

Finding hidden order in spin models with persistent homology

Bart Olsthoorn^{1,*}, Johan Hellsvik^{1,†} and Alexander V. Balatsky^{1,2‡}

¹*Nordita, KTH Royal Institute of Technology and Stockholm University,
Roslagstullsbacken 23, SE-106 91 Stockholm, Sweden*

²*Department of Physics, University of Connecticut, Storrs, CT 06269, USA*

(Dated: November 20, 2020)

Persistent homology (PH) is a relatively new field in applied mathematics that studies the components and shapes of discrete data. In this work, we demonstrate that PH can be used as a universal framework to identify phases of classical spins on a lattice. This demonstration includes hidden order such as spin nematic ordering and spin liquids. By converting a small number of spin configurations to barcodes we obtain a descriptive picture of configuration space. Using dimensionality reduction to reduce the barcode space to color space leads to a visualization of the phase diagram.

I. INTRODUCTION

In condensed matter physics we often deal with materials where collective degrees of freedom such as magnetization, density modulations, coherent condensate formation, etc. can undergo order to disorder transitions as a function of external parameters. The distinctly different orderings of these degrees of freedom constitute different phases of the system, characterized by observables known as order parameters.

In the Landau theory of phase transitions, the transition between phases involves the breaking of a symmetry of the system. A classical example is the breaking of rotational symmetry in ferromagnetic materials. In the paramagnetic (high temperature) phase the spins are fully disordered and the system has the symmetry of the crystalline structure. At the onset of ferromagnetic order (low temperature) the symmetry is reduced. Order parameters reveal key symmetries of the phases and serve as indicators for the phase transitions where the condensed matter system undergoes a qualitative change as a result of change in control parameters such as temperature, pressure, and electromagnetic fields. Conventionally, in the Landau theory, these order parameters are constructed by hand.

In the presence of frustration between atomic spins, magnetic materials feature rich phase diagrams with different forms of collinear and noncollinear antiferromagnetic orderings, as well as more unconventional phases that lack long range magnetic orderings. At the same time we know that both the easily identifiable phases, such as the ferromagnet, and more complicated and not easily visualized orders, like spin nematic, do describe strong correlations between spins. Of relevance for the present work are classical and quantum spin liquids [1–5], the spin-ice phase [6–8], and nematic ordering that involves breaking of spin rotational symmetry while in the absence of magnetic ordering [7, 9–12]. The spin nematic

ordering is an example of a hidden order parameter which can be difficult to observe directly in an experiment, but that still reveals itself in the heat capacity signature and other observables at the phase transition into nematic state [4, 13].

The commonly used indicators of phase transitions are heat capacity, susceptibility of the relevant order parameter(s), and derived quantities such as the reduced fourth-order Binder cumulant [14]. While a peak in the heat capacity as a function of temperature can reveal the presence of a phase transition, it can not alone be used to characterize the phases. Understanding of the order parameter with concomitant susceptibility is a required step.

A new paradigm to investigate phase transitions based on neural networks recently emerged. It has been shown that neural networks can classify spin configurations sampled from a Monte Carlo simulation and thus can function as order parameter identifiers [15, 16]. The neural network is trained in a supervised fashion, i.e. it is given labeled microstates, where the labels are determined by a known conventional order parameter. Intriguingly, recent work has shown that it is also possible to identify phases without prior information of the phase diagram, including phases without conventional order parameters [17–19].

In this work, we use a qualitatively different approach to *automatically construct* order parameters, based on a persistent homology (PH), a recently developed field in applied mathematics. In this approach the objective is to identify components and shapes given discrete data and a distance metric. These features are captured in a persistence diagram or barcode, two different graphical representations of the same information. Roughly speaking, these diagrams show at what scale different features are present in the data, and over what range of scale they persist.

PH is used in many fields including image processing [20], complex networks [21] and natural language processing [22]. However, it has not seen much use in condensed matter yet. In cosmology, it has been applied to study the cosmic web and its structure [23]. In materials science, it is used to study nanoporous materials [24] and other larger structures such as silica glass and polymers

* bartol@kth.se

† hellsvik@kth.se

‡ balatsky@hotmail.com

[25, 26]. A recent work proposes persistent homology observables to study dynamics of a quantum many-body system, using the two-dimensional Bose gas as an example [27]. There are three works that discuss the application of persistent homology for the detection of phase transitions. Firstly, Donato *et al.* [28] study the mean-field XY model and a classical Φ^4 model. The phase transitions are detected by computing the persistent homology of the configuration space with molecular dynamics simulations. Secondly, the recent work by Tran *et al.* [29] demonstrates the detection of phase transitions in the classical XY model and in quantum models. Thirdly, while this manuscript was under review we learned of related work by Cole *et al.* [30], where the phase transition in a number of two-dimensional lattice spin models is detected.

