

Learning the Ising model with generative neural networks

Francesco D'Angelo^{1,*} and Lucas Böttcher^{2,3,4,†}

¹*Institute of Neuroinformatics, University of Zurich and ETH Zurich, 8057, Zurich, Switzerland*

²*Computational Medicine, UCLA, Los Angeles, California 90095-1766, USA*

³*Institute for Theoretical Physics, ETH Zurich, 8093, Zurich, Switzerland*

⁴*Center of Economic Research, ETH Zurich, 8092, Zurich, Switzerland*



(Received 8 December 2019; accepted 8 May 2020; published 2 June 2020)

Recent advances in deep learning and neural networks have led to an increased interest in the application of generative models in statistical and condensed matter physics. In particular, restricted Boltzmann machines (RBMs) and variational autoencoders (VAEs) as specific classes of neural networks have been successfully applied in the context of physical feature extraction and representation learning. Despite these successes, however, there is only limited understanding of their representational properties and limitations. To better understand the representational characteristics of RBMs and VAEs, we study their ability to capture physical features of the Ising model at different temperatures. This approach allows us to quantitatively assess learned representations by comparing sample features with corresponding theoretical predictions. Our results suggest that the considered RBMs and convolutional VAEs are able to capture the temperature dependence of magnetization, energy, and spin-spin correlations. The samples generated by RBMs are more evenly distributed across temperature than those generated by VAEs. We also find that convolutional layers in VAEs are important to model spin correlations whereas RBMs achieve similar or even better performances without convolutional filters.

DOI: [10.1103/PhysRevResearch.2.023266](https://doi.org/10.1103/PhysRevResearch.2.023266)

I. INTRODUCTION

After the successful application of deep learning and neural networks in speech and pattern recognition [1], there is an increased interest in applying generative models in condensed matter physics and other fields [2,3]. For example, restricted Boltzmann machines (RBMs) [4–6], as one class of stochastic neural networks, were used to study phase transitions [7,8], represent wave functions [9,10], and extract features from physical systems [11,12]. Furthermore, variational autoencoders (VAEs) [13] have been applied to different physical representation learning problems [14]. In addition to the application of generative neural networks to physical systems, connections between statistical physics and the theoretical description of certain neural network models helped to gain insights into their learning dynamics [11,15–17].

Despite these developments, the representational characteristics and limitations of neural networks have been only partially explored. To better understand the representational properties of RBMs and VAEs, we study their ability to capture the temperature dependence of physical features of the Ising model. This allows us to quantitatively compare learned distributions with corresponding theoretical predictions of

a well-characterized physical system. We train both neural networks with realizations of Ising configurations and use the trained neural networks to generate new samples. Previous studies utilized single RBMs for each temperature to learn the distribution of Ising configurations that consist of a few dozen spins [15,18,19]. Our work complements Refs. [15,18,19] in three ways. First, we use system sizes that are about one order of magnitude larger than in previous studies. This has profound implications for monitoring the learning progress of RBMs. For the system sizes we consider, it is computationally not feasible anymore to directly evaluate the partition function as outlined in Ref. [18], and we therefore provide an overview of loss-function approximations that allow us to monitor RBMs during training. Second, instead of using single machines for each temperature, we focus on the training of one generative neural network for all temperatures and use classification networks to determine the temperatures of generated Ising configurations. Our results indicate that this type of training can improve the quality of VAE samples substantially. Third, in addition to RBMs, we also consider nonconvolutional and convolutional VAEs and quantitatively compare different neural network architectures in terms of their ability to capture magnetization, energy, and spin-spin correlations.

Our paper proceeds as follows. In Sec. II we give a brief introduction to generative models and summarize the concepts that are necessary to train RBMs and VAEs. To keep track of the learning progress of both models, we provide an overview of different monitoring techniques in Sec. III. In Sec. IV we apply RBMs and VAEs to learn the distribution of Ising spin configurations for different temperatures. We conclude our paper and discuss our results in Sec. V.

*fdangelo@ethz.ch

†lucasb@ethz.ch

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. *Restricted Boltzmann machine.* An RBM is composed of one visible layer (blue) and one hidden layer (red). In this example, the respective layers consist of six visible units $\{v_i\}_{i \in \{1, \dots, 6\}}$ and four hidden units $\{h_i\}_{i \in \{1, \dots, 4\}}$. The network structure underlying an RBM is bipartite.

II. GENERATIVE MODELS

Generative models are used to approximate the distribution of a data set $\mathcal{D} = \{x_1, \dots, x_m\}$ whose entities are samples $x_i \sim \hat{p}$ drawn from a distribution \hat{p} ($1 \leq i \leq m$). Usually, the distribution \hat{p} is unknown and thus approximated by a generative model distribution $p(\theta)$, where θ is a corresponding parameter set. We can think of generative models as neural networks whose underlying weights and activation function parameters are described by θ . Once trained, generative models are used to generate new samples similar to the ones of the considered data set \mathcal{D} [11]. To quantitatively study the representational power of generative models, we consider two specific architectures: (1) RBMs and (2) VAEs. We introduce the basic concepts behind RBMs and VAEs in Secs. II A and II B.

A. Restricted Boltzmann machine

RBMs are a particular type of stochastic neural networks and were first introduced by Smolensky in 1986 under the name *Harmonium* [20]. In the early 2000s, Hinton developed some efficient training algorithms that made it possible to apply RBMs in different learning tasks [6,21]. The network structure of an RBM is composed of one visible and one hidden layer. In contrast to Boltzmann machines [4], no intralayer connections are present in RBMs (i.e., the network structure of RBMs is bipartite). We show an illustration of an RBM network in Fig. 1.

Mathematically, we describe the visible layer by a vector $\mathbf{v} = (v_1, \dots, v_m)^T$, which consists of m visible units. Each visible unit represents one element of the data set \mathcal{D} . Similarly, we represent the hidden layer by a vector $\mathbf{h} = (h_1, \dots, h_n)^T$, which is composed of n hidden units. To model the distribution of a certain data set, hidden units serve as additional degrees of freedom to capture the complex interaction between the original variables. Visible and hidden units are binary [i.e., $v_i \in \{0, 1\}$ ($1 \leq i \leq m$) and $h_j \in \{0, 1\}$ ($1 \leq j \leq n$)]. We note that there also exist other formulations with continuous degrees of freedom [22]. In an RBM, connections between units v_i and h_j are undirected, and we describe their weights by a weight matrix $W \in \mathbb{R}^{m \times n}$ with elements $w_{ij} \in \mathbb{R}$. Furthermore, we use $\mathbf{b} \in \mathbb{R}^m$ and $\mathbf{c} \in \mathbb{R}^n$ to account for biases in the visible and hidden layers. According to these definitions, there is an energy

$$E(\mathbf{v}, \mathbf{h}) = -\mathbf{b}^T \mathbf{v} - \mathbf{c}^T \mathbf{h} - \mathbf{v}^T W \mathbf{h} \quad (1)$$

associated with every configuration (\mathbf{v}, \mathbf{h}) . The probability of the system to be found in a certain configuration (\mathbf{v}, \mathbf{h}) is described by the Boltzmann distribution [4,11]

$$p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-E(\mathbf{v}, \mathbf{h})}, \quad (2)$$

where $Z = \sum_{\{\mathbf{v}, \mathbf{h}\}} e^{-E(\mathbf{v}, \mathbf{h})}$ is the canonical partition function. The sum $\sum_{\{\mathbf{v}, \mathbf{h}\}}$ is taken over all possible configurations $\{\mathbf{v}, \mathbf{h}\}$.

Due to the absence of intralayer connections in RBMs, hidden (visible) variables are mutually independent given the visible (hidden) ones. Therefore, the corresponding conditional probabilities are [23]

$$p(\mathbf{v}|\mathbf{h}) = \prod_{i=1}^m p(v_i|\mathbf{h}) \quad \text{and} \quad p(\mathbf{h}|\mathbf{v}) = \prod_{j=1}^n p(h_j|\mathbf{v}), \quad (3)$$

where

$$p(h_j = 1|\mathbf{v}) = \sigma \left(c_j + \sum_{i=1}^m w_{ij} v_i \right), \quad (4)$$

$$p(v_i = 1|\mathbf{h}) = \sigma \left(b_i + \sum_{j=1}^n w_{ij} h_j \right), \quad (5)$$

and $\sigma(x) = 1/[1 + \exp(-x)]$ denotes the sigmoid function.

We now consider a data set \mathcal{D} that consists of N i.i.d. realizations $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$. Each realization $\mathbf{x}_k = (x_{k1}, \dots, x_{km})^T$ ($1 \leq k \leq N$) in turn consists of m binary elements. We associate these elements with the visible units of an RBM. The corresponding log-likelihood function is [23]

$$\begin{aligned} \log \mathcal{L}(\mathbf{x}, \theta) &= \log p(\mathbf{x}_1, \dots, \mathbf{x}_N; \theta) = \log \prod_{k=1}^N p(\mathbf{x}_k; \theta) \\ &= \sum_{k=1}^N \log p(\mathbf{x}_k; \theta) = \sum_{k=1}^N \log \sum_{\{\mathbf{h}\}} \frac{e^{-E(\mathbf{x}_k, \mathbf{h})}}{Z} \\ &= \sum_{k=1}^N \log \left[\frac{1}{Z} e^{\mathbf{b}^T \mathbf{x}_k} \prod_{j=1}^n (1 + e^{(c_j + \sum_{i=1}^m w_{ij} x_{ki})}) \right] \\ &= -N \log(Z) + \sum_{k=1}^N \mathbf{b}^T \mathbf{x}_k \\ &\quad + \sum_{k=1}^N \sum_{j=1}^n \log (1 + e^{(c_j + \sum_{i=1}^m w_{ij} x_{ki})}), \end{aligned} \quad (6)$$

where the parameter set θ describes weights and biases and the partition function is

$$Z = \sum_{\{\mathbf{h}\}} e^{\mathbf{c}^T \mathbf{h}} \prod_{i=1}^m (1 + e^{b_i + \sum_{j=1}^n w_{ij} h_j}). \quad (7)$$

Based on Eq. (6), we perform a maximum-likelihood estimation of RBM parameters. We may express the log-likelihood function in terms of the free energy $\mathcal{F} = -\log Z$ [24]. This yields $\log \mathcal{L} = N\mathcal{F} - \sum_{k=1}^N \mathcal{F}_c(\mathbf{x}_k)$, where $\mathcal{F}_c(\mathbf{x}_k) = -\log(\sum_{\{\mathbf{h}\}} e^{-E(\mathbf{x}_k, \mathbf{h})})$ is the (clamped) free energy of an RBM whose visible states are clamped to the elements of \mathbf{x}_k [25].




