

Symmetries and phase diagrams with real-space mutual information neural estimation


Doruk Efe Gökmen¹, Zohar Ringel², Sebastian D. Huber¹ and Maciej Koch-Janusz^{1,3,4}

¹*Institute for Theoretical Physics, ETH Zurich, 8093 Zurich, Switzerland*

²*Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel*

³*Department of Physics, University of Zurich, 8057 Zurich, Switzerland*

⁴*James Franck Institute, The University of Chicago, Chicago, Illinois 60637, USA*

 (Received 2 April 2021; revised 10 August 2021; accepted 5 October 2021; published 6 December 2021)

Real-space mutual information (RSMI) was shown to be an important quantity, formally and from a numerical standpoint, in finding coarse-grained descriptions of physical systems. It very generally quantifies spatial correlations and can give rise to *constructive* algorithms extracting relevant degrees of freedom. Efficient and reliable estimation or maximization of RSMI is, however, numerically challenging. A recent breakthrough in theoretical machine learning has been the introduction of variational lower bounds for mutual information, parametrized by neural networks. Here we describe in detail how these results can be combined with differentiable coarse-graining operations to develop a single unsupervised neural-network-based algorithm, the RSMI-NE, efficiently extracting the relevant degrees of freedom in the form of the operators of effective field theories, directly from real-space configurations. We study the information contained in the *statistical ensemble* of constructed coarse-graining transformations and its recovery from partial input data using a secondary machine learning analysis applied to this ensemble. In particular, we show how symmetries, also emergent, can be identified. We demonstrate the extraction of the phase diagram and the order parameters for equilibrium systems and consider also an example of a nonequilibrium problem.

DOI: [10.1103/PhysRevE.104.064106](https://doi.org/10.1103/PhysRevE.104.064106)

I. INTRODUCTION

Constructing coarse-grained descriptions of physical systems is a fundamental operation, both practically and from the foundational theory viewpoint. It is often difficult to simulate a complex system from first principles, even if a microscopic model exists, necessitating layers of descriptions and independent simulations of properties at differing scales [1]. This is the practical counterpart to the powerful methodology of deriving effective theories providing a summary of physics at this scale, which may involve qualitatively new emergent properties.

The above idea finds its full realization within the framework of the renormalization group (RG) [2–5], where the coarse-graining transformation of local degrees of freedom (DOFs) gives rise to the RG-flow in the space of theories. While momentum-space methods had a profound impact on physics [6,7], real-space renormalization, despite some veritable successes [8–11], has not yet reached a similar status. In practice, real-space procedures are nontrivial to design, involve ad hoc choices, and can rarely be executed exactly, amplifying any approximation as they are iterated. Moreover, analytical understanding is often lacking. At the same time disordered and complex systems, where the notion of momentum may not be available, are naturally amenable to real-space approaches and thus provide a strong motivation for development of improved methods.

The arbitrariness of the real-space RG transformation choice can be drastically reduced: an optimal coarse graining rule *for a given system* exists [12,13], theoretically defined

by maximizing the information the coarse DOFs retain about distant parts of the system, i.e., the real-space mutual information (RSMI). Such coarse-grained representation is an optimal (lossy) compression of information about long-distance properties of the system. This observation can, however, be made much more powerful. Since long-range information is due to the scaling operators, its optimal compression not only defines a better iterative RG rule, but, performed *for each point in the phase diagram*, should allow us to *directly* extract all the operators themselves, without explicitly executing the RG flow. This was formally proven, at least for critical systems in equilibrium [14]. Conceptually, this opens the possibility of identifying formal components of the effective theory directly from real-space data, using only its statistical properties.

The core quantity of interest, the RSMI, is however challenging to maximize or estimate, as is mutual information (MI) in general [15], which in practice limits the applicability of such general approach to but the simplest of systems.

Here we overcome this limitation by developing a highly efficient algorithm, the RSMI neural estimator (RSMI-NE), computing the optimal coarse-grainings. Our approach is based on state-of-art machine learning techniques of estimating mutual information by maximizing its rigorous lower-bounds [15,16]. A key algorithmic idea, introduced in Ref. [16], is to parametrize these bounds by a neural network f_{Θ} , optimizing over its parameters Θ . We use this differentiable variational *ansatz* for RSMI to optimize the RG transformation, expressed by another neural network *ansatz* with parameters Λ . Crucially, both networks are combined and trained together using stochastic gradient descent

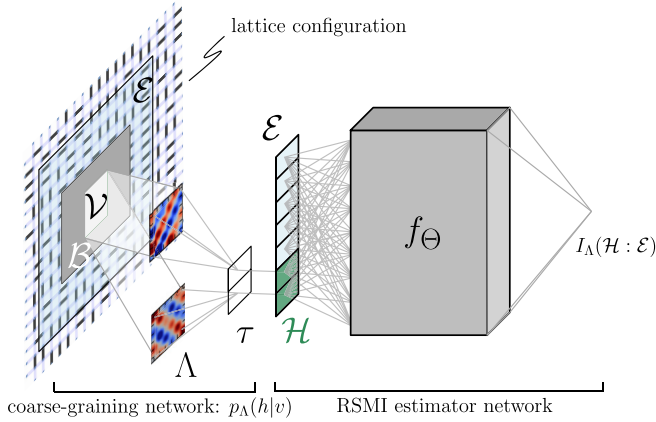


FIG. 1. The architecture of RSMI-NE. The coarse-graining transformation extracts the order parameters and relevant operators via the transformations Λ and discretizing step τ . The long-range information $I_\Lambda(\mathcal{H} : \mathcal{E})$ which Λ maximize is estimated by f_Θ , all of which are parametrized by neural networks and co-trained together.

backpropagating through the whole structure (see Fig. 1), even in the presence of discrete-valued coarse-graining. The algorithm is several orders of magnitude faster than RSMI estimation in Ref. [12], allowing us to explore large systems and length-scales, demonstrating superior convergence and stability.

The methods presented here were advertised in a companion work [17], where the emphasis was on the extraction of relevant operators on the lattice. Here, in contrast, we have two other areas of focus. First, we present and examine in detail the algorithmic aspects of the method, and the ML/statistics tools it uses. Second, we study the properties of a novel object: the *ensemble* of the coarse-graining transformations. Taking a step beyond identifying relevant operators in a single instance of RSMI optimization, we demonstrate that their distribution, generated by multiple independent optimization runs, reveals information about the symmetries of the system, broken and *emergent*. We consider application of the RSMI formalism to further systems, including nonequilibrium ones.

The manuscript is organized as follows: In Sec. II we describe the main theoretical and algorithmic components of the RSMI-NE. Specifically, in Sec. II A the general idea of the RSMI approach is briefly described, Sec. II B reviews the neural-network-based variational lower bounds on MI which are used and implemented in the RSMI-NE package [18] we make available. Section II C describes the parametrization of the coarse-graining as a neural network, and ensuring its *differentiability* for *discrete* latent variables; Sec. II D combines these elements into the complete algorithm and discusses convergence properties. Turning to physics, in Sec. III we detail the physical information (phase diagram, operators, symmetries) contained in the algorithm outputs, testing it in equilibrium statistical systems. Particularly, in Sec. III C, we show that the ensemble of coarse-graining transformations allows us to identify the symmetries of the system, also *emergent* ones, and retrieve them from incomplete data using a secondary ML analysis of the ensemble. The pos-

sible extension to nonequilibrium problems is discussed in Sec. IV, and the algorithm is validated on the example of the nonequilibrium chipping model. We conclude in Sec. V with a brief discussion of the scope of the method, and its possible extensions and applications. Short appendices give technical details related to the code and data generation.

II. THE RSMI-NE ALGORITHM

A. The RSMI variational principle

RG is rooted in the observation that most microscopic details are irrelevant for large-scale behavior of physical systems. It is, however, necessary to define a firm basis for determining exactly which short-scale details are projected out in a *coarse-grained* description. This has proven difficult in real space [19]. To address this issue, Ref. [12] proposed that the optimal real-space RG transformation maximizes an information theoretical quantity, the real-space mutual information (RSMI), measuring the information shared between a coarse-grained degree of freedom and its distant environment at the original fine level. This constitutes a universal principle for determining the coarse-grained description for any statistical system. Moreover, it gives *direct* access to the relevant operators on the lattice [14,17].

Consider a system of classical DOFs in any dimension denoted by a multivariate random variable \mathcal{X} , whose physics is encoded in a probability measure $p(x)$, either Gibbsian, i.e., $p(x) \propto e^{-\beta H(x)}$, or a generic nonequilibrium distribution. Here we denote by $x \sim p(x)$ an instance of the random variable \mathcal{X} drawn from the distribution $p(x)$. A coarse-graining rule $\mathcal{X} \rightarrow \mathcal{X}'$ is defined as a conditional distribution $p_\Lambda(x'|x)$, determined by a set of parameters Λ to be optimized. It is a probabilistic map generating a particular compressed representation of the original DOFs.

A coarse-graining is typically carried out on disjoint spatial blocks $\mathcal{V}_i \subset \mathcal{X}$, and it factorises: $p(x'|x) = \prod_i p_{\Lambda_i}(h_i|v_i)$, such that $\mathcal{X} = \bigcup_i \mathcal{V}_i$ and $\mathcal{X}' = \bigcup_i \mathcal{H}_i$, with $p_{\Lambda_i}(h_i|v_i)$ the coarse-graining rule applied to block i . If the system is translation invariant, then a fixed $\Lambda_i \equiv \Lambda$ suffices; otherwise, e.g., in disordered systems, it can be favorable to optimize each block individually.

The RSMI approach identifies coarse-graining filters extracting the most relevant long-range features as the ones retaining the most information shared by a block $\mathcal{V} \subset \mathcal{X}$ to be coarse-grained, and its distant environment \mathcal{E} [12,13], i.e., those that optimally *compress* this information. The environment is separated from \mathcal{V} by a shell of nonzero thickness constituting the buffer \mathcal{B} and forms the remainder of the system [see Fig. 1(a)]. The “shared information” between the random variables \mathcal{H} and \mathcal{E} is formalized by the Shannon mutual information:

$$I_\Lambda(\mathcal{H} : \mathcal{E}) = \sum_{h,e} p_\Lambda(e, h) \log \left[\frac{p_\Lambda(e, h)}{p_\Lambda(h)p(e)} \right], \quad (1)$$

where $p_\Lambda(e, h)$ and $p(h)$ are the marginal probability distributions of $p_\Lambda(h, x) = p_\Lambda(h|v)p(x)$ obtained by summing over the DOFs in $\{\mathcal{V}, \mathcal{B}\}$ and $\{\mathcal{V}, \mathcal{B}, \mathcal{E}\}$, respectively. The size of the buffer \mathcal{B} sets the RG scale, and acts a filter, only allowing the information about large-scale properties to

be compressed into the coarse-grained variables. Finding the optimal coarse-graining is thus formulated as a variational principle maximizing I_Λ as a function of parameters Λ .

Maximizing mutual information for high-dimensional random variables is known to be difficult [20,21], which limited the usefulness of the RSMI approach [12]. This challenge can now be overcome with the help of recent ML results combining mathematically rigorous variational bounds on mutual information [22–24] with deep learning [15,16]. In the following section we describe in detail the two components forming the core of the new fast RSMI-NE algorithm: efficient neural MI estimation and differentially parametrizing the coarse-graining operation.

B. Differentiable lower-bounds of RSMI

We follow the recent approach of Refs. [15,16]. Given possibly high-dimensional random variables \mathcal{X} , \mathcal{Y} , a variational upper or lower bound for $I(\mathcal{X} : \mathcal{Y})$ is constructed, and parametrized by a sufficiently expressive nonlinear neural network (NN) *ansatz* $f(x, y)$ modeling the statistical dependence of \mathcal{X} and \mathcal{Y} . The weights of the network are updated in an unsupervised learning scheme using the joint samples of \mathcal{X} , \mathcal{Y} (in our case: \mathcal{H} , \mathcal{E}), producing a sequence of differentiable bounds I_Λ , asymptotically exact.

It is possible to either minimize a variational upper-bound or to maximize a lower-bound of MI. Given our central aim of maximizing RSMI with respect to the parameters Λ of some coarse-graining network, we focus on the latter. We mainly use the noise-contrastive lower-bound of MI (InfoNCE), a multisample bound characterized by lower variance, but we also review the single-sample bounds it is the extension of. As we shall see, the general form of these bounds is motivated by the interpretation of MI as distinguishing between independently and jointly distributed random variables.