We demonstrate how PH provides an unsupervised method for constructing phase diagrams, capturing both conventional and unconventional phases with a single framework. Each microstate is converted into a persistence diagram (or barcode), after which two different microstates can be compared with a similarity metric on their corresponding persistence diagram. Defining a similarity metric (or kernel) for persistence diagrams is still an open problem and we examine the capability of a sliced Wasserstein distance for spin models. This provides a unbiased way to construct a phase diagram.

The paper is organized as follows. We introduce persistent homology as an *order parameter* in Sec. II, describe our method to construct phase diagrams in Sec. III and the XXZ model on pyrochlore lattice in Sec. IV. The results from the application of persistent homology to the XXZ model are presented in Sec. V followed by a discussion in Sec. VI and conclusion in Sec. VII.

II. PERSISTENT HOMOLOGY AS AN ORDER PARAMETER

Persistent homology (PH) is a method that identifies the qualitative features of a finite metric space (also called a point-cloud dataset) – for an introduction see [31, 32]. Given a finite metric space and a distance parameter r , the point-cloud is converted into a graph where edges are added between vertices with distance smaller than r . A simplicial complex K_1 is built from the graph, and the homology of K_1 can be computed efficiently by linear algebra. Performing this construction for increasing r yields a sequence of nested subcomplexes,

$$K_1 \subset K_2 \subset \dots \subset K_l = K, \quad (1)$$

where the sequence K is called a filtered simplicial complex. The computed homology group $H_k(K_i)$ identifies k -dimensional holes in the simplicial complex. For example, $k = 0$ relates to connected components (clusters) and $k = 1$ to 1-dimensional holes and so on. The rank of H_k counts the number of k -dimensional holes, also referred to as the Betti number β_k . The persistence (or lifetime) of

k -dimensional holes over the distance parameter r is visualized as a barcode, persistence diagram (PD), or lifetime diagram, all of which represent the same information.

Common types of complexes are the Vietoris-Rips (VR), Čech and Alpha complex [31, 32]. The complexes differ in number of simplices, affecting the computational complexity. We examine the capability of Alpha complex (also α -complex) that is often used when the input data is 2D or 3D, in our case spins on a 3D lattice. The α -complex is equivalent to the Čech complex when studying persistent homology, but it contains fewer simplices.

The α -complex is constructed from a Delaunay triangulation on the input point-cloud dataset [33]. For a 3-dimensional point cloud, the highest dimensional simplex is the 3-simplex (i.e. tetrahedron). Due to the input space being 3-dimensional, there are three homology groups of interest: H_0, H_1 and H_2 .

With this brief introduction we now move to outline how we apply PH to spin models. The spin configurations (microstates) from classical Monte Carlo simulations form a metric space in multiple ways. Firstly, a single spin configuration can be considered a point in state space (e.g. discrete in the case of an Ising model, continuous in case of a Heisenberg model). Given a distance metric between two states, the corresponding persistence diagram (PD) is describing the topology of the phase diagram, where the homology groups describe the shapes of the phases in parameter space. Secondly, a single site in the spin system can be considered a point, leading to a single PD per microstate. A phase diagram can be constructed by comparing PDs.

In this work, we focus on the latter concept, i.e. the homology groups describe a microstate (spin configuration) directly. Averaging over microstates is possible by simply merging barcodes and discarding the the information from which microstate each bar (or point in PD) originated.

The concept introduced above requires a distance metric between two sites, where we can use the information of the lattice and the spin. The symmetry of the barcode is dictated by the choice of distance metric and determines its capability of capturing distinct phases.

As a simple example, consider the ferromagnetic square lattice Ising model ($J = 1$) hosting a low-temperature symmetry-breaking ferromagnetic phase and a phase transition at $T_c \approx 2.269 J$ to a paramagnetic phase [34]. Choose a distance parameter r where aligned spins are closer than opposite spins, with the distance being an Euclidean distance between spin tops. This definition of a distance metric leads to domains forming connected components in homology group H_0 . Hence the Betti number $\beta_0(r)$ counts the number of domains present in the spin configuration, similar to the conventional order parameter of magnetization density $\langle M \rangle$. In general, any metric capturing the shapes present in the spin texture is able to identify the two phases.

The phase diagram of a 2D Ising model discussed above is a convenient example, but the phase diagram is too

simple to benefit from an approach with PH. In this work we study a more complex system, namely the XXZ model on a pyrochlore lattice, known to host a large number of competing phases [4].

III. PHASE DIAGRAM CONSTRUCTION

Traditionally, phase diagrams are constructed by observing a change in the order parameter that typically ranges from zero (phase A) to finite (phase B). In our proposed scheme, shown schematically in Fig. 1, the persistence diagram (PD) plays the role of the order parameter, and therefore we need to quantify the change in the PD.

