Applications of Topological Data Analysis to Statistical Physics and Quantum Field Theories

Prifysgol Abertawe
Swansea University

Nicholas Sale

Department of Mathematics
Swansea University

Submitted to Swansea University in fulfilment
of the requirements for the Degree of
Doctor of Philosophy

August 2022

Abstract

This thesis motivates and examines the use of methods from topological data analysis in detecting and analysing topological features relevant to models from statistical physics and particle physics.

In statistical physics, we use persistent homology as an observable of three different variants of the two-dimensional XY model in order to identify relevant topological features and study their relation to the phase transitions undergone by each model. We examine models with the classical XY action, a topological lattice action, and an action with an additional nematic term. In particular, we introduce a new way of computing the persistent homology of lattice spin model configurations and demonstrate its use in detecting topological defects called vortices. By considering the fluctuations in the output of logistic regression and k-nearest neighbours models trained on persistence images, we develop a methodology to extract estimates of the critical temperature and the critical exponent of the correlation length. We put particular emphasis on finite-size scaling behaviour and producing estimates with quantifiable error. For each model we successfully identify its phase transition(s) and are able to get an accurate determination of the critical temperatures and critical exponents of the correlation length.

In particle physics, we investigate the use of persistent homology as a means to detect and quantitatively describe center vortices in SU(2) lattice gauge theory in a gauge-invariant manner. The sensitivity of our method to vortices in the deconfined phase is confirmed by using twisted boundary conditions which inspires the definition of a new phase indicator for the deconfinement phase transition. We also construct a phase indicator without reference to twisted boundary conditions using a k-nearest neighbours classifier. Finite-size scaling analyses of both persistence-based indicators yield accurate estimates of the critical β and critical exponent of correlation length for the deconfinement phase transition. We also use persistent homology to study the stability of vortices under gradient flow and the classification of different vortex surface geometries.

Declaration

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree. This thesis is the result of my own investigations, except where otherwise stated. Any other sources are acknowledged by footnotes giving explicit references and a bibliography is appended. The University's ethical procedures have been followed.

The following people and institutions contributed to the publication of work undertaken as part of this thesis:

Candidate	Nicholas Sale, College of Science, Swansea University
Author 1	Jeffrey Giansiracusa, Durham University
Author 2	Biagio Lucini, College of Science, Swansea University

Paper 1 (Quantitative analysis of phase transitions in two-dimensional XY models using persistent homology) [1]
Located in Chapter 3
Candidate contributed the design of experiments and methodology. This includes designing the analysis pipeline, writing of all required new software, and performing the analysis and calculations themselves with said software. The candidate wrote the entire first draft of the paper which forms the basis for the vast majority of the text of the final paper after critique.
Author 1 co-designed the research question and the methodology, participated in the development of the investigation, critiqued the manuscript and advised on matters regarding topological data analysis. ($\sim 5\%$).
Author 2 co-designed the research question and the methodology, participated in the development of the investigation, critiqued the manuscript and advised on matters regarding statistical physics. ($\sim 5\%$).

Paper 2 (Probing center vortices and deconfinement in SU(2) lattice gauge theory with persistent homology) [2]
Located in Chapter 4
Candidate contributed the design of experiments and methodology. This includes designing the analysis pipeline, writing of all required new software, and performing the analysis and calculations themselves with said software. The candidate wrote the entire first draft

of the paper which forms the basis for the vast majority of the text of the final paper after critique.
Author 1 co-designed the research question and the methodology, participated in the development of the investigation, critiqued the manuscript and advised on matters regarding topological data analysis. ($\sim 5\%$).
Author 2 co-designed the research question and the methodology, participated in the development of the investigation, critiqued the manuscript and advised on matters regarding quantum field theory. ($\sim 5\%$).

We the undersigned agree with the above stated "proportion of work undertaken" for each of the above published peer-reviewed manuscripts and preprints contributing to this thesis: Signed Candidate [redacted]

Author 1 [redacted]
Author 2 [redacted]

Nicholas Sale
August 2022

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed .[redacted]..
Date ...16/08/2022..

This thesis is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed [redacted]..
Date ...16/08/2022..

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed .[redacted]..
Date ...16/08/2022..

