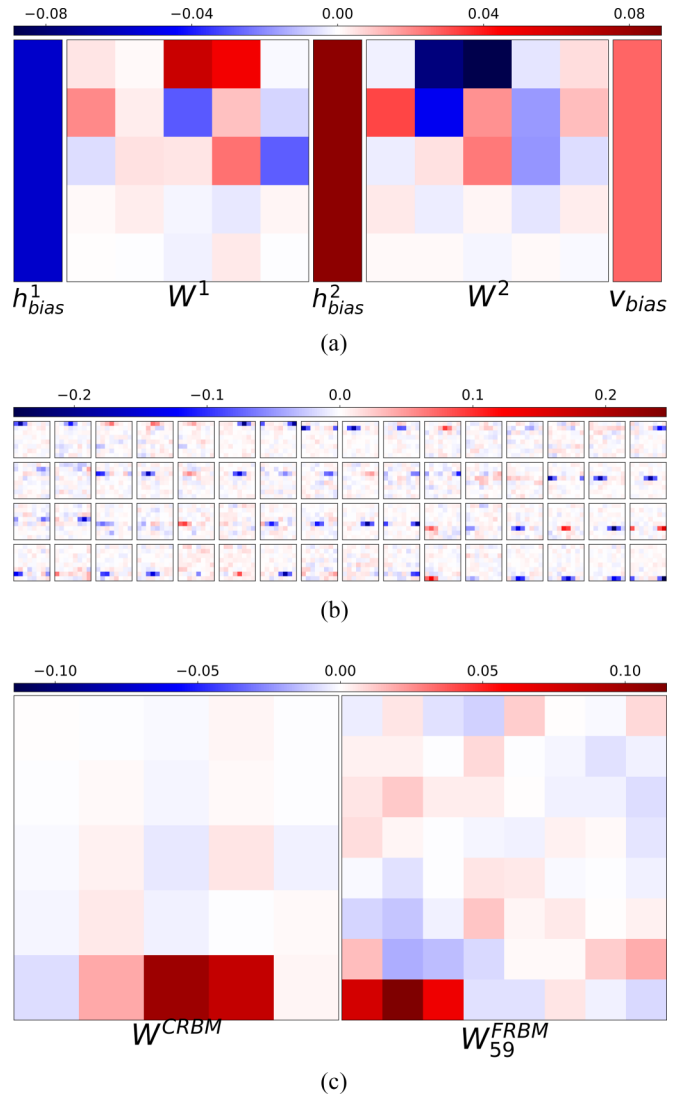


FIG. 5. FRBM and CRBM trained at $L = 8$ for the Kitaev model at $T = 0.018$ with periodic boundary conditions. (a) CRBM with two convolution kernels W^k with size 5×5 . The lattice points η_r interact with an effective interaction with range 4×3 . (b) Weight matrix W of the FRBM, which is reshaped $W_{i,j} \rightarrow W_{i,(j_1,j_2)}$ so that each panel represents an interaction between visible units. The panels were also sorted to show the translational invariance. An effective interaction of 3×1 remains. [(c), left] is the kernel of a CRBM trained with one kernel with size 5×5 . [(c), right] is the 59 component in the figure above.

A. Periodic boundary conditions

1. Comparing FRBM and CRBM with METROPOLIS

In this section, performance will be compared between CRBM, FRBM, and local METROPOLIS MC with periodic boundary conditions. First, METROPOLIS MC simulations using parallel tempering at $L = 8$ are performed to generate 4×10^4 states at each temperature. Then, for each temperature, a different CRBM with kernel size $2 \times 5 \times 5$ and an FRBM with a weight matrix with size $8^2 \times 8^2$ were trained at the lattice size $L = 8$.

The trained FRBM and CRBM kernel for $T = 0.018$ are plotted in Fig. 5. Before training the FRBM, all gauge field components can potentially interact, which stands in contrast to the CRBM where only a 5×5 field of neighboring η_i can interact. After training as expected, the trained weight matrix of the FRBM learns to ignore long-range interactions by setting them to 0 and only three gauge field interactions in the x - y -direction remain. Note that all weight matrix elements in the FRBM are almost the same just shifted around to obtain translation invariance, which is inherent in the CRBM. For the CRBM an effective interaction, where 4×3 field of lattice sites interact with each other, remains after training. Interestingly, the FRBM only captures a strong interaction in the x - y -direction but not in the z -direction, while the CRBM shows interactions in both directions. This is still the case if the amount of hidden variables for the FRBM is increased.

To further understand the origin of the difference between CRBM and FRBM, the CRBM was trained again with a smaller kernel with size $1 \times 5 \times 5$. This kernel is compared in Fig. 5(c) to one of the interaction terms of an FRBM trained with 8^2 hidden units. When restricting the CRBM to just one kernel it learns the same pattern as the FRBM and accordingly performs similarly. Since the reduced CRBM no longer can capture the full interaction, it focuses on the stronger interaction in the z -direction and neglects the interaction in the x - y -direction. This explains why the FRBM with 8^2 hidden units performs badly. The FRBM with 2×8^2 hidden units should be able to learn the interaction to the same degree as the CRBM since a CRBM with kernel size $2 \times 5 \times 5$ can be mapped to an FRBM with 2×8^2 hidden units. The question that remains is why in this case the FRBM cannot learn the interaction in the x - y -direction. To find an answer the learned

$1 \times 5 \times 5$ CRBM kernel is duplicated to yield a kernel of size $2 \times 5 \times 5$. When this kernel is trained it does not learn the interaction in the z -direction, as it gets trapped in the local minima where only the interaction in the x - y -direction is learned. We conclude that it is much easier for the FRBM to get stuck in local minima since it not only needs to learn the interaction but also needs to learn to be translationally invariant, which is already encoded in the CRBM.