Common distance metrics for PDs are the bottleneck distance and the Wasserstein distance [35]. Both essentially measure the similarity by attempting to match the points (i.e. persistent features) in the two diagrams. These metrics are too computationally expensive for our purposes. Instead, we use the sliced Wasserstein (SW) distance which is an approximation of the Wasserstein distance [36]. The distance is measured between two PDs of same k th homology group, i.e. $\text{SW}(\text{PD}_i^k, \text{PD}_j^k)$. For the total distance, we consider holes (H_1) and voids (H_2):

$$D_{ij} = \text{SW}(\text{PD}_i^1, \text{PD}_i^1) + \text{SW}(\text{PD}_i^2, \text{PD}_i^2) \quad (2)$$

The distance matrix D is established by calculating the SW distance between all pairs of system parameters (e.g. all combinations of interaction J and temperature T). Finally, the phase diagram is visualized by dimensionality reduction on D to a 3-dimensional color space (red, green and blue). There are obvious drawbacks to this simple approach such as the non-linearity of color perception, but we find it sufficient. When distances are small and more color space is required, it is possible to limit the visualization to a subregion of the phase diagram, as shown in Sec. V.

Dimensionality reduction algorithms aim to construct a low-dimensional image where distances of the original high-dimensional space are preserved as best as possible. Principal component analysis (PCA) is a common method based on matrix factorization, when given a set of input data points. In our case we only have a distance matrix D . Therefore PCA cannot be used directly, but multidimensional scaling (metric MDS) is the equivalent basic technique in this case [37]. In summary, metric MDS algorithm is used to reduce the distance matrix D to three RGB color channels that color-code the different phases in the phase diagram. Similar phases with small

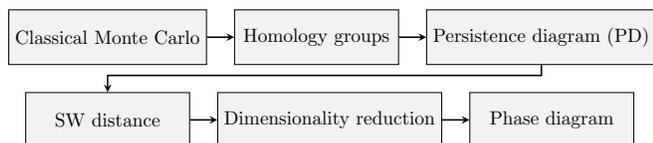


FIG. 1. Summary of phase diagram construction.

SW distances will appear close in color space and are therefore color-coded with similar colors.

IV. XXZ MODEL ON THE PYROCHLORE LATTICE

The pyrochlore lattice is a cornerstone in research on frustrated magnetism, and is having a lead role in experimental and theoretical explorations of spin-ice physics [1, 4, 6–8]. Motivated by the different chemistry of pyrochlore materials, a variety of spin Hamiltonians with long or short ranged interactions on the pyrochlore lattice have been investigated. The XXZ model on a pyrochlore lattice is a system with short ranged interactions having a rich phase diagram with competing antiferromagnetic, spin-ice, spin-liquid and spin-nematic phases [4, 6]. This model has also been used previously to test a machine learning model (support vector machine, SVM) that identifies phases [17]. The Hamiltonian of the XXZ model is given by

$$H_{\text{XXZ}} = \sum_{\langle i,j \rangle} J_{zz} S_{i,z} S_{j,z} - J_{\pm} (S_i^+ S_j^- - S_i^- S_j^+), \quad (3)$$

with $S_i^{\pm} = S_{i,x} \pm i S_{i,y}$ and $\mathbf{S}_i = (S_{i,x}, S_{i,y}, S_{i,z})$, $\|\mathbf{S}_i\| = 1$, and has the cubic symmetry of the pyrochlore lattice. In the following, energies and temperatures are expressed in terms of the antiferromagnetic $J_{zz} = 1$ exchange interaction.

Extensive classical Monte Carlo simulations, field theoretical analysis, and spin dynamics simulations by Taillefumier *et al.* [4] have established a $J_{\pm} - T$ phase diagram with six phases which here will be briefly recapitulated (numbering corresponding to [17] and naming to [4]):

- I Easy-plane antiferromagnet ($\text{AF}\perp$). Spins lie in the plane perpendicular to the local z axes. The phase occurs for large positive value of the ratio J_{\pm}/J_{zz} .
- II Paramagnetic (PM). No long range spin ordering. Spins point in random directions with exponential decay of spin correlations.
- III Easy-plane spin liquid ($\text{SL}\perp$) with algebraic spin-spin correlations, no long range magnetic ordering.
- IV Easy-plane spin-nematic ($\text{SN}\perp$) with algebraic spin-spin correlations, no long range spin dipole ordering. The rotational $U(1)$ symmetry of the local z axes is broken by the onset of a higher order multipolar ordering.
- V Spin ice (SI). Each tetrahedron on the lattice has the "two-in, two-out" spin configuration with the spins aligned along their local z axis (see Fig. 2 (a)). The phase has algebraic spin correlations, no long range spin ordering.

VI Pseudo-Heisenberg antiferromagnet (pHAF) with algebraic spin correlations distinct from phases III and V, no long range spin ordering.