1. Single-sample lower-bounds

With this motivation in mind, and MI defined as follows:

$$I(\mathcal{X} : \mathcal{Y}) = \mathbb{E}_{p(x,y)} \left[\log \frac{p(x|y)}{p(x)} \right] = \mathbb{E}_{p(x,y)} \left[\log \frac{p(y|x)}{p(y)} \right],$$

we introduce the conditional probability distribution $q(x|y)$ as a variational *ansatz* approximating $p(x|y)$. We shall first keep the form of $q(x|y)$ unconstrained, and derive a lower-bound of $I(\mathcal{X} : \mathcal{Y})$ in its terms. Our goal is to find the optimal *ansatz* that makes the bound tight. Subsequently, we explain how the form of $q(x|y)$ can be constrained at the onset to improve the corresponding lower-bound, yielding a more tractable estimator.

Since the Kullback-Leibler (KL) divergence between them:

$$D_{\text{KL}}[p(x|y)||q(x|y)] = \mathbb{E}_{p(x,y)} \left[\log \frac{p(x|y)}{q(x|y)} \right] \quad (2)$$

is nonnegative, we immediately obtain a lower-bound for $I(\mathcal{X} : \mathcal{Y})$, known as the Barber-Agakov (BA) bound [23]:

$$\begin{aligned} I(\mathcal{X} : \mathcal{Y}) &\geq \mathbb{E}_{p(x,y)} \left[\log \frac{q(x|y)}{p(x)} \right] \\ &= \mathbb{E}_{p(x,y)} [\log q(x|y)] + H(\mathcal{X}) =: I_{\text{BA}}(\mathcal{X} : \mathcal{Y}), \quad (3) \end{aligned}$$

where $H(\mathcal{X})$ is the entropy of \mathcal{X} . This bound is a functional of the *ansatz*: $I_{\text{BA}}(\mathcal{X} : \mathcal{Y}) = I_{\text{BA}}(\mathcal{X} : \mathcal{Y})[q(x|y)]$. Since $D_{\text{KL}} = 0$ if and only if $q(x|y) = p(x|y)$, the BA bound is tight only when the *ansatz* $q(x|y)$ equals $p(x|y)$.

The observation that mutual information measures the correlations between variables motivates the idea that in modeling $p(x|y)$ the *ansatz* $q(x|y)$ should focus on the dependencies between the variables \mathcal{X} and \mathcal{Y} . Consider thus the following *ansatz* family:

$$q(x|y) := \frac{p(x)}{Z(y)} e^{f(x,y)}, \quad (4)$$

with $Z(y) := \mathbb{E}_{p(x)}[e^{f(x,y)}]$. In the above energy-based form, the complex correlations within the possibly high-dimensional data \mathcal{X} are contained in the marginal distribution $p(x)$. The resulting lower-bounds are sensitive mainly to the variables' interdependency. In other words, maximizing the lower-bound of MI is rephrased as a search for a ‘‘critic’’ [15] function $f(x, y)$ modeling the relationships, between \mathcal{X} and \mathcal{Y} very well. The critic function, distinguishing the ‘‘positive’’ samples from the joint distribution, from the ‘‘negative’’ ones generated by the product of marginals, will be approximated by a neural network.

Substituting the energy-based *ansatz* into the BA bound, we obtain the *unnormalized* BA bound (UBA):

$$I_{\text{UBA}}(\mathcal{X} : \mathcal{Y}) := \mathbb{E}_{p(x,y)}[f(x, y)] - \mathbb{E}_{p(y)}[\log Z(y)]. \quad (5)$$

By the same arguments as above, the UBA bound is tight when $\frac{p(x)}{Z(y)} \exp f(x, y) = p(x|y)$.¹ Taking advantage of the strict concavity of the log function, one arrives at the *tractable* version of UBA bound [15]:

$$\begin{aligned} I_{\text{TUBA}}(\mathcal{X} : \mathcal{Y}) &:= \mathbb{E}_{p(x,y)}[f(x, y)] \\ &\quad - \mathbb{E}_{p(x)p(y)} \left[\frac{e^{f(x,y)}}{a(y)} \right] - \mathbb{E}_{p(y)} \left[\log \frac{a(y)}{e} \right]. \end{aligned}$$

In several studies, see, e.g., Nguyen, Wainwright, and Jordan (NWJ) [24], f -GAN by Nowozin *et al.* [25], and MINE- f by Belghazi *et al.* [16], the *baseline* function $a(y)$ is fixed to be the constant e . This choice simplifies the TUBA bound:

$$I_{\text{NWJ}}(\mathcal{X} : \mathcal{Y}) := \mathbb{E}_{p(x,y)}[f(x, y)] - e^{-1} \mathbb{E}_{p(x)p(y)}[e^{f(x,y)}]. \quad (6)$$

Note that in this case $f(x, y)$ should be optimized under the constraint that $q(x|y)$ is normalized. The MINE approach [16] has also recently been used to estimate entropy in physical systems [26].

2. Replica lower bounds

Despite the improvement due to the energy-based *ansatz*, the above ‘‘single-sample’’ bounds are known to suffer from a large variance [15,27]. An improved approach is to divide a single batch of samples (e.g., Monte Carlo) for the pair of random variables (\mathcal{X} , \mathcal{Y}) into minibatches of K -fold ‘‘replicated’’ random variables $(\mathcal{X}_i, \mathcal{Y}_i)_{i=1}^K$, and to derive the corresponding

¹This condition is equivalent to $f(x, y) = \log p(y|x) + c(y)$, because $\exp f(x, y) = p(x|y) \frac{Z(y)}{p(x)} = p(y|x) \frac{Z(y)}{p(y)}$ with $c(y) := \log Z(y) - \log p(y)$ being a constant in x .

“multisample” lower-bounds (the confusing dual usage of the term “sample” is standard). These are obtained by taking the average of the single-sample bounds, and address the issue of large variance by means of noise-contrastive estimation (NCE) [28], first proposed in the context of MI estimation in Ref. [27].

A multisample bound estimates $I(\mathcal{X}_1, \mathcal{Y})$, where $(\mathcal{X}_1, \mathcal{Y}) \sim p(x_1, y)$, given $K - 1$ additional independent “replicas” for one of the random variables, say \mathcal{X} (drawn from the marginal distribution). We denote them by $\mathcal{X}_{2:K} \sim \prod_{j=2}^K p(x_j)$. All of the K independent replicas of the random variable \mathcal{X} are treated as a single K -dimensional random variable $\mathcal{X}_{1:K}$, and it is easily seen that $I(\mathcal{X}_1 : \mathcal{Y}) = I(\mathcal{X}_{1:K} : \mathcal{Y})$ since only \mathcal{X}_1 was drawn jointly with \mathcal{Y} . Thus, we can apply the “single-sample” bounds to $I(\mathcal{X}_{1:K} : \mathcal{Y})$.

For example, for the NWJ bound Eq. (6) the optimal *ansatz* for the critic f is given by (see Sec. I B in the Supplemental Material [29]):

$$f^*(x_{1:K}, y) = 1 + \log \frac{p(y|x_{1:L})}{p(y)} = 1 + \log \frac{p(y|x_1)}{p(y)}. \quad (7)$$

This critic function can be made to take advantage of the additional replicas. Observing that by Eq. (4):

$$\frac{p(x_1|y)}{p(x_1)} = \frac{e^{f^*(x_1, y)}}{Z(y)},$$

we can take the following modified critic function:

$$g(x_1, y) := 1 + \log \frac{e^{f(x_1, y)}}{m(y; x_{1:K})}, \quad (8)$$

where $m(y; x_{1:K})$ is the K sample Monte Carlo estimate of the partition function $Z(y)$:

$$m(y; x_{1:K}) := \frac{1}{K} \sum_{i=1}^K e^{f(x_i, y)} \approx Z(y). \quad (9)$$

Substituting this critic in the NWJ bound for $I(\mathcal{X}_1 : \mathcal{Y})$, and averaging over K replica random variables such that $(\mathcal{X}_i, \mathcal{Y}_i) \sim p(x_i, y_i)$, i.e., using the NWJ bound with each \mathcal{Y}_i playing the role of \mathcal{Y} in turn, we arrive after some simple but tedious algebra at the InfoNCE lower-bound of MI [15]:

$$\begin{aligned} I(\mathcal{X} : \mathcal{Y}) &\geq I_{\text{NCE}}(\mathcal{X} : \mathcal{Y}) := (I_{\text{NWJ}}(\mathcal{X} : \mathcal{Y})) \\ &= \frac{1}{K} \mathbb{E}_{\prod_{k=1}^K p(x_k, y_k)} \left[\sum_{j=1}^K \log \frac{e^{f(x_j, y_j)}}{\frac{1}{K} \sum_{i=1}^K e^{f(x_i, y_j)}} \right]. \quad (10) \end{aligned}$$

This completes derivation of the InfoNCE lower-bound for MI, which is the default one used in our implementation. In Sec. I C in the Supplemental Material [29] we discuss some of its further properties, including the expression in terms of the categorical cross-entropy and conditions on its upper bound, which should be taken into account to avoid biased estimates.

3. Neural network architectures for the RSMI lower-bound estimator

A key idea which made the variational bounds for MI introduced above computationally relevant was to parametrize

the critic functions by neural networks f_{Θ} [16]. Their parameters Θ can then be optimized using standard methods, e.g., stochastic gradient descent, to maximize the lower bounds.

Multiple multilayer perceptron (MLP) architectures for the critic function $f \equiv f_{\Theta}(h, e)$ have been considered [15]. Here, we opt for a *separable* form, such that

$$f_{\Theta}(h, e) = v^T(h)u(e), \quad (11)$$

where v and u are array-valued functions (here, neural networks, whose weights constitute Θ) that depend only on hidden variables and the environment, respectively. The networks v and u independently map \mathcal{H} and \mathcal{E} to a so-called embedding space. This choice allows construct the scores matrix F_{ij} (see below), storing the values of f_{Θ} for all pairs of jointly and independently drawn samples, in N passes of the MLP (N passes for both v and u networks) for a sample dataset of size N . This is in contrast to N^2 passes for all $N(N - 1)$ independent and N joint samples in a concatenated architecture $f_{\Theta}(h, e) = f_{\Theta}([h, e])$.

We opted for two hidden layers each with 32 neurons fully connected to the layer containing the $(\mathcal{H}, \mathcal{E})$ data. The embedding dimension is 8. The neurons are activated by the rectified linear unit (ReLU) function (see, e.g., Ref. [30]). We note that the results of RSMI-NE are not sensitive to these architectural details.

C. The coarse-graining network

The second key element of the algorithm is the coarse-graining probability distribution $p_{\Lambda}(h|v)$. To take advantage of the differentiable nature of the RSMI estimators described above, and the possibility of efficient gradient descent training, we consider *ansatz* parametrized by neural networks, as well. In particular, we use the following composite architecture (see Fig. 1):

$$h = \tau \circ (\Lambda \cdot v). \quad (12)$$

Combinations of the local DOFs are selected by an inner product with the parameters Λ , which can be understood in terms of a generalized Kadanoff block-spin transformation, [2] before being mapped to a discrete variable by the map τ . In practice the first operation can be represented by a single layer network with parameters Λ , and the number of kernels can be varied according to the symmetries of the system. We emphasize that the RSMI approach does not rely on the specific type of the variational *ansatz* for coarse-graining; the inner product form is a choice of convenience here. We briefly discuss the possibility of a more general coarse-graining Λ network *ansatz*, comprising multiple layers, in Sec. III D. The final step is a nonlinear stochastic mapping τ into a state h of the coarse-grained variable with a predetermined type (e.g., binary spins). This embedding is both crucial [31] and algorithmically nontrivial, as the discretization operation needs to be differentiable [32].

1. Gumbel-softmax reparametrization trick for discretization of coarse-grained variables

In the RSMI-NE the coarse-grained variables h are inputs to the MI estimator. Since the value of MI depends on what kind of distribution h belongs to, we need to ensure that this

estimation step is not falsified by, e.g., neglecting to force the output of the coarse-grainer into a discrete binary variable form, rather than a real number, if we decided h to be Ising spins. The apparent problem is that generating stochastic discrete h seems to spoil the differentiability of the whole setup. This is in fact somewhat similar to the problem encountered in variational autoencoders (VAEs), which is solved there using the so-called *reparametrization trick*, effectively allowing us to only differentiate w.r.t. to the parameters of the latent space probability distribution. With this intuition in mind, we discuss the solution to the issue in RSMI-NE [33].