The University's ethical procedures have been followed and, where appropriate, that ethical approval has been granted.

Signed .[redacted]..
Date ...16/08/2022..

Contents

Acknowledgements

The undertaking of my PhD over the last three years has at times been a source of great frustration. Thankfully, these times have been in the small minority compared to the most part which has been an utter joy and a distinctly rewarding experience. This is due, in no small part, to the efforts of my supervisors Jeff Giansiracusa and Biagio Lucini whom I would like to thank very much; Jeff for supporting me right from the start, inspiring a continued love of all forms of mathematics, and because his continued supervision was not even a question after his move to Durham University. Biagio for introducing me to the fascinating area in which I have worked these last three years, for agreeing to become one of my supervisors part way through my PhD, and for his endless patience in explaining physics to me. I must also thank my fellow students: James Maxwell for his support and the time he spent with me on a fascinating project on hyperfields, and Dimitrios Bachtis for his excellent advice on the machine learning of phase transitions at a crucial time in my PhD. I thank my family for their unfailing love and for always wanting to try to understand what I have been spending my time on. Finally but by no means least, I thank my girlfriend Marleen. Without her thousand small acts of love and support over the last three years, none of this would have been possible.

List of Figures

List of Tables

1 Introduction

1.1 Topological Objects in Physics

Topology is ubiquitous in modern physics, and even a summary of all the areas of physics where topological arguments crop up is far beyond the scope of this introduction. In this thesis we are interested in certain phenomena involving topology that appear in the study of condensed matter and in particle physics. On the condensed matter side we shall investigate defects in materials or fields which are classified using topology. Moving to particle physics, we will be interested in field singularities which lie along closed surfaces, the topology of these surfaces, and how they embed into spacetime. However, the focus of this thesis is on computational physics and simulations on discrete models of spacetime. The usual topological tools relying on continuous maps break down here: the topology of discrete space does not support many interesting phenomena. But the topology in the underlying physics has not gone away, we simply need to look for it using a new lens. An argument of this thesis is that topological data analysis is an appropriate such lens for studying the phenomena mentioned above.

As an example, in [3], Mermin lays out the use of winding numbers and the fundamental group π_1 in classifying defects in media consisting of planar spins. Assign each point in the plane a unit-length vector lying in the same plane. We could imagine that this describes the orientation of molecules in a material, or of small magnetic moments in a larger magnetic film. Such a field is specified by a map

$$f : \mathbb{R}^2 \to S^1$$

where S^1 is the circle and f describes the angle the vector makes with the positive x-axis at each point in the plane. Now suppose this map is continuous everywhere except potentially at some point $P \in \mathbb{R}^2$, the location of a potential defect. Then given a simple closed curve in the plane $\gamma : S^1 \to \mathbb{R}^2$ that encloses P and is traversed anticlockwise, we can calculate the winding number

$$w(\gamma) = \frac{1}{2\pi} \oint_\gamma df$$

which sums up the changes in the angle f. Since γ is a closed path, the integral must be an integer multiple of 2π and therefore $w : (S^1 \to \mathbb{R}^2) \to \mathbb{Z}$ gives a signed count of how many times the vector field winds round the circle, and in which direction, while traversing γ. Note that this is a continuous map: so long as we avoid P, continuously deforming γ should continuously change w. But w is discrete and must therefore remain constant. Suppose that we have a vector field f such that this winding number is non-zero. Then f must turn through a full circle no matter how small we shrink γ around P. The derivative of f at P therefore diverges and we do indeed have a singularity, or defect. The point is that we did not need to shrink γ. Calculating the winding number around a huge loop far away from P, but still enclosing it, would give the same result and lead us to the same conclusion that there is some singularity enclosed by γ. The simple topological notions of continuity and continuous deformation allow us to detect a local defect at P from its global effect. The same kind of argument shows that the nature of the singularity is classified up to continuous deformation by the winding number too. Continuously deforming the vector field f instead of γ will also leave the winding number invariant. Composing γ and f gives us a map $f \circ \gamma : S^1 \to S^1$ from the circle to itself and allows us to skip the middle man of having a vector field in the plane. The winding number of f around γ simply becomes the degree of the map $f \circ \gamma$ and helps make it clear that what we are really trying to get at is the topological fact that $\pi_1(S^1) \cong \mathbb{Z}$.