Next, we analyze the specific heat at different temperatures. We focus only on the low-temperature crossover since the high-temperature crossover does not pose any problem for the MC. The METROPOLIS sampling was performed using parallel tempering. For the RBMs, a modified version of parallel tempering is used. Between parallel tempering corrections steps, $k = 4$ Gibbs steps were performed with the RBMs (see Appendix B 2). Note that k is adapted so that the k Gibbs steps take as much time as one correction step $k = \frac{t_{\text{Kitaev}}}{t_{\text{CRBM}}}$ (see Appendix D).

For the lattice size $L = 8$, the specific heat values [Fig. 6(a)] for both the CRBM, FRBM, and METROPOLIS MC agree. Nevertheless, METROPOLIS MC shows a larger error in the crossover regime. Observe that for the METROPOLIS algorithm, each update flips one gauge field component at most. In contrast, each sample for the RBMs is almost completely uncorrelated with the previous one. This is also reflected in the integrated autocorrelation time [Fig. 6(b)]. The CRBM achieves an autocorrelation time that is 100 times smaller than that by METROPOLIS MC. The decreased autocorrelation will make it possible to sample effectively at large system sizes. The CRBM autocorrelation is also two times smaller than the one achieved by the FRBM. This is still the case if the amount of hidden units is increased by a factor of 2. The CRBM has smaller autocorrelation than the FRBM due to a smaller loss during training, which can be traced back to the inability of the FRBM to learn the interactions in the z -direction. This is also reflected in the acceptance rate, which is larger for the CRBM. One can also see that the FRBM with more hidden units achieves a larger acceptance rate than the other FRBM at low temperatures.

For $L = 16$, METROPOLIS acquires problems in converging at temperatures close to the crossover region as can be seen in Fig. 7(a). The temperature points have to be increased by a factor of 2, to take advantage of parallel tempering so

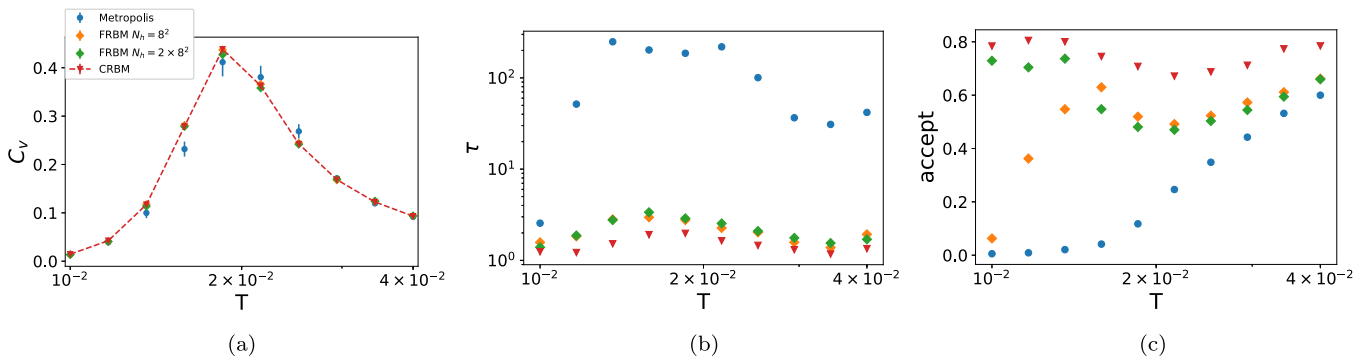


FIG. 6. Numerical results for the Kitaev model at $L = 8$ with periodic boundary conditions zoomed around the low temperature crossover generating 4×10^4 samples. (a) Specific heat, (b) autocorrelation time, and (c) acceptance rate for the METROPOLIS, FRBM, and CRBM methods.

