

The Many-Body Physics of Deep Learning

Zhe Wei Kho

Department of Physics, University of California, Davis

Deep learning has received considerable attention in recent times owing to its success in a vast variety of applications. An exciting cross-fertilization between deep learning architectures and many-body physics is also occurring, the reason being that the two are intricately connected to each other. In this paper, we briefly review the relationship between these two concepts, as well as explore how this blossoming field of research can lead to exciting new developments in both domains.

I. INTRODUCTION

More is different. In this aptly titled article [1], Anderson highlighted the notion of emergent phenomena in complex and many-body systems, noting that as the size of systems are scaled up, the interactions between the constituents of the system result in the emergence of various intriguing phenomena. Many of these are almost impossible to predict intuitively based only upon an elementary understanding of the properties of the individual constituents. Rather, “the understanding of these new behaviours requires research which is as fundamental in its nature as any other”.

In recent years, there has been an growing appreciation and recognition of the fact that such emergent phenomena occur throughout numerous hierarchical scales and contexts in nature – ranging from physics to biology and even the social sciences, leading to considerable interest in the study of complex systems.

In particular, the study of quantum many-body systems is currently one of the most exciting and challenging problems in contemporary physics research, with the immense potential to yield not only novel insights into fundamental physics but also open the door towards the development promising and innovative new technologies and applications. Quantum many-body systems predominantly occur in the domain of condensed matter physics, which in itself is a rapidly expanding field, where the complex many-body interactions are responsible for the rich variety of interesting and important physical phenomena occurring in mesoscopic systems, which include but are not limited to superconductivity, superfluidity, quantum magnetism, quantum transport, disorder and localisation, as well as topological effects

Unfortunately, more is not only *different*, but also *exponential*. Although most many-body systems can be characterised by relatively simple-looking model Hamiltonians, they cannot be solved analytically in general, and must be tackled via approximation schemes or numerical approaches. The key obstacle to direct numerical methods such as Quantum Monte Carlo schemes lies in the exponential increase in the dimensionality of the Hilbert space of the system with the size of the system – a system composed of N binary spins would require 2^N bits of information to describe, which makes it computa-

tionally expensive in terms of both resources and time, and practically constrains us to systems of fairly limited and unrealistically small sizes.

A promising and novel solution for getting around this problem, which has gained traction only in the preceding few years, lies in the form of applying machine learning architectures to reformulate the many-body problem. Naturally, one may question whether this is merely just part of the passing fad involving the attempted application of machine learning to just about any imaginable situation, inspired rightly or wrongly by the general successes of machine learning. While this approach appears to be rather unorthodox at first glance, it turns out that deep learning and many-body physics actually share an intricate connection with each other – and this could pave the way to not only a powerful means of tackling many-body systems, but also shed light into the black box of why deep learning works so well in so many different situations.

II. APPLICATIONS OF MACHINE LEARNING TO MANY-BODY PHYSICS

There are numerous areas in which machine learning has had, or promises to have a significant impact in advancing the study of many-body systems. These can be broadly classified into the four major areas below:

- (1) *Quantum machine learning* [2]. Whilst not exactly an ‘application’ of machine learning to many-body systems, quantum machine learning is the idea of speeding up classical machine learning algorithms such as support vector machines [3, 4] by employing many-body quantum systems to implement quantum computing protocols such as Grover’s search algorithm. It is also possible that potential new quantum algorithms may be devised in the near future.
- (2) *Learning to distinguish phases of matter*. Given datasets representing typical configurations of the system, the machine learns on-the-fly to identify and distinguish any phases present [5–7]. This is promising for the study of systems which exhibit complex phases that do not lend themselves to a

ready description in terms of clearly defined order parameters or structure factors, such as topological phases.