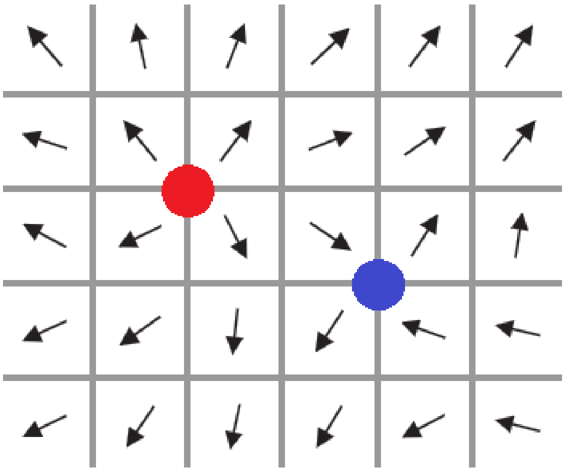

Figure 1: A vortex (red) and antivortex (blue) in a configuration of the 2D XY model – a discrete version of a unit-length vector field in the plane.

Now consider vortices and antivortices in the 2D XY model as illustrated in Figure 1 (we will look at these in more detail in Chapter 3). Instead of a vector field on the plane, we can think of configurations of the XY model as unit-length vector fields $f : \mathbb{Z}^2 \to S^1$ on the lattice. We might intuitively associate a winding number of 1 to the vortex (in red) and -1 to the antivortex (in blue). But if we try and define winding numbers and fundamental groups rigorously here, the arguments fall apart. The discrete nature of the space means that an 'unwinding' of the field becomes a perfectly admissible continuous deformation. Yet the physics in this case ensures that vortices remain intact and behave as we would expect in the continuum. There is a broad and fascinating underlying problem here of how to model and describe such topological objects when the usual definitions in terms of continuity and homotopy do not apply. Moreover, we are now looking at actual configurations, i.e. data. Data is messy and the vortices we see in such data may not look as well-formed as those in Figure 1. In Chapter 3 we will use topological data analysis to characterise vortices and antivortices.

Such examples are not constrained to condensed matter physics. In quantum field theory one might study certain Yang-Mills gauge fields that are classical solutions to the equations of motion, called instantons. If we assume that the field strength decays going outwards towards infinity, then such gauge fields are classified up to continuous deformation by the third homotopy group of the gauge group $\pi_3(G)$ [4] (isomorphic to $\mathbb{Z}$ when $G = \mathrm{SU}(n)$ for any $n > 1$ [5]). In this thesis we will look at topological defects from Yang-Mills theory called center vortices which form along closed surfaces. For SU(2), the parity of the number of such vortices is a topological invariant. Once again, working on the lattice and with real data ruins any chance of characterising these objects through homotopy groups. In Chapter 4 we use topological data analysis to detect center vortices in SU(2) lattice gauge theory, the lattice incarnation of SU(2) Yang-Mills.

1.2 Machine Learning of Phase Transitions

Many of the tools of topological data analysis are quantitative in nature. They output a numerical summary of the input data which can then be used in further data analysis. Therefore in this thesis we will not only use topological data analysis to detect topological defects, but also investigate its ability to quantitatively analyse phase transitions associated with those defects. This work lies in the same vein as the emerging body of work exploring the use of machine learning and other data analysis methods to detect and classify phase transitions in statistical physics systems. One of the motivations of this approach is to develop methodologies which require minimal *a priori* knowledge about the systems in question. The hope then is that these data-centric methods will be able to offer new insights into those models at the forefront of physics which seem to defy analytical methods. Much of the work in this area makes use of neural network models which, while unparalleled in machine learning tasks, are generally hard to interpret. This thesis explores the idea that topological data analysis, along with simple machine learning models, may provide a more interpretable approach by allowing us to relate certain phase transitions to the topological defects that may have a hand in driving them.

1.3 Thesis Outline

Chapter 2 covers the necessary background on the techniques we employ in this thesis. It will cover the basic definitions and theorems behind persistent homology, the main topological data analysis tool we use, along with some of the ways to analyse its output. It goes on to introduce two simple machine learning classifiers, the finite-size scaling method for quantitative analysis of phase transitions, and some of the statistical techniques we make use of. The chapter ends with a summary of some of the existing work in the application of topological data analysis to statistical physics and quantum field theory.