FIG. 2. *Block Gibbs sampling in an RBM.* We show an example of an RBM with a visible layer that consists of six visible units and a hidden layer that consists of four hidden units. Because of the bipartite network structure of RBMs, units within one layer can be grouped together and updated in parallel (*block Gibbs sampling*). Initially visible units (green) are determined by the data set. Hidden (red) and visible (blue) units are then updated in an alternating manner.

To train an RBM, we have to find a set of parameters θ that minimizes the difference between the distribution p_θ of the machine and the distribution \hat{p} of the data. One possibility is to quantify the dissimilarity of these two distributions in terms of the Kullback-Leibler (KL) divergence (i.e., relative entropy) [4]

$$D_{\text{KL}}(\hat{p}(\mathbf{x})||p_\theta(\mathbf{v})) = \sum_{\mathbf{x}} \hat{p}(\mathbf{x}) \log \frac{\hat{p}(\mathbf{x})}{p_\theta(\mathbf{v})} \geq 0. \quad (8)$$

For continuous distributions, the sum in Eq. (8) has to be replaced by a corresponding integral. As an alternative to KL-divergence minimization, it is possible to minimize the negative log likelihood. This optimization problem can be solved using a stochastic gradient descent algorithm [26], so that the model parameters are iteratively updated:

$$\theta_i^{t+1} = \theta_i^t - \eta \frac{\partial \mathcal{R}(\mathcal{D}, \theta^t)}{\partial \theta_i^t}, \quad (9)$$

where $\mathcal{R}(\mathcal{D}, \theta)$ denotes the loss function (i.e., negative log likelihood or KL divergence) and η is the corresponding learning rate. In the case of an RBM, the derivatives of the loss function are

$$\frac{\partial \mathcal{R}(\mathcal{D}, \theta)}{\partial \theta} = \mathbb{E}_{\mathbf{x} \sim \hat{p}(\mathbf{x})} \left[\frac{\partial E}{\partial \theta} \right] - \mathbb{E}_{\mathbf{v} \sim p(\mathbf{v})} \left[\frac{\partial E}{\partial \theta} \right], \quad (10)$$

$$\begin{aligned} \frac{\partial \mathcal{R}(\mathcal{D}, \theta)}{\partial w_{ij}} &= \langle \hat{p}(h_j = 1 | \mathbf{x}) x_i \rangle_{\mathbf{x} \sim \hat{p}(\mathbf{x})} \\ &\quad - \langle p(h_j = 1 | \mathbf{v}) v_i \rangle_{\mathbf{v} \sim p(\mathbf{v})}, \end{aligned} \quad (11)$$

$$\frac{\partial \mathcal{R}(\mathcal{D}, \theta)}{\partial b_j} = \langle x_i \rangle_{\mathbf{x} \sim \hat{p}(\mathbf{x})} - \langle v_i \rangle_{\mathbf{v} \sim p(\mathbf{v})}, \quad (12)$$

$$\frac{\partial \mathcal{R}(\mathcal{D}, \theta)}{\partial c_i} = \langle \hat{p}(h_j = 1 | \mathbf{x}) \rangle_{\mathbf{x} \sim \hat{p}(\mathbf{x})} - \langle p(h_j = 1 | \mathbf{v}) \rangle_{\mathbf{v} \sim p(\mathbf{v})}, \quad (13)$$

where $\langle \cdot \rangle$ indicates the ensemble average over different realizations of the quantities of interest. According to Eqs. (11) to (13), weights and biases are updated to minimize the differences between data and model distributions.

The bipartite network structure of RBMs allows us to update units within one layer in parallel [see Eqs. (4) and (5)]. We illustrate this so-called *block Gibbs sampling* procedure in Fig. 2. After a certain number of updates of the visible and hidden layers, the model distribution p_θ equilibrates. It has been shown empirically [6] that a small number of updates may be sufficient to train an RBM according to Eqs. (11) to (13). This technique is also known as *k-contrastive divergence* (CD- k), where k denotes the number of Gibbs sampling

steps. An alternative to CD- k methods is *persistent contrastive divergence* (PCD) [27]. Instead of using the environment data as initial condition for the sampling process (see Fig. 2), PCD uses the final state of the model from the previous sampling process to initialize the current one. The idea behind PCD is that the model distribution changes only slightly after updating the model parameters θ according to Eq. (9) when using a small learning rate. For further details on the RBM training process, see Sec. IV A.

B. Variational autoencoder

A variational autoencoder (VAE) is a graphical model for variational inference and consists of an encoder and a decoder network. It belongs to the class of latent variable models, which are based on the assumption that the unknown data distribution \hat{p} can be described by random vectors $\mathbf{z} \in \mathbb{R}^d$ and an appropriate prior $p(\mathbf{z})$ in a low-dimensional and unobserved (latent) space. For a given prior $p(\mathbf{z})$, we need a corresponding probabilistic model $p_\theta(\mathbf{x}, \mathbf{z})$ with parameter set θ to describe the data-generation process. To infer $p_\theta(\mathbf{x}, \mathbf{z})$, we need to maximize the marginal log likelihood $\log p_\theta(\mathbf{x})$ that we can rewrite as follows [28]:

$$\begin{aligned} \log p_\theta(\mathbf{x}) &= \log \int p_\theta(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \log \int p_\theta(\mathbf{x} | \mathbf{z}) p(\mathbf{z}) d\mathbf{z} \\ &= \log \int p_\theta(\mathbf{x} | \mathbf{z}) p(\mathbf{z}) \frac{q_\phi(\mathbf{z} | \mathbf{x})}{q_\phi(\mathbf{z} | \mathbf{x})} d\mathbf{z} \\ &\geq \mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x})} [\log p_\theta(\mathbf{x} | \mathbf{z})] - D_{\text{KL}}[q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})] \\ &:= \mathcal{E}(\mathbf{x}, \theta, \phi), \end{aligned} \quad (14)$$

where we applied Jensen's inequality in the fourth step. The integration over the latent variable \mathbf{z} in the second step is usually intractable, and we therefore introduced the approximate posterior distribution $q_\phi(\mathbf{z} | \mathbf{x})$ with parameter set ϕ and use the variational-inference principle to obtain a tractable bound of the marginal log likelihood, the so-called *evidence lower bound* (ELBO) [13]. We use $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbb{I})$ as the latent variable prior and $\mathcal{E}(\mathbf{x}, \theta, \phi)$. The ELBO is a tractable lower bound of the log likelihood and can thus be maximized to infer $p_\theta(\mathbf{x} | \mathbf{z})$. The term $\mathbb{E}[\log p_\theta(\mathbf{x} | \mathbf{z})]$ may be interpreted as a reconstruction error, because maximizing it makes the output of the decoder more similar to the input of the encoder and the term $D_{\text{KL}}[q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})]$ ensures that the latent representation is Gaussian such that data points with similar features have a similar Gaussian representation.

So far, we outlined the general idea behind latent variable models, but we still need to specify the approximate




FIG. 3. *Variational autoencoder.* A VAE is composed of two neural networks: the encoder (yellow) and decoder (green). In this example, each network consists of an input and an output layer (multiple layers are also possible). The dotted black arrow between the two networks represents the reparameterization trick [see Eq. (17)]. Each unit in the input layer of the decoder requires a corresponding mean μ and variance σ as input.

(variational) posterior $q_\phi(\mathbf{z}|\mathbf{x})$ and model likelihood $\log p_\theta(\mathbf{x}|\mathbf{z})$. A common choice for the approximate posterior is a factorized Gaussian distribution

$$q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2)) = \prod_{j=1}^d \mathcal{N}(z_j|\mu_j, \sigma_j^2), \quad (15)$$

where $\boldsymbol{\mu} = \boldsymbol{\mu}(\mathbf{x})$ and $\boldsymbol{\sigma} = \boldsymbol{\sigma}(\mathbf{x})$ denote mean and standard deviation vectors of the model. Given the binary nature of our data, we describe the model likelihood $p_\theta(\mathbf{x}|\mathbf{z})$ by a Bernoulli distribution with $\Pr(x_i = 1) = p_i$:

$$\begin{aligned} \log p_\theta(\mathbf{x}|\mathbf{z}) &= \log \prod_{i=1}^m p_i^{x_i} (1 - p_i)^{1-x_i} \\ &= \sum_{i=1}^m x_i \log p_i + (1 - x_i) \log(1 - p_i). \end{aligned} \quad (16)$$

Standard inference frameworks would determine the variational parameters per sample $\{\boldsymbol{\mu}_k, \boldsymbol{\sigma}_k, \mathbf{p}_k\}$ ($1 \leq k \leq N$). Instead, variational autoencoders utilize encoder and decoder networks to infer $q_\phi(\mathbf{z}|\mathbf{x})$ and $p_\theta(\mathbf{x}|\mathbf{z})$ (see Fig. 3). The encoder maps the input data to the variational parameters of the approximate posterior [i.e., $\mathbf{x}_k \rightarrow \{\boldsymbol{\mu}_k(\mathbf{x}), \boldsymbol{\sigma}_k(\mathbf{x})\}$], and the decoder maps a certain point in the latent parameter space to the corresponding likelihood [i.e., $\mathbf{z} \rightarrow \mathbf{p}(\mathbf{z})$]. This procedure allows us to amortize the cost of inference by learning only the parameters of the two neural networks.

To summarize, a VAE is based on an encoder network that outputs the parameters $(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\sigma}(\mathbf{x}))$ of the latent Gaussian representation $q_\phi(\mathbf{z}|\mathbf{x})$ of $p(\mathbf{z}|\mathbf{x})$. The second neural network in a VAE is a decoder that uses samples of the Gaussian $q_\phi(\mathbf{z}|\mathbf{x})$ as input to generate new samples according to the distribution $p_\theta(\mathbf{x}|\mathbf{z})$ (see Fig. 3). We estimate all parameters by maximizing the ELBO using backpropagation [13] and therefore need a deterministic and differentiable map between the output of the encoder and the input of the decoder w.r.t.

$\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$. To calculate the corresponding derivatives in the ELBO, we need to express the random variable \mathbf{z} as some differentiable and invertible transformation g of another auxiliary random variable $\boldsymbol{\epsilon}$ [i.e., $\mathbf{z} = g(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\epsilon})$]. We employ the so-called *reparameterization trick* [13] that uses $g(\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\epsilon}) = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}$ such that

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}, \quad (17)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ and \odot is the elementwise product.