- (3) *Deep learning as a recommender system.* In designing Monte Carlo schemes, one needs to maximise the acceptance rate whilst maintain the correct detailed balance conditions. Notably, the standard Metropolis algorithm suffers from critical slowing down near the critical point, and numerous algorithms such as the cluster update algorithms of Swendsen-Wang and Wolff have been proposed to overcome this problem. Analogously, there are the ‘worm’ algorithms for the quantum Monte Carlo simulation of quantum many-body systems. Rather than relying on the ingenuity of physicists in coming up with more efficient schemes, Liu et al. [8] discovered that, without any prior guidance, a deep learning architecture could learn how to implement cluster updates, possibly paving the way for machine-learning-designed Monte Carlo algorithms that update and adapt on-the-fly.
- (4) *Applying deep learning architectures directly to characterise the many-body problem.* Beginning with seminal papers by Bény [9] and Mehta and Schwab [10] that established a direct mapping between deep learning (restricted Boltzmann machines) and the renormalization group, various others such as Carleo and Troyer [11] and Gao and Duan [12] have gone on to employ deep learning networks to efficiently represent many-body states. In line with this spirit, novel approaches like the neural network renormalization group [13] have also been proposed.

However, in the context of this short review, we will primarily focus on the last sub-field, for that is where a deep and non-trivial link between the deep learning and many-body physics exists, as opposed to what one may consider merely a generic application of machine learning.

III. MACHINE LEARNING

The field of machine learning generically involves the construction of algorithms that enable the network or architecture of the ‘machine’ to learn from and perform predictive analysis on data. Another way of framing it is that machine learning tries to extract the relevant or key features of the system from the given underlying (microscopic) data. In particular, a major subcategory of machine learning, deep learning, is typically composed of layers of nonlinear processing units that are coupled together (Fig. 1), not unlike a lattice of coupled spins.




FIG. 1. Graphical representation of a typical neural network with several hidden layers. Image adapted from <http://neuralnetworksanddeeplearning.com/>.

Deep learning has received considerable attention from the scientific and engineering communities in recent years, bolstered in no small part by the resounding victory of the computer program AlphaGo over the top human players in the strategy board game Go. The game Go was widely considered to be one of the cornerstones for artificial intelligence owing to its complexity and large branching factor, and so the success of AlphaGo firmly established the power and potential capabilities of deep learning. Technical advances such as the advent of GPU-based computing, which greatly sped up processing times, as well as the release of the open-source machine learning library, *TensorFlow*, further made machine learning viable and readily accessible to the general community at large.

The advent of this machine learning revolution has generated significant advances in a wide range of fields as diverse as computer vision, natural language processing and even pharmaceutical drug design. Although we know that machine learning works – and works amazingly well in a variety of situations, the entire architecture is only understood on a heuristic level, and the reason why it works so well is not yet fully understood.

Contrary to the common misbelief that the power of deep learning networks fundamentally stems from nothing more than having a sufficiently large set of adjustable parameters to fit the possible space spanned by the dataset (which brings to mind John von Neumann’s critique that ‘with four parameters, I can fit an elephant, and with five I can make him wiggle his trunk’), the fact is that deep learning is not just a matter of overfitting the data. The architecture and design of the features of the network seem to play a highly nontrivial role in determining how efficiently the machine can be trained and how effective it is. Furthermore, on careful examination, the number of neurons or nodes and parameters

required to realize a successful deep learning scheme is significantly smaller than the space spanned by the possible functions that could describe the data (which unsurprisingly, scales exponentially with the dimensionality of the dataset). This leads to the conundrum: *why does deep and cheap learning work so well?* The answer, or at least a possible route to it, as we shall see later, stems from notions of symmetry and locality that are deeply-rooted in physics [14].

Restricted Boltzmann Machines

Of particular interest is the class of *stochastic energy-based models* known as *restricted Boltzmann machines*. The restricted Boltzmann machine has a relatively straightforward architecture, comprising nodes that are arranged in two layers – the *visible layer* $\{X_i\}$ and the *hidden layer* $\{h_i\}$. Fig. 2 depicts a schematic of this architecture.

Each of these nodes (sometimes referred to as spins or bits) is a binary-valued random variable. The term *restricted* in the name of these machines refers to us imposing the condition that no direct couplings exist between nodes belonging to the same layer. Only nodes belonging to different layers can be directly coupled to each other, as represented by the lines between the nodes in Fig. 2. The strength or weight of the coupling between nodes X_i and h_j is denoted by the (symmetric) coupling matrix W_{ij} . Each node can also be influenced by a local bias factor, denoted by $\{a_i\}$ for those in the visible layer and $\{b_j\}$ for those in the hidden layer.