Chapter 3 covers a project based on using persistent homology to detect topological defects in several variants of the 2D XY model and study the phase transitions driven by these defects. It begins by introducing the models being studied and their phase transitions, as well as the previous work on using TDA to study them. It then introduces the methodology we use to apply persistent homology, before going on to cover several analyses using that methodology. The main part is a quantitative analysis of the phase transitions where the critical temperatures and a critical exponent of the transition are accurately estimated, which was published as a paper [1].

Chapter 4 covers a project based on using persistent homology to detect center vortices in SU(2) lattice gauge theory and study the deconfinement phase transition undergone by this model. It begins by introducing the model being studied, its phase transition, and the proposed mechanism for the transition that we consider. It then introduces the methodology we use to apply persistent homology, before going on to cover several analyses using that methodology. The main part is a quantitative analysis of the phase transition where the critical couplings and a critical exponent of the transition are accurately estimated, which has been released as a preprint [2].

Chapter 5 summarises the contributions of this thesis and discusses potential future work.

2 Background

2.1 Topological Data Analysis

Topological data analysis (TDA) is a relatively recent field of mathematics based on applying the ideas and tools of algebraic topology to the analysis of data. The motivation behind it is to provide a means to robustly describe shape in data. Initially the idea was to identify purely topological structure such as the number of connected components or holes, but more recent perspectives acknowledge the ability of its tools to summarise the geometry of data too (e.g., [6]). It has found applications in a wide variety of fields. A small list of references that attempts to capture the breadth of fields of application is [7–16]. The principal tool of TDA is persistent homology which is what we will be using in this thesis.

2.1.1 Persistent Homology

Persistent homology was introduced in its modern form in [17] and popularised in [18]. We shall give an overview here, but other useful references are [19–22]. In this section we assume a basic knowledge of simplicial/cubical complexes and their homology. A brief review of these can be found in Appendix B.

Given a simplicial or cubical complex X, the k-dimensional holes in it are described by the k-th homology $H_k(X)$. In particular, we compute homology with respect to some fixed field $\mathbb{F}$ so that $H_k(X)$ is a vector space. However, data coming from applications (including physics simulations) may not take the form of such a complex. Even if it does, it may not have the kind of structure we want to consider built in. For example, data may take the form of a set of points P in a metric space, sampled from some underlying manifold. A discrete set can be considered to be a simplicial or cubical complex, but its k-th homology is trivial for $k \geq 1$ and its 0-th homology simply counts the cardinality of P. However, we probably wanted to know something about the homology of the manifold that the points were sampled from. In order to insightfully apply homology then, we must use the data to construct a complex which (ideally) captures the topological features we are interested in. A common example of how to do this is given by the Vietoris-Rips simplicial complex of a point cloud.

Definition 2.1 (Vietoris-Rips Complex)**.** Given a finite set of points P with a distance function $d : P^2 \to \mathbb{R}$ and a choice of radius $\epsilon \geq 0$, the Vietoris-Rips complex $VR_\epsilon(P)$ is the simplicial complex with vertices P and simplices

$$\Sigma = \{\, \sigma \in \mathcal{P}(P) \mid \ \forall x, y \in \sigma \, : \, d(x,y) \leq \epsilon \,\}.$$

Rather than viewing P as a discrete set of points, the Vietoris-Rips complex gives us a way to view it as a piecewise linear object "at the scale of" ϵ. However this leaves us to find a single good value of ϵ which may be difficult to find, or may not even exist at all as we see in Figure 2.

