

Quantitative and interpretable order parameters for phase transitions from persistent homology

Alex Cole*

*GRAPPA and ITFA, Institute of Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, the Netherlands*Gregory J. Loges[†] and Gary Shiu[‡]*Department of Physics, University of Wisconsin—Madison, Madison, Wisconsin 53706, USA*

(Received 25 November 2020; revised 29 July 2021; accepted 7 September 2021; published 28 September 2021)

We apply modern methods in computational topology to the task of discovering and characterizing phase transitions. As illustrations, we apply our method to four two-dimensional lattice spin models: the Ising, square ice, XY, and fully frustrated XY models. In particular, we use persistent homology, which computes the births and deaths of individual topological features as a coarse-graining scale or sublevel threshold is increased, to summarize multiscale and high-point correlations in a spin configuration. We employ vector representations of this information called persistence images to formulate and perform the statistical task of distinguishing phases. For the models we consider, a simple logistic regression on these images is sufficient to identify the phase transition. Interpretable order parameters are then read from the weights of the regression. This method suffices to identify magnetization, frustration, and vortex-antivortex structure as relevant features for phase transitions in our models. We also define “persistence” critical exponents and study how they are related to those critical exponents usually considered.

DOI: [10.1103/PhysRevB.104.104426](https://doi.org/10.1103/PhysRevB.104.104426)**I. INTRODUCTION**

Given an unknown condensed matter system sitting in front of you, the zeroth order question you may ask is as follows: What is its phase structure? With sufficient technical ability, one may vary the various coupling constants, external temperature, etc., and measure its ensuing equilibrium configurations. One way to understand the phase structure is to carefully search through the entire parameter space and deduce for which parameter regimes the system looks similar (i.e., the system remains in the same phase). In doing so, one may occasionally encounter boundaries where some symmetry is broken or some specific heat diverges, indicating a new phase. Having identified these phases, a natural next question is how to distinguish them in practice, i.e., what order parameters describe the various phase transitions. These questions are naturally phrased in the language of machine learning (ML). Namely, the question “How many phases are there?” is an exercise in *unsupervised learning*, while the question “How are different phases distinguished?” is an exercise in *supervised learning*. Note that this is an exercise in distinguishing statistical ensembles and incurs some amount of uncertainty.

Recently, ML techniques have been applied to these very tasks. Unsupervised methods such as principal component analysis, clustering algorithms, and autoencoders have been used to identify phase transitions (see, e.g., Refs. [1–8]).

Support vector machines have been shown to be a useful tool in quantifying characteristics of phase transitions [9–11]. Supervised learning with neural networks has proven useful in this classification task (see, e.g., Refs. [12–20]) but often lacks the desired level of interpretability.

In this paper we use *persistent homology* [21,22] (see Refs. [23–25] for reviews) as a tool for detecting and characterizing phase transitions using a supervised learning approach, although is amenable to unsupervised learning as well. As illustrations, we apply our method to study two-dimensional lattice spin systems. Persistent homology is a technique from topological data analysis (TDA) that identifies the births and deaths of topological features throughout a family of discrete complexes. This family often corresponds to the data set at various coarse-graining scales. By now, persistent homology has been fruitfully applied in a wide variety of fields, including sensor networks [26], image processing [27], genomics [28], protein structure [29,30], neuroscience [31,32], cosmology [33,34], and string theory [35,36], to name only a few. In the context of spin systems, persistent homology encodes multiscale and high-order correlations in a data set. The main takeaway from our work is that this representation of a spin system configuration is not only sufficient to distinguish phases in spin systems but additionally provides interpretable order parameters for the phase transitions. For example, we find that persistent homology identifies such varied phenomena as magnetization, frustration, and (anti)vortices in spin systems. Additionally, as a multiscale technique, persistent homology can capture a system’s approach toward scale invariance, i.e., its critical behavior. We work with *persistence images* [37], which are vectorized representations of persistent homology informa-

* a.e.cole@uva.nl

† gloges@wisc.edu

‡ shiu@physics.wisc.edu

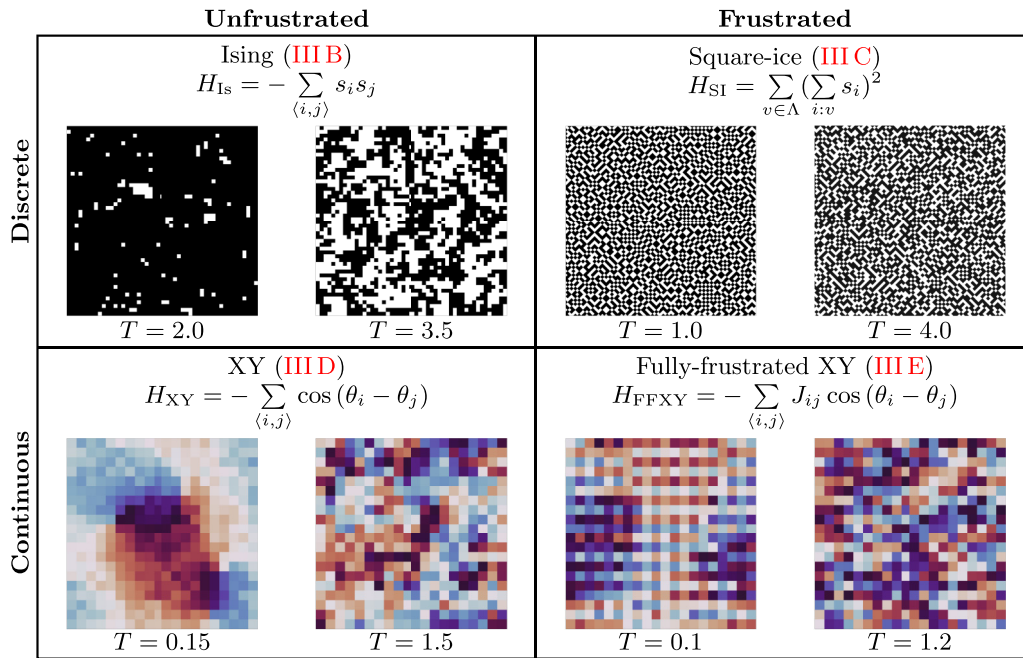


FIG. 1. Overview of the considered models. Example spin configurations in the low- and high-temperature phases are shown for each model.

tion. This framework allows us to define quantitative order parameters and quantify the uncertainty that a particular spin configuration belongs to a particular phase.

Persistent topological methods have been applied to statistical mechanics in a few cases, but so far these applications have been largely qualitative in terms of statistics. Reference [38] studied the relationship between phase transitions and topology changes in configuration space. More recently, Ref. [39] studied the relaxation dynamics of a two-dimensional Bose gas with persistent homology, and Ref. [40] performed unsupervised learning on persistence diagrams to visualize their phase structure. Reference [41] explores the properties of a lattice model by computing pairwise distances under a particular metric on the persistence diagrams and visualizes the phase diagram via a dimensional reduction. However, the use of persistent topological methods in obtaining *quantitative* information about statistical mechanics systems is in our view an underdeveloped subject. The purpose of this paper is to provide an early foray in this important direction.

In this manuscript we use persistent homology to quantitatively characterize phase transitions in four different lattice spin systems. We consider discrete and continuous spin models with and without frustration in the ground state (see Fig. 1). Each example contains a distinct lesson. We begin with an obligatory analysis of the two-dimensional Ising model (Is). We are able to easily identify the model’s phase transition relying only on training data far from the critical temperature. The magnetization as order parameter is immediately extracted from the weights of the corresponding logistic regression. Additionally, we examine the multiscale nature of the information probed by persistence images. In particular, we define “persistence” critical exponents that capture the model’s approach toward criticality, finding interesting

connections to the critical exponents usually considered. We then turn to the square-ice model (SI), for which there is no local order parameter due to frustrated low-energy dynamics. We again find a successful classification and are able to identify an order parameter associated with the low-temperature phase’s “scale of frustration.” Our technique quickly picks up on the scale of this feature. We subsequently turn to continuous spin models, beginning with the XY model, where a simple logistic regression on the persistence images discovers the Kosterlitz-Thouless (KT) phase transition and corresponding vortex-antivortex structure in the low-temperature phase. Vortex-antivortex pairs are shown to give a distinctive signature in the persistence images that the logistic regression discovers and decides to use on its own. Finally, we consider the fully frustrated XY model (FFXY), where frustration prevents the formation of (anti)vortices in the low-temperature phase. In this case, our method identifies small scale correlations between next-to-nearest neighbors that reflect the system’s attempt to satisfy competing constraints.