The solution has three steps. The first result needed is the *Gumbel-max reparametrization*: Let h be a categorical random variable which can be in one of the states $\{i\}_{i=1}^N$ with the set of probabilities $\{\pi_i\}_{i=1}^N$. It can be shown that

$$k^* = \operatorname{argmax}_{k \in \{1:N\}} \{g_k + \log \pi_k\}_{i=1}^N \quad (13)$$

is a categorical random sample drawn from the distribution defined by $\{\pi_i\}_{i=1}^N$, where $\{g_i\}_{i=1}^N$ are N *parameterless* random variables drawn from the Gumbel distribution [33,34] centered at the origin. All the parametric dependence is therefore in the constants $\{\pi_i\}_{i=1}^N$, separated from the source of randomness, which in principle allows differentiation of the distribution.

Since argmax is not differentiable itself, in the second step it is *smoothed*, in a controlled and reversible fashion. Given $\{g_i\}_{i=1}^N$ we define a vector-valued random variable utilizing the softmax function Eq. (S12) [29], whose j th component takes the form

$$\operatorname{softmax}_{j,\epsilon}(\{g_i + \log \pi_i\}_{i=1}^N) = \frac{\exp[(\log \pi_j + g_j)/\epsilon]}{\sum_{i=1}^N \exp[(\log \pi_i + g_i)/\epsilon]}, \quad (14)$$

where ϵ is the smearing parameter. For $\epsilon \rightarrow 0$ the softmax becomes the argmax , mapping the argument vector $y = \{g_i + \log \pi_i\}_{i=1}^N$ into a N -component one-hot vector (one-hot encoding maps each of N possible states i of a discrete variable into a N -dimensional vector, with 1 on i th position, and zeros elsewhere) with some k^* th entry taking the value 1, thereby marking $y_{k^*} = \max y$. The result is a Gumbel-softmax random variable [33], approximately (or pseudo-)discrete, for small enough ϵ (do not confuse with a discrete random variable defined by taking the maximum component of the softmax function). For $\epsilon \approx 0$, a sample vector $h \sim \operatorname{softmax}_{\epsilon}(\{g_i + \log \pi_i\}_{i=1}^N)$ has a single component very close to 1 and all other components take very small values, comparable to machine precision. Conversely, in the limit $\epsilon \rightarrow \infty$ the distribution becomes uniform over all components.

This is used in the third step, where we anneal the smoothing parameter. There is a trade-off between small ϵ which leads to very noisy gradient estimates, and large ϵ at which the gradients have low variance but the samples h are far from being discrete. To reconcile this, we start the training at a high value of ϵ and anneal it exponentially toward a small positive value during training and thus stiffen the pseudodiscrete variable into an increasingly better approximation of a discrete one. The annealing procedure is described in more detail in Sec. II A 3 in the Supplemental Material [29].

D. Unsupervised learning scheme for the combined network

The results of the preceding section enable us to construct a variational *ansatz* $\tilde{I}_{\Lambda,\Theta}(\mathcal{H} : \mathcal{E})$, differentiable with respect to the parameters of the coarse-graining filter Λ and the estimator Θ . We stress that it is upper-bounded by the exact value of RSMI:

$$\max_{\Lambda} \tilde{I}_{\Lambda,\Theta}(\mathcal{H} : \mathcal{E}) \leq I_{\Lambda^*}(\mathcal{H} : \mathcal{E}), \quad \forall \Theta, \quad (15)$$

where Λ^* stands for the optimal solution. The equality holds if and only if the estimator becomes exact, i.e., for the optimal parameters $\Theta = \Theta^*$ of the energy-based *ansatz* f of InfoNCE. Thus, the search for the optimal RSMI coarse-graining, for any tuning parameters of the system, becomes a well-defined and tractable variational problem. It can be solved by simultaneously optimizing both set of trainable parameters $\{\Lambda, \Theta\}$ toward the same objective in an unsupervised learning scheme, which we now describe.

The inputs of the RSMI-NE can be e.g., the Monte Carlo (MC) samples from the desired model, for example as in Sec. III, but the algorithm can also be run on measured data. Since we estimate RSMI using the InfoNCE bound, the sampling is divided into minibatches, each containing K samples. We separate in each sample the visible patch \mathcal{V} and its environment \mathcal{E} , dismissing a finite buffer that separates them. Then a single minibatch is denoted by the multidimensional random variable $(v_{1:K}, e_{1:K}) = (v_1, \dots, v_K, e_1, \dots, e_K)$. As usual, ensuring good quality MC sampling is important.

Let Λ^s and Θ^s denote the network parameters for the coarse-graining, and the critic f , respectively, at training step s . We initialize them as tensors containing random numbers. At each step s , in the samples in the minibatch v_i are coarse-grained into $h_i[\Lambda^s]$ and the scores matrix $F_{ij}(\Theta^s, \Lambda^s) = f(h_i[\Lambda^s], e_j; \Theta^s)$ is computed for the InfoNCE at current values of the network parameters. In F_{ij} the entries with $i = j$ denote the jointly drawn samples and the rest denote independently drawn samples for the coarse-grained degree of freedom and the environment. As described above discrete h are generated by a layer τ .

The InfoNCE prediction (that of $p(h, e)$ being equal to $p(h)p(e)$ or not, as defined in Eq. (S13) [30]) for the minibatch is computed using the scores matrix as

$$Q(h_{1:K}, e_{1:K}; \Theta^s, \Lambda^s) = \sum_{j=1}^K \frac{\exp F_{jj}(\Theta^s, \Lambda^s)}{\sum_{i=1}^K \exp F_{ij}(\Theta^s, \Lambda^s)}. \quad (16)$$

Then $\log Q(h_{1:K}, e_{1:K}; \Theta^s, \Lambda^s) + \log K$ gives our single minibatch estimate of RSMI.

The gradients of the minibatch estimate of RSMI with respect to Λ and Θ are used to update the network parameters. We use the adam optimizer [35] to perform stochastic gradient-ascent. We found that using the same learning rate for both Λ and Θ leads to efficient training. We repeat the above over all minibatches, until all samples are fed to the network once. This constitutes one epoch of training. In Algorithm 1 the training procedure is given in pseudocode.

We train for multiple epochs until convergence criteria are satisfied (see Sec. II C in the Supplemental Material [29]). For illustration, we plot in Fig. 2 the time series of the RSMI estimates and the coarse-graining filters during the training for 2D

Algorithm 1. One epoch for the unsupervised learning procedure for the RSMI-net using InfoNCE lower-bound.

```

1:  $\eta =$  learning rate
2:  $\epsilon =$  relaxation parameter for Gumbel-softmax distribution
3:  $\Theta^0 \leftarrow$  random hyperparameter tensor
4:  $\Lambda^0 \leftarrow$  random hyperparameter tensor
5: for  $s$  in  $1 : n$  do
6:    $\epsilon^s \leftarrow$  reduce Gumbel-softmax relaxation parameter
7:    $\tau^s \leftarrow \tau(\epsilon^s)$ 
8:   for  $i$  in  $1 : K$  do
9:     for  $j$  in  $1 : K$  do
10:       $h_{ij}^s[\Lambda^s] \leftarrow \tau^s(\Lambda^s \cdot v_{ij}^s)$ 
11:       $F_{ij}(\Theta^s, \Lambda^s) \leftarrow f(h_{ij}^s[\Lambda^s], e_j^s; \Theta^s)$ 
12:    end for
13:  end for
14:   $Q(x_{1:K}, y_{1:K}; \Theta^s, \Lambda^s) \leftarrow \sum_{j=1}^K \frac{F_{ij}(\Theta^s, \Lambda^s)}{\sum_{i,j=1}^K \exp F_{ij}(\Theta^s, \Lambda^s)}$ 
15:
16:   $\Delta \Theta^s \leftarrow \eta \nabla_{\Theta} [\log Q(\Theta, \Lambda^s)] \Big|_{\Theta=\Theta^s}$ 
17:   $\Theta^s \leftarrow \Theta^s + \Delta \Theta^s$ 
18:
19:   $\Delta \Lambda^s \leftarrow \eta \nabla_{\Lambda} [\log Q(\Theta^s, \Lambda)] \Big|_{\Lambda=\Lambda^s}$ 
20:   $\Lambda^s \leftarrow \Lambda^s + \Delta \Lambda^s$ 
21: end for
22:  $\tilde{I}_{\Lambda}(\mathcal{H} : \mathcal{E}) = \frac{1}{n} \sum_{t=1}^n \log Q(x_{1:K}, y_{1:K}; \Theta^t, \Lambda^t) + \log K$ 
23: return  $\tilde{I}_{\Lambda}(\mathcal{H} : \mathcal{E}), \Lambda^n$ 

```

\triangleright initialise InfoNCE *ansatz* $f(h, e)$
 \triangleright initialise coarse-graining filter
 \triangleright loop over all n K -replica samples for $(\mathcal{V}, \mathcal{E})$
 \triangleright Anneal Gumbel-softmax layer
 \triangleright Coarse-grain visible degrees of freedom
 \triangleright ij 'th element of scores matrix
 \triangleright InfoNCE "prediction"
 \triangleright Update parameters of the RSMI estimator network:
 \triangleright automatic differentiation
 \triangleright stochastic gradient-ascent
 \triangleright Update parameters of the coarse-grainer network:
 \triangleright average over n samples

critical Ising antiferromagnet and interacting dimer models below and above the BKT transition point (see Sec. III). Upon convergence, we are left with an optimized coarse-graining represented by the final Λ -parameters, and an estimate of the RSMI given by a moving average of the time-series of minibatch estimates.

III. RSMI-NE IN EQUILIBRIUM SYSTEMS

The RSMI-NE algorithm yields a comprehensive characterization of long-distance properties of an equilibrium statistical system: its phase diagram, correlations, symme-

tries. The companion work [17] discusses in detail the construction of order parameters, or, more generally, relevant operators. Here we demonstrate how the extracted quantities, and their dependence on the tuning parameters of the system and the buffer length-scale, reveal the critical points and the nature of correlations in the phases. We illustrate this on the examples of a dimer model with aligning interactions and the 2D Ising model. We examine the information, particularly on symmetries, also *emergent*, contained in the statistical *ensemble* of coarse-graining filters, and show its retrieval with ML techniques, which

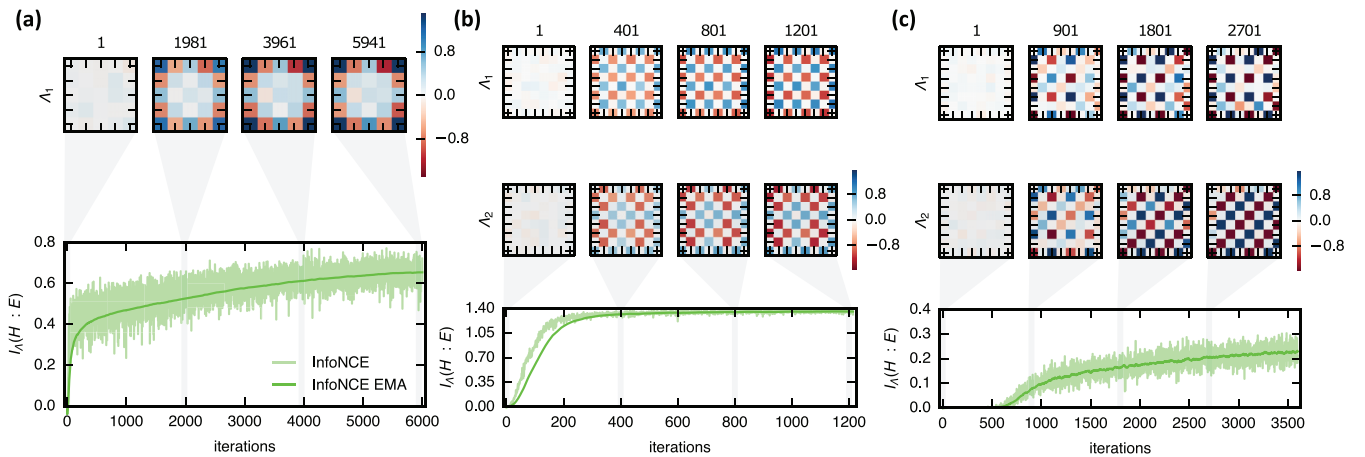


FIG. 2. Convergence of real-space mutual information value and the coarse-graining filters. The light-green curve shows the the time series of RSMI and dark green its exponential moving average. In the top panel, the time series of the coarse-graining filters are given. (a) The antiferromagnetic Ising model on a 2D square lattice at the critical point. The RSMI converges to $\log 2$ and the optimal filter couples to the boundary degrees of freedom in \mathcal{V} with an alternating sign pattern, due to the onset of antiferromagnetic order. (b) The same for the interacting dimer model at $T = 0.4 < T_{\text{BKT}}$. (c) For the interacting dimer model at $T = 15.0 \gg T_{\text{BKT}}$. See Sec. III for details on the Ising and dimer models and the interpretation of these results.