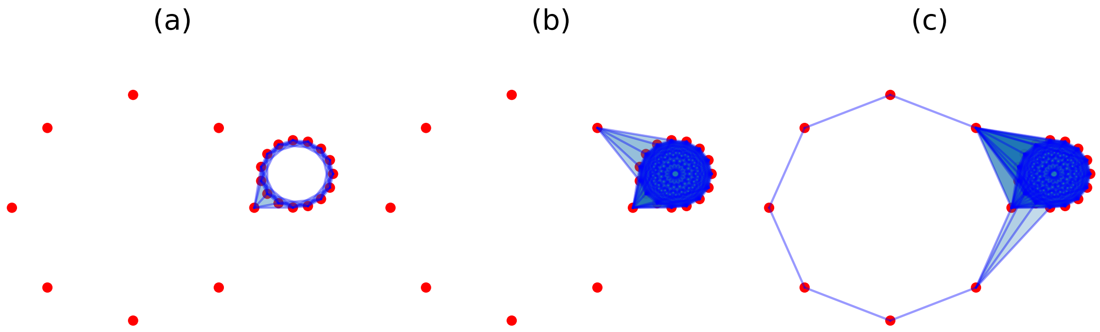

Figure 2: An example data set (red) showing that the Vietoris-Rips complex (blue) may not capture all relevant scales in data at once. The data has two obvious 1-dimensional holes. (a) a low radius may capture the smaller circle but not the larger one. (b) an intermediate radius already fills in the hole in the small circle before even detecting the larger circle. (c) only at higher choices of radius do we detect the larger circle.

Persistent homology is an attempt to address this problem by considering all possible scales, in this case all possible values of ϵ, and relating the homology of each complex making use of the functoriality of homology (see Appendix B for more details).

Proposition 2.1. *Let X, Y be simplicial (cubical) complexes and $f : X \to Y$ be a simplicial (cubical) map. Then $\forall k \in \mathbb{N}$ there is an induced linear map $f_k : H_k(X) \to H_k(Y)$ on k-th homology. The identity map $id : X \to X$ induces the identity map* *on homology. Moreover, if $g : Y \to Z$ is a simplicial (cubical) map, then*

$$(g \circ f)_k = g_k \circ f_k.$$

The maps we want to consider in the case of Vietoris-Rips complexes (and indeed most cases) are simply inclusion maps. Note that for $\epsilon \leq \epsilon'$ we have that $VR_\epsilon(P) \subseteq VR_{\epsilon'}(P)$, so in particular we can consider the inclusion map $VR_\epsilon(P) \hookrightarrow VR_{\epsilon'}(P)$. Let $\epsilon_1, \ldots, \epsilon_N$ be the values at which $VR_\epsilon(P)$ undergoes a change (or simply a change in homology). Then we can consider the sequence

$$VR_{\epsilon_1}(P) \hookrightarrow VR_{\epsilon_2}(P) \hookrightarrow \ldots \hookrightarrow VR_{\epsilon_N}(P)$$

known as a filtration of simplicial complexes. Applying the k-th homology functor yields the sequence

$$H_k(VR_{\epsilon_1}(P)) \to H_k(VR_{\epsilon_2}(P)) \to \cdots \to H_k(VR_{\epsilon_N}(P)) \tag{1}$$

of vector spaces known as a persistence module. This allows us not only to record the homology of the complex at different steps of the sequence, but also match up the homology classes at one step to those at later steps. See Figure 3 for an example.

The Vietoris-Rips complex is not the only way we might end up with a nested collection of complexes. Most of the filtrations we introduce in this thesis take the following more general form.

Definition 2.2 (Sublevel Set Filtration)**.** Let X be a simplicial (cubical) complex. A function $f : \Sigma(X) \to \mathbb{R}$ on the simplices (cubes) of X is a monotonic map if it maps each simplex (cube) to a value in $\mathbb{R}$ in such a way that if τ is a face of σ, then $f(\tau) \leq f(\sigma)$. Given such a map f, the sublevel set filtration of f is the collection of simplicial (cubical) complexes

$$\{ f^{-1}(-\infty, t] \mid t \in \mathbb{R} \}$$

along with inclusion maps

$$f^{-1}(-\infty, s] \hookrightarrow f^{-1}(-\infty, t]$$

for all $s \leq t \in \mathbb{R}$. Here f is known as a filter, and t is the filtration index.