To learn the distributions of physical features of the Ising model for different temperatures (i.e., classes), we use a conditional VAE (cVAE) [29] (see Sec. IV B for further details). The advantage of cVAEs over standard VAEs is that they allow to condition the encoder and decoder on the label of the data. We use l to denote the total number of classes and $c \in \{1, \dots, l\}$ is the label of a certain class. The data in a certain class is then distributed according to $p(\mathbf{z}|c)$ and the model distributions of the cVAE are given by $q_\phi(\mathbf{z}|\mathbf{x})$ and $p_\theta(\mathbf{x}|\mathbf{z}, c)$. These modified distributions are used in the ELBO of a cVAE. From a practical point of view, we have to provide extra dimensions in the input layers of the decoder of cVAEs to account for the label of a certain class.

C. Comparison of the two models

In Secs. II A and II B, we outlined the basic ideas behind RBMs and VAEs and discussed how they can be used as generative neural networks that are able to extract and learn features of a data set. Both models have in common that they create a low-dimensional representation of a given data set and approximate the underlying distribution minimizing the error in the reconstruction [see Eqs. (8) and (14)].

One major difference between the two models is that the compression and decompression is performed by the same network in the case of an RBM. For a VAE, however, two networks (encoder and decoder) are used to perform these tasks. A second important difference is that the latent distribution is a product of Bernoulli distributions for RBMs and a product of Gaussians for VAEs. Therefore, the representational power of RBMs is restricted to 2^m , where m is the number of hidden units. In the case of VAEs, it is \mathbb{R}^d where d is the dimension of the latent space. The greater representational power of a VAE may be an advantage, but it could also result in overfitting the data.

III. MONITORING LEARNING

To monitor the learning progress of RBMs and VAEs, we can keep track of the *reconstruction error* (i.e., the squared Euclidean distance between the original data and its reconstruction). Another monitoring approach uses the *binary cross-entropy* (BCE) [30]

$$R(\mathbf{x}) = - \sum_{i=1}^m x_i \log p_i(\mathbf{x}) + (1 - x_i) \log[1 - p_i(\mathbf{x})], \quad (18)$$

where $x_i \in \{0, 1\}$ and $p_i(\mathbf{x})$ is the reconstruction probability of bit i . In the following subsections, we describe additional methods that allow us to monitor approximations of log likelihood [see Eq. (6)] and KL divergence [see Eq. (8)]. Approximations are important since log likelihood and KL

divergence can be calculated exactly only for small system sizes. In Ref. [18], the partition function of RBMs has been computed directly since the considered system sizes contained only a few dozen spins.

A. Pseudo-log likelihood

One possibility to monitor the learning progress of an RBM is the so-called pseudo-log likelihood [31]. We consider a general approximation of an m -dimensional parametric distribution:

$$\begin{aligned} p(\mathbf{x}; \theta) &= \prod_{i=1}^m p\left(x_i \middle| \bigcap_{j=1}^{i-1} x_j; \theta\right) \\ &\approx \prod_{i=1}^m p(x_i | x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_m; \theta) \\ &= \prod_{i=1}^m p(x_i | \mathbf{x}_{-i}; \theta) := \mathcal{P}(\mathbf{x}; \theta), \end{aligned} \quad (19)$$

where \mathbf{x}_{-i} is the set of all variables apart from x_i . We now take the logarithm of $\mathcal{P}(\mathbf{x}; \theta)$ and obtain the *pseudo-log likelihood*

$$\log \mathcal{P}(\mathbf{x}; \theta) = \sum_{i=1}^m \log p(x_i | \mathbf{x}_{-i}). \quad (20)$$

Based on Eq. (20), we conclude that the pseudo-log likelihood is the sum of the log probabilities of each x_i conditioned on all other states apart from x_i . The summation over all states x_i can be computationally expensive, so we approximate $\log \mathcal{P}(\mathbf{x}; \theta)$ by the log-likelihood estimate [32]

$$\log \tilde{\mathcal{P}}(\mathbf{x}; \theta) = m \log p(x_i | \mathbf{x}_{-i}), \quad i \sim \mathcal{U}(0, m), \quad (21)$$

where we sample uniformly to efficiently approximate the pseudo-log likelihood. For RBMs, the approximated pseudo-log likelihood is

$$\log \tilde{\mathcal{P}}(\mathbf{x}; \theta) = m \log \left(\frac{e^{-\mathcal{F}_c(\mathbf{x})}}{e^{-\mathcal{F}_c(\tilde{\mathbf{x}})} + e^{-\mathcal{F}_c(\mathbf{x})}} \right), \quad (22)$$

where $\tilde{\mathbf{x}}$ corresponds to the vector \mathbf{x} but with a flipped i th component. Equation (22) follows from the identity

$$\begin{aligned} p(x_i | \mathbf{x}_{-i}) &= \frac{p(x_i, \mathbf{x}_{-i})}{p(\mathbf{x}_{-i})} = \frac{p(x_i, \mathbf{x}_{-i})}{\sum_{x_i \in \{0,1\}} p(x_i, \mathbf{x}_{-i})} \\ &= \frac{p(\mathbf{x})}{p(\tilde{\mathbf{x}}) + p(\mathbf{x})} = \frac{e^{-\mathcal{F}_c(\mathbf{x})}}{e^{-\mathcal{F}_c(\tilde{\mathbf{x}})} + e^{-\mathcal{F}_c(\mathbf{x})}}. \end{aligned} \quad (23)$$

B. Estimating KL divergence

In most unsupervised learning problems, except some synthetic problems such as Gaussian mixture models [33], we do not know the distribution underlying a certain data set. We are given only samples and can therefore not directly determine the KL divergence [see Equation 8]. Instead, we have to use some alternative methods. According to Ref. [34], we use the nearest-neighbor (NN) estimation of the KL divergence. Given two continuous distributions p and q and N and \tilde{N} corresponding i.i.d. m -dimensional samples $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and

$\{\mathbf{y}_1, \dots, \mathbf{y}_{\tilde{N}}\}$, we want to estimate $D_{\text{KL}}(p||q)$. For two consistent estimators [35] p' and q' , the KL divergence can be approximated by

$$\begin{aligned} D_{\text{KL}}(p||q) &= \sum_{\mathbf{x} \sim p(\mathbf{x})} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} \\ &\approx \frac{1}{N} \sum_{i=1}^N \log \frac{p'(\mathbf{x}_i)}{q'(\mathbf{x}_i)}. \end{aligned} \quad (24)$$

We construct the consistent estimators as follows. Let $\rho_k(\mathbf{x}_i)$ be the Jaccard distance [36] between \mathbf{x}_i and its k -NN in the subset $\{\mathbf{x}_j\}$ for $j \neq i$ and $\nu_k(\mathbf{x}_i)$ the Jaccard distance between \mathbf{x}_i and its k -NN in the set $\{\mathbf{y}_l\}$. We obtain the KL divergence estimate

$$\hat{D}(p||q) = \frac{m}{N} \sum_{i=1}^N \log \frac{\nu_k(\mathbf{x}_i)}{\rho_k(\mathbf{x}_i)} + \log \frac{\tilde{N}}{N-1}. \quad (25)$$

In the context of generative models this means that we can calculate the KL divergence in terms of the Jaccard distances $\rho_k(\mathbf{x}_i)$ and $\nu_k(\mathbf{x}_i)$.

IV. LEARNING THE ISING MODEL

We now use the generative models of Sec. II to learn the distribution of Ising configurations at different temperatures. This enables us to examine the ability of RBMs and cVAEs to capture the temperature dependence of physical features such as energy, magnetization, and spin-spin correlations. We first focus on the training of RBMs in Sec. IV A and in Sec. IV B we describe the training of cVAEs. However, before focusing on the training, we give a brief overview of some relevant properties of the Ising model.

The Ising model is a mathematical model of magnetism and describes the interactions of spins $\sigma_i \in \{-1, 1\}$ ($1 \leq i \leq m$) according to the Hamiltonian [37]

$$\mathcal{H}(\{\sigma\}) = -J \sum_{\langle i, j \rangle} \sigma_i \sigma_j - H \sum_i \sigma_i, \quad (26)$$

where $\langle i, j \rangle$ denotes the summation over nearest neighbors and H is an external magnetic field. We consider the ferromagnetic case with $J = 1$ on a two-dimensional lattice with 32×32 sites. In an infinite two-dimensional lattice, the Ising model exhibits a second-order phase transition at $T = T_c = 2/\log(1 + \sqrt{2}) \approx 2.269$ [37,38]. To train RBMs and cVAEs, we generate 20×10^4 realizations of spin configurations at temperatures $T \in \{1.5, 2, 2.5, 2.75, 3, 4\}$ using the M(RT)² algorithm [37,39]. In previous studies, RBMs were applied to smaller systems with only about 100 spins to learn the distribution of spin configurations with one RBM per temperature (see Refs. [15,18] for further details).

We also consider the training of one RBM per temperature and compare the results with those obtained by a single RBM and cVAE that were trained for all temperatures. If RBMs and cVAEs are able to learn the distribution of Ising configurations $\{\sigma\}$ at a certain temperature, the model-generated samples should have the same (spontaneous) magnetization, energy, and spin-spin correlations as the M(RT)² samples. We compute the spontaneous magnetization

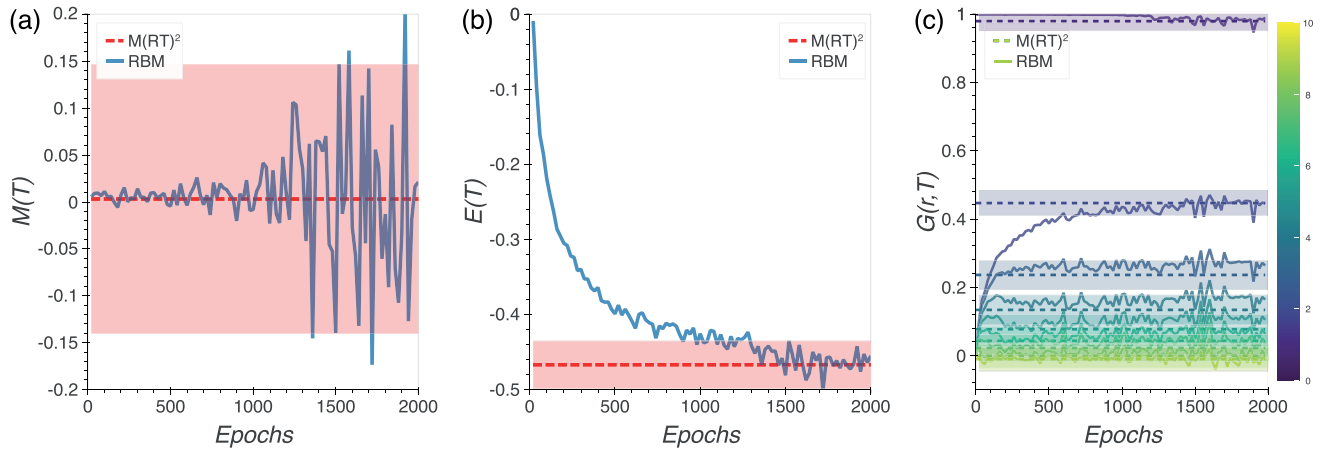


FIG. 4. Evolution of physical features during training. We show the evolution of the magnetization $M(T)$ (a), energy $E(T)$ (b), and correlation function $G(r, T)$ (c) for the training of a single RBM at $T = 2.75$. Dashed lines and shaded regions indicate the mean and standard deviation of the corresponding $M(RT)^2$ samples and different colors in the plot of $G(r, T)$ represent different radii r .

according to

$$M_S(T) = \lim_{H \rightarrow 0^+} \left\langle \frac{1}{m} \sum_{i=1}^m \sigma_i \right\rangle. \quad (27)$$

Furthermore, we determine the spin-spin correlations between two spins σ_i and σ_j at positions r_i and r_j :

$$G(r_i, r_j) = \langle (\sigma_i - \langle \sigma_i \rangle)(\sigma_j - \langle \sigma_j \rangle) \rangle = \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle. \quad (28)$$

We use translational-invariance and rewrite Eq. (28) as

$$G(r_i, r_j, T) = G(r, T) = \langle s_i s_{i+r} \rangle - M_S^2(T), \quad (29)$$

where $r = |r_i - r_j|$ is the Euclidean distance between spins σ_i and σ_j . The correlation function $G(r, T)$ quantifies the influence of one spin on another at distance r .