The central defining characteristic of the restricted Boltzmann machine is then captured in the so-called *energy function*

$$E(\mathbf{X}, \mathbf{h}) = - \sum_{\langle i,j \rangle} w_{ij} X_i h_j - \sum_i a_i X_i - \sum_i b_i h_i \quad (1)$$

associated with a particular configuration of spins $\{\mathbf{X}, \mathbf{h}\}$, with the probability of the system existing in this configuration given by

$$P(\mathbf{X}, \mathbf{h}) = \frac{1}{\mathcal{Z}} e^{-E(\mathbf{X}, \mathbf{h})}, \quad (2)$$

where \mathcal{Z} is the *partition function*

$$\mathcal{Z} \equiv \sum_{\{\mathbf{X}, \mathbf{h}\}} e^{-E(\mathbf{X}, \mathbf{h})}. \quad (3)$$

The probability distribution represented by Eq. 2 is none other than the Boltzmann distribution commonly encountered in statistical physics, hence the name given to this particular architecture.

At this point, it is probably fruitful to pause and examine the blatantly obvious analogy between the description

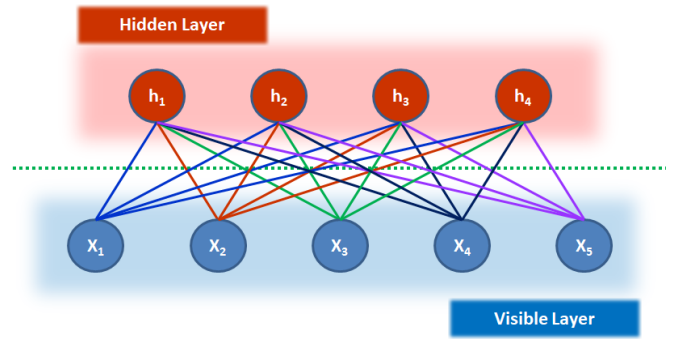


FIG. 2. Schematic of a restricted Boltzmann machine. The machine comprises two layers – the visible layer and the hidden layer. Only nodes in different layers are coupled directly to each other; nodes within the same layer are uncoupled. Each node can also experience a local bias that varies from node to node.

of the restricted Boltzmann machine and a typical spin system characterised by an Ising-type Hamiltonian

$$H_{\text{Ising}} = - \sum_{\langle i,j \rangle} J_{ij} S_i S_j - \sum_i h_i S_i, \quad (4)$$

where $\{S_i\}$ are binary variables describing the spin at the lattice sites $\{i\}$ (up or down), $\{J_{ij}\}$ is the coupling matrix between the spins and $\{h_i\}$ is the local external magnetic field. The key remaining difference between the two lies in the stratification of the spins in the restricted Boltzmann machine to two different layers, which we will expound upon in great detail later.

From the architecture of the restricted Boltzmann machine, we observe that each layer is *conditionally independent* (Markov shielding) – that is to say, given the configuration of the hidden layer $\{\mathbf{h}\}$, any two nodes in the visible layer, X_i and X_j , are independent with respect to each other, and vice versa with the roles of the hidden and visible layers swapped. One can also derive the activation probabilities of each node or spin to be

$$P(X_i = 1 | \mathbf{h}) = \sigma \left[a_i + \sum_j w_{ij} h_j \right] \quad (5)$$

$$P(h_i = 1 | \mathbf{X}) = \sigma \left[b_i + \sum_j w_{ij} X_j \right], \quad (6)$$

where

$$\sigma(x) \equiv \frac{1}{1 + e^{-x}} \quad (7)$$

is the *sigmoid* function.

The values and significance of the local biases $\{a_i\}$ and $\{b_i\}$, as well as the coupling matrix w_{ij} , have hitherto been undiscussed; in fact determining the values of these parameters is the key task of the learning procedure. These usually involve an iterative sequence of steps that

seeks to minimise the difference between the input distribution and the output distribution predicted by the machine (typically characterised by the *Kullen-Leibler divergence*). However, we shall not delve into these technical details here since these bear no direct relation to the physics that we intend to discuss.

The restricted Boltzmann machine has received much attention in the physics community owing to some of its attractive features in comparison to other more general models: (1) it has in-built notions of spatial locality (in that only adjacent layers are coupled); (2) it shares a close connection with the *renormalization group* [10] and (3) the existence of efficient algorithms for training.