In summary, one phase has long range antiferromagnetic ordering (I AF \perp), one phase is disordered (II PM), one phase has spin-nematic ordering (IV SN \perp), and three of the phases (III SL \perp , V SI, and VI pHAF) are classical spin liquids. Relative to the paramagnetic phase, in which the system has the symmetries of the pyrochlore lattice, breaking of symmetry occurs only on entering the easy-plane antiferromagnetic phase (I AF \perp) or on entering the easy-plane spin-nematic phase (IV SN \perp), hence only when crossing in and or out from one of these two phases does the system go through a phase transition, with characteristic peak structure in the heat capacity and order parameter susceptibilities [4]. Transitions into or from one of the three spin liquid phases constitute a gradual evolution (crossover) of spin configurations, involving no breaking of symmetry, and lacks sharp peak structures in the heat capacity. The latter quantity nevertheless contains information that can be used to define a distinct criteria for where the crossover occurs [4]. The three classical spin liquid phases can be distinguished by the different form of algebraic spin correlation reported for each of the phases [4], correlations that can be measured in neutron scattering experiments. Rather remarkably, the easy-plane spin-nematic phase have the same algebraic spin correlation as the easy-plane spin liquid phase, yet the phases differ in the regard that in the latter the rotational symmetry in the local z axes is broken by the nematic order parameter.

To test the proposal of PH for detection of the phases, we use representative spin configurations as an input. Spin configurations to construct the phase diagram were sampled within classical Monte Carlo simulations. In the Metropolis-Hastings algorithm, a spin flip is accepted if energy E is lowered, otherwise accepted with probability $\exp(-\beta\Delta E)$ [38], where $\beta = 1/T$ is the inverse temperature. The number of sites is given by $N = 16L^3$ with L^3 the number of cubic unit cells, each containing 16 sites, with each site having 6 nearest neighbors. The simulations are carried out with $L = 4$, i.e. $N = 1024$ spins. Spin configurations are sampled on a parameter space grid $(T/J_{zz}, J_{\pm}/J_{zz})$. Similar to [17], we sample 17 temperatures T/J_{zz} with logarithmically spaced points between 10^1 and 10^{-3} . Exchange interactions J_{\pm}/J_{zz} are sampled with 29 linearly spaced points $[-1, 0.4]$. Appendix A lists the details of the implementation and the open-source code.

V. PERSISTENT HOMOLOGY ON THE XXZ MODEL

A single spin configuration is converted into a point cloud by placing a point at each spin tip, where the spin length is set to $1/4$ of the lattice tetrahedron edge length

$d = \frac{a}{2\sqrt{2}}$. Euclidean distance is used once the point cloud is established using the distance metric

$$\mathcal{D}(i, j) = r(i, j) + \frac{a}{2\sqrt{2}} \frac{\|\mathbf{S}_i - \mathbf{S}_j\|}{4}, \quad (4)$$

where $r(i, j)$ is the Euclidean distance between two sites. This corresponds to the distance of tips of the spin arrows (see Fig. 2). Note that \mathbf{S}_i is in the global crystal frame, not in the local frame of Equation 3 (see Appendix A for coordinate frame conversion). For nearest neighbors, this means $\mathcal{D}(i, j)$ is in the range $(d/2, 3d/2)$.

Figure 2 shows the exact barcode for a single unit cell in the spin ice (V SI) phase, where spins are set according to the "two-in, two-out" rule. The connected components (H_0) reveal that there are two length scales present. The smaller length scale ($\alpha^2 \leq 0.010$) corresponds to neighboring spins forming 1-simplices (e.g. Fig. 2 (a)). The larger length scale ($\alpha^2 \geq 0.034$) leads to a single connected component, as shown in Fig. 2 (b). The 1-dimensional holes (H_1) appear at $\alpha^2 \geq 0.034$ as