A. Training RBMs

We consider all 20×10^4 samples for $T \in \{1.50, 2.0, 2.5, 2.75, 3.0, 4.0\}$ and first train one RBM per temperature using the ADAM optimizer [40] with a five-step PCD approach (see Sec. II A). For all machines, we use the same network structure and training parameters and set the number of hidden units to $n = 900$. The number of visible units is equal to the number of spins (i.e., $m = 32^2 = 1024$). We initialize each Markov chain with uniformly at random distributed binary states $\{0, 1\}$, train over 200 epochs, and repeat this procedure 10 times. We set the learning rate to $\eta = 10^{-4}$ and consider minibatches of size 128. In addition, we encourage sparsity of the weight matrix W by applying L1 regularization with a regularization coefficient of $\lambda = 10^{-4}$. That is, we consider an additional contribution $-\lambda \|W\|_1$ in the log likelihood and use

$$\log \mathcal{L}'(\mathbf{x}, \theta, \lambda) = \log \mathcal{L}(\mathbf{x}, \theta) - \lambda \|W\|_1 \quad (30)$$

instead of Eq. (6). Thus, minimizing $-\log \mathcal{L}'(\mathbf{x}, \theta, \lambda)$ means that we are also minimizing the 1-norm $\|W\|_1$ (i.e., the maximum absolute column sum) of the weight matrix W . We initialize all weights in W according to a Gaussian distribution

with the *glorot* initializer [41] and biases \mathbf{b} and \mathbf{c} with uniformly distributed numbers in the interval $[0, 0.1]$. In Fig. 4 we show the evolution of magnetization, energy, and correlations during training of an RBM for Ising samples at $T = 2.75$. We observe that it takes about 1500 epochs of training to properly capture magnetization fluctuations, energy, and spin correlations.

In addition to using one RBM per temperature, we also train one RBM for all temperatures. We use the same samples and training parameters as before. However, we add additional visible units to encode the temperature information of each training sample. For the training of this RBM, we start from a random initial configuration, train over 400 epochs, and repeat this procedure 10 times.

To monitor the training progress, we tested the considered models every 20 epochs on 10% of the data by monitoring reconstruction error, binary cross-entropy, pseudo-log likelihood, and KL divergence and its inverse (see Sec. III). We compute reconstruction error and binary cross-entropy based on a single Gibbs sampling step. For the KL divergence, we generated 12×10^3 samples by performing 10^3 sampling steps and keeping the last 10^2 samples. We repeated this procedure 20 times for each temperature and averaged over minibatches of 256 samples to have better estimators.

B. Training cVAEs

We train one cVAE for all temperatures and consider an encoder that consists of three convolutional layers with 32, 64, and 64 kernels (or filters) of size 3 [42] and a *rectified linear unit* (ReLU) as activation function. Each convolutional layer is followed by a *maxpool* layer to downsample the “pictures.” In a *maxpool* layer, only the maximum values of subregions of an initial representation are considered. To regularize the cVAE, we use *batchNormalization* [43] and a dropout rate of 0.5 [44]. The last layer is a fully connected flat layer with 400 units (200 mean and 200 variance) to represent the 200 dimensional Gaussian latent variable \mathbf{z} .

For the decoder, we use an input layer that consists of 200 units to represent the latent variable \mathbf{z} and concatenate

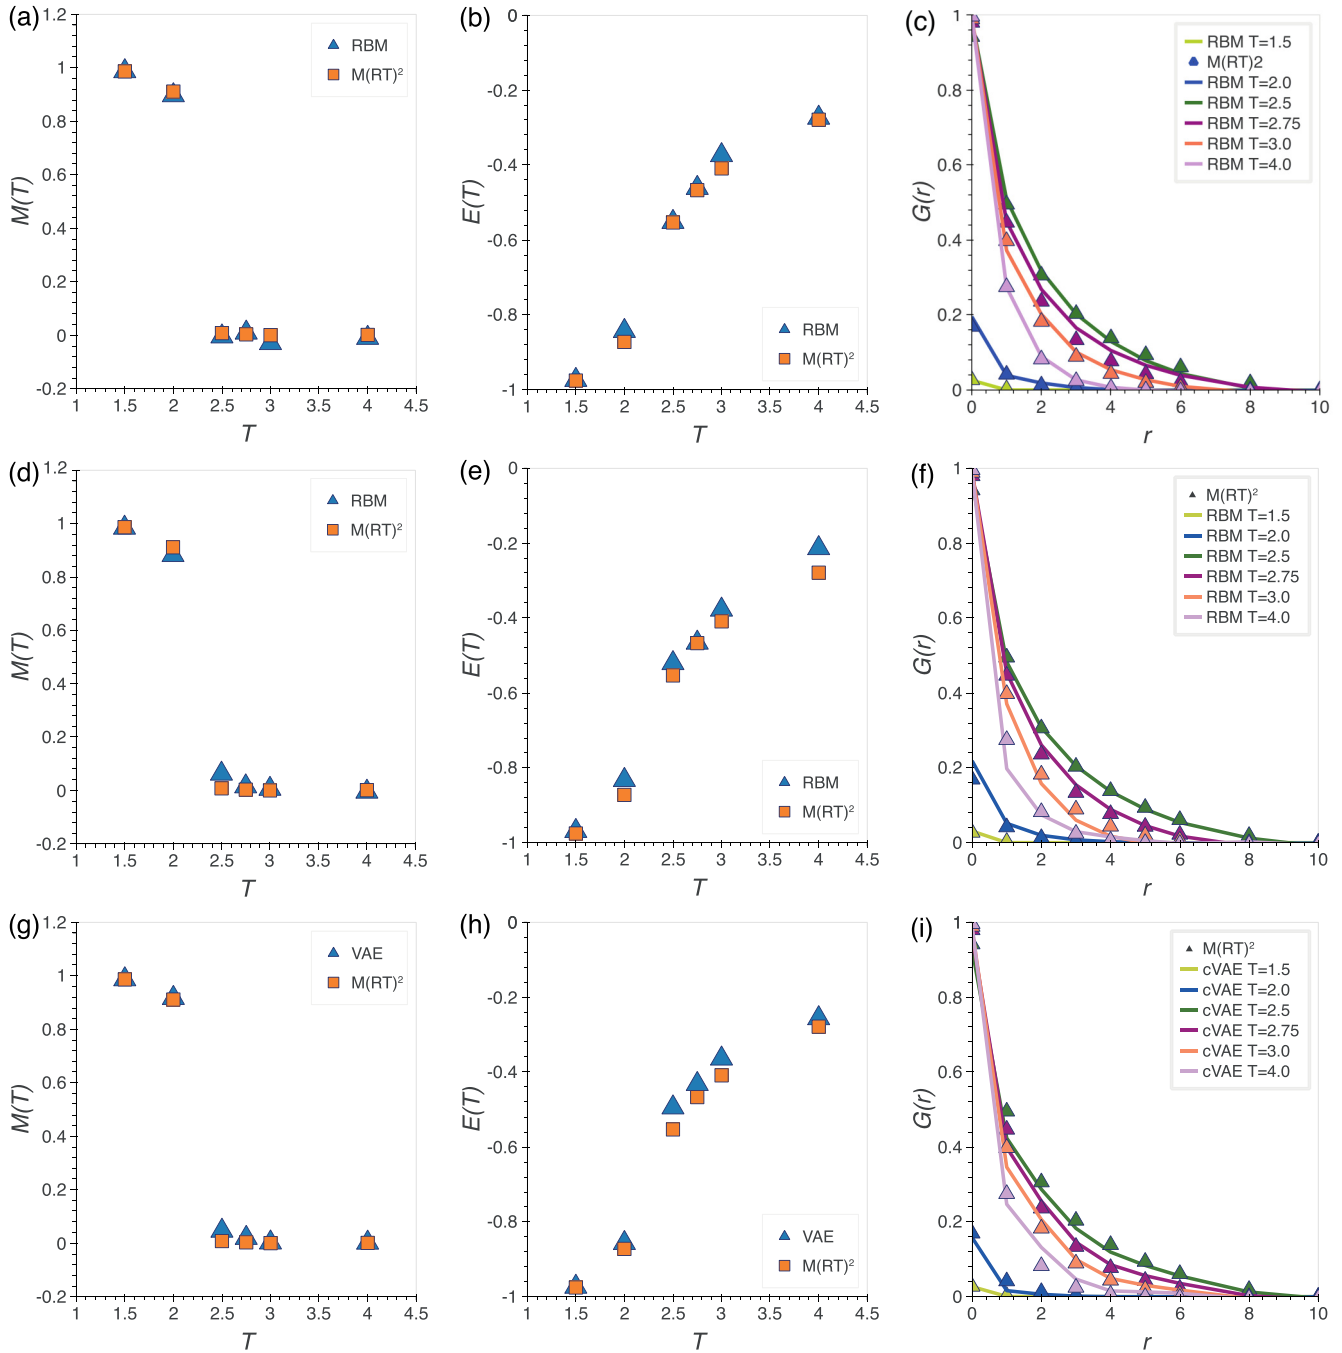


FIG. 5. Physical features of RBM and convolutional VAE Ising samples. We use RBMs and convolutional cVAEs to generate 20×10^4 samples of Ising configurations with 32×32 spins for temperatures $T \in \{1.5, 2, 2.5, 2.75, 3, 4\}$. We show the magnetization $M(T)$, energy $E(T)$, and correlation function $G(r, T)$ for neural network and corresponding $M(RT)^2$ samples. In panels (a)–(c), we separately trained one RBM per temperature, whereas we used a single RBM for all temperatures in panels (d)–(f). In panels (g)–(i), we show the behavior of $M(T)$, $E(T)$, and $G(r, T)$ for samples that are generated with a convolutional cVAE that was trained for all temperatures. Error bars are smaller than the markers.

it with the additional temperature labels. We upsample their dimension with a fully connected layer of 2048 units that are reshaped to be fed into three deconvolutional layers. The number of filters in these deconvolutional layers are 64, 64, and 32 with filter size 3 and each deconvolutional layer is followed by an upsampling layer to upsample the dimension of the “pictures.” The last layer of the encoder is a deconvolutional

layer with one single filter to have a single sample as output. We train the complete architecture over 10^3 epochs using the ADAM optimizer [40] and a learning rate of $\eta = 10^{-4}$. During training, we keep track of the ELBO [see Eq. (14)] and the KL divergence and its inverse (see Sec. III) for 10% of the data.