The representative power of restricted Boltzmann machines can be further enhanced by stacking layers, thereby organizing the network of nodes into a *deep architecture*. It can be proven that such an organization can be faithfully mapped to a *deep Boltzmann machine* consisting of an additional third ‘deep’ layer on top of the usual restricted Boltzmann machine.

IV RENORMALIZATION GROUP

Before moving on to establish the explicit link between the architecture of the restricted (or deep) Boltzmann machine and the physics of many-body systems, we need to first discuss another key concept – that of the *renormalization group*. The renormalization group [15, 16] is one one of the most significant, if not the most, concepts underlying the modern understanding of many-body systems.

The central inspiration for the renormalization group stems from the curious phenomena of *universality* in the critical behaviour of many body systems: systems possessing very distinct microscopic descriptions can exhibit the same scaling laws in the vicinity of the critical point (near a phase transition); these systems are said to belong to the same *universality class*. A canonical example would be the correspondence between the liquid-gas phase transition of simple fluid and a single-axis ferromagnet (the Ising model).

These ideas led Wilson [15] to formulate the notion of a *renormalization group flow* in which the critical points are mapped to fixed points of some transformation of the Hamiltonian or system. Under such a transformation, which we will denote by \mathbf{R}_λ , the new parameters \mathbf{p} of the Hamiltonian are transformed as

$$\mathbf{p}' = \mathbf{R}_\lambda(\mathbf{p}). \quad (8)$$

In particular, at the critical point, we have

$$\mathbf{p}^* = \mathbf{R}_\lambda(\mathbf{p}^*). \quad (9)$$

Put simply, at the critical point, the parameters and functional form of the renormalized Hamiltonian remain invariant under the transformation.




FIG. 3. The effect of block spin decimation on a 2D Ising lattice for three different regimes: (i) below the critical temperature, (ii) at the critical temperature and (iii) above the critical temperature. The black and white regions denote the up and down spins at each site respectively.

An example of such a transformation is the *block-spin decimation* first proposed by Kadanoff. In this scheme, neighbouring spins are grouped into blocks, each block of spins is then replaced with a single spin whose new spin orientation is determined using a majority-rule on the constituent spins. In essence, we are tracing out some degrees-of-freedom to form a smeared-out and smaller lattice (in loose terms, we are steadily zooming out in how we view the system).

This procedure is qualitatively summed up in Fig. 3, where the block-spin decimation is performed on a standard 2D lattice of Ising spins for representative configurations corresponding to three different parameter regimes: (i) below the critical temperature T_C , (ii) close to T_C and (iii) above T_C . As the renormalization group transformations occur, regime (i) gets progressively colder (T flows or decreases towards the fixed point $T_1^* = 0$), while regime (iii) gets progressively hotter (T flows or increases towards the fixed point $T_3^* = \infty$). Regime (ii) meanwhile remains relatively invariant under the transformation since it is already at the non-trivial fixed point $T_2^* = T_C$.

The spins that are traced out (or summed over) generate new effective correlations and couplings between the new spins; in this sense these traced-out spins play a role akin to the hidden layer in restricted Boltzmann machines.

V DEEP LEARNING & MANY-BODY PHYSICS

Although much of the work exploring the connection between deep learning and many-body physics or applying deep learning techniques to tackling many-body problems only proliferated in the past few years or so, these two domains of study actually share a deep relation dating back decades ago when they were still being

formulated.

In 1967, Ricciardi and Umezawa [17] proposed a many-body model attempting to characterize the interactions between neurons in the brain, demonstrating how this model could possibly support learning and memory processing. Most deep learning networks also drew heavy inspiration from such many-body representations of neurons (hence the name neural networks), and the Hopfield network class of random-field Ising spin models is one of the early successful models of how memory and information might be encoded nonlocally in the structure of the network rather than at local individual sites or neurons.

Why Deep Learning Works so Well

The central problem experienced in both scenarios (many-body physics and machine learning) is similar in nature – given a dataset (microscopic details or spin configuration) that is possibility exponential in the space that it spans, how can we describe or extract the key features embodied (macroscopic observables and characteristics) using only a *polynomial* amount of resources and time?