An important feature of our analysis is the simplicity of our machine learning architecture. Once the relevant spin configurations are reduced to persistence images, the phase classification and extraction of order parameters can be achieved via a simple logistic regression. This reflects the fact that persistent homology condenses these data sets into their most relevant (and interpretable) features.

The code and data used in our analysis are made available in Ref. [42]. The organization of this manuscript is as follows. In Sec. II we give a brief introduction to persistent homology, persistence images, and our computational choices for applying these techniques to spin models. In Sec. III we apply our methods to spin models of increasing complexity. We conclude in Sec. IV.

II. PERSISTENT HOMOLOGY AND PERSISTENCE IMAGES

We are interested in developing general interpretable order parameters for phase transitions in spin systems. Some inspiration can be drawn from the hallowed two-dimensional (ferromagnetic) Ising model. In this case, spontaneous magnetization in the low-temperature phase leads to large, continuous domains where all spins are aligned. As the temperature is decreased toward $T = 0$, these domains grow, so that at sufficiently low temperature, the entire system is aligned. On the other hand, in the high-temperature phase, spins receive enough thermal energy to be randomly oriented. In the language of ML, the (non)existence and scale of magnetic domains manifests as a pattern in the hierarchical clustering of aligned spins. In other words, if we consider the set of aligned spins and perform successive coarse-graining transformations, then we would be able to distinguish these two phases by the number of domains at different coarse-graining scales. Note that this is a *multiscale* concept that probes high-order correlation functions.

In fact, clustering can be viewed as the most basic topological information about a data set, giving the total number of “connected components.” We may then consider the hierarchical (i.e., multiscale) topologies corresponding to higher-dimensional features as well, for example loops. A unified description of topological features of all dimensions is given by algebraic topology, and the hierarchical or multiscale version of algebraic topology is persistent homology [21–25].

We now give a brief description of simplicial homology, referring the reader to Refs. [23,24] for details. We begin by embedding our data in a discrete complex. We use both simplicial and cubical complexes in this work. In a simplicial complex, points (0-simplices) may be connected in pairs by edges (1-simplices), in triples by triangular faces (2-simplices), and so on. Simplicial complexes must be closed under taking faces: For example, if a 2-simplex is in the complex, then so, too, must be its three edges and three vertices. A cubical complex is similar, but it consists of points (0-cubes), line segments (1-cubes), squares (2-cubes), and so on. Topological aspects of the simplicial or cubical complex are then captured by its homology groups. These groups, denoted H_p ($p = 0, 1, 2, \dots$), consist of equivalence classes of p -cycles, where two p -cycles are in the same equivalence class if they can be smoothly deformed into one another. H_0 consists of connected components, H_1 consists of noncontractible loops, and so on, with the Betti numbers b_p giving the number of inequivalent, nontrivial p -cycles.

The core insight of *persistent* homology is that such a procedure can be significantly enhanced in its stability and information content if instead of a single complex, a monotonically growing family, called a *filtration*, is considered. We will generally use ν to denote the scale parametrizing the filtration. Often the growing of the filtration corresponds to the increasing of a coarse-graining scale, so that multiscale information is captured. See Fig. 2. As this coarse-graining scale increases, p -cycles are created (for example, loops form) and destroyed (for example, loops are “filled in”). The mathematics of persistent homology allows us to track the births and deaths of *individual* topological features, where *birth*

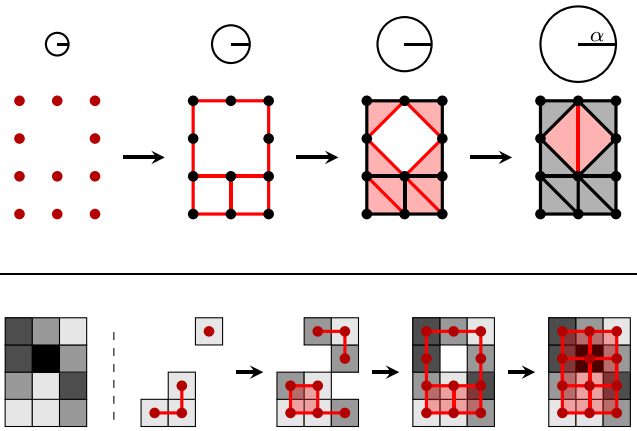


FIG. 2. Top: Four steps in the α filtration for a grid of points, such as appears in our discrete-spin models. The filtration parameter when a p -simplex is included is α^2 , where α is the radius of the simplex’s circumsphere. The α -complexes are pictured in black/red, with the most recently added p -simplices being shown in red. The nontrivial 1-cycle around the “gap” in the grid of points is born in the second step and dies in the last step. Bottom: Four steps in the sublevel filtration for a scalar function (represented by grayscale, with darker squares corresponding to larger values) defined on a 4×3 grid, such as appears in our continuous-spin models. The filtration parameter is the threshold, ν , for the sublevel sets and the cubical complexes are pictured in red. A nontrivial 1-cycle around the largest value is born and then dies in the last two steps.

refers to the value of the filtration parameter when the cycle appears and similarly for death. This information is usually summarized via a *persistence diagram* (see Fig. 3), which is a scatter plot of these births and deaths.

While persistence diagrams are often suitable for visualization, they are not very well suited for statistical analysis. In the end, we are interested in the statistical task of quantifying the probability that a given spin configuration belongs to a particular phase of the system. Therefore, rather than scatter plots, we might prefer a summary statistic that lives in a vector space. These also aid us in quantitatively characterizing the change in the system’s persistent homology as some parameter is varied. We therefore use persistence images for our analysis, which are formed by appropriately smoothing the persistence diagram and binning so as to have a low(er)-dimensional representation of the persistence data. That is, for a persistence diagram consisting of a number of points $\{(b_k, d_k)\}$ one chooses a number of bins $\{\text{bin}_i\}$ and forms

$$PI_i = \int_{\text{bin}_i} db dp \sum_k \frac{w(p_k)}{2\pi\sigma^2} \exp \left[-\frac{(b - b_k)^2 + (p - p_k)^2}{2\sigma^2} \right], \tag{1}$$

where the sum in k runs over all points in the persistence diagram, and the *persistence*, p_k , for each point is its “lifetime,” $d_k - b_k$. The weight $w(p)$ should be chosen to vanish at zero persistence in order to highlight those more important features which are longer lived. In what follows we use $w(p) = \log(1 + p)$. (See Ref. [37] for more details on the stability properties of persistence images.)

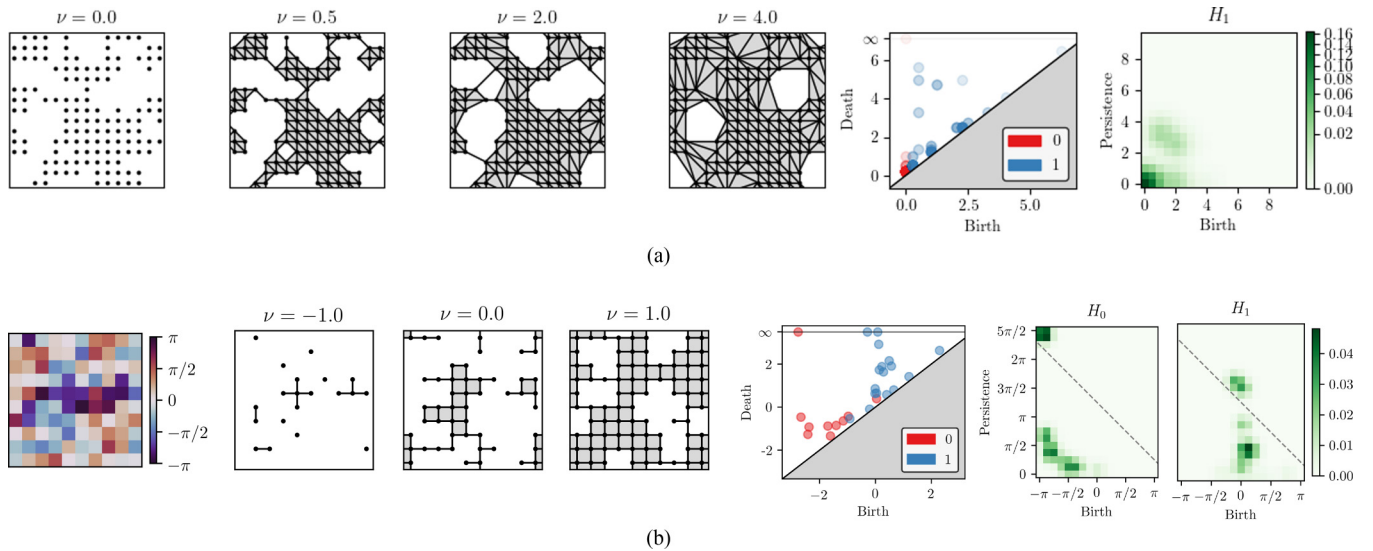


FIG. 3. Example filtrations along with their corresponding persistence diagrams, showing the individual births and deaths of topological features, and persistence images. (a) Ising model on 16×16 square lattice, $T = 3.0$. Simplices of radius less than $\sqrt{\nu}$ are included in the complex. At $\nu = 4.0$ there remain nontrivial 1-cycles around the holes in the simplicial complex—these correspond to the long-lived blue points in the upper left region of the persistence diagram and the shaded region around a persistence of three in the H_1 persistence image and (b) XY model on 16×16 square lattice, $T = 1.4$. At low values of the threshold ν , added spins are largely isolated, leading to a number of 0-cycles (connected components) in the cubical complex which then die as they are connected to other components via intermediate spins being added. These 0-cycles correspond to the red points in the lower left region of the persistence diagram and the shaded region in the H_0 persistence image.