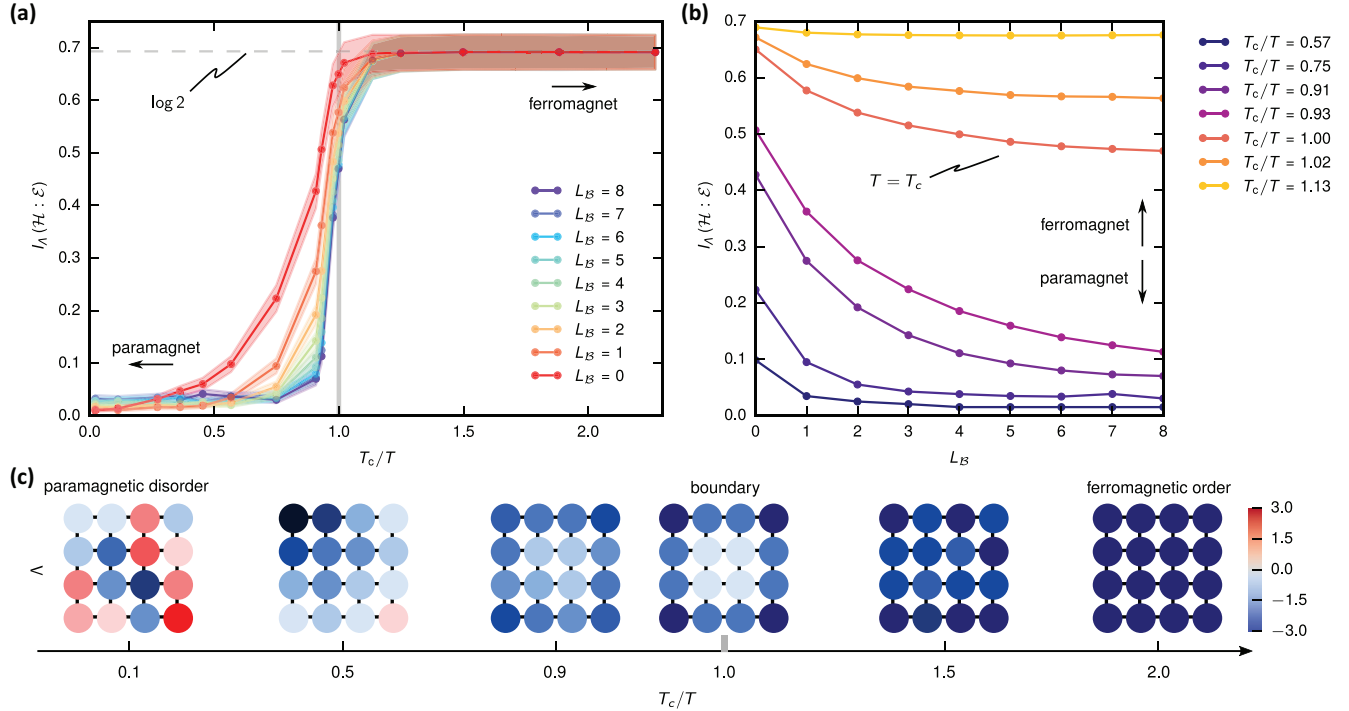


FIG. 3. Maximal RSMI as a function of temperature and its scaling with L_B . (a) The dependence of the maximal RSMI on temperature for different buffer widths L_B . (b) The scaling of the maximal RSMI with L_B at different temperatures. It is found that the RSMI decays exponentially in the paramagnetic phase, whereas the decay is slower at $T \leq T_c$. (c) The evolution of the RSMI-optimal filters with temperature at $L_B = 4$.

is of practical importance when faced with incomplete inputs.

We emphasize that, in contrast to many applications of ML (see Refs. [36,37] for a recent review), the success of RSMI-NE in extracting physical data *is not* serendipitous or due to a particular choice of architecture, but a consequence of RSMI being a well-defined *physical quantity* [14], which the ML methods used approximate numerically.

A. The phase diagram from the parameter dependence of RSMI and its scaling with buffer size

Real-space mutual information quantifies the totality of spatial correlations in the system, and thus their changing structure, especially due to phase transitions, should be reflected in its value. This is indeed the case, as shown below. The nature of these correlations (power-law versus exponential) further determines the decay properties of RSMI as a function of the length scale set by the buffer width.

We use the example of the classical 2D Ising model:

$$K[x = \{x_i\}] = \beta J \sum_{\langle i,j \rangle} x_i x_j, \quad (17)$$

with $x_i = \pm 1$, as the simplest test case for RSMI-NE. It undergoes a second order phase transition between a ferromagnetic for $J < 0$ or antiferromagnetic order for $J > 0$, and a disordered paramagnetic phase at inverse temperature $\beta = \ln(1 + \sqrt{2})/2 \approx 0.44$ [38]. We investigate this model in the temperature range $T_c/T \in [0, 2.5]$ by optimizing RSMI at buffer widths $L_B \in [0, 8]$.

As shown in Fig. 3(a), the temperature dependence of the maximal information $I_\Lambda(T)$, i.e., the amount of long-range information attained *with the (at given T) optimal Λ* , is a clear indicator of the second order phase transition, and of the existence of two phases. At $T < T_c$, independent of the buffer width, exactly 1 bit of information is recovered. This precise quantization is due to RSMI effectively counting the (two) segregated phase space sectors corresponding to the ferromagnetic ground states, and reveals the long-range order.

Phase transitions are reflected by nonanalyticities in $I_\Lambda(T)$ (cf. the behavior of the mutual information in the absence of buffer in Refs. [39,40]). At $T = T_c$ we find that the RSMI has a steplike decay which becomes sharper at larger buffer width L_B . At larger temperatures, the long-range order is destroyed by the thermal fluctuations. The short-range nature of the paramagnetic phase results in an exponential decay of the RSMI with L_B , see Figs. 3(b) and 6(a). This is to be contrasted with the critical phase of the dimer model with power-law correlations, where the maximal information decays only algebraically with L_B ; see Fig. 6(b).

We next turn to the more complex example of the interacting dimer model, defined by the partition function:

$$Z(T) = \sum_{\{C\}} \exp(-E_C/T), \quad (18)$$

with T the temperature and C denoting dimer configurations on the square lattice obeying the constraint of exactly one dimer at every vertex; see Fig. 4(a). The energy $E_C = N_C(\parallel) + N_C(=)$ counts plaquettes covered by parallel dimers favored by the interaction.




FIG. 4. (a) A generic valid dimer covering on the square lattice. (b) The four ground states of the dimer model with aligning nearest neighbor interactions break C_4 and lattice translation symmetries.

The essence of this system is in the interplay of aligning interaction energy and entropic effects due to the nonlocal cooperation of local dimer covering constraints. At low- T , the former facilitates long-range order (LRO), crystallizing the system into one of four translation symmetry breaking *columnar* states; see Fig. 4(b). With increasing T the system undergoes a Berezinskii-Kosterlitz-Thouless (BKT) transition at $T_{\text{BKT}} = 0.65(1)$ [41], entering a critical phase characterized by algebraic decay of correlations with exponents continuously changing with T . The effective theory of the system is given by a sine-Gordon field theory [41–43]. In particular, for $T \rightarrow \infty$ the aligning interactions are irrelevant and this description reduces to a free Gaussian field theory.

To test our method on the dimer model, we generate its Monte Carlo samples across the whole temperature range, using the directed loop algorithm [41] (see Sec. III in the Supplemental Material [29] for implementation details) for 64×64 systems. These are used as inputs to RSMI-NE. We restrict the coarse-grained variables \mathcal{H} to a two-component binary vector $\{\pm 1, \pm 1\}$, a choice suggested by the systematic procedure in Sec. III D. Hence, we are looking for a two-component vector of filters Λ_1, Λ_2 determining how the visible region \mathcal{V} is mapped onto \mathcal{H} .

Though the BKT transition is of entirely different nature to the Ising example considered above, we find that optimizing the filters Λ_1, Λ_2 for all T readily reveals the structure of the phase diagram [see Fig. 5(a)]. To wit, for $T < T_{\text{BKT}}$ its value is constant and equal to $\log 4$, or 2 bits. The information shared between distant parts of the system in the ordered phase is precisely which of the four columnar states they are in. This is analogous to the ferromagnetic order of the Ising model, i.e., the optimal RSMI counts the number of segregated phase space sectors in long-range ordered phases. Moreover, the algebraic decay of $I_\Lambda(T)$ with the buffer size for $T > T_{\text{BKT}}$, as seen in Fig. 6, is indicative of a critical phase with power-law decaying correlations. In particular, we have found that the RSMI scales as $I_\Lambda(T \rightarrow \infty) \sim L_B^{-\nu}$, with the exponent obtained from the best fit [see Fig. 6(c)] to be $\nu \approx 1.16$.

We conclude that the value of the optimal RSMI as a function of the system's parameters provides us with information about position of the critical points, type of phase transition, the nature of correlation decay in the phases, as well as the number of sectors in the long-range ordered phases.

B. Correlations from the parameter dependence and flow of the optimal coarse-graining filters

Much more can be learned about spatial correlations upon examining the coarse-graining filters $\Lambda(T)$. First, the optimal filters, with which the highest RSMI value was attained, them-

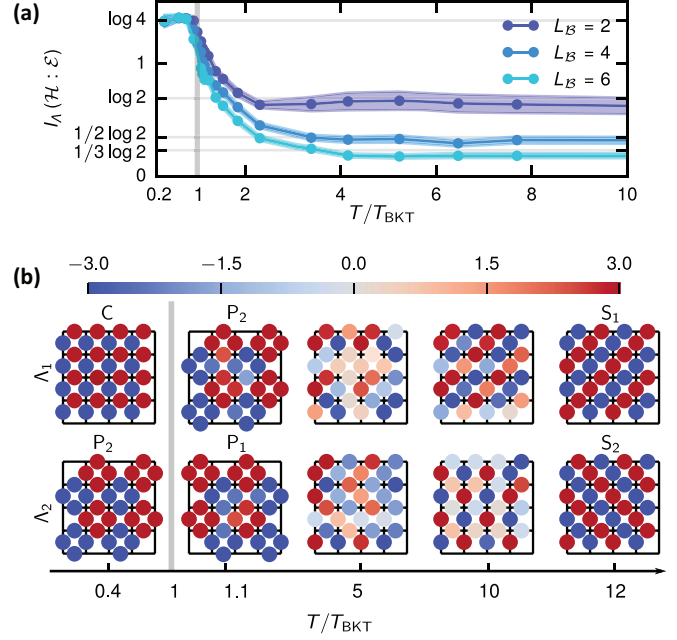


FIG. 5. RSMI analysis of the interacting dimer model. (a) Total RSMI extracted with the optimal filters as a function of T and its scaling with the buffer size. (b) Samples of optimal filters obtained with RSMI-NE for different T [columnar (C), plaquette (P_1, P_2), and staggered (S_1, S_2)].

selves depend on the tuning parameters of the physical system and in fact carry the information about the phase diagram. Particularly, they reflect the symmetries of the system (see also Sec. III C). Moreover, the filters depend on the length scale L_B , reflecting an RG flow. In Ref. [17] we further show that they correspond to the (lattice discretization of) relevant operators in the field theory describing the system, in light of which observation the intriguing results of this subsection become natural.

1. The optimal coarse-graining filters of the 2D Ising model

The temperature dependence and the relation of the optimal filters to the phase diagram are clear in the Ising model results, as seen in Fig. 3(c). Here we used a fixed buffer width




FIG. 6. The scaling of RSMI with the buffer size L_B . (a) The RSMI decays exponentially in the Ising paramagnet (at $T = 4$). For $L_B \geq 2$, the RSMI numerically drops to 0. (b) At the Ising critical point, the RSMI scales algebraically; the line gives the best fit for the exponent $\nu \approx 0.09$. (c) Same as in panel (b) but for the high-temperature free dimers ($T = 15$). The fit gives the RSMI scaling exponent $\nu \approx 1.16$ for the free dimers.