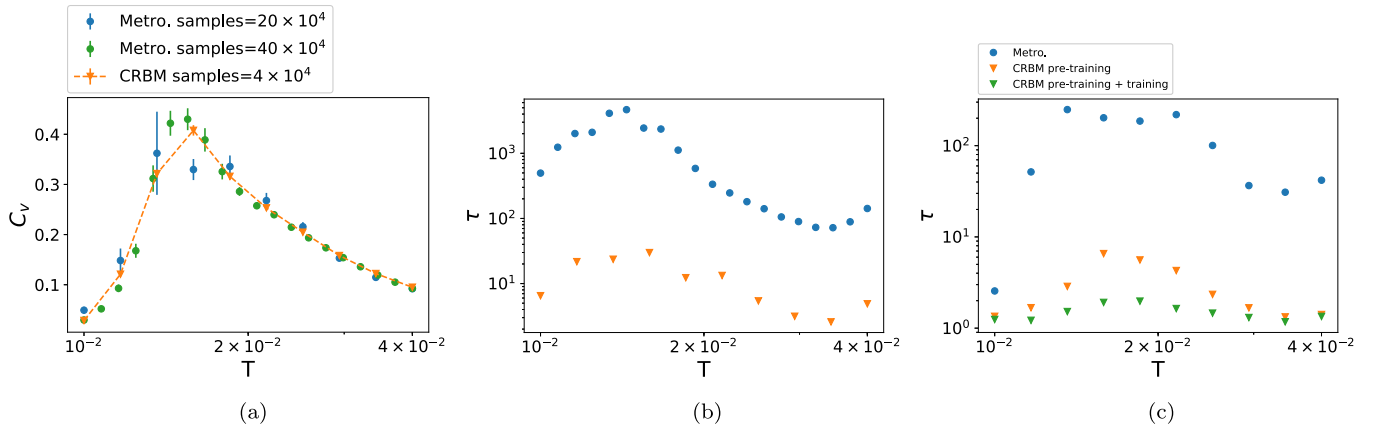


FIG. 7. (a) Specific heat and (b) autocorrelation time at $L = 16$ for the Kitaev model with periodic boundary conditions zoomed around the low temperature crossover. (Blue) METROPOLIS parallel tempering with ten temperature points. (Orange) CRBM parallel tempering with ten temperature points. (Green) METROPOLIS parallel tempering with 20 temperature points. (c) Autocorrelation time with 4×10^4 samples at $L = 8$. Metropolis is compared to both the CRBM only trained with pre-training and the CRBM completely trained.

that the results are more stable and also increase the sample size to 40×10^4 . In contrast, the CRBM works well with ten times fewer samples and half as many temperature points, even though it was trained at $L = 8$. This is explained by the autocorrelation time, which, close to the crossover region, is about 150 times larger for METROPOLIS. Note that the use of a fully connected RBM would not be possible here due to the increasing need for samples due to the larger lattice size. Besides, it would need to learn from the poorly performing METROPOLIS.

2. Pretraining

We discuss the effect of the pretraining step where the CRBM is first trained with 10^4 states generated at random. The cheap pretraining mainly speeds up the convergence of training. To train the CRBM without the pretraining step, the model needs to be trained on average 200 epochs to reach the minimum of the loss. In contrast, this is reduced to an average of 30 epochs if pretraining is applied beforehand. In addition, if only the pretraining is done, the CRBM still performs better than the METROPOLIS algorithm, see Fig. 7(c). Note that to perform the pretraining step, no METROPOLIS MC simulation

needs to be carried out beforehand. This gives us the possibility to sample the physical states from the pre-trained CRBM and then use those same states for finalizing its training. This was tested and gives similar results to training with states sampled with local METROPOLIS.

B. Open boundary conditions

To treat the open boundary conditions we adapt the CRBM by using the locally connected CRBM (LCRBM) to learn the free energy of a system with open boundary conditions in x - y -direction and periodic boundary conditions in the z -direction. The LCRBM is similar to the CRBM in that in the bulk of the lattice a translation invariance is assumed, i.e., the weights are shared. The two models are, however, different in that for the LCRBM new unshared kernels are introduced at the edges. For this purpose, three unshared kernels are introduced on the left and right sides to learn boundary effects. Instead of training a new LCRBM, we use the values of the CRBM used for periodic boundaries and then retrain it at $L = 12$, which converges after only a few epochs. The kernels before training and after are very similar. For $L = 16$, the specific heat is

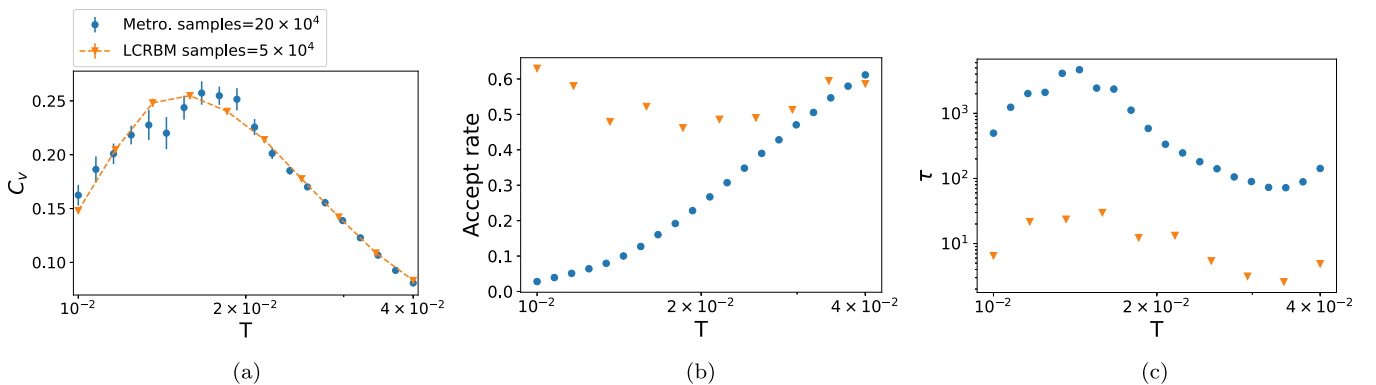


FIG. 8. Numerical results for the Kitaev model with $\alpha = 1$, $L = 16$, and open boundary conditions for the (a) specific heat, (b) acceptance rate of the parallel tempering exchange between the physical distribution and the LCRBM, and (c) autocorrelation time of the energy. Blue points refer to the METROPOLIS parallel tempering with 20 temperature points and orange triangles are CRBM parallel tempering with ten temperature points.