In the domain of many-body physics, this has been achieved to a certain degree of success by building upon ideas from the renormalization group, as well as the *special properties* of the states-of-interest. In most systems, one is mainly concerned with studying the ground and low-lying excited states. These states occupy only a small region of the entire Hilbert space of states, and typically satisfy the *entanglement entropy scaling law*, as opposed to typical thermal states that satisfy the usual extensive volume scaling behaviour. The relatively low amount of entanglement present enables these states to be efficiently represented using tensor network or matrix product states, which constitute the basis of some of the more popular numerical schemes today, such as the density matrix renormalization group (DMRG) in 1D and multi-scale entanglement renormalization ansatz (MERA) in 2D.

Building on these suggestive analogies between deep learning architectures, Lin et al. [14] argue that, just as physical considerations such as symmetry and scaling behaviour enable physics to be so effective, deep learning too works so unreasonably well because of constraints imposed by physics on the properties and nature of real datasets, “hinging on symmetry, locality, and polynomial log-probability in data from or inspired by the natural world.” They further go on to establish a dictionary-of-sorts that draws direct parallels between various concepts in Physics and their corresponding counterparts in Machine Learning (Table I).

Physics	Machine Learning
Hamiltonian	Surprisal (Self-Information)
Quadratic Hamiltonian	Gaussian Distribution
Partition Function	Softmax Function
Spin	Node (Bit)
Relevant Operator	Feature
Irrelevant Operator	Noise

TABLE I. ‘Dictionary’ demonstrating between the correspondence between concepts from Physics and their counterparts from Machine Learning, as put forth by Lin et al. [14]

Direct Correspondence Between RBMs and RG

These observations and somewhat qualitative conjectures put forth by Lin et al. and various others were put on firm ground by Mehta and Schwab, who were able to go beyond intuition and establish a direct mapping between the architecture of restricted Boltzmann machines and the renormalization group procedure [10]. Notably, they derived a one-to-one correspondence between the variation renormalization group transformation operator \mathbf{R}_λ acting on the many-body Hamiltonian and the *conditional probability* of the hidden spins given the configuration of the visible spins in the corresponding restricted Boltzmann machine representation.

As an explicit illustration of this correspondence, let us consider the block-spin renormalization of a 1D Ising model, where the renormalization group transformation can be performed exactly [18]. The isotropic 1D Ising model in the absence of an external magnetic field is described by the Hamiltonian

$$H_{\text{Ising}} = -J \sum_i S_i S_{i+1}, \quad (10)$$

with the associated canonical partition function

$$\mathcal{Z} = \sum_{\{S_i=\pm 1\}} \prod_i \exp[K S_i S_{i+1}], \quad (11)$$

where $K \equiv \beta J$. To perform the block-spin decimation, we trace over all odd spins to get

$$\mathcal{Z} = \sum_{\{S_j=\pm 1\}} \prod_j 2 \cosh[K(S_j + S_{j+1})], \quad (12)$$

where we have relabelled the spins. In the next step, we attempt to rewrite this in the same form as the original partition function and so demand that

$$2 \cosh[K(S_j + S_{j+1})] = [g(K)]^{N/2} \exp[K' S_j S_{j+1}], \quad (13)$$




FIG. 4. Block-spin renormalization on a 1D Ising model. Alternate spins are successively traced out, leading to a new model that is still described by the same Ising Hamiltonian but with a different effective coupling constant J (or K).

where K' is the renormalized coupling constant. Solving the set of recursion relations obtained from Eq. (13) leads to the result

$$K' = \frac{1}{2} \log [\cosh(2K)]. \quad (14)$$

This process is pictorially depicted in Fig. 4: after each block-spin decimation step, we recover the same 1D Ising model, but with a different renormalized effective coupling constant $K^{(n)}$ that satisfies the recursive relation

$$K^{n+1} = \frac{1}{2} \log [\cosh(2K^{(n)})]. \quad (15)$$

At the same time, let us consider the restricted Boltzmann architecture depicted in Fig. 5. When we trace out the hidden layers successively, it is straightforward to see that this generates effective couplings between the nodes in the same layer that obey exactly the same relation as Eq. (15). Figs. 4 and 5 therefore establish a direct mapping between the two scenarios.

Similar analysis can also be performed for the 2D Ising model, but since the renormalization group transformation cannot be performed exactly in closed form, it has to be performed numerically, as was done by Mehta and Schwab [10], who also discovered that the structure of the couplings in the restricted Boltzmann machine representation displayed features reminiscent of block-spin grouping.