A. Filtrations

The data we consider come from square-lattice spin models, some with discrete, \pm , Ising spins and others with continuous angles. In all models we employ periodic boundary conditions. We now describe the filtrations we use for these two cases, and give examples of how features of the spin configurations are captured by the persistence data.

1. Filtrations for discrete spins

With discrete spins on a square lattice (Ising and square-ice) we choose to represent our data via a point cloud, taking the locations of all spins aligned with a predetermined direction as the data. We choose to take all spins which are aligned with the total magnetization (no matter how small). After creating the point cloud from a given spin configuration, we then use an α filtration to create the persistence diagram/image. The filtration corresponds to a coarse-graining of the point cloud, parameterized by the areas of balls enclosing the simplices: See Fig. 2 for a small example.

As an example, Fig. 3(a) shows several steps in the α filtration for an example Ising model spin configuration in the disordered phase, along with the derived persistence diagram and H_1 persistence image. To compute the persistence of α -complexes we use the GUDHI class `AlphaComplex` [43].

2. Filtrations for continuous spins

We also consider models where the spins are continuous (XY and fully frustrated XY). In these cases, a spin configuration is a function $f : \Lambda \rightarrow S^1$ from the lattice, Λ , of N spin sites to their angles. We consider models with global

$O(2)$ symmetry so that we may place the total magnetization (no matter how small) at angle $\theta = 0$ and think of the function f as mapping Λ into $(-\pi, \pi]$. The sublevel sets with threshold $\nu \in (-\pi, \pi]$, $f^{-1}(-\pi, \nu] \subseteq \Lambda$, consisting of all lattice sites with angles less than ν then give a filtration of (periodic) cubical complexes. (Here cubical complexes are natural because of the underlying cubical lattice.) These sublevel sets experience topology change when the threshold ν passes a critical point of f , as is familiar from Morse theory [44]. In this case 0-cycles have nontrivial births, corresponding to spin values where f has a local minimum. As such we include both 0- and 1-cycles in the derived persistence images.

Figure 3(b) shows an example spin configuration for the XY model in the disordered phase, along with its corresponding persistence diagram and the combined H_0, H_1 persistence image. To compute the persistence of cubical complexes we use the GUDHI class `PeriodicCubicalComplex` [45].

III. PHASE CLASSIFICATION AND CRITICAL PHENOMENA

In this section we apply our methods to the task of phase classification in simple two-dimensional lattice spin models. We consider four such models: The Ising and square-ice models have discrete, \pm , spins and the XY and fully frustrated XY models have continuously varying spins. Sample spin configurations for each model are generated at a number of temperatures using standard Monte Carlo sampling techniques. Example spin configurations at low and high temperatures for each model are shown in Fig. 1.

For each model considered, classification into two phases is performed using only the persistence images. A subset of samples with extreme temperatures are used to train a logistic regression and then the accuracy of the regression is evaluated using the known temperatures of all samples. We normalize our persistence images using the ℓ_1 -norm, so they may be interpreted as probability densities for finding cycles with particular births/deaths for a given system. Unnormalized persistence images contain information about the total number of p -cycles and also lead to a successful classification.

A. Logistic regression

In the following sections we will be classifying spin configurations based on their persistence images. Since the persistence images are information rich, we are able to use perhaps the simplest classification scheme, logistic regression, to great effect. Here we quickly recall the procedure of logistic regression. One benefit of logistic regression is that it is easy to tell what aspects of the data are used by the classification algorithm. We will use these to extract order parameters for the phase transitions under consideration.

Persistence images $PI \in \mathbb{R}^n$ [$n \sim \mathcal{O}(400)$ in our examples] are vectors of positive numbers representing the distribution of cycles at different values of birth and persistence. A logistic regression depends on parameters λ_0 and $\lambda \in \mathbb{R}^n$ and the sigmoid function $\sigma : \mathbb{R} \rightarrow (0, 1)$, given by

$$\sigma(z) = \frac{1}{1 + e^{-z}}. \tag{2}$$

The sigmoid interpolates between $\sigma(-\infty) = 0$ and $\sigma(\infty) = 1$. A persistence image is declared to be in “category 0” if $\sigma(\lambda_0 + \lambda PI) < \frac{1}{2}$ and in “category 1” if $\sigma(\lambda_0 + \lambda PI) > \frac{1}{2}$. In our examples “category 0” will correspond to a low-temperature phase and “category 1” will correspond to a high-temperature phase. The parameters $\lambda_{i=0, \dots, n}$ are learned by training on a subset of the persistence images, $PI^{(k)}$, which are labeled into the two categories (i.e., phases) with $y^{(k)} \in \{0, 1\}$. Training amounts to maximizing the log-likelihood,

$$\sum_k (y^{(k)} \log \{\sigma[\lambda_0 + \lambda PI^{(k)}]\} + (1 - y^{(k)}) \times \log \{1 - \sigma[\lambda_0 + \lambda PI^{(k)}]\}) - C \sum_i \lambda_i^2, \tag{3}$$

with respect to λ_0 and λ , where the constant $C = 0.1$ controls the ℓ_2 regularization used to prevent overfitting. By training on extreme temperatures, we incur some inaccuracy due to our extrapolating to intermediate temperatures; these will not concern us too much, as we will find successful classification regardless.

On training, the regression can be applied to the rest of the persistence images to give an “average classification” at each temperature. This can be interpreted as quantifying the regression’s certainty that a particular temperature belongs to a particular phase. The temperature at which the average classification on testing data is 0.5 gives an estimate of the critical temperature. In addition, the learned coefficients λ_i may be investigated to learn which bins (i.e., regions) of the persistence images are most discerning when it comes

to distinguishing the low- and high-temperature data. Bins where $\lambda_i \gg 0$ will identify features prevalent in the high-temperature phase, while $\lambda_i \ll 0$ will identify features that are prevalent in the low-temperature phase. Bins where $\lambda_i \gg 0$ ($\lambda_i \ll 0$) will identify the dimensionality, size, and longevity of features which are characteristic of the high-temperature (low-temperature) phase. These will constitute our order parameters.

B. Ising model

The Ising model on a two-dimensional square lattice is very well understood, largely in part to Onsager’s exact solution [46]. Spins $s_i \in \{-1, 1\}$ live at the vertices of the lattice with ferromagnetic interactions governed by the local Hamiltonian

$$H_{\text{Is}} = - \sum_{\langle i, j \rangle} s_i s_j, \tag{4}$$

where the sum is over nearest-neighbor pairs. In the thermodynamic limit there is a second-order phase transition at $T_{\text{Is}} = \frac{2}{\log(1+\sqrt{2})} \approx 2.27$. At low temperatures there is spontaneous magnetization, while there is a disordered phase at high temperatures. While this model is well understood, it provides a good first application of our method. We are able to easily extract the magnetization as order parameter from a simple logistic regression. We additionally study the relationship of new “persistence” critical exponents to those usually studied.

1. Logistic regression and order parameter

For temperatures $T \in \{1.00, 1.05, \dots, 3.50\}$ we generate 1000 sample spin configurations for a $N \times N$ square lattice of N^2 spins for $N = 15, 25, 50$. For each sample we construct the persistence image using a weight $\log(1+p)$ and $\sigma = 0.5$ [see Eq. (1)]. Training of a logistic regression on the persistence images is conducted only on a subset of samples with extreme temperatures, well within the expected phases (see the left-hand side of Fig. 4). The classification extrapolates very well to the intermediate temperatures and for $N = 50$ gives an estimate of $T \approx 2.37$ for the critical temperature. The discrepancy from the known critical temperature may be attributed to finite-size effects.