We also consider nonconvolutional VAE architectures in Appendix A. However, in the absence of convolutional layers,

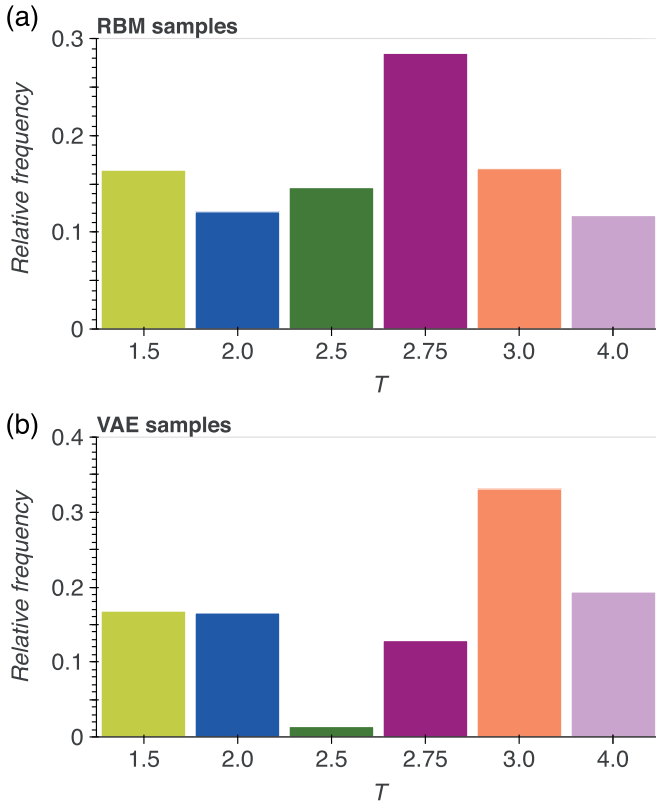


FIG. 6. Relative frequencies of samples at different temperatures. We show the relative frequencies of samples that are obtained by starting trained RBMs (a) and convolutional cVAEs (b) from random initial configurations. The data are based on 20×10^4 samples for each temperature.

we found that some physical features are not well captured anymore (see Appendix A for further details).

C. Training classifier

In the case of single RBMs and cVAEs that we train for all temperatures, we have to use a classifier to determine the temperature of the generated samples. We consider a classifier network that consists of two convolutional layers with, respectively, 32 and 64 filters and four fully connected layers of [256,128,64,32] units with a *ReLU* activation function and stride 2 convolution. Similar to Eq. (30), we regularize all layers using a L2 regularizer with a regularization coefficient of 0.01. The last layer is a softmax fully connected layer that outputs the probability of a sample to belong to each temperature class, and we determine the sample temperature by identifying the highest probability class. We trained the model for 100 epochs with the ADAM optimizer [40] and a learning rate of 10^{-4} . We monitor both the test loss (i.e., categorical cross-entropy) and accuracy. With this model, we achieve a test accuracy of 89% for the considered six temperature classes.

D. Generating Ising samples

After training both generative neural networks with 20×10^4 realizations of Ising configurations at different

temperatures [45], we can now study the ability of RBMs and cVAEs to capture physical features of the Ising model. To generate samples from an RBM, we start 200 Markov chains from random initial configurations whose mean corresponds to the mean of spins in the data at the desired temperature. This specific initialization together with the encoding of corresponding temperature labels helps sampling from single RBMs that were trained for all temperatures. We perform 10^3 sampling steps, keep the last 10^2 samples, and repeat this procedure 2×10^2 times for each temperature. The total number of samples is therefore 20×10^4 . For the trained cVAE, we sample from the decoder using 20×10^3 200-dimensional vectors with normally distributed components for each temperature. We concatenate these vectors with the respective temperature labels and use them as input of the decoder. For the single RBM and cVAE that we trained for all temperatures, we use the classifier (see Sec. IV C) to identify the temperature of the newly generated patterns.

We show the distribution of RBM and cVAE samples over different temperature classes in Fig. 6. For both neural networks, the relative frequency of samples with temperature $T \geq 2.5$ is larger than for lower temperatures. In addition, the results of Fig. 6 suggest that the considered convolutional cVAE has difficulties to generate samples for $T \approx T_c \approx 2.269$. Nonconvolutional VAEs suffer from the same problems and also have difficulties to properly capture physical features (see Appendix A). Interestingly, we do not observe this issue for the RBM samples in Fig. 6.

In general, patterns at criticality are difficult to learn and generate since they exhibit long-range correlation effects. One possibility to enforce the generation of VAE samples at T_c is to determine the latent representation of samples at the critical temperature and use slightly perturbed versions of these samples as input in the decoder.

Next, we determine magnetization, energy, and correlation functions of the generated samples at different temperatures (see Fig. 5). In the top panel of Fig. 5, we observe that RBMs that were trained for single temperatures are able to capture the temperature dependence of magnetization, energy, and spin correlations. In the case of single RBMs and convolutional cVAEs that were trained for all temperatures, we also find that the aforementioned physical features are captured well. However, our results indicate that nonconvolutional VAEs have difficulties in capturing magnetization, energy, and spin correlations (see Appendix A for further details).

Convolutional layers have been proven effective in image classification [46–48] and generation [49]. The advantage of convolutional architectures is their ability to model coarse-grained (“global”) and fine-grained (“local”) features of a given data set. The application of convolutional layers allows us to account for the two-dimensional structure of the input data and corresponding (local) interactions of spatially neighboring spins by dividing the original 32×32 lattice into smaller overlapping regions. These regions provide the basis for a partitioned feature extraction/reconstruction that preserves local spatial properties of the data [50]. The max-pooling operations that we use in our convolutional VAE compress the input data and preserve their two-dimensional structure. For the outlined reasons, convolutional VAEs have a clear advantage over nonconvolutional VAEs in terms of

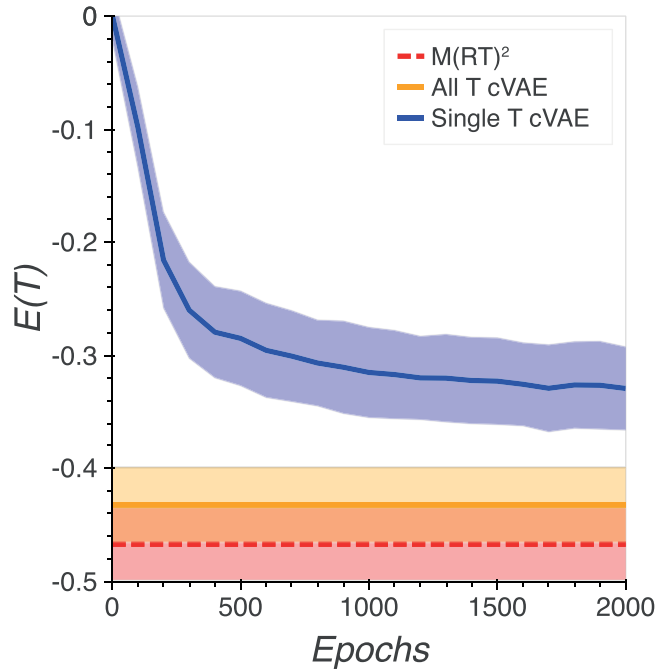


FIG. 7. Evolution of energy during training of a convolutional VAE. We show the evolution of the energy $E(T)$ for the training of a single convolutional VAE at $T = 2.75$ (blue line). The dashed red line indicates the mean energy of the corresponding $M(RT)^2$ samples. The yellow line shows the mean energy of the convolutional cVAE trained on all temperatures for 1000 epochs.

local feature extraction. This is particularly evident when comparing their ability to capture spin-spin correlations and energy [see Fig. 5 (bottom) and Appendix A].

Interestingly, the RBMs we consider are able to capture correlation effects without additional convolutional layers. In general, such performance differences originate from several factors including: (1) the intrinsic properties of the chosen model (e.g., network structure, data representation, etc.) and (2) the underlying optimization algorithm. Thus, even if convolutional layers provide an efficient way to capture local information, this does not imply that nonconvolutional models cannot achieve the same results. In our case, both models rely on a maximum-likelihood parameter estimation. However, in the case of VAEs, we can maximize only the ELBO [see Eq. (14)] without any guarantee that this function shares the




FIG. 8. Snapshots of Ising configurations. We show snapshots of Ising configurations for $T \in \{1.5, 2.5, 4\}$. The configurations in the top, middle, and bottom panels are based on $M(RT)^2$, RBM, and convolutional cVAE samples, respectively.

same local maxima with the actual likelihood. For RBMs, our results suggest that contrastive-divergence learning can find solutions that preserve spatial information without using convolutional layers.

In our analysis, we also noticed another important difference between the learning dynamics of RBMs and VAEs. Restricted Boltzmann machines produce similar results independent of whether they were trained on single or multiple temperatures (see Fig. 5) whereas cVAEs exhibit a better performance if trained on all temperatures (see Fig. 7). Patterns at low and high temperatures are easier to learn and seem to serve as reference points for the encoder during the training of VAEs on all temperatures. This type of training may be useful in other VAE-learning applications.

For a more visual comparison, we show snapshots of Ising configurations in the original data and compare them to RBM and VAE-generated spin configurations in Fig. 8. We also visualize the differences between original data and model-generated Ising samples in terms of two-dimensional

TABLE I. Wasserstein distance between neural-network and training data magnetization and energy distributions. Wasserstein distance \mathcal{W} [51] between the magnetization and energy distributions of $M(RT)^2$ samples and corresponding RBM and cVAE samples. Smaller values of \mathcal{W} indicate a higher similarity of distributions.