Given that finding the ‘correct’ or good renormalization group transformation \mathbf{R}_λ to use is not an exact science, but rather an art. One has to identify which are the *irrelevant degrees-of-freedom* and also attempt to perform the rescaling of couplings in a *tractable* manner so as not to generate an exponentially increasing number of couplings between the new spins. This naturally leads to the question of whether the deep learning architecture




FIG. 5. Restricted Boltzmann machine hierarchy corresponding to the RG process depicted in Fig. 4. The bold lines indicate actual direct couplings between the nodes, while the dashed lines indicate the effective couplings between nodes in the same layer should the hidden layer above it be traced over.

can learn the best renormalization group procedure to perform.

It turns out that the answer is a resounding ‘yes’: Koch-Janusz and Riegler proposed and implemented a deep learning algorithm [19] based on analysing the real-space mutual information (RMSI) between different parts of the system, that was able to identify the relevant degrees-of-freedom in the system to keep while tracing out the rest in an *unsupervised learning context*.

VI OUTLOOK

Much work in this nascent field is still ongoing in areas such as attempting to represent a variational expression of the many-body wavefunction in terms of restricted Boltzmann machines [11] or neural networks [12], as well as further work investigating the full extent of the relationship between quantum many-body systems and machine learning paradigms.

In particular, in a recent paper published just a few months ago, Chen et al. [20] extended Mehta and Schwab’s result in demonstrating an equivalence between restricted Boltzmann machines and tensor network states. They proved that the former can always be represented in terms of tensor network states, while tensor network states can be mapped to a restricted Boltzmann machine representation under certain conditions, and further provided explicit algorithms for performing these mappings.

The significance of this lies in the fact that restricted Boltzmann machines can be parametrized with far fewer parameters and also lend themselves to training with established pre-existing machine learning algorithms, which may be more efficient as compared to current complicated many-body simulation schemes. The represen-

tation in Boltzmann machine architecture also opens up new windows for more ready analysis of the many-body state in terms of information measures such as complexity.

Indeed, the cross-fertilization between many-body physics and machine learning is an exciting endeavour that promises to bring about revolutionary advances in the way that we look at both machine learning architectures and many-body systems alike.

-
- [1] P. W. Anderson, *Science* **177**, 393 (1972).
 - [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature* **549**, 195 (2017).
 - [3] P. Rebentrost, M. Mohseni, and S. Lloyd, *Physical Review Letters* **113** (2014).
 - [4] Z. Li, X. Liu, N. Xu, and J. Du, *Physical Review Letters* **114** (2015).
 - [5] J. Carrasquilla and R. G. Melko, *Nature Physics* **13**, 431 (2017).
 - [6] K. Chng, J. Carrasquilla, R. G. Melko, and E. Khatami, *Physical Review X* **7** (2017).
 - [7] L. Wang, *Physical Review B* **94** (2016).
 - [8] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Physical Review B* **95** (2017).
 - [9] C. Bény, arXiv: 1301.3124 (2013).
 - [10] P. Mehta and D. J. Schwab, arXiv: 1410.3831 (2014).
 - [11] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
 - [12] X. Gao and L.-M. Duan, *Nature Communications* **8** (2017).
 - [13] L. Wang and S.-H. Li, arXiv: 1802.02840 (2018).
 - [14] H. W. Lin, M. Tegmark, and D. Rolnick, *Journal of Statistical Physics* **168**, 1223 (2017).
 - [15] K. G. Wilson, *Physical Review B* **4**, 3174 (1971).
 - [16] L. P. Kadanoff, A. Houghton, and M. C. Yalabik, *Journal of Statistical Physics* **14**, 171 (1976).
 - [17] L. M. Ricciardi and H. Umezawa, *Kybernetik* **4**, 44 (1967).
 - [18] N. Goldenfeld, *Lectures on phase transitions and the renormalization group*, *Frontiers in physics* No. v. 85 (Addison-Wesley, Advanced Book Program, Reading, Mass, 1992).
 - [19] M. Koch-Janusz and Z. Ringel, *Nature Physics* **14**, 578 (2018).
 - [20] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, *Physical Review B* **97** (2018).
 - [21] I. Frérot and T. Roscilde, *Physical Review Letters* **116** (2016).
 - [22] H. E. Stanley, *Reviews of Modern Physics* **71**, S358 (1999).