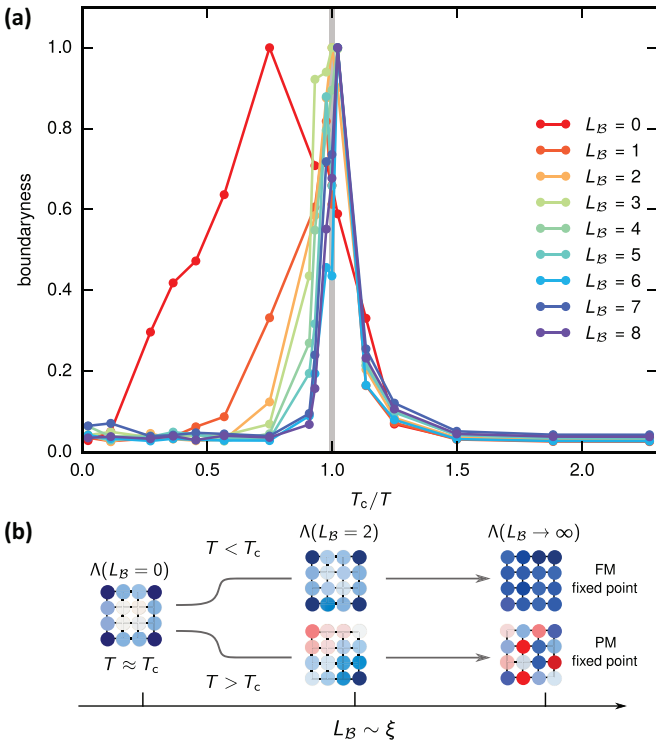


FIG. 7. Relative strength of coupling to boundary vs bulk spins in \mathcal{V} , as a function of temperature. (a) The empirical “boundaryness” measure [see Eq. (19)] of the optimal coarse-graining filter at different temperatures and buffer sizes. (b) Near the critical point T_c the optimal Λ averages the boundary spins in \mathcal{V} . The separation L_B of \mathcal{V} from its environment \mathcal{E} effectively sets the RG scale ξ . Growing L_B increasingly differentiates the filters Λ slightly below and above T_c , which ultimately flow to the ferromagnetic (FM) and paramagnetic (PM) fixed points. Note the transformations are *not* iterated.

$L_B = 4$, a visible region of size 4×4 and a 64×64 grid for the whole system.

In the high- and low- T limits the paramagnetic and ferromagnetic phases result in optimal filters, which are respectively random and uniform. The uniform filter acts as the ferromagnetic order parameter, labeling the configurations by their magnetization. Consistently, for the antiferromagnetic Ising model, we found that the optimal filter at low- T is the staggered magnetization (see Fig. 2). At the second-order critical point, the filters exhibit a boundary behavior [12], i.e., they correspond to the magnetization on the boundary of the visible block. This is because in the *critical* Ising system the information shared between \mathcal{V} and \mathcal{E} is proportional to the surface of their interface, and not to the volume of \mathcal{V} [39,44]. The boundary filter thus signals the presence of the critical point.

While the occurrence of the boundary filter is associated with the critical T_c , the range of temperatures where this happens depends on L_B . Put differently, the accuracy to which the critical fixed point can be resolved depends on the length scale set by the size of the buffer. To visualize this in Fig. 7(a) we plot an empirical measure of the relative strength of the boundary versus bulk couplings in $\Lambda(T)$, at different values

of L_B , which we call the “boundaryness”:

$$\text{boundaryness} := \frac{|\sum_{i \in \text{boundary}} \Lambda_i|}{|\sum_{i \in \text{bulk}} \Lambda_i|}, \quad (19)$$

and we rescale by the maximum over T in Fig. 7(a) for clarity (note that the maximum value was found approximately the same for all L_B except for 0, where it was much greater). This ratio peaks around $T = T_c$, becoming increasingly sharp as L_B grows.

This behavior is readily understood, since L_B (together with the total finite system size, cf. Eq. (5) in Ref. [14]) effectively controls the RG scale. Indeed, a corresponding flow of the optimal filters can be constructed. As shown in Fig. 7(b), for small L_B at $T \approx T_c$ the optimal filter is a boundary one, both above or below T_c . At this scale the critical point is not resolved very well. As L_B is increased, however, the $T < T_c$ and $T > T_c$ cases are increasingly differentiated, and they eventually flow to the ferromagnetic and paramagnetic fixed points, respectively. This, of course, is consistent with the presence of a repulsive fixed-point in the RG flow of the 2D Ising model.

2. The optimal coarse-graining filters of the interacting dimer model

The optimal coarse-grainings of the dimer model Eq. (18) likewise depend on the tuning parameters of the system (i.e., the temperature); see Fig. 5(b). In contrast to the 2D Ising model example, however, this dependence is continuous for $T > T_{\text{BKT}}$. This, in fact, provides another indication that the transition is of the BKT type (in addition to the algebraic scaling of the RSMI curve in the critical phase).

More concretely, in the high- and low- T limits, three classes of filters emerge: independent optimizations (see discussion of the filter *ensemble* in Sec. III C) return exclusively sets of filters $\Lambda_{1,2}$ that correspond to columnar and plaquette at low temperatures and staggered ones at high temperatures. They are denoted as \mathbf{C} , $\mathbf{P}_{1,2}$, and $\mathbf{S}_{1,2}$ in Fig. 5(b). We call these filters “pristine” as they reflect limiting cases. They also reveal information about the symmetries. In particular, the pristine plaquette and columnar filters at $T \rightarrow 0$ break the discrete translation or rotation symmetry of the lattice, respectively. Any pair of $\Lambda_{1,2}$ drawn out of these classes of filters defines a bijection between the four columnar states in Fig. 4 and the four distinct states $(\pm 1, \pm 1)$ of the compressed degrees of freedom in \mathcal{H} . They thus label uniquely the ordered states, (which is the reason the recovered RSMI is exactly 2 bits for $T < T_{\text{BKT}}$), and correspond precisely to the dimer symmetry breaking order parameter of Ref. [41]. In Ref. [17] we show the columnar and plaquette filters correspond to the electrical charge operators of the sine-Gordon field theory, i.e., the operators with the lowest scaling dimensions, and so the most relevant in the RG sense.

The degeneracy of plaquette $\mathbf{P}_{1,2}$ and columnar \mathbf{C} filters in their RSMI value is lifted when the rotation symmetry is restored at BKT transition: the pristine columnar filter, which breaks the lattice rotation symmetry explicitly, is not found above T_{BKT} .

In the limit of $T \rightarrow \infty$ the optimal filters are the staggered $\mathbf{S}_{1,2}$. These can be shown [17] to exactly correspond to the

spatial gradients of the height field in the sine-Gordon description of the system, or equivalently to the electrical fields. At $T \rightarrow \infty$ these are in fact the only terms in the field theory, which is then that of a free Gaussian field.

In the critical phase $T > T_{\text{BKT}}$, where the system is characterized by power-law correlations with temperature-dependent exponents, the resulting filters continuously interpolate between pristine plaquette and staggered ones. This is due to the competition between the electric field operator and plaquette correlations, or in other words the gradient and the cosine terms in the sine-Gordon field theory [41,43]. In a finite system the correlations due to these operators of slightly differing (for $T \gtrsim T_{\text{BKT}}$) scaling dimensions both contribute to RSMI, though in the thermodynamic limit the more relevant gradient term would dominate (which indeed happens for larger T).

C. The ensemble of coarse-graining filters and its analysis: Operators and symmetries

1. Extracting the pristine filters (lattice operators)

The above described mixing of the pristine filters in intermediate parameter regimes and in finite-size systems may seem troublesome, but we find it can in fact be resolved, and the solution to this problem is useful in itself.

The key observation is that due to the RSMI-NE being a stochastic algorithm it produces in independent runs a *distribution* of RSMI-optimal transformations [thus, Fig. 5(b) shows a sample of filters at each T]. This distribution defines *the ensemble of filters*, a novel concept we introduce.

The ensemble contains physical information, particularly about the symmetries, and, in contrast to individual filters, also *emergent* ones (see below). Crucially, the ensemble also allows us to address the problem of filter mixing due to competing correlations. The pristine filters, which correspond to the (lattice representation of) scaling operators [17], can be identified not only at the limiting temperatures, but also through data analysis of the ensemble in a window of intermediate temperatures.

To show this we perform a principal component analysis (PCA) of the ensemble of RSMI-optimal filters for an intermediate temperature range $0.7 < T < 3.7$ above the BKT transition, where the pristine components do not explicitly appear. The goal is to find the most distinctive features of the ensemble which vary with the changing system parameters controlling the location in the phase diagram (here: temperature), while filtering out variations due to a specific realization of statistical noise or the random initial conditions of the training. As we coarse-grain the dimer model using 8×8 two-component filters, we consider a 64-dimensional vector space where each coarse-graining filter component is a point. The input to the PCA consists of the ensemble of coarse-graining filters, flattened into 1D arrays. The resulting principal components are reshaped back into 8×8 arrays, so that each defines a coarse-graining. To visualize the results we then project the full space of coarse-graining filters onto the hyperplane given by the most important principal components.

The results are shown in Fig. 8. An important observation is that the highest principal components are in fact given

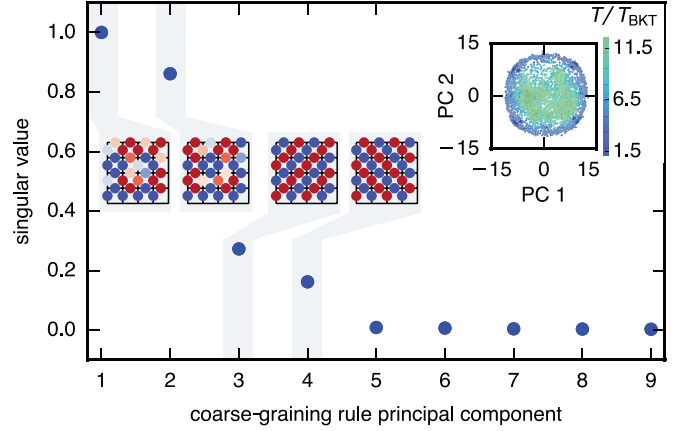


FIG. 8. Analysis of the ensemble of coarse-graining rules. PCA spectrum of the ensemble of filters $\Lambda(T)$ for a restricted temperature window $0.7 < T < 3.7$ above the BKT point. Top-right inset: projection of the full ensemble on the two highest PCA components. The overlap with those “plaquette” filters falls with T .

by the pristine filters. This justifies describing the filters in intermediate regimes as “mixtures,” as suggested by the intuitive physical picture of the competing correlations. We can thus identify the relevant operators, i.e., the plaquette (electric charge) and staggered (electric field) filters while never seeing data from parameter regimes where they entirely dominate. This is important, as MC simulations may be costly, or we may be dealing with experimental data whose range we do not have full control over.

2. Discovering broken and emergent symmetries

Here we show how the different symmetries of the model, including emergent ones, can be observed in the statistical properties of the filter ensemble. To this end recall first (Sec. III B2) that a pair of plaquette filters, or a plaquette and a columnar one are degenerate in RSMI for $T < T_{\text{BKT}}$, and label uniquely the symmetry-broken states. This degeneracy is reflected in the equal frequency with which they appear as the optimal solutions in individual RSMI-NE runs. Likewise, the disappearance of the rotation-symmetry-breaking columnar filter from the ensemble above T_{BKT} signals the lifting of the columnar and plaquette degeneracy and restoration of the rotation symmetry. These observations can be pushed further.

To wit, though each individual instance of RSMI optimization returns an object which we can interpret as an operator, more generally there often exists an RSMI-degenerate subspace for the kernels of these operators, i.e., the filters Λ . In fact, the ensemble of many optimization instances (at the same physical parameters) allows us to map out this subspace, and the resulting distribution contains systematic information about symmetries. To show this, we project the ensemble of filters generated in the dimer model from 500 optimization runs for each temperature onto the orthogonal bases formed by the pristine staggered and plaquette filters introduced above.