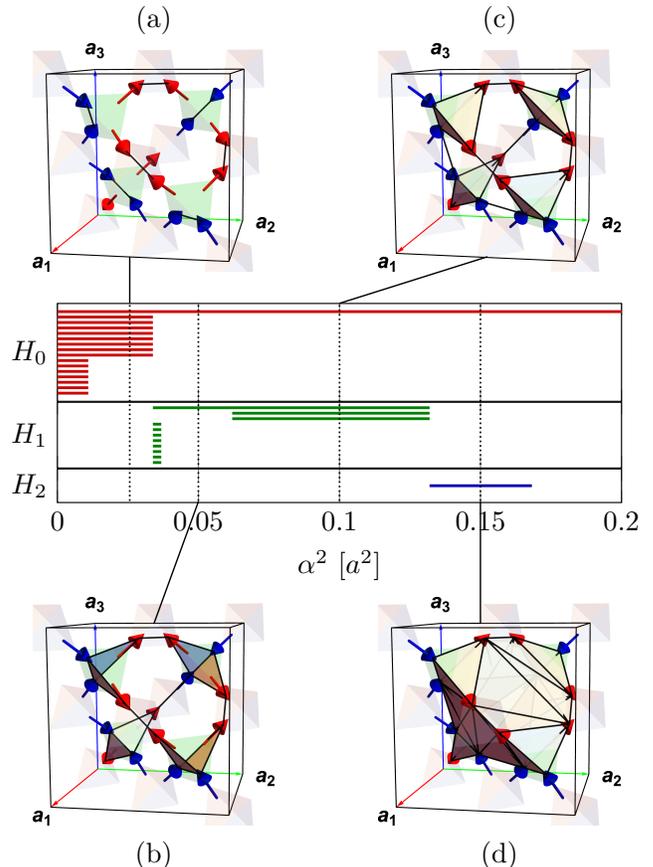


FIG. 2. Barcode for a single unit cell with an exact spin ice configuration "two-in, two-out". The spins are all aligned to the local z axis and there is no long range order. Shaded tetrahedrons form crystal lattice. The α -complex at certain α^2 filtration values is displayed. The characteristic features in H_1 and H_2 are captured by the Wasserstein distance to distinguish the spin ice phase from other phases.

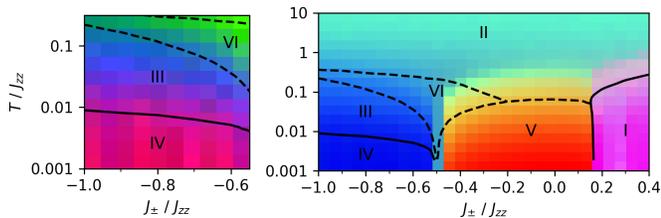


FIG. 3. Phase diagram of the XXZ model on a pyrochlore lattice. Constructed by using dimensionality reduction (MDS) on the distance matrix of $(T/J_{zz}, J_{\pm}/J_{zz})$ parameter space to the (R,G,B) color space. Sliced Wasserstein distance on homology groups is used. Crossover (transition) boundaries are drawn with solid (dashed) lines [4].

9 bars. This includes 8 empty triangles (i.e. complex of 1-simplices rather than a single 2-simplices) and one large hole (also visible at $\alpha^2 = 0.05$, Fig. 2 (b)). The empty triangles are filled at $\alpha^2 = 0.037$ and become 2-simplices, shown as shaded triangles in Fig. 2 (b). At $\alpha^2 = 0.062$, the 1-simplices form a large tetrahedron with 2-simplices at its corners. A tetrahedron of 1-simplices has $\beta_1 = \text{rank}(H_1) = 3$, which corresponds to 3 bars for H_1 in Fig. 2 (c). All 1-dimensional holes are filled at $\alpha^2 \geq 0.131$. At this point, a single 2-dimensional void (H_2) appears for $\alpha^2 = [0.132, 0.168]$, shown in Fig. 2 (d).

Generally, each phase has persistent features at different characteristic length scales. We sample 32 statistically independent spin configurations for a single $(T/J_{zz}, J_{\pm}/J_{zz})$ parameter combination. The α -complexes are calculated by GUDHI [33]. The resulting barcodes are merged into a single barcode, discarding the information of the microstate origin from each bar. The similarity between barcodes is measured by the sliced Wasserstein (SW) distance, implemented in persim [39], see Equation 2.

The distance matrix D is dimensionality reduced with MDS to a 3-dimensional space which is interpreted as red, green and blue channels. Figure 3 shows the phase diagram. The paramagnetic phase II PM and the phases V SI and I AF $_{\perp}$ are clearly distinct due to large pairwise sliced Wasserstein distances. The small distances between III SL $_{\perp}$ and IV SN $_{\perp}$ require a rescaling of the colormap to be visible, in the separate left pane. Also, the VI pHAF phase has small SW distances to the paramagnetic phase and is therefore difficult to detect, but a small color gradient is present.

Figure 4 shows the lifetime diagrams and confirms that the phases II PM, III SL $_{\perp}$, IV SN $_{\perp}$ and VI SN $_{\perp}$ are similar, leading to smaller SW distances. A more detailed view of the lifetime diagrams as temperature decreases is shown in the Appendix Fig. 6. Note that the spin ice phase V SI exhibits two H_1 features (with lifetimes around 0.07 and 0.1), corresponding to the two longer H_1 bar lengths in Fig. 2. In general, each phase has its own characteristic features and differences in its fingerprint are captured by the sliced Wasserstein distance. The qualitative similarities of phases and significant dif-

ferences between lifetime diagrams, I AF vs V SI vs IV SN $_{\perp}$ vs III SL $_{\perp}$ vs VI pHAF phase, leads to the proposal that these lifetime diagrams can serve as respective order parameters (i.e. lifetime PH portraits of the phases).