FIG. 9. (a) Specific heat around the first low temperature crossover with $\alpha = 1$ and $L = 24$ at different temperatures with open boundary conditions. (b) Energies of samples during the MC, for METROPOLIS and the LCRBM, at $T = 10^{-2}$. Panel (c) compares specific heat obtained with the LCRBM at $L = \{12, 20, 30\}$. Note that four times more samples were generated for the METROPOLIS algorithm to reduce errors to obtain agreement between the two methods.

compared between LCRBM and METROPOLIS [see Fig. 8(a)]. The LCRBM has an extremely low error even though eight times fewer computations were performed, which is explained by the lower autocorrelation time. This is especially noticeable at lower temperatures and close to the crossover.

Another aspect of interest is the acceptance rate and the autocorrelation time, Figs. 8(b) and 8(c). The acceptance rate for METROPOLIS increases from almost 0 to 0.6 with increasing temperatures. This is expected as smaller T translate directly to smaller acceptance rates. For the LCRBM the acceptance rate is not strongly temperature-dependent, as it is mostly influenced by how well the LCRBM fits the physical probability distribution. One has to keep in mind that each time a state is accepted by METROPOLIS a maximum of one gauge component has been flipped. In contrast, each state accepted by the physical distribution from the LCRBM is almost independent of its predecessor. This is confirmed by the autocorrelation time. The autocorrelation of the LCRBM at the crossover is an average of ~ 100 times smaller than METROPOLIS. Taking into account that more temperature points are needed, this is equivalent to a ~ 200 times faster simulation. Note that, to obtain error values that are small enough using METROPOLIS, four times more samples had to be generated.

Increasing the lattice size to $L = 24$, the METROPOLIS method does not converge, whereas the LCRBM trained at $L = 8$ converges easily [Fig. 9(a)]. For METROPOLIS the warm-up phase had to be increased to 20×10^4 and 100×10^4 states had to be sampled. This can be traced back to the high autocorrelation time. In Fig. 9(b) the energies during the MC at low a temperature are shown. The METROPOLIS method takes a long time until equilibrium is reached and one can see strong correlations between samples. In contrast, the LCRBM is in equilibrium after the first LCRBM parallel tempering step, as one LCRBM parallel tempering step consists of 140 Gibbs steps ($L = 24$) and one corrections step.

The specific heat for $L = 12, 20, 30$ as shown in Fig. 9(c) only show small differences between each other in agreement with Motome *et al.*[27]. For two dimensions there is no phase transition just a crossover since the specific heat is not singular at $L \rightarrow \infty$.

V. CONCLUSION

To conclude, we employ convolutional restricted Boltzmann machines to learn an effective energy for the 2D Ising and Kitaev models. In contrast to the fully connected RBMs, which suffer from long training times when the lattice size is increased, CRBM was shown to be more efficient because it can be trained for smaller lattice sizes before applying to the larger lattices using translation invariance of the model. It was also shown that for the Kitaev model the CRBM better captures the physical interaction than the FRBM since it does not need to learn to be translational invariant. We showed that, not only can a CRBM reproduce thermodynamic observables accurately and give results with smaller errors for both periodic and open boundary conditions for the Kitaev model, but that it is also able to simulate lattice sizes up to $L = 30$, which is not possible using the local METROPOLIS algorithm in any reasonable time. Our results also confirm that the Kitaev model does not possess a phase transition but only a crossover. It is of interest to explore in future Kitaev-like model systems with a phase transition, as the CRBM could enable more accurate finite-size scaling.

ACKNOWLEDGMENT

We thank N.B. Perkins for the stimulating discussions.

APPENDIX A: TRAINING

The physical model with energy $E_{\text{phys}}(x)$ will be approximated by the RBM. States in the physical model are encoded as -1 and $+1$. The RBM encodes its states as, 0 and 1. While sampling and training, a conversion between the two is necessary.

Training is done supervised similarly to Ref. [19] by minimizing the loss function

$$\text{loss}(W) = \frac{1}{M} \sum_{i=1}^M [E_{\text{phys}}(x_i) - F_{\text{RBM}}(x_i; W) - C(W)]^2 \quad (\text{A1})$$