We begin in the ordered phase. In Fig. 9(a), the distribution of the optimal filters [pairs of $\pm P_{1,2}$, or $(\pm C, \pm P_{1,2})$] is shown projected onto the plaquette filters, exhibiting four distinct peaks corresponding to $\pm P_{1,2}$, and a central peak

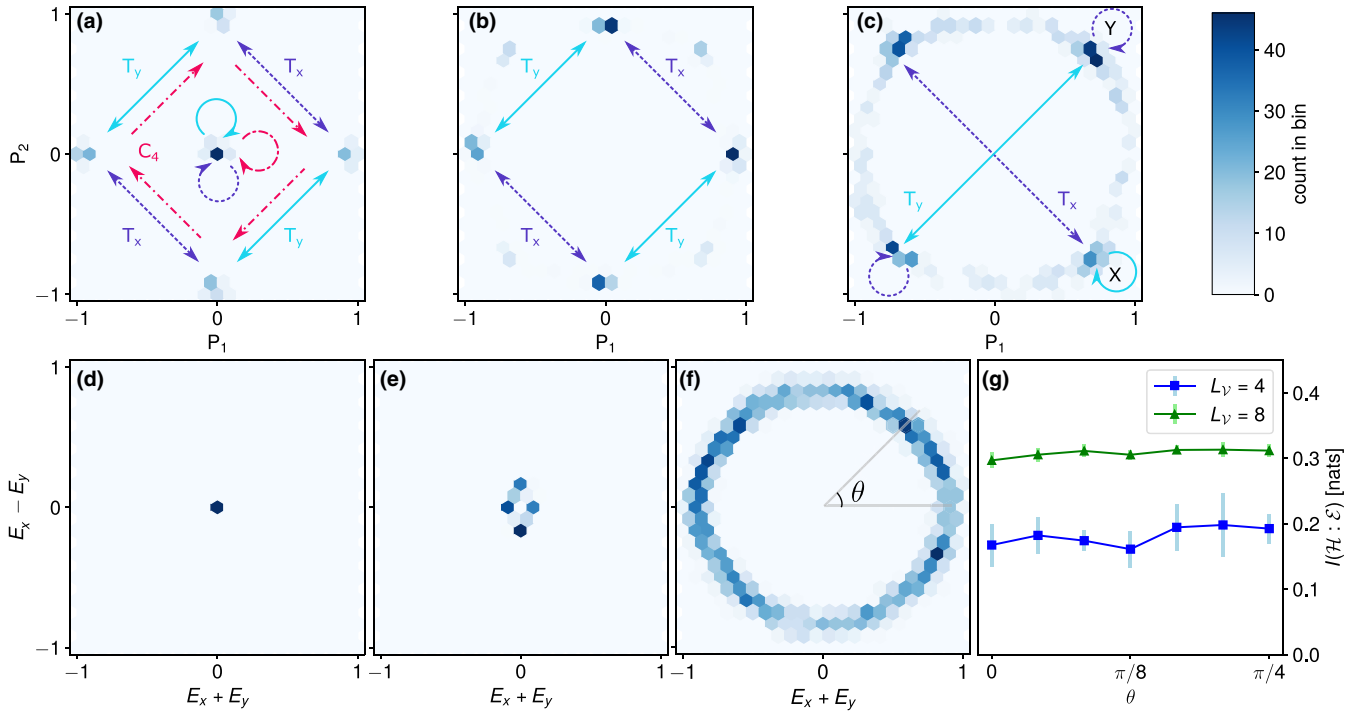


FIG. 9. Identifying symmetries in the ensemble of optimal coarse-grainings. (a)–(c) The projection of the ensemble of optimal filters onto the orthogonal basis (P_1, P_2) of *pristine* plaquette filters at temperatures $\{0.2 < T_{\text{BKT}}, 0.7 \gtrsim T_{\text{BKT}}, 1.5\}$, respectively. The induced action of the symmetries on the filter components are marked with red arrows (dashed-dotted line) for C_4 rotation, and with purple (dashed line) and light blue (solid line) for horizontal and vertical translations (T_x, T_y) , respectively. The central peak in panel (a) corresponds to the columnar subspace $\pm C$, closed under the action of both translation and C_4 . We see in panel (b) that there appear two independent representations of the translation symmetry, namely one with the basis (P_1, P_2) and also another with $(X, Y) := (P_1 - P_2, P_1 + P_2)$, see discussion in the text. For clarity, we show the $T_{x,y}$ action in these two representations separately in panels (b) and (c). (d)–(f) The projection onto the orthogonal basis $(E_x + E_y, E_x - E_y)$ at temperatures $\{0.2 < T_{\text{BKT}}, 0.7 \gtrsim T_{\text{BKT}}, 15.0\}$, respectively. (f) The ensemble for high- T is approximately invariant under *continuous* rotations in the space of filters. (g) The RSMI is constant as a function of the rotation angle of the electrical field. This degeneracy in RSMI value is the origin of the *emergent* $U(1)$ symmetric ensemble.

corresponding to $\pm C$. It can easily be checked that these form closed subspaces under the action of discrete lattice translations $T_{x,y}$ in x and y directions, and C_4 rotations. Indeed, since both families of filter pairs are bijections of the four broken-symmetry states, they give simultaneous representations of the two broken symmetries.

More concretely, the action of C_4 around a lattice vertex leads to a \mathbb{Z}_4 -cycle between the peaks, as shown by the red arrows (dashed-dotted line) in Fig. 9(a). Under $T_{x,y}$, on the other hand, $\pm P_{1,2}$ gives a representation of $\mathbb{Z}_2 \times \mathbb{Z}_2$, as shown by the light-blue (solid) and purple (dashed) arrows in Fig. 9(a). The columnar subspace $\pm C$ leads to a trivial representation as it is left invariant under both symmetries.

Above the BKT transition, a more complex picture arises [41,42]. Especially in finite systems, the plaquette correlations remain strong and though the C_4 symmetry is restored, the low-lying excitations consisting of exchanging horizontally and vertically aligned dimer pairs around plaquettes do not immediately restore the translation symmetry [41]. Since the two broken translation symmetry states differ by the “site-parity” [e.g., the second and third configurations in Fig. 4(b)], using one of the coarse-graining components to label the site-parity the filters can still recover at least one bit of long-range information (cf. Fig. 5) in a finite-sized system close to the transition point.

In particular, the pairs $(C, P_{1,2})$ become immediately sub-optimal above BKT [cf. the disappeared central peak in Fig. 9(b)], as they cannot produce a bijective labeling due to the restored C_4 symmetry. In contrast, the pairs consisting of $\pm P_{1,2}$ can still do the identification faithfully if T sufficiently close to T_{BKT} , so that \mathcal{V} contains at most one flipped plaquette. This is reflected in the persistence of the four $\pm P_{1,2}$ peaks in Fig. 9(b).

Note, however, the development of four additional peaks on the diagonals in Fig. 9(b). They are readily understood: the information-rich “site-parity” identification is insensitive to multiple plaquette flips if the filters $(X, Y) := [(P_1 - P_2), (P_1 + P_2)]^2$ are used instead. This is because they exclusively couple to horizontal or vertical dimers in each component, hence the notation.

As shown explicitly in Fig. 9(c), T_x (T_y) inverts the sign of X (Y), whereas T_y (T_x) stabilises it, resulting in a 1D representation of \mathbb{Z}_2 . Thus, slightly above the BKT transition there are two sets of peaks giving two independent representations

²This corresponds to an alternative (but equivalent) form $(\cos \varphi, \sin \varphi)$ of the \mathcal{O}_1 electric charge operator in the continuum sine-Gordon theory with the height field shifted by $\pi/4$ with respect to the expression in terms of $P_{1,2}$.

of the broken translation symmetry, corresponding to the two labelings above. As T is increased, the $\pm P_{1,2}$ peaks broaden and disappear; see Fig. 9(c).

For completeness, we note that (X, Y) is suboptimal in the ordered phase as it cannot deterministically map *both* the site-parity *and* the orientation into binary variables.³ Thus, these diagonal peaks are absent in Fig. 9(a).

Yet, still a further important question can be asked. The above discussion was about lattice symmetries, we know though that in the effective field theory of the dimer model this discrete C_4 symmetry is in fact enlarged to a full $U(1)$. Can the ensemble of the RSMI-optimal filters provide a hint of this *emergent symmetry*?

Surprisingly, the answer is affirmative. This is despite the fact that the filters are trained on the lattice, and even despite the imposed discrete-valuedness of the coarse-grainings. The emergent continuous symmetry is represented on the space of kernels Λ of the coarse-graining maps, and reflected in the ensemble.

To show this we project the filters in the ensemble onto the orthogonal basis of the pristine filters $(E_x + E_y, E_x - E_y)$. As shown in Fig. 9(f), at high T the kernel distribution of the electrical fields returned in individual RSMI runs (all of which exhibit a, at first glance “similar,” diagonal pattern) is indeed invariant under continuous rotations! In fact, as plotted in Fig. 9(g), the RSMI is constant as a function of the angle ϑ of the *rotated* electrical fields $R(\vartheta)(\Lambda_{S_1}, \Lambda_{S_2}) \sim [\cos(\vartheta)E_x - \sin(\vartheta)E_y, \sin(\vartheta)E_x + \cos(\vartheta)E_y]$, thus demonstrating the connection between the (emergent) operator symmetry and the RSMI-degeneracy. For temperatures below the BKT transition, the filters do not overlap with the electrical fields [Fig. 9(d)], while for those closely above [Fig. 9(e)] the continuous $U(1)$ symmetry is not yet emerged, due to the effect of the plaquette operators in this finite system [contrast with Figs. 9(b) and 9(c)].

The above analysis of the ensemble can be improved by e.g., using more sophisticated methods to disentangle the mixtures more data-efficiently. Nevertheless, we conclude that even simple data analysis allows us to extract the most important operators, and the symmetries (see also Ref. [45]).

Finally, we remark that the filter ensemble may be a very useful and natural notion in disordered systems, where the RSMI approach also applies [13], and the filters may depend on the quenched disorder realization. We leave this intriguing possibility to future work.

D. Coarse-graining filters: Type and number of components

In the previous sections we demonstrated among others that the RSMI-optimal filters label ordered states, and how, more generally, the symmetries of the system, broken or emergent, even continuous, are manifested on the space of filters Λ . An important question is whether and how these results depend on the constraints imposed on the image of Λ , i.e., the type of the coarse-grained variable \mathcal{H} .

³Explicitly, it maps the two horizontal (vertical) columnar configurations to tuples of probabilities (for getting +1 for the given component) $\Lambda \cdot \mathcal{V} = (1 \pm 1, 2)/2 \ ((2, 1 \pm 1)/2)$.




FIG. 10. Coarse-graining smears the discrete lattice height field into a sine-Gordon field. (a) The histogram of values taken by the height field $\varphi(\mathbf{r}_i) \bmod 2\pi$ (defined on vertices i) for T across the BKT transition; it is independent of temperature. (b), (c) The histogram of values of the block-averaged field $\langle \varphi \rangle_{\mathbf{r}_i \in \mathcal{V}}$, respectively, for 4×4 and 8×8 blocks (corresponding to $L_{\mathcal{V}} = 8, 16$ in dimer configurations). (d) The range of the continuous sine-Gordon field $\varphi(\mathbf{r})$.

There are, in fact, two questions: that of continuous- versus discrete-valued \mathcal{H} , and that of the number of \mathcal{H} components, which we address in this section. We begin with the former, in particular explaining the surprising emergence of a continuous symmetry in ensemble of filters (as in the high T dimers) even as \mathcal{H} is discrete.

1. Discrete and continuous variables

To this end, let us examine the role of coarse-graining (in a broad sense) in constructing a continuum low-energy effective field theory starting from a discrete lattice model, for which the dimer model is a good example.

On any bipartite planar lattice, a dimer configuration $C(\mathbf{r}_i)$ can be mapped onto a unique height profile $\varphi(\mathbf{r}_i)$ [43]. In 2D square lattice, the height at a given site is a four-state variable, invariant under a shift by 2π , giving it a natural interpretation as the orientation of the dimer connected to the given site. Equivalently, the distribution of $\varphi(\mathbf{r}_i)$ is uniform over four *discrete* values [see Fig. 10(a)].

The *continuum* effective theory of the interacting dimer model is given by a sine-Gordon (SG) action:

$$S[\varphi] = \int d^2\mathbf{r} \left\{ \frac{g(T)}{2} |\nabla\varphi(\mathbf{r})|^2 + \cos[4\varphi(\mathbf{r})] \right\}, \quad (20)$$

where the lattice height $\varphi(\mathbf{r}_i) \in \mathbb{Z}_4$ is continued into a *real-valued* $\varphi(\mathbf{r}) \in [0, 2\pi)$. The first term in Eq. (20) is the energy density associated to the electrical field $\mathbf{E} := \nabla\varphi$, which accounts for the entropy of the dimers and dominates at high T . The cosine *locking* potential orders $\varphi(\mathbf{r})$ into a flat profile with one of the four values $\{\pi/4, 3\pi/4, 5\pi/4, 7\pi/4\}$, corresponding to the four ground states of the lattice model. In other words, the BKT transition happens precisely at the temperature at which the range of $\varphi(\mathbf{r})$ fragments into four sectors, which eventually collapse into discrete delta-peaks at $T = 0$, as shown in Fig. 10(d).