VI. DISCUSSION

The PH method outlined here is using only 32 spin configurations for each point in parameter space $(T/J_{zz}, J_{\pm}/J_{zz})$ and does not require data on energies or heat capacities, making it an efficient technique. The α -complexes are faster to calculate than the Čech and Vietoris-Rips complex and capture the holes and voids created by spins locally ordering. The computational bottleneck is the calculation of sliced Wasserstein distances between persistence diagrams. This could be improved by a faster distance metric for persistence diagrams, which is still an open problem in persistent homology. Finally, we note that persistent homology has very few free parameters in comparison to neural networks, which facilitates the interpretation of results.

In comparison to the work by Greitemann *et al.* [17] our method has a couple of key differences. Their work relies on the support vector machine (SVM) and the construction of monomials of spin components. We use persistent homology instead, involving no regression coefficients, and operate directly on the spin configurations. For the phase diagram construction, Greitemann *et al.* use the Fiedler vector to partition the phase diagram given a distance matrix D . We instead use the simpler technique of multidimensional scaling (metric MDS) di-

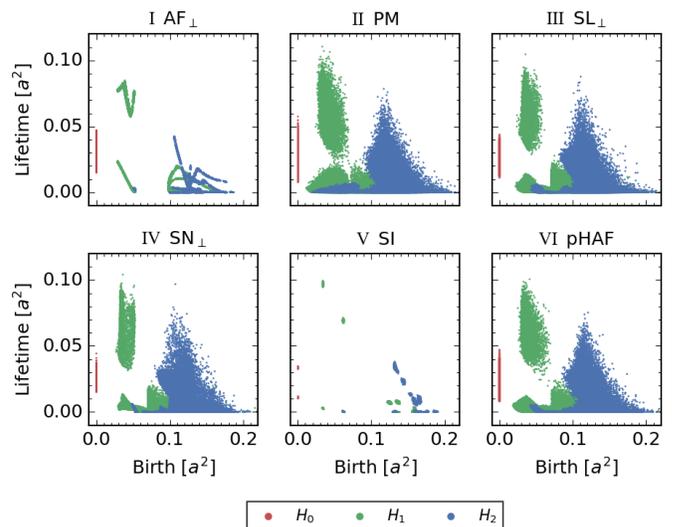


FIG. 4. Lifetime diagram of the α -complexes constructed from spin configurations for different phases. The lifetime diagram shows the same information as a persistence diagram, with the diagonal along the x-axis, i.e. lifetime is death - birth. The corresponding $(J_{\pm}/J_{zz}, T/J_{zz})$ phase space parameters are I = $(0.4, 0.001)$, II = $(0, 10)$, III = $(-0.8, 0.056)$, IV = $(-0.8, 0.001)$, V = $(0, 0.001)$ and VI = $(-0.5, 0.001)$.

rectly on the distance matrix D .

Recent work has shown that PH can detect phase transitions in a variety of classical and quantum models [28, 29]. We show the universality of the approach by identifying 6 different phases at once, including three classical spin liquids and construct a full two-dimensional phase diagram. Even though the phases are detected by a change in persistent homology, a deeper understanding of the phase (e.g. symmetry breaking condition) is not immediately clear. However, the barcode reveals the length scales where patterns are present, which can be a starting point for a specific theory constructed by hand.

The detection of phases is possible by studying persistence diagrams (Fig. 4) and changes are quantified by the sliced Wasserstein distance. We reduce the persistence diagrams space to color space with matrix factorization, which proves to work best when SW distances are large and of similar magnitude. Alternatively, dimensionality reduction based on the neighbor graph approach (e.g. Uniform Manifold Approximation and Projection, UMAP [40]) could highlight local structure better at the expense of representing global structure of the original high dimensional barcode space.

VII. CONCLUSION

We demonstrate that persistent homology (PH) can be used to capture different types of order in spin models. The application of our method to the XXZ model on the pyrochlore lattice with its six phases demonstrates the versatility of this approach. Both the phase with long range order and the phases with only local ordering are identified.

The barcodes reveal the characteristic length scales present in the spin model. Using sliced Wasserstein distance, we obtain a distance matrix for all pairwise system parameters. The distance matrix can be visualized using dimensionality reduction, revealing similar regions in parameter space. Alternatively, the persistence diagrams can be inspected directly to observe changes as the system temperature decreases.

In summary, persistent homology provides a new general computational framework to study both long range and local spin ordering. Extending the persistent homology framework to quantum spin models will be the topic of future work. Moreover, the phase transition of spin glasses (e.g. in the Edwards–Anderson model [41]) could also be an interesting test-case for this framework [42].

VIII. ACKNOWLEDGEMENTS

The authors are grateful to R. Matthias Geilhufe, Jens H. Bardarson, and Qian Yang for discussions. We acknowledge funding from the VILLUM FONDEN via the Centre of Excellence for Dirac Materials (Grant No. 11744), the European Research Council ERC HERO grant, University of Connecticut, and the Swedish Research Council (VR) through a neutron project grant (BIFROST, Dnr. 2016-06955). The authors acknowledge computational resources from the Swedish National Infrastructure for Computing (SNIC) at the National Supercomputer Centre at Linköping University, the Centre for High Performance Computing (PDC), the High Performance Computing Centre North (HPC2N), and the Uppsala Multidisciplinary Centre for Advanced Computational Science (UPPMAX).