Temperature	$\mathcal{W}(M_{\text{RBM}} M_{\text{M(RT)}^2})$	$\mathcal{W}(M_{\text{cVAE}} M_{\text{M(RT)}^2})$	$\mathcal{W}(E_{\text{RBM}} E_{\text{M(RT)}^2})$	$\mathcal{W}(E_{\text{cVAE}} E_{\text{M(RT)}^2})$
$T = 1.50$	0.002	0.0003	0.006	0.002
$T = 2.00$	0.014	0.0054	0.06	0.02
$T = 2.50$	0.099	0.0471	0.02	0.06
$T = 2.75$	0.042	0.0402	0.01	0.04
$T = 3.00$	0.024	0.0225	0.04	0.04
$T = 4.00$	0.021	0.0243	0.13	0.02
Mean	0.034	0.023	0.045	0.029

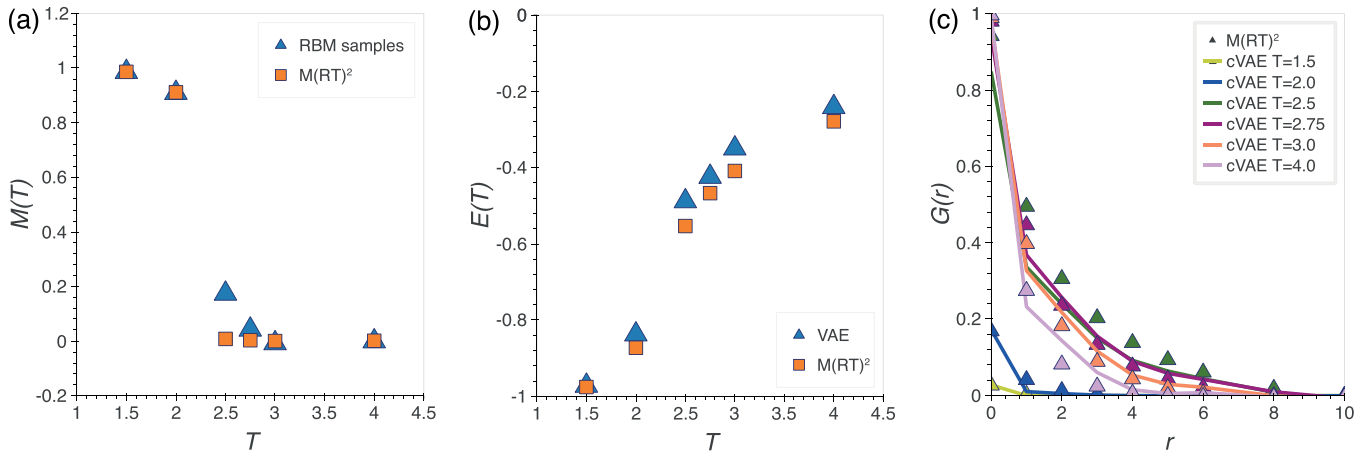


FIG. 9. *Physical features of nonconvolutional VAE Ising samples.* We use nonconvolutional cVAEs to generate 20×10^4 samples of Ising configurations with 32×32 spins for temperatures $T \in \{1.5, 2, 2.5, 2.75, 3, 4\}$. We show the magnetization $M(T)$ (a), energy $E(T)$ (b), and correlation function $G(r, T)$ (c) for neural-network and corresponding $M(RT)^2$ samples. The nonconvolutional cVAE was trained for all temperatures. Error bars are smaller than the markers.

energy-magnetization scatter plots and magnetization distributions in Appendix B. We observe that the magnetization distributions of the training data are qualitatively well captured by the corresponding neural-network-generated magnetization distributions. To make a quantitative comparison between model and training data distributions, we determine their corresponding Wasserstein distances \mathcal{W} [51]. Smaller values of \mathcal{W} indicate a higher similarity of two distributions. Based on the results that we show in Table I, we conclude that, in terms of a smaller Wasserstein distance, the considered cVAEs are able to capture the magnetization (and energy) distributions of the training data better than the considered RBMs. We, however, note that the mean values may be better captured by RBMs than cVAEs for certain temperatures (see Fig. 5).

As a further extension of our framework, we also study the ability of RBMs and cVAEs to capture magnetization, energy, and spin correlations of Ising systems with heterogeneous coupling distributions (see Appendix C). Using the example of an Ising model with a Gaussian coupling distribution, we

find that both neural networks are also able to qualitatively reproduce the temperature dependence of the mentioned thermodynamic quantities. Our results indicate that RBMs perform slightly better than cVAEs in this task for the considered network architecture and training parameters. We note that alternative approaches such as teacher-student frameworks also can be used to generate spin configurations after inferring spin couplings from a given training data set (the inverse Ising problem) [2,52]. Teacher-student frameworks are useful tools if the model family is known and, in the case of the Ising model, system sizes are small enough if spin couplings are heterogeneous. In contrast, the RBMs and cVAEs that we consider in this paper are applicable to homogeneous and heterogeneous Ising systems.

V. CONCLUSIONS AND DISCUSSION

In comparison with other generative neural networks, the training of RBMs is challenging and computationally expensive [53]. Moreover, conventional RBMs can handle only




FIG. 10. *Two-dimensional visualization of Ising samples.* We show the distribution of 20×10^4 Ising samples for temperatures $T \in \{1.5, 2, 2.5, 2.75, 3, 4\}$. The data in panels (a)–(c) are based on $M(RT)^2$, RBM, and convolutional cVAE samples, respectively.