Algorithm 1 Train the RBM

```

 $N \leftarrow$  train size


---


 $M \leftarrow$  batch size
 $K \leftarrow$  amount of used kernels
 $L_W \leftarrow$  kernel size
 $ub \leftarrow$  Buffer update size
initialize  $W \sim \mathcal{N}(m = 0, \sigma = \frac{2}{KL_W^2})$  with size  $K \times L_W \times L_W$ 
initialize  $v_{\text{bias}} = 0$  with size 1
initialize  $h_{\text{bias}} = 0$  with size  $K$ 
 $param = [W, v_{\text{bias}}, h_{\text{bias}}]$ 
sample  $x_{\text{phys}} \sim P_{\text{phys}}$  with METROPOLIS
calculate  $E_{\text{phys}}(x_{\text{phys}})$ 
initialize  $\bar{x}_{\text{RBM}} \sim \text{Binomial}(\text{size} = (ub, L, L), p = 0.5)$ 
function UPDATE BUFFER
  update  $\bar{x}_{\text{RBM}}$  with  $k$  steps of Gibbs sampling
   $x_{\text{RBM}} \leftarrow x_{\text{RBM}} \cup \bar{x}_{\text{RBM}}$ 
  calculate  $E_{\text{phys}}(\bar{x}_{\text{RBM}})$ 
  if  $\text{size}(x_{\text{RBM}}) > N$  then
    remove last elements of  $x_{\text{RBM}}$ 
  end if
end function
iterate UPDATE BUFFER until  $\text{size}(x_{\text{RBM}}) = N$ 
while  $\text{loss} > \epsilon$  do
  UPDATE BUFFER
   $x_{\text{tot}} = x_{\text{RBM}} \cup x_{\text{phys}}$ 
  for  $x = (M \text{ elements of } x_{\text{tot}})$  do
     $\text{diff}^i = E_{\text{phys}}(x^{(i)})\beta - F_{\text{RBM}}(x^{(i)}; param)$ 
     $C = \frac{1}{M} \sum_{i=1}^M \text{diff}^i$ 
     $\text{loss} = \frac{1}{M} \sum_{i=1}^M (\text{diff}^i - C)^2$ 
     $param \leftarrow param - \text{ADAM}(\frac{\partial \text{loss}}{\partial param})$ 
  end for
end while

```

where L is the lattice size, β is the inverse temperature, and C is a value that can be chosen freely because the probability function is invariant under the addition of a constant to the energy. C is chosen such that the loss is minimal:

$$C(W) = \frac{1}{M} \sum_{i=1}^M E_{\text{phys}}(x_i)\beta - F_{\text{RBM}}(x_i; W). \quad (\text{A2})$$

In practice, to minimize the loss the ADAM batch-gradient descent is used. A combination of samples from both the physical distribution P_{phys} and the RBM's distribution P_{RBM} are used for training. These are expensive, as after each training step P_{RBM} changes. This means that new samples need to be generated for training. In addition, in each training step, $E_{\text{phys}}(x)$ for these new samples has to be computed, which is numerically expensive. To alleviate the computational load, we perform two additional steps. First, a pretraining step is performed in which completely random states are generated and trained until the loss of the physical states no longer decreases. This reduces the training time because $E_{\text{phys}}(x)$ only needs to be computed once for each random state, and the random states can be used to pretrain the RBM at all temperatures. The second step is to use a buffer of past RBM samples where instead of generating new samples in each training step, only a portion of the used states are new and the rest are reused

from previous training steps. In each epoch, ub new samples are added to the buffer. In practice, a value of $ub = 200$ is chosen. Below we give the details of the algorithms involved.

APPENDIX B: SAMPLING**1. Gibbs sampling**

The standard way to sample from an RBM is Gibbs sampling. The conditional probability $P(v|h) = \frac{P(v,h)}{P(h)}$ is computed to

$$P(h_i = 1|v) = \frac{e^{h_{\text{bias}} + \sum_j W_{ij} v_j}}{1 + e^{h_{\text{bias}} + \sum_j W_{ij} v_j}} \quad (\text{B1})$$

$$= \sigma \left(h_{\text{bias}} + \sum_j W_{ij} v_j \right), \quad (\text{B2})$$

$$P(v_j = 1|h) = \sigma \left(v_{\text{bias}}^j + \sum_i h_i W_{ij} \right). \quad (\text{B3})$$

The conditional probabilities describe a binomial distribution that depends either on h or on v . A bipartite Markov chain is constructed. First, a sample $h^{(1)}$ is drawn from the conditional distribution $P(h^{(1)}|v^{(1)})$. Second, the new visible sample is drawn from $P(v^{(2)}|h^{(1)})$.

$$v^1 \xrightarrow{P(h^1|v^1)} h^1 \xrightarrow{P(v^2|h^1)} v^2.$$

The only difference between this and a normal Markov chain is the middle step where h is sampled. The RBM can be reformulated so that it is equivalent to a normal Markov chain, as the transition probability from $v^{(1)}$ to $v^{(2)}$ is

$$P(v^{(2)}|v^{(1)}) = \sum_h P(v^{(2)}|h)P(h|v^{(1)}). \quad (\text{B4})$$

Note that this quantity cannot be computed analytically, but it can be proven that this conditional probability fulfills the detailed balance equation. This means that the RBM describes an ergodic Markov chain.

2. Parallel tempering

Often problems plague METROPOLIS MC simulations at small temperatures because the acceptance rate is very low. A method that tackles this is parallel tempering. Parallel tempering takes advantage of the fact that probability distributions at similar temperatures overlap. It introduces an exchange between high and low temperatures so that the low-temperature dynamics can take advantage of the high acceptance rates at high temperatures. States are exchanged at different temperatures with METROPOLIS acceptance probability

$$A(x_1, x_2) = \min(1, e^{F_1(x_1)\beta_1 + F_2(x_2)\beta_2 - F_1(x_2)\beta_1 - F_2(x_1)\beta_2}). \quad (\text{B5})$$

For the Kitaev model, exchanges are computationally cheap because the eigenvalue problem for both states has already been solved. In Fig. 10(a), the full sampling cycle can be seen. First, a METROPOLIS step and then a state exchange with the neighboring temperature is performed.