-
- [1] C. Lacroix, F. Mila, and P. Mendels, eds., *Introduction to frustrated magnetism : materials, experiments, theory* (Springer, Berlin, 2010).
 - [2] L. Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199 (2010).
 - [3] L. Savary and L. Balents, Quantum spin liquids: a review, *Reports on Progress in Physics* **80**, 016502 (2017).
 - [4] M. Taillefumier, O. Benton, H. Yan, L. D. C. Jaubert, and N. Shannon, Competing Spin Liquids and Hidden Spin-Nematic Order in Spin Ice with Frustrated Transverse Exchange, *Physical Review X* **7**, 041057 (2017).
 - [5] C. Broholm, R. J. Cava, S. A. Kivelson, D. G. Nocera, M. R. Norman, and T. Senthil, Quantum spin liquids, *Science* **367**, eaay0668 (2020).
 - [6] S. T. Bramwell and M. J. P. Gingras, Spin Ice State in Frustrated Magnetic Pyrochlore Materials, *Science* **294**, 1495 (2001).
 - [7] N. Shannon, K. Penc, and Y. Motome, Nematic, vector-multipole, and plateau-liquid states in the classical $O(3)$ pyrochlore antiferromagnet with biquadratic interactions in applied magnetic field, *Physical Review B* **81**, 184409 (2010).
 - [8] R. Edberg, L. Ø. Sandberg, I. M. B. Bakke, M. L. Haubro, L. C. Folkers, L. Mangin-Thro, A. Wildes, O. Zaharko, M. Guthrie, A. T. Holmes, M. H. Sørby, K. Lefmann, P. P. Deen, and P. Henelius, Dipolar spin ice under uniaxial pressure, *Physical Review B* **100**, 144436 (2019).
 - [9] P. Chandra, P. Coleman, and A. I. Larkin, Ising transition in frustrated Heisenberg models, *Physical Review Letters* **64**, 88 (1990).
 - [10] M. Capati, M. Grilli, and J. Lorenzana, Nematic phase without Heisenberg physics in FeAs planes, *Physical Review B* **84**, 214520 (2011).
 - [11] A. Völl and S. Wessel, Spin dynamics of the bilinear-biquadratic $S = 1$ Heisenberg model on the triangular lattice: A quantum Monte Carlo study, *Physical Review B* **91**, 165128 (2015).
 - [12] Y. Kohama, H. Ishikawa, A. Matsuo, K. Kindo, N. Shannon, and Z. Hiroi, Possible observation of quantum spin-nematic phase in a frustrated magnet, *Proceedings of the National Academy of Sciences* **116**, 10686 (2019).
 - [13] G. Aeppli, A. V. Balatsky, H. M. Rønnow, and N. A. Spaldin, Hidden, entangled and resonating order, *Nature*