Numerically, the lattice and the SG pictures are bridged by mapping the discrete lattice heights onto *block-averaged* values $\langle \varphi \rangle_{\mathbf{r}_i \in \mathcal{V}}$ on the coarse grid. As seen in Figs. 10(b) and 10(c), this simple procedure indeed smears out the discrete heights into a quasi-continuous variable of the same

character as the SG field. The larger block size L_V , the closer the coarse-grained variable to a continuous SG field.

This observation is crucial in clarifying the role of discretization in RSMI-NE, and sheds light on the scale L_V . Particularly, for the case of free dimers the RSMI filter computes the block average of the electric field:

$$\Lambda_\alpha \cdot \mathcal{V} = \sum_i \Lambda_\alpha^i \mathcal{V}_i = \sum_I \underbrace{\sum_{i \in I} \Lambda_\alpha^i \mathcal{V}_i}_{=: E_\alpha^I} = \langle E_\alpha \rangle_{\mathcal{V}}, \quad (21)$$

where i iterates over bonds, and I over disjoint smallest unit blocks (here this is $L_V = 4$) on which the electric field can be defined; α denotes the electric field (or filter) component. On the other hand, in terms of the lattice height field, this means that (see Supplemental Material of Ref. [17]):

$$\Lambda_\alpha \cdot \mathcal{V} = \langle \partial_\alpha \varphi \rangle_{\mathcal{V}} = \partial_\alpha \langle \varphi \rangle_{\mathcal{V}}, \quad (22)$$

where ∂_α is the lattice gradient. In other words, the inner product computes the electrical field averages using the block-averaged heights (e.g., using $L_V = 8$, as we had in the discussion of the dimer model), which approximate the continuous ones of the SG theory. This is the reason why the emergent continuous $U(1)$ symmetry can manifest itself in the filter ensemble when the gradient part of the SG action dominates (in the finite system as $T \rightarrow \infty$), despite the discrete input data.

Note that composing $\Lambda_\alpha \cdot \mathcal{V}$ with a binary discretization map τ does not interfere with the $U(1)$ symmetry in the filter ensemble. We emphasize the same filters are obtained if no discretization is assumed. Discretization becomes important, though, when the support of φ starts to get locked into four values at low T . The information to be retained precisely comes from identifying these four peaks, and τ increases the efficiency of finding the filters bijectively mapping the symmetry-broken states by dramatically restricting the search space. In this sense it acts as an *a posteriori* entropy cutoff and regularization.

We thus conclude that discrete-valuedness of coarse-graining maps does not interfere with or preclude continuous symmetries manifesting, and further serves as a regularization, which is very useful when the number of configuration samples is limited.

2. Number of components

Next, we explain how the necessary number of hidden variable components can be *discovered* systematically, providing information about the system. This particularly will justify using a single component variable in the Ising example, and two for the dimer model.

The essence of the RSMI-NE approach is the efficient compression of the long-range information. Any compression method contains a trade-off (explicit or implicit) between the compression rate, e.g., given by the total number of bits retained, and the preservation of relevant information. Ideally one should compress to preserve just sufficiently enough information relevant for the downstream task the compressed representation is used in—but not more than that. In RSMI-NE the compression rate is effectively controlled by the number and the alphabet of allowed values of the

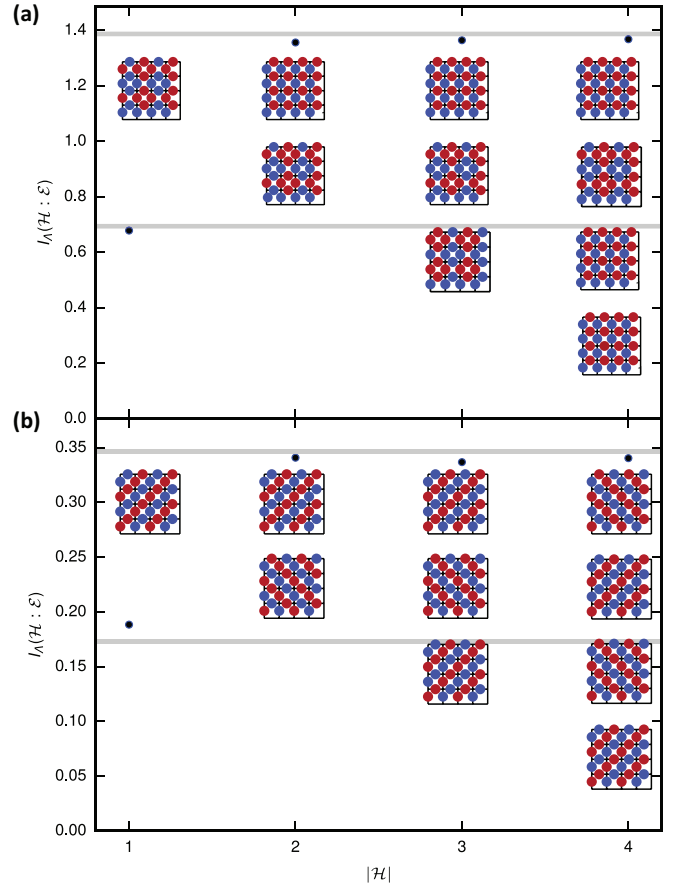


FIG. 11. The maximal RSMI versus number of components of \mathcal{H} in the dimer model. (a) For $T < T_{\text{BKT}}$, the only long-range information is about the type of ground-state the system crystallizes into. Using two binary components for \mathcal{H} suffices to encode this information and adding further components does not improve this result, as reflected by the value of $I_\Lambda(\mathcal{H} : \mathcal{E})$ attained. (b) At $T \rightarrow \infty$, two electric field components recover the maximal long-range information $\approx \frac{1}{2} \log 2$ (for $L_B = 4$). This is not improved by additional components, which converge to filters linearly dependent on the first two.

coarse degrees of freedom \mathcal{H} (for example, $\{\pm 1\}$, $\{\uparrow, \downarrow\}$, $\{|0\rangle, |1\rangle, |2\rangle\}$, etc.). The relevant long-range information, on the other hand, is a property of the physical system, which we do not have control over.

Since *a priori* we do not know (though in practice we may often anticipate) what the amount of long-range information in the system is, a simple practical procedure to determine the optimal compression rate is to find the maximal number of compressed variables $|\mathcal{H}|$, above which the retained RSMI does not further increase significantly. Figure 11 shows that this happens at $|\mathcal{H}| = 2$ for the dimer model both in the high- T limit, and the columnar ordered phase at low T . In both cases we use a fixed buffer width of $L_B = 4$. We find that using a single binary component leads to only half the RSMI value attained with $|\mathcal{H}| = 2$, and is obtained by an optimal filter equal to one of the components of the two-component rule. For $|\mathcal{H}| > 2$, the RSMI saturates into either $\frac{1}{2} \log 2$ at $T \rightarrow \infty$ or $\log 4$ at $T < T_{\text{BKT}}$. Thus, we verify that at most two of the components are linearly independent, and additional

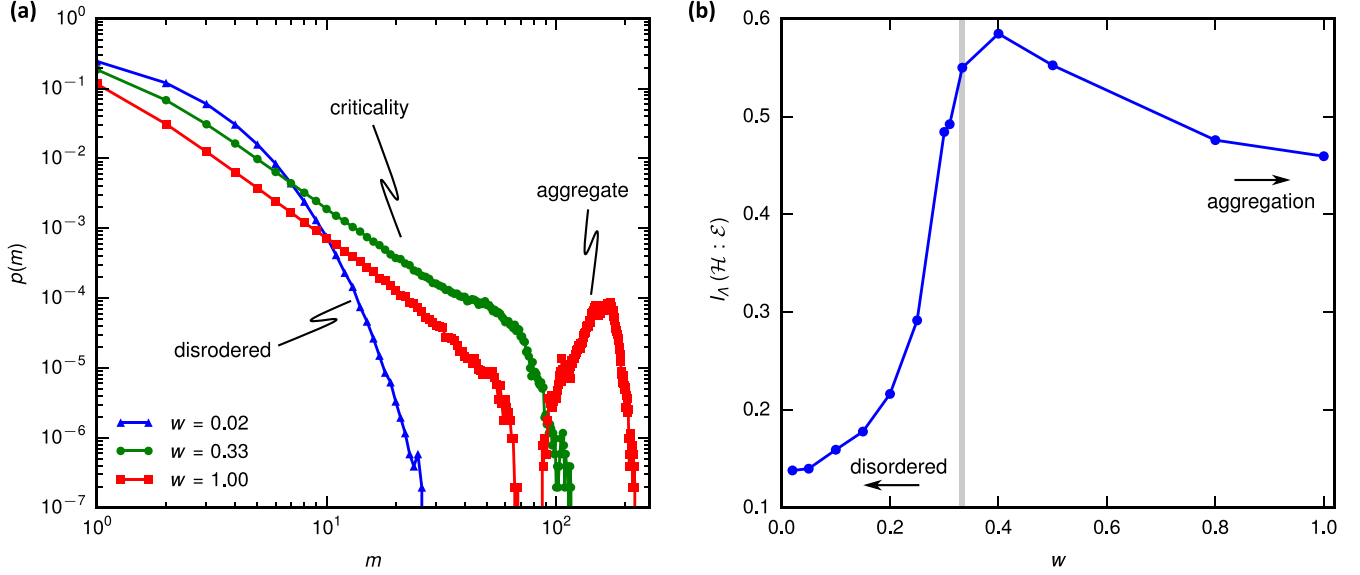


FIG. 12. Phase transition in the chipping and aggregation model. (a) Observing the decay profile of the marginal mass distribution for a given site at different values of w is one way of qualitatively assessing the different phases of the chipping model. (b) By providing merely the real-space sample configurations, the peaking maximal value of RSMI signals the nonequilibrium critical point. Furthermore, the aggregated and low- w phases can be distinguished by the distinct saturated values.

filters do not extract distinct information from \mathcal{V} . In the 2D Ising model example, in contrast, only a single filter suffices (and it corresponds to magnetization). The physical intuition behind the above procedure is clear: it finds the number of relevant operators whose correlations explain the total information shared between distant parts of the system.

Finally, we note that while the variational *ansatz* used for the coarse-graining was a shallow network dotted into the configurations (before the nonlinear Gumbel-softmax step), more general or deeper network architectures can be considered. RSMI maximization can be performed over any class of variational functions. For specific systems and inputs a more expressive *ansatz* could possibly recover larger RSMI. In more abstract terms, since the coarse-graining rules are related to the RG-relevant operators [17], such choices would be able to extract operators which cannot be written as linear functions of the local degrees of freedom, should these be important. Though for such complex multilayered architectures the patterns of the weights themselves may not be directly interpretable, the extracted Λ filters can still be used as operators, as we have done in computing the correlation

functions in Ref. [17]. We again emphasize that the physical interpretability is not a consequence of the architectural choices but of the physical nature of the RSMI quantity [14].

IV. POSSIBLE EXTENSION TO NONEQUILIBRIUM: A MODEL WITH CHIPPING AND AGGREGATION

The RSMI-NE algorithm we described does not in any way use or rely on the existence of a Hamiltonian generating the probability distribution. It can thus be directly applied to general nonequilibrium distributions, though the formal understanding of the optimal filters in this situation, analogous to results of Refs. [14,17], is currently missing. This is an exciting research direction, whose development we leave to future work. Here, however, we provide a short validation of the idea.

To this end we consider the nonequilibrium example of the 1D chipping and aggregation model of Ref. [46]. Its stochastic dynamics is defined by the update rules given below. At any time increment Δt , masses m_i on site i in the chain are modified according to the following moves:

$$\text{chipping, at rate } p = \Delta t : \begin{cases} \text{if } m_i > 0 & : m_i \mapsto m_i - 1, \quad m_{i\pm 1} \mapsto m_{i\pm 1} + 1, \\ \text{if } m_i = 0 & : \text{do nothing,} \end{cases} \quad (23)$$

$$\text{aggregation, at rate } p = w\Delta t : m_i \mapsto 0, \quad m_{i\pm 1} \mapsto m_{i\pm 1} + m_i. \quad (24)$$

We consider the case with mass density $\rho = 1$, so that $\sum_i m_i = L$, where L is the length of the chain for which we impose periodic boundary conditions.

This system undergoes a nonequilibrium transition between phases, whose qualitative difference is reflected in the marginal probability distribution $p(m)$ of masses at each site,

as shown in Fig. 12(a). For low aggregation w , the distribution $p(m)$ is an exponentially decaying function with increasing mass per site. At high w , a macroscopic proportion of the masses aggregate (or condense) to a single site, corresponding to a delta-function peak around $m \approx L$, with an algebraically decaying part for the remainder of the masses. Criticality of

the intermediate w region is reflected by a purely algebraically decaying distribution $p(m)$.