A small modification leads to parallel tempering with an exchange correction. Instead of using METROPOLIS MC for the updates, we introduce the RBM as a different probability

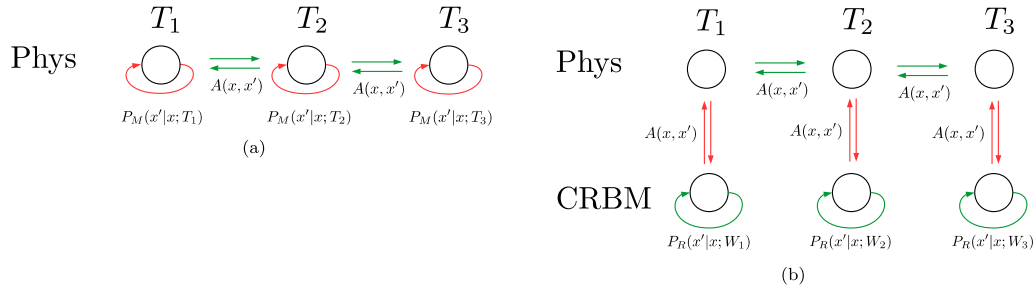


FIG. 10. (a) Parallel tempering step where $P_M(x'|x)$ represents a local METROPOLIS MC step and $A(x', x)$ [Eq. (B5)] is the probability that the two temperatures exchange states. Green arrows represent a cheap operation. Red arrows represent an expensive operation. (b) Parallel tempering exchange correction where instead of updating with METROPOLIS states, exchanges with CRBMs are introduced. A P_R step represents k Gibbs steps.

distribution that the physical states can exchange with. The more the approximate RBM distribution and physical distribution overlap, the larger the acceptance rate is [see Eq. (B5)]. Sampling entails three steps as seen in Fig. 10(b). First, the RBM states are updated through Gibbs sampling a certain amount of times (in practice we use 60 steps). The RBM states are then exchanged with the physical distribution with probability $A(x_1, x_2)$. Lastly, a normal parallel tempering exchange is performed. Each time a state is accepted, it is almost uncorrelated to the previous one. This is in stark contrast with local METROPOLIS MC simulations where only one lattice point is updated each time as a new state is accepted. Note that physical exchanges are cheap because one does not need to recalculate the eigenvalue problem to know what the free energy at a different temperature is. Note that this method was not used for the Ising model.

It is interesting to see what would happen if the RBM is not statistically corrected. We show the specific heat for both the CRBM and FRBM at $L = 8$ in Fig. 11(a) and the LCRBM at $L = 24$ in Fig. 11(b). Observe that the CRBM and LCRBM without corrections are close to the corrected versions while the FRBM results lie quite far. This is due to the fact that the FRBM was no able to learn the interaction in the z -direction. Furthermore, note that the increase in lattice size

did not affect the closeness between corrected and uncorrected CRBM.

APPENDIX C: ANALYTICAL SOLUTION

As an alternative to learning the 2D Ising model, it is possible to compute the values of the convolutional kernels analytically. This is done similarly to Ref. [23]. The energy associated with an interaction between two neighboring spins in the Ising model is $E(s_1, s_2) = s_1 s_2$. For the CRBM if the kernel is chosen to be

$$W_1 = \begin{pmatrix} W & 0 \\ W & 0 \end{pmatrix},$$

$$W_2 = \begin{pmatrix} W & W \\ 0 & 0 \end{pmatrix},$$

the interaction term between s_1 and s_2 is $F(v_1, v_2) = \frac{v_{\text{bias}}(v_1 + v_2)}{4} + \log(1 + e^{h_{\text{bias}} + W(v_1 + v_2)})$. Note that for each v_i , the $v_{\text{bias}} v_i$ term will appear four times in the free energy. This produces the $1/4$ factor in front of v_{bias} . Also, note that $s = \pm 1$ and $v = \{0, 1\}$. If $F(v_1, v_2) = \beta E(s_1, s_2) - C$, then the CRBM will be equivalent to Ising model. s_1, s_2 can take four different states, which yield three equations. Equation (C1) is obtained for the case where both spin are up. Equation (C2)