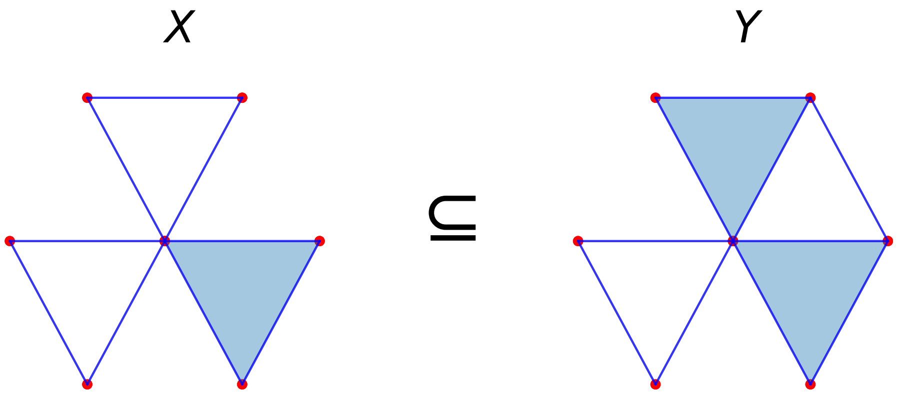

Figure 3: An example of the map induced on homology by an inclusion of simplicial complexes. Note that a shaded triangle represents a 2-simplex. If there is a triangle with no shading, we should think of this as simply three 1-simplices forming a hole that is not filled in by a 2-simplex. Suppose a basis of $H_1(X)$ is given as $\{[1,0]^\top, [0,1]^\top\}$ where the first vector corresponds to the lower left cycle and the second corresponds to the upper one, and suppose a basis of $H_1(Y)$ for the right space is also given as $\{[1,0]^\top, [0,1]^\top\}$ where the first vector corresponds to the lower left cycle and the second corresponds to the upper right one. Then the inclusion $X \hookrightarrow Y$ induces a linear map $H_1(X) \to H_1(Y)$ given by the matrix $\left(\begin{smallmatrix} 1,0 \\ 0,0 \end{smallmatrix}\right)$. Although both spaces have two 1-dimensional homology classes, only the lower left hole in X persisted after mapping into Y. The upper hole in X was killed off, and a new unrelated one was born in Y.

For example, we see that the Vietoris-Rips filtration for a point set P is simply the sublevel set filtration of the function on the complete simplicial complex $f : \mathcal{P}(P) \to \mathbb{R}$ that sends $\sigma \mapsto \mathrm{diam}(\sigma) = \max\{d(x,y) \,|\, x, y \in \sigma\}$. The functoriality of homology still allows us to describe the topology of this more general nested collection of complexes and inclusion maps. However, a sequence of vector spaces with linear maps between them is hardly a concise summary. In order to state the results that turn persistent homology into a practical tool, it will help to present some more precise definitions. The following exposition is inspired by [23]. It is worth noting that the following is just an overview of the theory in which there is much ongoing research. Definitions and theorems will be general enough for our purposes, but more restrictive than what is possible in an effort to simplify the picture.

Definition 2.3 (Persistence Module)**.** Fix some indexing set $T \subseteq \mathbb{R}$. A persistence module $\mathbb{V}$ over T is a family of vector spaces $\{V_t\}_{t \in T}$ along with linear maps $v_{s,t} : V_s \to V_t$ for all $s \leq t \in T$ such that $v_{t,t} = id$ and $v_{t,u} \circ v_{s,t} = v_{s,u}$ when $s \leq t \leq u$.

The sequence (1) was an example of a persistence module. Applying the k-th homology functor to the sublevel set filtration of a function f is also a persistence module which we denote by $H_k(f)$. In general, if $X = \{X_t\}_{t \in T}$ is a filtered complex, then we may write $H_k(X)$ for the persistence module coming from applying H_k to each X_t and inclusion map $X_s \hookrightarrow X_t$. The following is a particularly simple example of a persistence module.

Definition 2.4 (Interval Module)**.** Let $b, d \in \overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. The interval module $\mathbb{I}_{[b,d)}$ is defined

$$\mathbb{I}_{[b,d)} \quad = \quad \underbrace{0 \to \cdots \to 0}_{t<b} \to \underbrace{\mathbb{F} \xrightarrow{id} \ldots \xrightarrow{id} \mathbb{F}}_{b \leq t < d} \to \underbrace{0 \cdots \to 0}_{d \leq t}$$

and we may define the interval modules $\mathbb{I}_{(b,d)}$, $\mathbb{I}_{[b,d]}$, $\mathbb{I}_{(b,d]}$ similarly.