- Reviews Materials **5**, 477 (2020).
- [14] K. Binder, Finite size scaling analysis of ising model block distribution functions, *Zeitschrift für Physik B Condensed Matter* **43**, 119 (1981).
- [15] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nature Physics* **13**, 431 (2017).
- [16] I. A. Iakovlev, O. M. Sotnikov, and V. V. Mazurenko, Profile approach for recognition of three-dimensional magnetic structures, *Physical Review B* **99**, 024430 (2019).
- [17] 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, *Physical Review B* **100**, 174408 (2019).
- [18] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nature Physics* **13**, 435 (2017).
- [19] O. Balabanov and M. Granath, Unsupervised learning using topological data augmentation, *Physical Review Research* **2**, 013354 (2020).
- [20] H. Edelsbrunner, Persistent homology in image processing, in *Graph-Based Representations in Pattern Recognition*, edited by W. G. Kropatsch, N. M. Artner, Y. Haxhimusa, and X. Jiang (Springer, Berlin Heidelberg, 2013) pp. 182–183.
- [21] D. Horak, S. Maletić, and M. Rajković, Persistent homology of complex networks, *Journal of Statistical Mechanics: Theory and Experiment* **2009**, P03034 (2009).
- [22] X. Zhu, Persistent homology: An introduction and a new text representation for natural language processing, in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13* (AAAI Press, 2013) p. 1953–1959.
- [23] T. Sousbie, The persistent cosmic web and its filamentary structure - I. theory and implementation, *Monthly Notices of the Royal Astronomical Society* **414**, 350 (2011).
- [24] Y. Lee, S. D. Barthel, P. Dłotko, S. M. Moosavi, K. Hess, and B. Smit, Quantifying similarity of pore-geometry in nanoporous materials, *Nature Communications* **8**, 15396 (2017).
- [25] M. Buchet, Y. Hiraoka, and I. Obayashi, Persistent homology and materials informatics, in *Nanoinformatics*, edited by I. Tanaka (Springer, Singapore, 2018) pp. 75–95.
- [26] G. Kusano, Y. Hiraoka, and K. Fukumizu, Persistence weighted gaussian kernel for topological data analysis, *Proceedings of Machine Learning Research* **48**, 2004 (2016).
- [27] D. Spitz, J. Berges, M. K. Oberthaler, and A. Wienhard, Finding universal structures in quantum many-body dynamics via persistent homology (2020), arXiv:2001.02616 [cond-mat.quant-gas].
- [28] I. Donato, M. Gori, M. Pettini, G. Petri, S. De Nigris, R. Franzosi, and F. Vaccarino, Persistent homology analysis of phase transitions, *Physical Review E* **93**, 052138 (2016).
- [29] Q. H. Tran, M. Chen, and Y. Hasegawa, Topological Persistence Machine of Phase Transitions (2020), arXiv:2004.03169.
- [30] A. Cole, G. J. Loges, and G. Shiu, Quantitative and interpretable order parameters for phase transitions from persistent homology (2020), arXiv:2009.14231 [cond-mat.stat-mech].
- [31] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington, A roadmap for the computation of persistent homology, *EPJ Data Science* **6**, 17 (2017).
- [32] R. Ghrist, Barcodes: The persistent topology of data, *Bulletin of the American Mathematical Society* **45**, 61 (2007).
- [33] V. Ruvvrau, Alpha complex, in *GUDHI User and Reference Manual* (GUDHI Editorial Board, 2020) 3.3.0 ed.
- [34] L. Onsager, Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition, *Physical Review* **65**, 117 (1944).
- [35] H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction* (American Mathematical Society, Providence, Rhode Island, 2009).
- [36] M. Carrière, M. Cuturi, and S. Oudot, Sliced wasserstein kernel for persistence diagrams, in *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17* (JMLR.org, 2017) p. 664–673.
- [37] C. K. Williams, On a connection between kernel PCA and metric multidimensional scaling, *Machine Learning* **46**, 11 (2002).
- [38] A. W. Sandvik, A. Avella, and F. Mancini, Computational Studies of Quantum Spin Systems, in *AIP Conference Proceedings*, Vol. 1297 (2010) pp. 135–338.
- [39] Persim 0.0.9 documentation.
- [40] L. McInnes, J. Healy, N. Saul, and L. Grossberger, Umap: Uniform manifold approximation and projection, *The Journal of Open Source Software* **3**, 861 (2018).
- [41] S. F. Edwards and P. W. Anderson, Theory of spin glasses, *Journal of Physics F: Metal Physics* **5**, 965 (1975).
- [42] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, *Reviews of Modern Physics* **58**, 801 (1986).

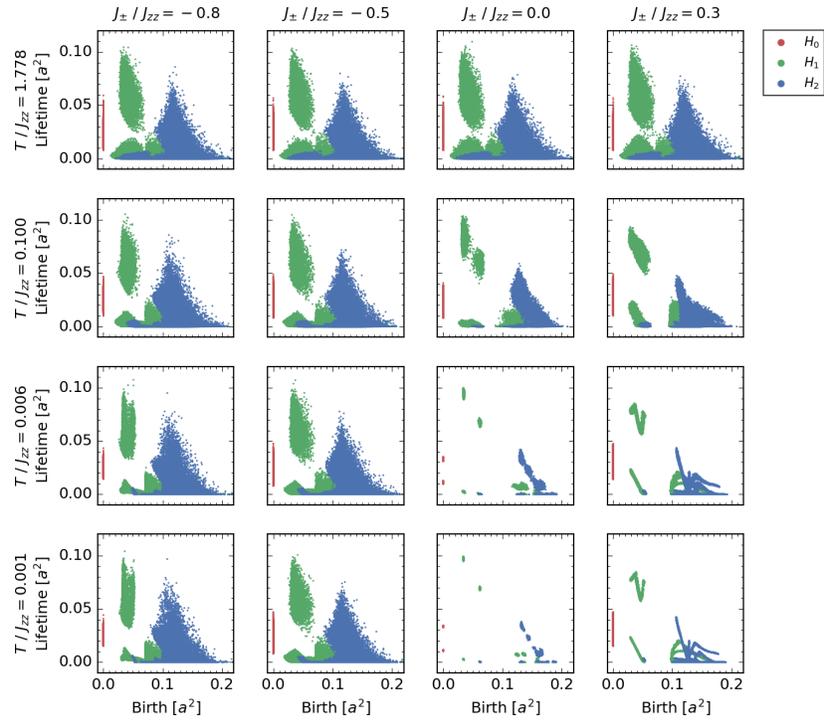


FIG. 6. Lifetime diagram of the α -complexes constructed from spin configurations at different phase space parameters. At the lowest temperature $T/J_{zz} = 0.001$ (bottom row), persistent features at characteristic length are visible.