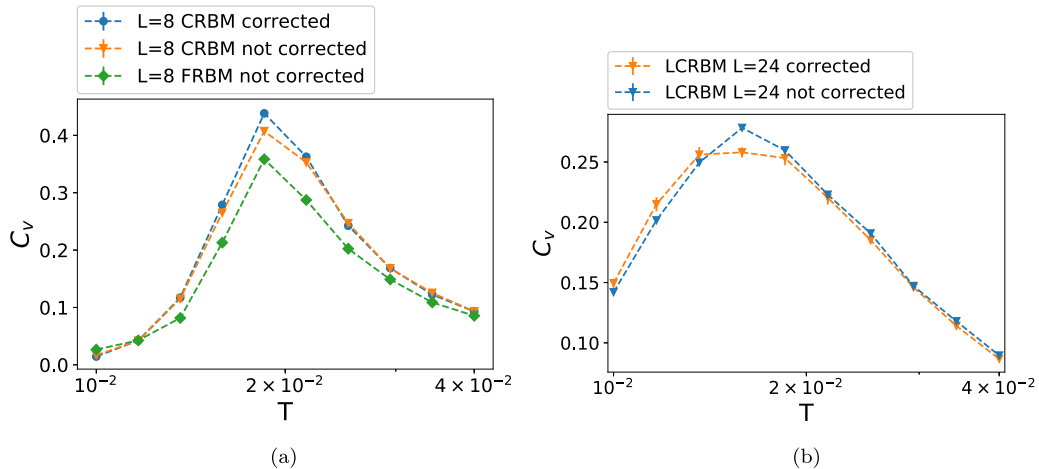


FIG. 11. Specific heat results compared between various RBMs that were statistically corrected with the parallel tempering exchange correction and ones without correction procedure. Here, 4×10^4 samples were generated. Panel (a) shows a comparison of FRBM and CRBM at $L = 8$ with periodic boundary conditions, and (b) refers to LCRBM at $L = 16$ with open boundary conditions.

is obtained for the case where one spin is up and the other is down, and Eq. (C3) for the case when both spins are down

$$\frac{v_{\text{bias}}}{2} + \log 1 + e^{h_{\text{bias}}+2W} = \beta + C, \quad (\text{C1})$$

$$\frac{v_{\text{bias}}}{4} + \log 1 + e^{h_{\text{bias}}+W} = -\beta + C, \quad (\text{C2})$$

$$\log 1 + e^{h_{\text{bias}}} = \beta + C, \quad (\text{C3})$$

they can be reduced to

$$v_{\text{bias}} = 2 \log \frac{1 + e^{h_{\text{bias}}}}{1 + e^{h_{\text{bias}}+2W}}, \quad (\text{C4})$$

$$2\beta = \log \frac{\sqrt{(1 + e^{h_{\text{bias}}})(1 + e^{h_{\text{bias}}+2W})}}{1 + e^{h_{\text{bias}}+W}}, \quad (\text{C5})$$

and Eq. (C5) can be re-expressed as

$$h_{\text{bias}} = -2W + \log \left(\frac{(e^W - 1)(\pm \sqrt{(e^W + 1)^2 - 4e^{4\beta+W}}) - 2e^{4\beta+W} + e^{2W} + 1}{2(e^{4\beta} - 1)} \right). \quad (\text{C6})$$

It only has a real solution if the term in the square root is positive, which means that

$$|W| > W_{\text{min}} = \log(2e^{2\beta}(\sqrt{e^{4\beta} - 1} + e^{2\beta}) - 1). \quad (\text{C7})$$

We expect that the minimal W would yield the smallest autocorrelation time. For this case

$$\begin{aligned} W &= \pm W_{\text{min}}, \\ h_{\text{bias}} &= -W, \\ v_{\text{bias}} &= -2W. \end{aligned} \quad (\text{C8})$$

Using this analytical result for performing MC yields similar autocorrelation times and the same expectation values as obtained from a trained kernel.

APPENDIX D: TRAINING AND SAMPLING PARAMETERS

In this Appendix, we list all our training and sampling parameters.

1. Training

All models were trained with ADAM with a learning rate of $\lambda = 10^{-3}$ and a batch-size of 20. The Ising model was trained at $L = 3$. First, the energy $E_{\text{phys}}(x)$ of all possible $2^{3^2} = 512$ states was computed. Then the loss in Eq. (5) was minimized




FIG. 12. Ratio $k = \frac{t_{\text{CRB}}}{t_{\text{metro}}}$ between the time it takes to perform one local METROPOLIS step in the Ising model and the time to perform one CRBM Gibbs update step with a kernel of size $2 \times 2 \times 2$.

until a loss of 10^{-7} is reached. This procedure is demonstrated in a JUPYTER NOTEBOOK in [25].

The Kitaev model was trained at $L = 8$ in two steps. First, the pretraining step described in Appendix A was performed until the loss did no longer decrease, this happened, on average, after 800 epochs. Second, 10^4 states are sampled with METROPOLIS MC. These states are then combined with a buffer of states sampled from the CRBM. The buffer update size is set to $ub = 200$ as described in Appendix A. This step is iterated until the loss no longer decreases, which happens on average after 30 epochs.

2. Sampling

For the Ising model sampling for local METROPOLIS was done on the CPU and sampling for the CRBM was done on the GPU. A local METROPOLIS step takes $t_{\text{metro}} = 55$ ns on a Intel(R) Core(TM) i7-9700K CPU. The time it takes to perform a CRBM Gibbs update step is lattice size-dependent, the ratio between a METROPOLIS and a CRBM update step $k = \frac{t_{\text{CRB}}}{t_{\text{metro}}}$ is computed for a GeForce RTX 2070 Super and is shown in Fig. 12. To compare the speed of both methods k METROPOLIS steps are performed between each recorded state, whereas for the CRBM only one Gibbs update step is performed. The different temperatures are simulated after each other starting with the highest temperature. For both CRBM and METROPOLIS MC a warm-up phase is performed before the recording of the observables. For consistency, this warm-up phase was chosen to be 10% of the total simulation.