The coefficients of the trained logistic regression (see the right-hand side of Fig. 4) show that the low-temperature configurations are identified by their having many small, short-lived cycles. These may be understood as arising both from 2×2 blocks of aligned spins (which lead to very short-lived 1-cycles) as well as 1-cycles wrapping small groups of isolated spins which are flipped relative to the large domains of aligned spins: The latter become more and more important as the temperature is increased. In the high-temperature phase, spins are oriented randomly, leading to a more uniform distribution of 1-cycle sizes. Using persistent homology we are able to easily identify the magnetization as the order parameter, as is well known.

2. Multiscale behavior near criticality

Since persistent homology contains multiscale information about a spin configuration, it seems reasonable that one should

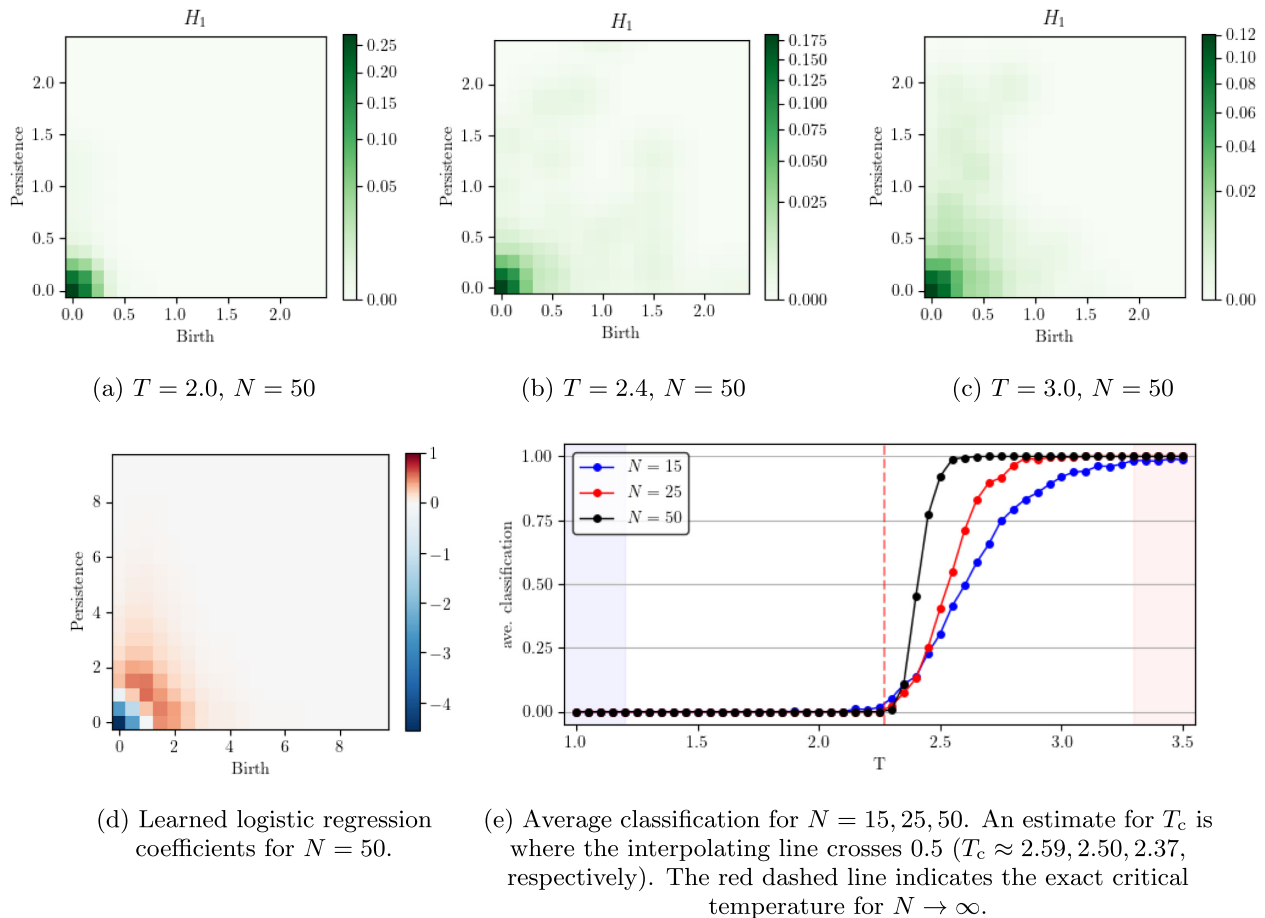


FIG. 4. Phase classification for the Ising model on an $N \times N$ square lattice. [(a) (c)] Example persistence images for the first homology group below, near, and above the critical temperature. [(d) and (e)] The results of classifying these persistence images into two phases, having trained a logistic regression only on a subset of persistence image with temperatures in the highlighted regions in (e). The learned coefficients in (d) indicate that the low-temperature phase is characterized by small, short-lived features.

be able to probe a model’s approach toward scale invariance via critical exponents. Indeed, we are able to see aspects of scale invariance appearing at criticality by looking at statistics derived from the persistence diagram. One-dimensional statistics such as the Betti numbers, births, and deaths can be found by counting points in different regions of the persistence diagrams. In this way we may compute the 1-cycle death probability density, $D_T(d)$, at each temperature, which quantifies the distribution of feature sizes in the spins. In Fig. 5

we see that deaths are exponentially distributed with a long tail forming at criticality, indicative of a diverging correlation length and the emergence of power-law behavior.

To be more quantitative, we may fit each $D_T(d)$ to a function of the form

$$D_T(d) = A d^{-\mu} e^{-d/a_{\text{death}}}. \quad (5)$$

Here d is the filtration parameter at the death scale of a cycle, and A is a numerical constant. There are two critical exponents to be extracted: μ gives the power-law behavior at criticality, while the correlation area a_{death} diverges at criticality according to $a_{\text{death}} \sim |T - T_c|^{-\nu_{\text{death}}}$. We are limited by the IR cutoff of the model, namely the finite area of the lattice, but we may still estimate these exponents. As a consistency check, we ask how these might be related to previously studied critical exponents. Using scaling arguments, one can show that at criticality the proportion of clusters of k aligned spins goes as

$$P(\text{cluster of size } k) \sim k^{-\tau}, \quad (6)$$

where the critical exponent is $\tau \approx 2.032$ [47,48]. The function $D_T(d)$ is not directly measuring the size of clusters, since the death of a 1-cycle around an island of spins is influenced non-trivially by the shape and “nesting” of clusters. Nevertheless,

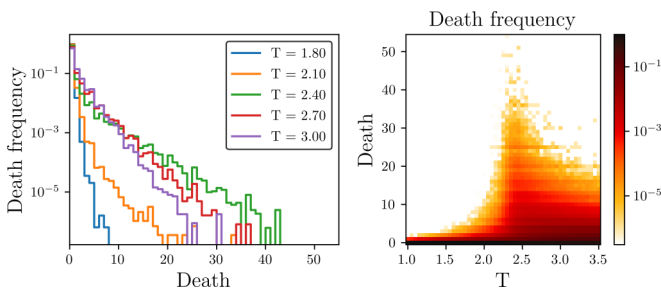


FIG. 5. Ising model death distributions, $N = 50$. The slight horizontal stripes in the figure on the right (e.g., at death = 25) are symptomatic of the underlying lattice.




FIG. 6. μ and a_{death} for Ising death distributions, $N = 50$. The red dashed lines indicate the previously estimated critical temperature $T \approx 2.37$. Error estimates are derived from fitting multiple simulations.

it seems reasonable to expect that at criticality the distribution of 1-cycle deaths should follow a similar power-law distribution. Recall that the value of the filtration parameter at the death of a 1-cycle is the *area* of the disks placed on each point in the point cloud, and so roughly corresponds to the number, k , of spins enclosed by the 1-cycle.

The fit parameters μ and a_{death} are shown in Fig. 6, where we see clearly the diverging correlation area as criticality is approached from above. The value of μ at our previously

estimated critical temperature, $T \approx 2.37$, is consistent with $\mu \approx \tau = 2.032$ as anticipated above, although a more detailed study would be needed to determine the value of μ more exactly. We see also the linear behavior of a_{death}^{-1} with temperature, indicating $\nu_{\text{death}} \approx 1$. That this is the same degree of divergence as the correlation length of the spin-spin correlation function $\langle s(0)s(r) \rangle \sim e^{-r/\xi}$, $\xi \sim |T - T_c|^{-1}$ can be understood by the following *rough* argument.