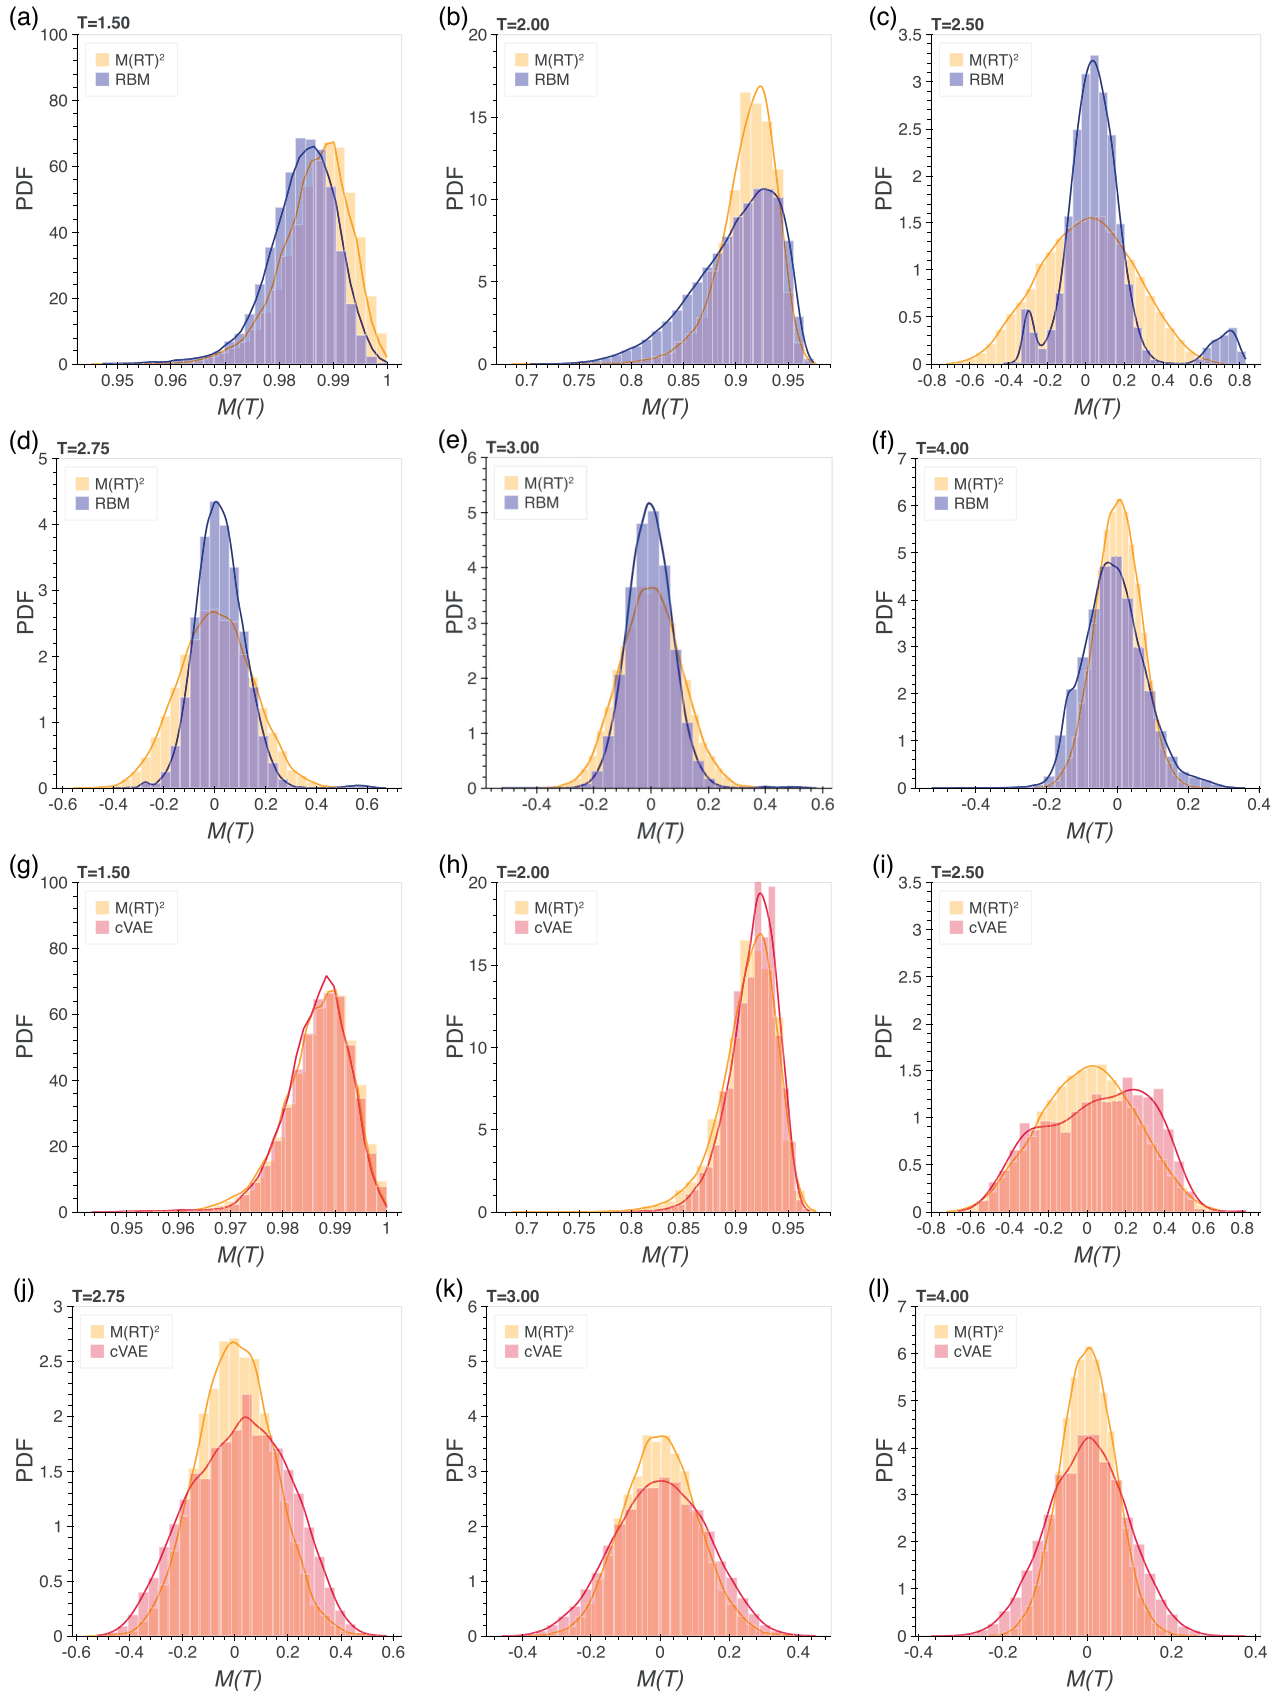


FIG. 11. *Distribution of magnetization at different temperatures.* We show the distribution of 20×10^4 Ising samples for temperatures $T \in \{1.5, 2.0, 2.5, 2.75, 3, 4\}$. The data are based on $M(RT)^2$ (yellow), RBM (blue), and convolutional cVAE (red) samples, respectively.

binary data and generalized architectures are necessary to describe continuous input variables [22]. For these reasons and the good performance that convolutional models have demonstrated in image classification [46–48] and generation [49], RBM-based representation learning is being more and more replaced by models like VAEs and generative adversarial networks (GANs) [54–56]. However, performance assessments of RBMs and VAEs are mainly based on a few common test data sets (e.g., MNIST) that are frequently used in the deep learning community [11]. This approach can provide only limited insights into the representational properties of neural network models.

To better understand the representational characteristics of RBMs and VAEs, we studied their ability to learn the distribution of physical features of Ising configurations at different temperatures. This approach allowed us to quantitatively compare the distributions learned by RBMs and VAEs in terms of features of an exactly solvable physical system. We found that physical features are well captured by the considered RBM and VAE models. However, samples generated by the employed VAE are less evenly distributed across temperatures than is the case for samples that we generated with an RBM. In particular, close to the critical point of the Ising model, the considered VAEs has difficulties to generate corresponding samples. In addition, our results also showed that convolutional layers improve the performance of VAEs in terms of their ability to capture magnetization, energy, and spin correlations (see Appendix A for more information on nonconvolutional VAEs). Such layers are useful to model and preserve local information [57]. Interestingly, RBMs achieve a similar or even better performance without additional convolutional layers. One possible reason for the observed behavior is that the training of VAEs relies on maximizing the ELBO [see Eq. (14)], which may not share the same local maxima with the actual log likelihood. Another distinctive feature between the considered models is that RBMs show a similar performance independent of whether they were trained on one or all temperatures, whereas the performance of VAEs improves if trained on all temperatures. Due to long-range correlation effects in the vicinity of the critical temperature, features of Ising configurations at criticality are more difficult to capture than at low and high temperatures. In the case of VAEs that were trained on all temperatures, low- and high-temperature samples seem to provide reference points that facilitate feature extraction at criticality. To summarize, in the context of physical feature extraction, our results suggest that RBMs are still able to compete with the more recently developed VAEs.

Possible directions for future research include the study of representational characteristics of RBMs, VAEs, and other generative neural networks by applying them to alternative physical systems (e.g., disordered magnetic systems [58]) [59]. Furthermore, it might be useful to explore connections between the learning capabilities of generative neural networks and corresponding phase-space approximation methods for Ising-like systems such as cluster variational approximation methods (CVAMs) [60–62].

Future studies may also focus on a comparison between the sampling performance of neural networks and conventional

methods. It could be of interest to compare the training and sampling performance between maximum-entropy models and generative neural networks [2]. In the case of one RBM (or cVAE) that was trained for all temperatures, the sampling always involves a classification, and it might be the case that one has to discard many samples until one has generated a sample with the desired temperature. Even if the physical features are well captured, this kind of sampling is not as efficient as $M(RT)^2$, cluster algorithms, or histogram methods. If one uses single machines for each temperature, the sampling process is more efficient. For cVAEs, one just has to generate different vectors with normally distributed components, and for RBMs the sampling is based on parallel updates of the visible units.

All codes and further detailed descriptions are available on GitHub [63].

ACKNOWLEDGMENTS

We thank Gonzalo Ruz and Lenka Zdeborova for helpful discussions and comments. Numerical calculations were performed on the ETH Euler cluster. L.B. acknowledges financial support from the SNF Early Postdoc Mobility fellowship on “Multispecies interacting stochastic systems in biology” and the Army Research Office (W911NF-18-1-0345).

APPENDIX A: NONCONVOLUTIONAL VARIATIONAL AUTOENCODER

In addition to the convolutional architecture of cVAEs that we described in Sec. IV B, we also considered their nonconvolutional counterparts. We used an encoder that is composed




FIG. 12. *Gaussian coupling distribution.* Gaussian distribution of Ising spin couplings [see Eq. (26)]. The mean and standard deviation of the Gaussian distribution are both equal to one.

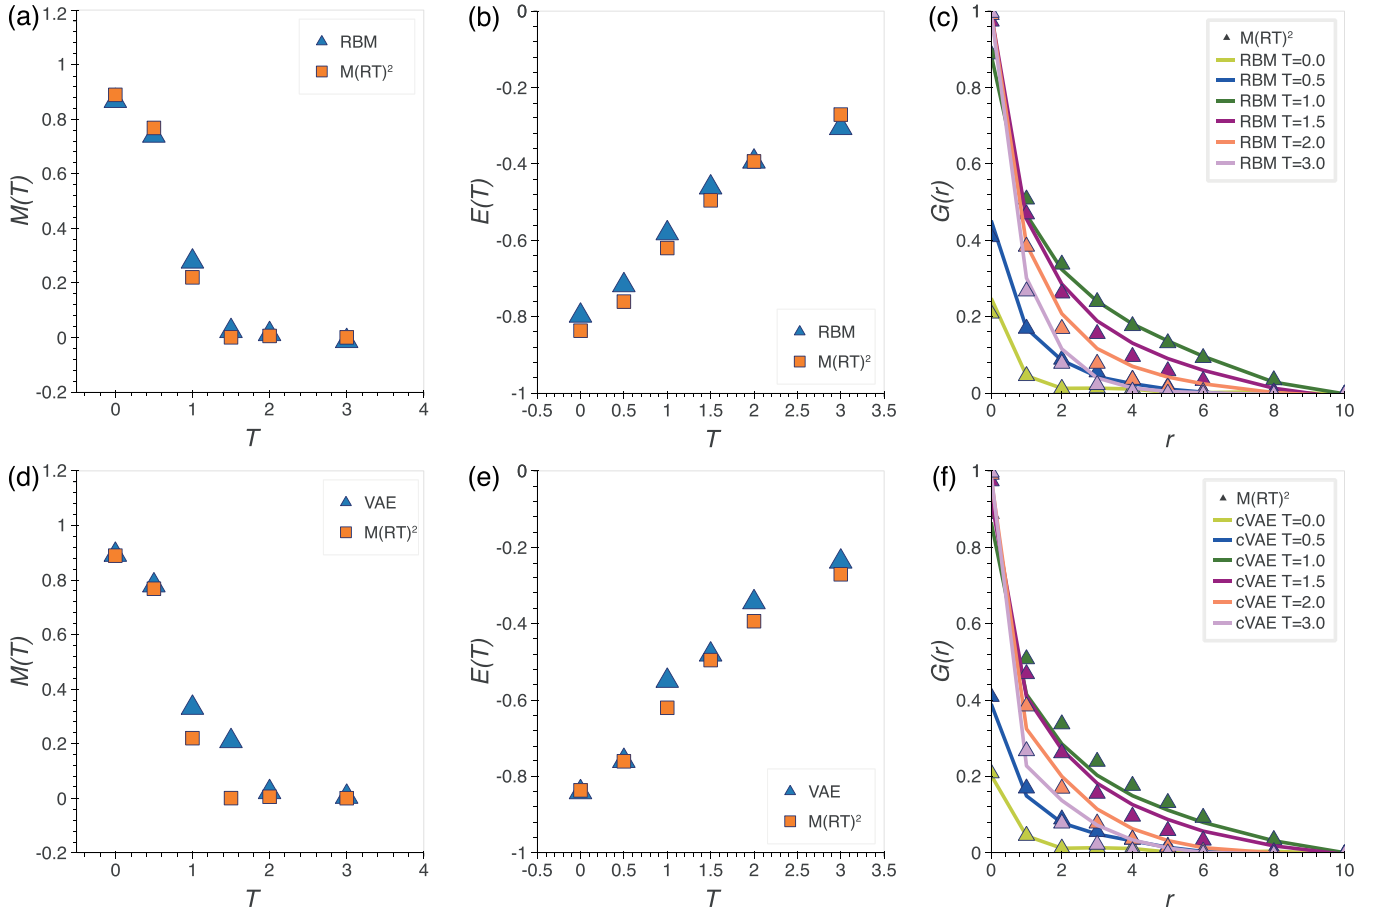


FIG. 13. *Physical features of RBM and convolutional VAE Ising samples for heterogeneous couplings.* We use RBMs and convolutional cVAEs to generate 20×10^4 samples of Ising configurations with 32×32 spins for temperatures $T \in \{10^{-6}, 0.5, 1, 1.5, 2, 3\}$ and a Gaussian coupling distribution whose mean and standard deviation are equal to one. We show the magnetization $M(T)$, energy $E(T)$, and correlation function $G(r, T)$ for neural network and corresponding $M(RT)^2$ samples. In panels (a)–(c), we used a single RBM for all temperatures; in panels (d)–(f), we show the behavior of $M(T)$, $E(T)$, and $G(r, T)$ for samples that are generated with a convolutional cVAE that was trained for all temperatures. Error bars are smaller than the markers.

of an input layer with 1024 units, one fully connected hidden layers with 700 units, and a sigmoid activation function. As for the convolutional VAEs, we also used dropout to avoid overfitting. The output layer has 400 units and represents a 200-dimensional Gaussian latent variable \mathbf{z} . We use a decoder with a symmetrical architecture composed of an input layer that consists of 200 units to represent the latent variable \mathbf{z} . There are six additional units that we use to encode the temperature of Ising samples. The remaining layers in the decoder are one fully connected hidden layer with 700 units and a sigmoid activation function. The output layer is sigmoidal and consists of 1024 units.