To perform the local METROPOLIS sampling for the Kitaev model, parallel tempering was used. A METROPOLIS update at each temperature and a temperature exchange are per-

TABLE I. Speed of one CRBM Gibbs update with a $2 \times 5 \times 5$ kernel and the computation of the free energy of a gauge configuration for the Kitaev model with an Intel(R) Xeon(R) X5670 at different lattice sizes. k is the amount of Gibbs updates that are performed between correction steps.

L	8	16	20	24	30
CRBM Gibbs	0.09 ms	0.14 ms	0.17 ms	0.21 ms	0.28 ms
Kitaev free energy	0.35 ms	4.5 ms	12 ms	30 ms	98 ms
k	4	30	70	140	350

formed iteratively. To sample using the CRBM three steps are performed repeatedly. First, at each temperature states are updated k times using the Gibbs update from the CRBM. Second, a correction step where the CRBM and physical distribution exchange states is performed. Note that for this, the free energy of the Kitaev model needs to be computed. Lastly, a normal temperature exchange is performed. The number of steps k is varied depending on L . k is chosen such that the

Gibbs sampling of the CRBM takes as much time as one correction step. Note that the most expensive operation in the correction step is the computation of the free energy. The timings and values for $k = \frac{t_{\text{Kitaev}}}{t_{\text{CRBM}}}$ can be found in Table I. Note that the warm-up for METROPOLIS MC was chosen to be 50% of the generated samples, and the one for the CRBM just 10%. The higher percentage for METROPOLIS was chosen due to the higher autocorrelation time.

-
- [1] J. Carrasquilla, *Adv. Phys.: X* **5**, 1797528 (2020).
 - [2] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
 - [3] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
 - [4] D. Sehayeck, A. Golubeva, M. S. Albergo, B. Kulchytskyy, G. Torlai, and R. G. Melko, *Phys. Rev. B* **100**, 195125 (2019).
 - [5] G. Carleo, Y. Nomura, and M. Imada, *Nat. Commun.* **9**, 1 (2018).
 - [6] J. Carrasquilla and R. G. Melko, *Nat. Phys.* **13**, 431 (2017).
 - [7] Y. Zhang, R. G. Melko, and E.-A. Kim, *Phys. Rev. B* **96**, 245119 (2017).
 - [8] J. Greitemann, K. Liu, and L. Pollet, *Phys. Rev. B* **99**, 060404(R) (2019).
 - [9] J. F. Rodriguez-Nieva and M. S. Scheurer, *Nat. Phys.* **15**, 790 (2019).
 - [10] S. Efthymiou, M. J. S. Beach, and R. G. Melko, *Phys. Rev. B* **99**, 075113 (2019).
 - [11] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 041101(R) (2017).
 - [12] X. Y. Xu, Y. Qi, J. Liu, L. Fu, and Z. Y. Meng, *Phys. Rev. B* **96**, 041119(R) (2017).
 - [13] H. Shen, J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 241104(R) (2017).
 - [14] Y. Nagai, H. Shen, Y. Qi, J. Liu, and L. Fu, *Phys. Rev. B* **96**, 161102(R) (2017).
 - [15] C. Chen, X. Y. Xu, J. Liu, G. Batrouni, R. Scalettar, and Z. Y. Meng, *Phys. Rev. B* **98**, 041102(R) (2018).
 - [16] H. Shen, J. Liu, and L. Fu, *Phys. Rev. B* **97**, 205140 (2018).
 - [17] S. Li, P. M. Dee, E. Khatami, and S. Johnston, *Phys. Rev. B* **100**, 020302(R) (2019).
 - [18] G. Torlai and R. G. Melko, *Phys. Rev. B* **94**, 165134 (2016).
 - [19] L. Huang and L. Wang, *Phys. Rev. B* **95**, 035105 (2017).
 - [20] M. Norouzi, M. Ranjbar, and G. Mori, in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2009), pp. 2735–2742.
 - [21] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, in *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09* (Association for Computing Machinery, New York, 2009) p. 609–616.
 - [22] D. P. Kingma and J. Ba, *arXiv:1412.6980*.
 - [23] N. Yoshioka, Y. Akagi, and H. Katsura, *Phys. Rev. E* **99**, 032113 (2019).
 - [24] R. Al-Rfou, G. Alain, A. Almahairi, C. Angermueller, D. Bahdanau, N. Ballas, F. Bastien, J. Bayer, A. Belikov, A. Belopolsky *et al.*, *arXiv:1605.02688*.
 - [25] Simulating the Ising model with the CRBM method - JUPYTER NOTEBOOK, https://github.com/danielalcalde/MCMC_CRBM/blob/master/Ising_CRBM.ipynb.
 - [26] A. Kitaev, *Ann. Phys. (NY)* **321**, 2 (2006).
 - [27] J. Nasu, M. Udagawa, and Y. Motome, *Phys. Rev. B* **92**, 115122 (2015).
 - [28] J. Nasu, M. Udagawa, and Y. Motome, *Phys. Rev. Lett.* **113**, 197205 (2014).