The death of a 1-cycle in the α filtration roughly corresponds to the *area* of the cluster of spins that it encloses and $D_T(d)$ roughly corresponds to the probability that a contiguous region of spins with area d is aligned. Consider for simplicity looking to estimate the probability that a disk of spins with radius R are all aligned. At infinite temperature where the spins are randomly aligned, this probability would simply be $2^{-\pi R^2/\ell^2}$, where ℓ is the lattice spacing. If we suppose that $P_{\text{disk}}(R) \sim e^{-R^2/a}$ for some “correlation area” a even at finite temperature, then how is a related to ξ as defined by $\langle s(R)s(0) \rangle \sim e^{-R/\xi}$ in the disordered phase? To estimate $P_{\text{disk}}(R + \ell) \sim e^{-(R+\ell)^2/a} \approx e^{-R^2/a - 2\ell R/a}$, imagine asking that a circle of $\sim R$ spins all be aligned with the disk of (aligned) spins of radius R that they encircle. For simplicity, we ignore conditional aspects of the probability and subleading terms. This should then take the form $e^{-R^2/a} e^{-\#R/\xi}$, from which we conclude that $a \sim \ell \xi$: In particular, $a \sim \xi \sim |T - T_c|^{-\nu}$

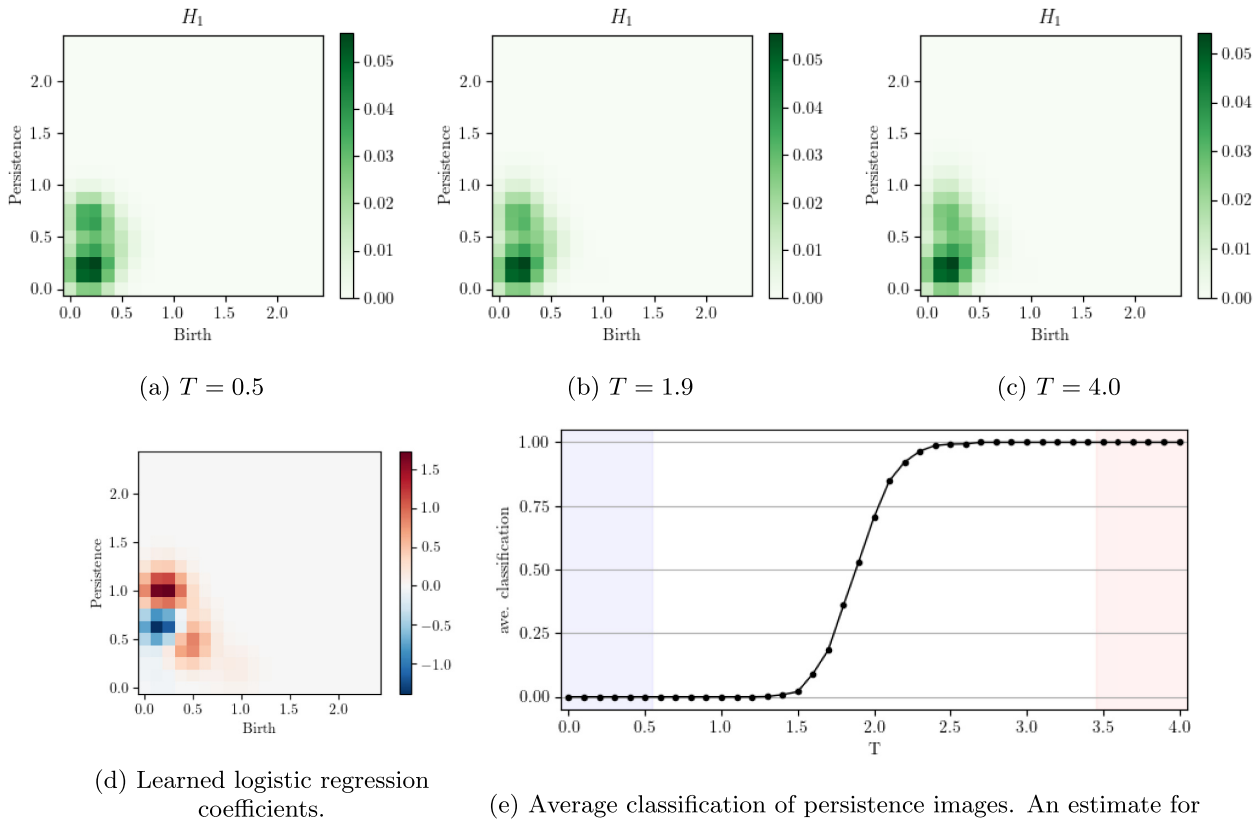


FIG. 7. Phase classification for the square ice model on a 50×50 lattice. [(a) (c)] Example persistence images for the first homology group below, near, and above the estimated critical temperature. [(d) and (e)] The results of classifying these persistence images into two phases, having trained a logistic regression only on a subset of persistence images with temperatures in the highlighted regions in (e). The learned coefficients in (d) indicate that the phase transition involves a subtle shift to the characteristic length scale in the spin configurations.

with the same critical exponent as the critical temperature is approached from above.

It would be interesting to further understand the relationship between the persistence critical exponents we defined and those typically studied.

C. Square-ice model

The square-ice model places spins, $s_i \in \{-1, 1\}$, on the edges rather than vertices of a square lattice and is governed by the local Hamiltonian

$$H_{\text{SI}} = \sum_{v \in \Lambda} \left(\sum_{i:v} s_i \right)^2, \quad (7)$$

where $i : v$ denotes those spins on edges adjacent to the vertex v . This is a particular instance of the 16-vertex model which in general has a rich phase structure (e.g., see Refs. [49,50]). In contrast to the Ising model there is no spontaneous magnetization at low temperatures. Rather, the ground state is highly degenerate: Any configuration with exactly two up and two down spins adjacent to every vertex has zero energy. This leads to frustration in the low-energy dynamics, as adjacent vertices v compete to minimize $(\sum_{i:v} s_i)^2$. This competition takes place at small scales, so that many 1-cycles die very quickly in the filtration. Nevertheless, we are still able to identify a shift in the distribution of p -cycle births and deaths and reliably classify samples into two phases. In

this case, the frustration introduces a particular length scale to the topological features in the low-temperature phase, while the distribution of sizes in the high-temperature phase is less restricted.

1. Logistic regression and order parameter

We generate 1000 sample spin configurations for a 50×50 lattice with 5000 spins at temperatures $T \in \{0.0, 0.1, \dots, 4.0\}$. Each sample gives a persistence image with a weight $\log(1+p)$. Again training a logistic regression only on those persistence images with extreme temperatures (Fig. 7), we find an estimate of $T \approx 1.9$ for the critical temperature. From the logistic regression coefficients presented Fig. 7(d) we see that as the temperature increases there is a tendency for 1-cycles to be born later or to be longer-lived. Both are indicative of a changing local structure in the spin configurations. In the low-temperature phase, it is energetically beneficial for neighboring vertices to coordinate, resulting in a regular patterns of alternating up and down spins. This regularity forces 1-cycles to live at smaller scales than in the high-temperature phase. More precisely, the domains of alternating spins result in many 1-simplices of radius $1/2$ which are filled in at $v = 1/\sqrt{2} \approx 0.7$ —this is reflected in the location of the negative logistic regression coefficients shown in Fig. 7(d).