We trained the nonconvolutional cVAE in the same way as the convolutional cVAE (see Sec. IV B). We then generated samples and determined their temperatures with the classifier that we describe in Sec. IV C. In Fig. 9 we show the resulting temperature dependence of magnetization, energy, and spin correlations. We observe that energy and correlations are not captured as well as for convolutional architectures and RBMs (see Fig. 5).

APPENDIX B: ENERGY AND MAGNETIZATION SCATTER PLOTS

We generated 20×10^4 samples with the RBMs and convolutional cVAEs that we trained for all temperatures $T \in \{1.5, 2, 2.5, 2.75, 3, 4\}$. After determining the temperatures of the generated samples (see Sec. IV C), we compute different physical features including energy $E(T)$ and magnetization $M(T)$. We show the two-dimensional scatter plots of $E(T)$ and $M(T)$ in Fig. 10. The left panel of Fig. 10 shows the distribution of $E(T)$ and $M(T)$ in the original data. We find that the RBMs we consider have problems to reproduce the behavior at high temperatures (e.g., for $T = 4$, see middle panel of Fig. 10). As also apparent in Fig. 6, the convolutional cVAE has difficulties to generate samples at temperatures close to T_c (e.g., for $T = 2.5$). These observations are also reflected in the magnetization distributions that we show in Fig. 11. As outlined in the main text and Table I, the considered cVAEs are able to capture the magnetization distributions better than the considered RBMs in terms of smaller Wasserstein distances \mathcal{W} [51]. Still, the magnetization mean values of the

considered RBMs are better aligned with the corresponding training data mean values for certain temperatures than is the case for the considered cVAEs (see Fig. 5).

APPENDIX C: HETEROGENEOUS COUPLINGS

We also apply RBMs and cVAEs to an Ising model with spin couplings that are distributed according to a Gaussian distribution whose mean and standard deviation are both equal to one (see Fig. 12). We use the RBM and cVAE

architectures of Sec. IV and generated 20×10^4 $M(\text{RT})^2$ samples for temperatures $T \in \{10^{-6}, 0.5, 1, 1.5, 2, 3\}$ to train the two neural networks. Training protocols are also as outlined in Sec. IV. Only in the case of the RBM, we train over 90 epochs instead of 200 epochs, as in Sec. IV, and repeat this procedure 10 times. After training both neural networks, we generated 20×10^4 Ising configurations. We show the corresponding results in Fig. 13 and observe that both neural networks are able to qualitatively capture the temperature dependence of magnetization, energy, and spin correlations.

-
- [1] Y. LeCun, Y. Bengio, and G. Hinton, *Nature (London)* **521**, 436 (2015).
- [2] H. C. Nguyen, R. Zecchina, and J. Berg, *Adv. Phys.* **66**, 197 (2017).
- [3] K. Davidsen, B. J. Olson, W. S. DeWitt, III, J. Feng, E. Harkins, P. Bradley, and F. A. Matsen, IV, *ELife* **8**, e46935 (2019).
- [4] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cogn. Sci.* **9**, 147 (1985).
- [5] G. E. Hinton, *Neural Comput.* **1**, 143 (1989).
- [6] G. E. Hinton, *Neural Comput.* **14**, 1771 (2002).
- [7] D. Kim and D.-H. Kim, *Phys. Rev. E* **98**, 022138 (2018).
- [8] S. Efthymiou, M. J. S. Beach, and R. G. Melko, *Phys. Rev. B* **99**, 075113 (2019).
- [9] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [10] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [11] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, *Phys. Rep.* **810**, 1 (2019).
- [12] T. Wu and M. Tegmark, *Phys. Rev. E* **100**, 033311 (2019).
- [13] D. P. Kingma and M. Welling, *arXiv:1312.6114*.
- [14] S. J. Wetzel, *Phys. Rev. E* **96**, 022140 (2017).
- [15] S. Iso, S. Shiba, and S. Yokoo, *Phys. Rev. E* **97**, 053304 (2018).
- [16] S.-H. Li and L. Wang, *Phys. Rev. Lett.* **121**, 260601 (2018).
- [17] M. Koch-Janusz and Z. Ringel, *Nat. Phys.* **14**, 578 (2018).
- [18] G. Torlai and R. G. Melko, *Phys. Rev. B* **94**, 165134 (2016).
- [19] A. Morningstar and R. G. Melko, *J. Mach. Learn. Res.* **18**, 5975 (2017).
- [20] P. Smolensky, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1 (MIT Press, 1987).
- [21] G. E. Hinton, in *Neural Networks: Tricks of the Trade* (Springer, Berlin, Heidelberg, 2012), pp. 599–619.
- [22] T. Schmah, G. E. Hinton, S. L. Small, S. Strother, and R. S. Zemel, in *Advances in Neural Information Processing Systems* (MIT Press, Boston, MA, 2009), pp. 1409–1416.
- [23] Y. Li and X. Zhu, in *2018 International Joint Conference on Neural Networks (IJCN)* (IEEE, New York, 2018), pp. 1–10.
- [24] K. Huang, *Introduction to Statistical Physics* (Chapman and Hall/CRC, London, 2009).
- [25] M. Gabriele, E. W. Tramel, and F. Krzakala, in *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, 2015), pp. 640–648.
- [26] L. Bottou, F. E. Curtis, and J. Nocedal, *SIAM Rev.* **60**, 223 (2018).
- [27] T. Tieleman, in *Proceedings of the 25th International Conference on Machine Learning* (ACM, 2008), pp. 1064–1071.
- [28] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, *J. Am. Stat. Assoc.* **112**, 859 (2017).
- [29] K. Sohn, H. Lee, and X. Yan, in *Advances in Neural Information Processing Systems*, Vol. 28, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, 2015), pp. 3483–3491.
- [30] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, in *Advances in Neural Information Processing Systems* (MIT Press, Boston, MA, 2007), pp. 153–160.
- [31] J. Besag, *J. R. Stat. Soc. D* **24**, 179 (1975).
- [32] DeepLearning 0.1 documentation, <http://deeplearning.net/tutorial/rbm.html> (2020).
- [33] C. E. Rasmussen, in *Advances in Neural Information Processing Systems* (MIT Press, Boston, MA, 2000), pp. 554–560.
- [34] Q. Wang, S. R. Kulkarni, and S. Verdú, *IEEE Trans. Inf. Theory* **55**, 2392 (2009).
- [35] D. O. Loftsgaarden and C. P. Quesenberry, *Ann. Math. Stat.* **36**, 1049 (1965).
- [36] The choice of the distance metric depends on the nature of the samples. We use the Jaccard distance for the binary data in our models.
- [37] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, Cambridge, 2009).
- [38] L. Onsager, *Phys. Rev.* **65**, 117 (1944).
- [39] M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods* (John Wiley & Sons, 2009).
- [40] D. P. Kingma and J. Ba, 3rd International Conference for Learning Representations ICLR, San Diego, CA (2015), *arXiv:1412.6980*.
- [41] X. Glorot and Y. Bengio, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Vol. 9 (PMLR, 2010), pp. 249–256.
- [42] For filter sizes that approach the size of the underlying system, we would recover a nonconvolutional network architecture. If the filters are too small, we can account for local connectivity only in very small neighborhoods. In accordance with Ref. [64], we use filters of size 3. We also performed simulations for filter sizes 2 and 4 and found similar performances of the considered convolutional VAEs.
- [43] S. Ioffe and C. Szegedy, *arXiv:1502.03167*.
- [44] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [45] Training the considered RBMs for 10^3 epochs required seven days of computing time on an Intel Xeon E7-8867v3 (2.5 GHz)

- CPU. For cVAEs, the training took 20 hr on a GeForce GTX 1080 Ti GPU.
- [46] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in *Advances in Neural Information Processing Systems* (MIT Press, Boston, MA, 2012), p. 1097.
- [47] M. D. Zeiler and R. Fergus, in *European Conference on Computer Vision* (Springer, Berlin, Heidelberg, 2014), p. 818.
- [48] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2015), p. 1.
- [49] A. Radford, L. Metz, and S. Chintala, [arXiv:1511.06434](https://arxiv.org/abs/1511.06434).
- [50] B. Mitchell and J. Sheppard, in *2012 Proceedings of the 11th International Conference on Machine Learning and Computing*, Vol. 1 (IEEE, New York, 2012), pp. 162–167.
- [51] L. Rueschendorff, Wasserstein metric, Encyclopedia of Mathematic, http://encyclopediaofmath.org/index.php?title=Wasserstein_metric&oldid=32292.
- [52] A. Abbara, Y. Kabashima, T. Obuchi, and Y. Xu, [arXiv:1912.11591](https://arxiv.org/abs/1912.11591).
- [53] H. Schulz, A. Müller, and S. Behnke, in NIPS 2010 Workshop on Deep Learning and Unsupervised Feature Learning, Whistler, Canada (2010), <https://deeplearningworkshopnips2010.files.wordpress.com/2010/11/paper.pdf>.
- [54] C. K. Fisher, A. M. Smith, and J. R. Walsh, [arXiv:1804.08682](https://arxiv.org/abs/1804.08682).
- [55] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, in *Advances in Neural Information Processing Systems* (MIT Press, Boston, MA, 2014), pp. 2672–2680.
- [56] Z. Liu, S. P. Rodrigues, and W. Cai, [arXiv:1710.04987](https://arxiv.org/abs/1710.04987).
- [57] J. Long, E. Shelhamer, and T. Darrell, in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2015), pp. 3431–3440.
- [58] R. Hügli, G. Duff, B. O’Conchuir, E. Mengotti, A. F. Rodríguez, F. Nolting, L. Heyderman, and H. Braun, *Philos. Trans. R. Soc. London A* **370**, 5767 (2012).
- [59] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. Katsnelson, and A. A. Bagrov, *Nat. Commun.* **11**, 1593 (2020).
- [60] E. N. Cirillo, G. Gonnella, M. Troccoli, and A. Maritan, *J. Stat. Phys.* **94**, 67 (1999).
- [61] T. Tanaka, *Methods of Statistical Physics* (Cambridge University Press, Cambridge, 2002).
- [62] A. Pelizzola, *J. Phys. A* **38**, R309 (2005).
- [63] <https://github.com/dngfra/Restricted-Boltzmann-Machines>.
- [64] K. Simonyan and A. Zisserman, [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).