The shape of the marginal distribution $p(m)$ clearly demonstrates the different phases and the transition between them. Instead of, however, investigating such specific quantities, which requires at least an intuitive understanding of the physics of the system, one can generically identify the nonequilibrium critical point using RSMI-NE directly applied to the full real-space configurations. As shown in Fig. 12(b) the dependence of the optimal-RSMI at w readily points to a transition between the phases with different spatial correlation characteristics. We performed RSMI optimization on the model with $L = 256$, for a range of values of w , and used a buffer of width $L_B = 8$. While at all values of w we found that the optimal coarse-graining rule averages the mass for the region \mathcal{V} , the value of the maximal information saturates into different values in the low- w and the aggregated phases. Moreover, the transition point is marked by a slight peak in the RSMI.

These results demonstrate the in principle applicability of the RSMI maximization to nonequilibrium problems. While this example reaches the steady-state distribution very quickly, in a more general scenario involving far-from-equilibrium systems the formalism can be extended to screen out short-time correlations by introducing a buffer in the temporal direction. This line of research merits a full development, which we leave for future study.

V. CONCLUSIONS AND OUTLOOK

In this work, accompanying and extending Ref. [17], we demonstrate how recent rigorous results in ML-based estimation of information theoretic quantities [15,16] can be combined with other algorithmic ingredients [33,34] to yield a formally interpretable and numerically efficient algorithm extracting information about long-range properties of statistical systems from its raw configurational samples. We dub this unsupervised algorithm RSMI-NE, or the real-space mutual information neural estimator, from the key physical quantity of interest [12–14], and we provide a detailed introduction to the method and background concepts, as well as an examination of its properties.

The optimal real-space coarse-graining filters extracted depend on the parameters of the physical system and in fact are lattice representations of its most relevant operators in the sense of renormalization group (RG) [17]. The filters, along with the RSMI value, and its dependence of systems' parameters and length-scales, provide a comprehensive physical picture. Particularly, the position and nature of the critical points, decay and structure of the spatial correlations, together with the type of order is revealed by these quantities, which we demonstrate explicitly in equilibrium examples. Though we focused on statistical models on regular lattices, the algorithm can be applied to continuum or graph models as well.

We further introduce the notion and examine the properties of the *ensemble of the optimal coarse-graining filters*. We show that it contains important physical information, most

notably about the symmetries, including *emergent* ones. The ensemble can be the object of statistical analysis itself, allowing us to extract the relevant operators from partially complete data, or from restricted parameter regimes.

We also examined and validated the possibility of extending the applicability of the algorithm to nonequilibrium systems on the example the chipping model with aggregation [46], for which the presence and position of a nonequilibrium phase transition was detected.

Motivated by the above example, the full extension of the framework to the case of nonequilibrium systems is among the most promising and exciting future research directions. This requires investigating the RSMI approach with spatiotemporal coarse-graining rules, and the information shared between temporally ordered states [47] and extending the theoretical results of Ref. [14]. The existing algorithm does, however, already allow the investigation of spatial correlations in complex real-world systems, such as, e.g., meteorological precipitation data exhibiting critical points, possibly related to self-organized criticality [48].

Previous formal results [13] and the numerical possibility of individually optimizing the coarse-graining rules for each block invite the application of RSMI-NE to disordered systems. In particular, the ensemble of coarse-graining rules would be inherited naturally from the disorder distribution and the further statistical analysis may identify certain equivalence classes within this ensemble. This would be especially helpful in identifying the relevant DOFs in these challenging systems.

Finally, we emphasize that the RSMI-NE provides an important step toward the goal of automating certain aspects of theory building. The constructed outputs (the coarse-graining transformations) are effectively black-box algorithmic objects, which, however, can be assigned the formal identity of order parameters or scaling operators of the physical theory. They can explicitly be used as such to compute correlations functions or scaling exponents, as shown in the companion work Ref. [17]. The above results clearly invite further work in this direction.

ACKNOWLEDGMENTS

M.K.-J. is grateful to F. Alet for his comments on the physics of the interacting dimer model. We acknowledge insights into the physics of the chipping and aggregation model coming from an earlier study of this system together with R. Thomale. D.E.G., S.D.H., and M.K.-J. gratefully acknowledge financial support from the Swiss National Science Foundation and the NCCR QSIT, and the European Research Council under the Grant Agreement No. 771503 (TopMechMat), as well as from European Union's Horizon 2020 programme under Marie Skłodowska-Curie Grant Agreement No. 896004 (COMPLEX ML). Z.R. acknowledges support from ISF Grant No. 2250/19. Some of the computations were performed using the Leonhard cluster at ETH Zurich. This work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under Project ID No. eth5b.

- [1] J. D. Schieber and M. Hütter, Multiscale modeling beyond equilibrium, *Phys. Today* **73**, 36 (2020).
- [2] L. P. Kadanoff, Scaling laws for Ising models near T_c , *Phys. Phys. Fiz.* **2**, 263 (1966).
- [3] K. G. Wilson and J. Kogut, The renormalization group and the ϵ expansion, *Phys. Rep.* **12**, 75 (1974).
- [4] K. G. Wilson, The renormalization group: Critical phenomena and the Kondo problem, *Rev. Mod. Phys.* **47**, 773 (1975).
- [5] M. E. Fisher, Renormalization group theory: Its basis and formulation in statistical physics, *Rev. Mod. Phys.* **70**, 653 (1998).
- [6] K. G. Wilson, Renormalization group and critical phenomena. II. Phase-space cell analysis of critical behavior, *Phys. Rev. B* **4**, 3184 (1971).
- [7] D. J. Gross and F. Wilczek, Ultraviolet Behavior of Non-Abelian Gauge Theories, *Phys. Rev. Lett.* **30**, 1343 (1973).
- [8] D. S. Fisher, Random Transverse Field Ising Spin Chains, *Phys. Rev. Lett.* **69**, 534 (1992).
- [9] D. S. Fisher, Critical behavior of random transverse-field Ising spin chains, *Phys. Rev. B* **51**, 6411 (1995).
- [10] D. S. Fisher, Random antiferromagnetic quantum spin chains, *Phys. Rev. B* **50**, 3799 (1994).
- [11] F. Iglói and C. Monthus, Strong disorder RG approach of random systems, *Phys. Rep.* **412**, 277 (2005).
- [12] M. Koch-Janusz and Z. Ringel, Mutual information, neural networks, and the renormalization group, *Nat. Phys.* **14**, 578 (2018).
- [13] P. M. Lenggenhager, D. E. Gökmen, Z. Ringel, S. D. Huber, and M. Koch-Janusz, Optimal Renormalization Group Transformation from Information Theory, *Phys. Rev. X* **10**, 011037 (2020).
- [14] A. Gordon, A. Banerjee, M. Koch-Janusz, and Z. Ringel, Relevance in the Renormalization Group and in Information Theory, *Phys. Rev. Lett.* **126**, 240601 (2021).
- [15] B. Poole, S. Ozair, A. van den Oord, A. Alemi, and G. Tucker, On variational bounds of mutual information, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri, and R. Salakhutdinov, Proceedings of Machine Learning Research Vol. 97 (PMLR, 2019), pp. 5171–5180.
- [16] M. I. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, A. Courville, and D. Hjelm, Mutual information neural estimation, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause, Proceedings of Machine Learning Research Vol. 80 (PMLR, 2018), pp. 531–540.
- [17] D. E. Gökmen, Z. Ringel, S. D. Huber, and M. Koch-Janusz, Statistical physics through the lens of real-space mutual information, *Phys. Rev. Lett.* **127**, 240603 (2021).
- [18] Source code for the RSMI-NE is available online at <https://github.com/RSMI-NE/RSMI-NE>.
- [19] A. C. D. van Enter, R. Fernández, and A. D. Sokal, Regularity properties and pathologies of position-space renormalization-group transformations: Scope and limitations of Gibbsian theory, *J. Stat. Phys.* **72**, 879 (1993).
- [20] A. Kraskov, H. Stögbauer, and P. Grassberger, Estimating mutual information, *Phys. Rev. E* **69**, 066138 (2004).
- [21] D. H. Wolpert and D. R. Wolf, Estimating functions of probability distributions from a finite set of samples, *Phys. Rev. E* **52**, 6841 (1995).
- [22] M. D. Donsker and S. R. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time. IV, *Commun. Pure Appl. Math.* **36**, 183 (1983).
- [23] D. Barber and F. V. Agakov, Information maximization in noisy channels: A variational approach, in *Advances in Neural Information Processing Systems 16*, edited by S. Thrun, L. K. Saul, and B. Schölkopf (MIT Press, Cambridge, MA, 2004), pp. 201–208.
- [24] X. Nguyen, M. J. Wainwright, and M. I. Jordan, Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization, in *Advances in Neural Information Processing Systems 20*, edited by J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis (MIT Press, Cambridge, MA, 2008), pp. 1089–1096.
- [25] S. Nowozin, B. Cseke, and R. Tomioka, f-GAN: Training generative neural samplers using variational divergence minimization, in *Advances in Neural Information Processing Systems*, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, New York, 2016), Vol. 29.
- [26] A. Nir, E. Sela, R. Beck, and Y. Bar-Sinai, Machine-learning iterative calculation of entropy for physical systems, *Proc. Natl. Acad. Sci. USA* **117**, 30234 (2020).
- [27] A. van den Oord, Y. Li, and O. Vinyals, Representation learning with contrastive predictive coding, [arXiv:1807.03748](https://arxiv.org/abs/1807.03748).
- [28] M. Gutmann, and A. Hyvärinen, in *Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models*, Proceedings of Machine Learning Research Vol. 9 (PMLR, 2010), pp. 297–304.
- [29] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevE.104.064106> for technical details of simulation execution and data generation.
- [30] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [31] P. Y. Lu, S. Kim, and M. Soljačić, Extracting Interpretable Physical Parameters from Spatiotemporal Systems Using Unsupervised Learning, *Phys. Rev. X* **10**, 031056 (2020).
- [32] E. Jang, S. Gu, and B. Poole, Categorical reparameterization with Gumbel-softmax, [arXiv:1611.01144](https://arxiv.org/abs/1611.01144).
- [33] C. J. Maddison, D. Tarlow, and T. Minka, A* sampling, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, New York, 2014), pp. 3086–3094.
- [34] E. J. Gumbel, The maxima of the mean largest value and of the range, *Ann. Math. Stat.* **25**, 76 (1954).
- [35] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [36] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborova, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [37] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys.:* **X** **5**, 1797528 (2020).
- [38] L. Onsager, Crystal statistics. I. A two-dimensional model with an order-disorder transition, *Phys. Rev.* **65**, 117 (1944).
- [39] J. Wilms, M. Troyer, and F. Verstraete, Mutual information in classical spin models, *J. Stat. Mech.* (2011) P10011.
- [40] H. W. Lau and P. Grassberger, Information theoretic aspects of the two-dimensional Ising model, *Phys. Rev. E* **87**, 022128 (2013).
- [41] F. Alet, Y. Ikhlef, J. L. Jacobsen, G. Misguich, and V. Pasquier,

- Classical dimers with aligning interactions on the square lattice, [Phys. Rev. E **74**, 041124 \(2006\)](#).
- [42] F. Alet, J. L. Jacobsen, G. Misguich, V. Pasquier, F. Mila, and M. Troyer, Interacting Classical Dimers on the Square Lattice, [Phys. Rev. Lett. **94**, 235702 \(2005\)](#).
- [43] E. Fradkin, *Field Theories of Condensed Matter Physics*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2013).
- [44] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac, Area Laws in Quantum Systems: Mutual Information and Correlations, [Phys. Rev. Lett. **100**, 070502 \(2008\)](#).
- [45] R. Bondesan and A. Lamacraft, Learning symmetries of classical integrable systems, [arXiv:1906.04645](#).
- [46] R. Rajesh and S. N. Majumdar, Exact phase diagram of a model with aggregation and chipping, [Phys. Rev. E **63**, 036114 \(2001\)](#).
- [47] F. Creutzig, A. Globerson, and N. Tishby, Past-future information bottleneck in dynamical systems, [Phys. Rev. E **79**, 041925 \(2009\)](#).
- [48] O. Peters and J. D. Neelin, Critical phenomena in atmospheric precipitation, [Nat. Phys. **2**, 393 \(2006\)](#).