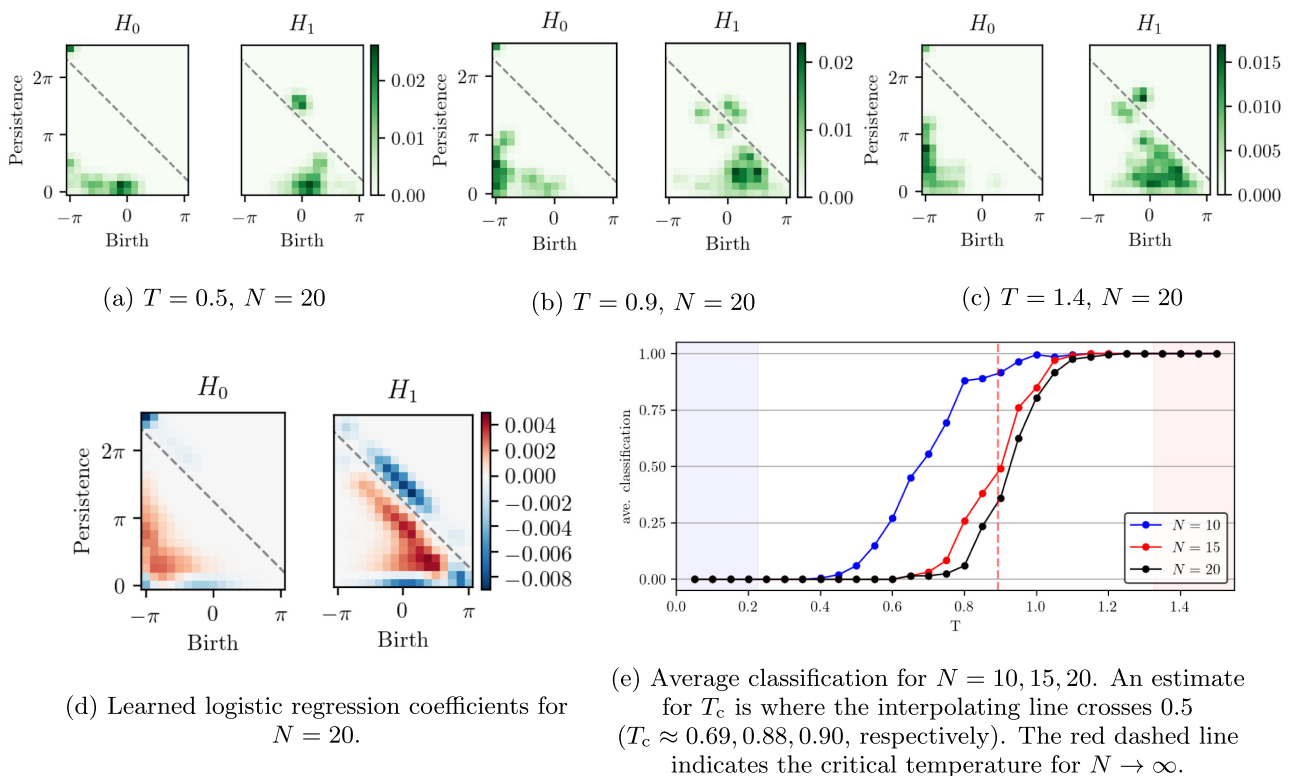


FIG. 8. Phase classification for the XY model on an $N \times N$ square lattice. [(a) (c)] Example persistence images for both the zeroth and first homology groups below, near, and above the critical temperature. The gray dashed lines visually separate the infinite persistence 0- and 1-cycles from those which have finite death. [(d) and (e)] The results of classifying these persistence images into two phases, having trained a logistic regression only on a subset of persistence images with temperatures in the highlighted regions in (e).




FIG. 9. Example spin configuration with three vortex-antivortex pairs, with seven steps in the sublevel filtration shown. Each vortex-antivortex pair corresponds to a number of 0-cycles which are born very early and 1-cycles which die late.

D. XY model

The XY model is a continuous-spin generalization of the Ising model. At each site of the square lattice spins take values in S^1 and are governed by

$$H_{XY} = - \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j). \tag{8}$$

There is a well-known KT phase transition at $T_{XY} \approx 0.892$ (see Refs. [51,52], among others). This is an infinite-order phase transition where at low temperatures there are bound vortex-antivortex pairs while at high temperatures free vortices proliferate and spins are randomly oriented.

1. Logistic regression and order parameter

With continuous spins each spin configuration implicitly contains much more information about the underlying dynamics. For temperatures $T \in \{0.05, 0.10, \dots, 1.50\}$ we generate 200 sample spin configurations on a $N \times N$ lattice with N^2 spins for $N = 10, 15, 20$. Persistence images are created for each sample, as in Fig. 3(b). The zeroth homology, in contrast to the α -complexes used for discrete spins, is very rich for the cubical complexes and we include both H_0 and H_1 persistence data in the persistence images. There is always a single 0-cycle and two 1-cycles which never die: These correspond to the p -cycles of the torus on which the lattice lives. We distinguish these immortal p -cycles from those cycles with late deaths ($d \approx \pi$) by giving the former a death of $d = \frac{5\pi}{2}$ by hand. Omitting these infinite persistence “torus cycles” results in a comparable phase classification.

Performing a logistic regression of the concatenated H_0 and H_1 persistence images by training on configurations with temperatures far away from the anticipated transition leads to the classification of Fig. 8. The critical temperature is estimated as $T_{XY} \approx 0.90$ for $N = 20$. We see that the low-temperature phase is characterized by p -cycles on the “boundary” of the persistence images. This we can understand in the following way. A (small-enough) loop around an isolated vortex has nontrivial winding number, which ensures that there are spins with angles close to both $-\pi$ and π if a vortex is present. This explains the strong blue regions in the corners of the logistic regression coefficients: For $\nu \approx -\pi$ a number of 0-cycles are born very early for each vortex and antivortex, giving the lower-left corner of the H_0 coefficients. One of these 0-cycles lives forever, giving the upper-left corner of the H_0 coefficients. In addition, there are 1-cycles which are born close to $\nu \approx \pi$, again corresponding to the extreme angles associated with the (anti)vortices. See Fig. 9 for an example of this interpretation in practice. When vortex-antivortex pairs happen to not be present at low temperatures, then all of the spins are

aligned close to $\theta = 0$, giving the short-lived features centered around a birth of zero along the bottom edges of both the H_0 and H_1 coefficients.

As before we may consider the distribution of p -cycle deaths as a function of temperature. In Fig. 10 we see that low temperatures there are two “populations” of both 0- and 1-cycles which merge into one as we pass into the high-temperature phase. This again can be attributed to the presence of vortex-antivortex pairs in the following way. Using the raw spin configurations we may count the number of (anti)vortices simply by looking for nontrivial winding in 2×2 blocks of the lattice. This can be compared with the number of 0-cycles with early death (e.g., $d \leq -\frac{3\pi}{4}$) and the number of 1-cycles with late death (e.g., $d \geq \frac{3\pi}{4}$). Averaging over samples with temperatures below 0.20 where the number of vortex-antivortex pairs is reasonably small on the 20×20 lattice leads to Fig. 11. There is a clear correlation between the number of extreme-death p -cycles and the number of vortex-antivortex pairs as determined directly from the spins. This topological signature of vortex-antivortex pairs should exist rather generally.

Previous investigations of the XY model and its KT phase transition using neural networks and PCA have faced difficulties in identifying vortices at low temperatures [4,53]. It is worth emphasizing the relative ease with which persistent homology identifies (anti)vortices as an important feature at low temperatures.

E. Fully frustrated XY model

A frustrated version of the XY model is obtained by changing some of the nearest-neighbor interactions to be antiferromagnetic. One such choice which is fully frustrated is

$$H_{FFXY} = - \sum_{\langle i,j \rangle} J_{ij} \cos(\theta_i - \theta_j), \tag{9}$$

where $J_{ij} = -1$ on every other row of horizontal edges and $J_{ij} = +1$ everywhere else. There are two phase transitions




FIG. 10. XY 0- and 1-cycle death distributions.

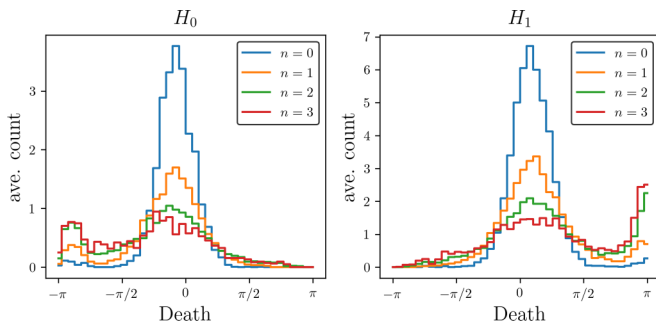


FIG. 11. Average counts of 0- and 1-cycle deaths for fixed number of vortex-antivortex pairs at low temperatures ($T_{\lesssim} 0.20$).

that occur at temperatures which are very close together: A phase transition at $T \approx 0.454$ is associated with the loss of \mathbb{Z}_2 symmetry, and a phase transition at $T \approx 0.446$ is associated with the loss of the $SO(2)$ rotational symmetry [54]. Because of their proximity we are unable to identify both transitions without an extensive set of simulations.

1. Logistic regression and order parameter

We generate 200 sample spin configurations on a 20×20 lattice with 400 spins for temperatures $T \in \{0.05, 0.10, \dots, 1.20\}$. As before with the XY model, the zeroth homology is quite rich and we include it in the persistence images. Training the logistic regression leads to the classification in Fig. 12, where the critical temperature is estimated as $T_{\text{FFXY}} \approx 0.39$. A more accurate estimation can be achieved by using training data closer to the phase tran-

sition. The learned coefficients show a strong tendency for both 0- and 1-cycles to shift to have persistence around $\frac{3\pi}{4}$ in the high-temperature phase. As in the square-ice model, our order parameter probes the small-scale structure of the frustration pattern. In particular, the low-temperature phase exhibits “pseudodomains” where many next-to-nearest neighbors take similar spin values. The alternating structure induced by the antiferromagnetic bands therefore leads to more isolated local minima (i.e., 0-cycles in the sublevel filtration) in the low-temperature phase. In the high-temperature phase, most of the local minima are born at $\theta \approx -\pi$, while in the low-temperature phase there are local minima at higher θ protected by these pseudodomains. This explains the blue band at the bottom of the H_0 logistic regression coefficients. The lack of vortices can be seen from the death distribution as a function of temperature in Fig. 13.

In our discussion we have used a sublevel filtration with cubical complex to quantify the homology of continuous-spin configurations. Another approach would be to construct point clouds by taking the locations of spins in a (sub)levelset and using an α filtration. By scanning through levelsets one can capture the topological features of $f : \Lambda \rightarrow S^1$ in a different way. For the fully frustrated XY model this leads to a comparable classification and estimate for the critical temperature.

IV. DISCUSSION

In this paper we have explored the use of persistent homology in quantitatively analyzing the phase structure and

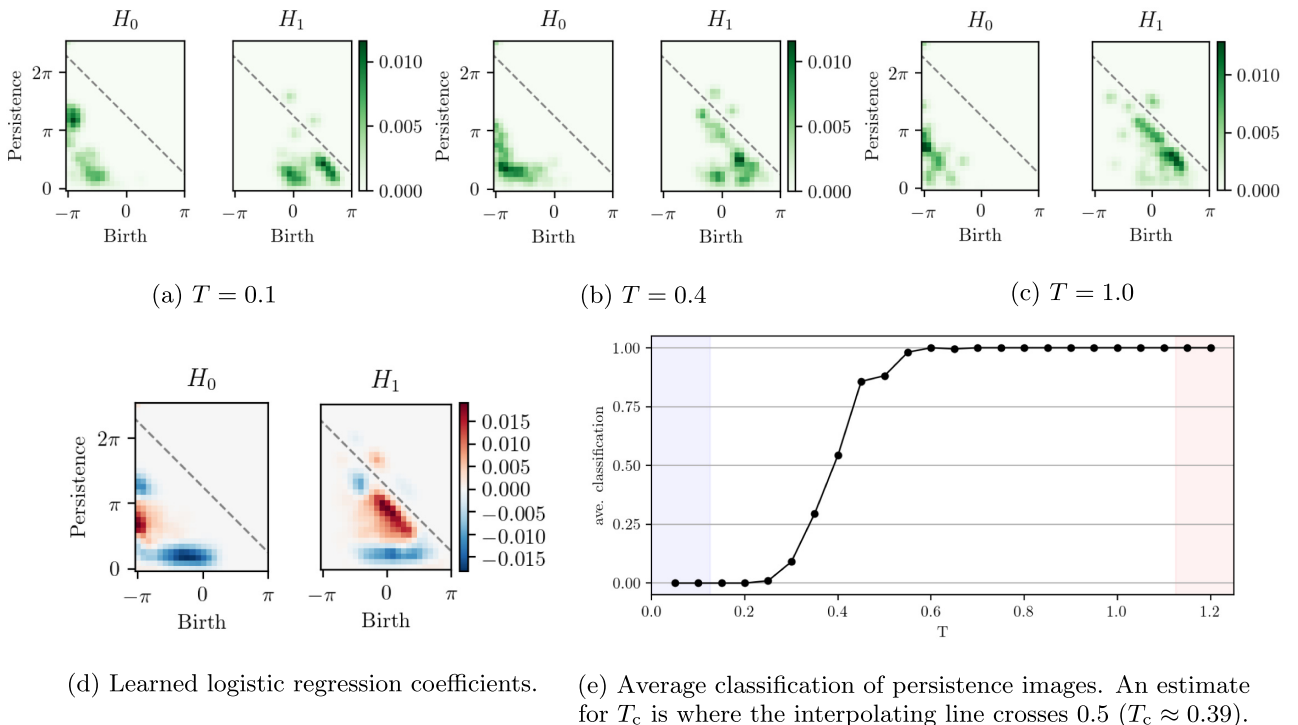


FIG. 12. Phase classification for the FFXY model on a 20×20 square lattice. [(a) (c)] Example persistence images for both the zeroth and first homology groups below, near, and above the critical temperature. The gray dashed lines visually separate the infinite persistence 0- and 1-cycles from those which have finite death. [(d) and (e)] The results of classifying these persistence images into two phases, having trained a logistic regression only on a subset of persistence images with temperatures in the highlighted regions in (e).




FIG. 13. FFXY 0- and 1-cycle death distributions.

critical behavior of lattice spin models. While the models we consider can be understood via other means, the use of persistent homology provides an interesting perspective into their statistical properties. Nonlocal features are naturally accounted for in this framework and could prove to be useful in more complex systems as well. As mentioned in the Introduction, the ideas exemplified in this work can be and have been applied to a wide variety of data sets: For other condensed matter systems one need only find a suitable representation of the data and choice of accompanying filtration. What we hope to have conveyed with the examples considered is that even the most simple-minded choice of data representation (spin locations or spin directions) can be successful, in part because the topological summary statistic which is the persistence image is quite rich. We are hopeful that TDA and persistent homology will continue to grow as tools in many areas of physics.

Phase classification using persistence images alone is accomplished successfully for the Ising, square-ice, XY, and fully frustrated XY models, providing a mix of examples with discrete/continuous spins and some with frustration. The resulting trained logistic regressions reveal those regions of the persistence image/diagram which are characteristic of low- and high-temperature phases. This allows for an easily interpretable classification, where (sometimes drastic) shifts in the distributions of p -cycle births and deaths are associated with a phase transition. In the case of the XY model,

there is a clear correlation between the number of early-death 0-cycles, late-death 1-cycles, and the number of bound vortex-antivortex pairs at low temperatures. In our classifications we relied on supervised learning, namely there being two phases and labeled samples with “extreme” temperatures on which to train. If the number of phases is unknown or labeled samples are unavailable, then unsupervised techniques will need to be used, but we expect the techniques utilized here to still prove useful in such cases. For example, clustering algorithms may be used in a straightforward way on the vector-valued persistence images.

The persistence data also display features of critical phenomena. For the Ising model one observes the emergence of power-law behavior in the distribution of 1-cycle deaths as the critical temperature is approached. We are able to estimate two critical exponents associated with this behavior: The correlation area diverges as $a_{\text{death}} \sim |T - T_c|^{-\nu_{\text{death}}}$ with $\nu_{\text{death}} \approx 1$ as expected for the 2D Ising model, and we estimate the critical exponent μ , introduced through $D_T(d) \sim d^{-\mu}$, to be $\mu \approx 2$, in agreement with expectations from the known power-law behavior of cluster sizes at criticality.

We have demonstrated the quantitative statistical capabilities of persistent homology for relatively simple 2D lattice spin systems. It would be interesting to apply these ideas and techniques to more complicated lattice spin models in higher dimensions or with no known order parameter. In more than two dimensions the higher homology groups may serve useful in quantifying nonlocal structures. We leave such work for the future.

ACKNOWLEDGMENTS

We thank Jeff Schmidt for useful discussions. We also thank the participants of the “Theoretical Physics for Machine Learning” and “Physics \cap ML” workshops at the Aspen Center for Physics and Microsoft Research, respectively, where partial results of this work were presented in early 2019, for discussion. G.J.L. and G.S. are supported in part by the DOE Grant No. DE-SC0017647 and the Kellett Award of the University of Wisconsin.

-
- [1] L. Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
 - [2] E. P. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nat. Phys.* **13**, 435 (2017).
 - [3] S. J. Wetzels and M. Scherzer, Machine learning of explicit order parameters: From the Ising model to $su(2)$ lattice gauge theory, *Phys. Rev. B* **96**, 184410 (2017).
 - [4] W. Hu, R. R. P. Singh, and R. T. Scalettar, Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination, *Phys. Rev. E* **95**, 062122 (2017).
 - [5] J. F. Rodriguez-Nieva and M. S. Scheurer, Identifying topological order through unsupervised machine learning, *Nat. Phys.* **15**, 790 (2019).
 - [6] R. Woloshyn, Learning phase transitions: Comparing PCA and SVM, [arXiv:1905.08220](https://arxiv.org/abs/1905.08220) [cond-mat.stat-mech].
 - [7] N. Rao, K. Liu, and L. Pollet, Inferring hidden symmetries of exotic magnets from learning explicit order parameters, *Phys. Rev. E* **104**, 015311 (2021).
 - [8] M. S. Scheurer and R.-J. Slager, Unsupervised Machine Learning and Band Topology, *Phys. Rev. Lett.* **124**, 226401 (2020).
 - [9] C. Giannetti, B. Lucini, and D. Vadacchino, Machine learning as a universal tool for quantitative investigations of phase transitions, *Nucl. Phys. B* **944**, 114639 (2019).
 - [10] K. Liu, J. Greitemann, and L. Pollet, Learning multiple order parameters with interpretable machines, *Phys. Rev. B* **99**, 104410 (2019).
 - [11] J. Greitemann, K. Liu, L. D. C. Jaubert, H. Yan, N. Shannon, and L. Pollet, Identification of emergent constraints and hidden order in frustrated magnets using tensorial kernel methods of machine learning, *Phys. Rev. B* **100**, 174408 (2019).

- [12] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
- [13] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine learning phases of strongly correlated fermions, *Phys. Rev. X* **7**, 031038 (2017).
- [14] P. Huembeli, A. Dauphin, and P. Wittek, Identifying quantum phase transitions with adversarial neural networks, *Phys. Rev. B* **97**, 134109 (2018).
- [15] A. Tanaka and A. Tomiya, Detection of phase transition via convolutional neural network, *J. Phys. Soc. Jpn.* **86**, 063001 (2016).
- [16] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys. X* **5**, 1797528 (2020).
- [17] A. Dawid, P. Huembeli, M. Tomza, M. Lewenstein, and A. Dauphin, Phase detection with neural networks: Interpreting the black box, *New J. Phys.* **22**, 115001 (2020).
- [18] Y. Zhang, P. Ginsparg, and E.-A. Kim, Interpreting machine learning of topological quantum phase transitions, *Phys. Rev. Res.* **2**, 023283 (2020).
- [19] S. Blücher, L. Kades, J. M. Pawłowski, N. Strodthoff, and J. M. Urban, Towards novel insights in lattice field theory with explainable machine learning, *Phys. Rev. D* **101**, 094507 (2020).
- [20] J. Arnold, F. Schäfer, M. Žonda, and A. U. J. Lode, Interpretable and unsupervised phase classification, *Phys. Rev. Research* **3**, 033052 (2021).
- [21] H. Edelsbrunner, D. Letscher, and A. Zomorodian, Topological persistence and simplification, in *Proceedings of the 41st Annual Symposium on Foundations of Computer Science* (IEEE, Los Alamitos, CA, 2000), pp. 454–463.
- [22] A. Zomorodian and G. Carlsson, Computing persistent homology, *Discr. Comput. Geom.* **33**, 249 (2005).
- [23] A. J. Zomorodian, *Topology for Computing*, Vol. 16 (Cambridge University Press, Cambridge, UK, 2005).
- [24] H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction* (American Mathematical Society, Providence, RI, 2010).
- [25] J. Murugan and D. Robertson, An introduction to topological data analysis for physicists: From lgm to frbs, [arXiv:1904.11044](https://arxiv.org/abs/1904.11044) [astro-ph.IM].
- [26] V. de Silva and R. Ghrist, Coverage in sensor networks via persistent homology, *Algebr. Geom. Topol.* **7**, 339 (2007).
- [27] G. Carlsson, T. Ishkhanov, V. De Silva, and A. Zomorodian, On the local behavior of spaces of natural images, *Int. J. Comput. Vis.* **76**, 1 (2008).
- [28] J. M. Chan, G. Carlsson, and R. Rabadan, Topology of viral evolution, *Proc. Natl. Acad. Sci. USA* **110**, 18566 (2013).
- [29] K. Xia and G.-W. Wei, Persistent homology analysis of protein structure, flexibility, and folding, *Int. J. Numer. Methods Biomed. Eng.* **30**, 814 (2014).
- [30] M. Gameiro, Y. Hiraoka, S. Izumi, M. Kramar, K. Mischaikow, and V. Nanda, A topological measurement of protein compressibility, *Jpn. J. Industr. Appl. Math.* **32**, 1 (2015).
- [31] A. E. Sizemore, J. Phillips-Cremins, R. Ghrist, and D. S. Bassett, The importance of the whole: topological data analysis for the network neuroscientist, *Network Neurosci.* **3**, 656 (2009).
- [32] F. A. N. Santos, E. P. Raposo, M. D. Coutinho-Filho, M. Copelli, C. J. Stam, and L. Douw, Topological phase transitions in functional brain networks, *Phys. Rev. E* **100**, 032414 (2019).
- [33] A. Cole and G. Shiu, Persistent homology and non-Gaussianity, *J. Cosmol. Astropart. Phys.* **2018**, 025 (2018).
- [34] M. Biagetti, A. Cole, and G. Shiu, The persistence of large scale structures I: Primordial non-Gaussianity, *JCAP* **04** (2021) 061.
- [35] M. Cirafici, Persistent homology and string vacua, *J. High Energy Phys.* **03** (2016) 045.
- [36] A. Cole and G. Shiu, Topological data analysis for the string landscape, *J. High Energy Phys.* **03** (2019) 054.
- [37] H. Adams, S. Chepushtanova, T. Emerson, E. Hanson, M. Kirby, F. Motta, R. Neville, C. Peterson, P. Shipman, and L. Ziegelmeier, Persistence images: A stable vector representation of persistent homology, *J. Mach. Learn. Res.* **18**, 218 (2017).
- [38] I. Donato, M. Gori, M. Pettini, G. Petri, S. De Nigris, R. Franzosi, and F. Vaccarino, Persistent homology analysis of phase transitions, *Phys. Rev. E* **93**, 052138 (2016).
- [39] D. Spitz, J. Berges, M. K. Oberthaler, and A. Wienhard, Finding universal structures in quantum many-body dynamics via persistent homology, *SciPost Phys.* **11**, 60 (2021).
- [40] Q. H. Tran, M. Chen, and Y. Hasegawa, Topological persistence machine of phase transitions, *Phys. Rev. E* **103**, 052127 (2021).
- [41] B. Olsthoorn, J. Hellsvik, and A. V. Balatsky, Finding hidden order in spin models with persistent homology, *Phys. Rev. Res.* **2**, 043308 (2020).
- [42] G. J. Loges and A. Cole, tda-spin-models, <https://github.com/gloges/TDA-Spin-Models> (2020).
- [43] V. Rourke, Alpha complex, in *GUDHI User and Reference Manual* (GUDHI Editorial Board, 2020), 3.3.0 ed.
- [44] J. Milnor, *Morse Theory (AM-51)*, Vol. 51 (Princeton University Press, Princeton, NJ, 2016).
- [45] P. Dlotko, Cubical complex, in *GUDHI User and Reference Manual* (GUDHI Editorial Board, 2020), 3.3.0 ed.
- [46] L. Onsager, Crystal statistics. i. A two-dimensional model with an order-disorder transition, *Phys. Rev.* **65**, 117 (1944).
- [47] A. D. Bruce and D. J. Wallace, Droplet theory in low dimensions: Ising systems in zero field, *J. Phys. A Math. Gen.* **16**, 1721 (1983).
- [48] R. Toral and C. Wall, Finite-size scaling study of the equilibrium cluster distribution of the two-dimensional Ising model, *J. Phys. A* **20**, 4949 (1987).
- [49] L. Foini, D. Levis, M. Tarzia, and L. F. Cugliandolo, Static properties of 2d spin-ice as a sixteen-vertex model, *J. Stat. Mech.* (2012) P02026.
- [50] D. Levis, Two-dimensional spin ice and the sixteen-vertex model, Ph.D. theses, Université Pierre et Marie Curie - Paris VI, 2012.
- [51] H. Weber and P. Minnhagen, Monte Carlo determination of the critical temperature for the two-dimensional xy model, *Phys. Rev. B* **37**, 5986 (1988).
- [52] M. Hasenbusch, The two-dimensional XY model at the transition temperature: A high-precision Monte Carlo study, *J. Phys. A* **38**, 5869 (2005).

- [53] M. J. S. Beach, A. Golubeva, and R. G. Melko, Machine learning vortices at the Kosterlitz-Thouless transition, [Phys. Rev. B](#) **97**, 045207 (2018).
- [54] M. Hasenbusch, A. Pelissetto, and E. Vicari, Multicritical behavior in the fully frustrated XY model and related systems, [J. Stat. Mech.](#) (2005) P12002.