The following result tells us that sufficiently tame persistence modules can be fully described in a succinct manner according to these interval modules.

Theorem 2.2 (Interval Decomposition of Persistence Modules [24, 25])**.** *Let $\mathbb{V}$ be a persistence module over index set T. If either T is finite, or V_t is finite-dimensional for all $t \in T$, then there exists a direct sum decomposition*

$$\mathbb{V} = \bigoplus_{I \in \mathcal{B}(\mathbb{V})} \mathbb{I}_I$$

where $\mathcal{B}(\mathbb{V})$ is some collection of intervals in $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$. Moreover, the decomposition is unique up to isomorphism. We will say that a persistence module is interval decomposable if it has such a direct sum decomposition into interval modules.

Note that the homology of a sublevel set filtration of a finite simplicial or cubical complex always satisfies the conditions of the theorem. The collection of intervals $\mathcal{B}(\mathbb{V})$ is called the barcode of $\mathbb{V}$. In this thesis we will only consider finite complexes, so that the index set T is also finite. All of our intervals will therefore be of the form $[b, d)$, closed on the left and open on the right. We say that a homological feature is born at index b and dies at index d. We often throw away the information about whether or not the endpoints of bars are open or closed anyway to form the persistence diagram.

Definition 2.5 (Persistence Diagram)**.** Given a persistence module $\mathbb{V}$, its persistence diagram $\mathcal{D}(\mathbb{V})$ is the multiset

$$\mathcal{D}(\mathbb{V}) = \{\, (b,d) \in \overline{\mathbb{R}}^2 \mid b \text{ and } d \text{ are the endpoints of an interval in } \mathcal{B}(\mathbb{V}) \,\}$$

where (b,d) is included with the appropriate multiplicity if there are multiple intervals with these endpoints. If we are looking at the persistence module corresponding to the k-th homology of a filtered complex and it is clear which complex we are talking about, in this thesis we shall often refer to the corresponding persistence diagram as PH_k.

Figure 4 shows an example of a filtration of simplicial complexes and the barcode and persistence diagram of its persistent homology.

Given some data and a way to summarise it via a sequence of simplicial or cubical complexes, we can now talk about the associated persistence diagram. We would now like to be able to compare the persistent homology obtained from different data. We start by defining a natural algebraic distance between persistence modules that measures the extent to which two modules fail to be isomorphic.

Definition 2.6 (Interleaving Distance)**.** Given persistence modules $\mathbb{V} = \{\, V_s \xrightarrow{v_{s,t}} V_t \,\}$ and $\mathbb{W} = \{\, W_s \xrightarrow{w_{s,t}} W_t \,\}$, an ϵ-interleaving between $\mathbb{V}$ and $\mathbb{W}$ is given by two collections of linear maps $\{\, \phi_t : V_t \to W_{t+\epsilon} \,\}$ and $\{\, \psi_t : W_t \to V_{t+\epsilon} \,\}$ such that the following diagrams commute for all s, t.

$$\begin{tikzcd}
V_s \arrow[rr] \arrow[dr, "\phi_s"] & & V_t \arrow[dr, "\phi_t"] & \\
& W_{s+\epsilon} \arrow[rr] & & W_{t+\epsilon}
\end{tikzcd}$$

$$\begin{tikzcd}
& V_{s+\epsilon} \arrow[rr] & & V_{t+\epsilon} \\
W_s \arrow[ur, "\psi_s"] \arrow[rr] & & W_t \arrow[ur, "\psi_t"] &
\end{tikzcd}$$

$$\begin{tikzcd}
V_t \arrow[rr] \arrow[dr, "\phi_t"] & & V_{t+2\epsilon} \\
& W_{t+\epsilon} \arrow[ur, "\psi_{t+\epsilon}"] &
\end{tikzcd}$$

$$\begin{tikzcd}
& V_{t+\epsilon} \arrow[dr, "\phi_{t+\epsilon}"] & \\
W_t \arrow[ur, "\psi_t"] \arrow[rr] & & W_{t+2\epsilon}
\end{tikzcd}$$

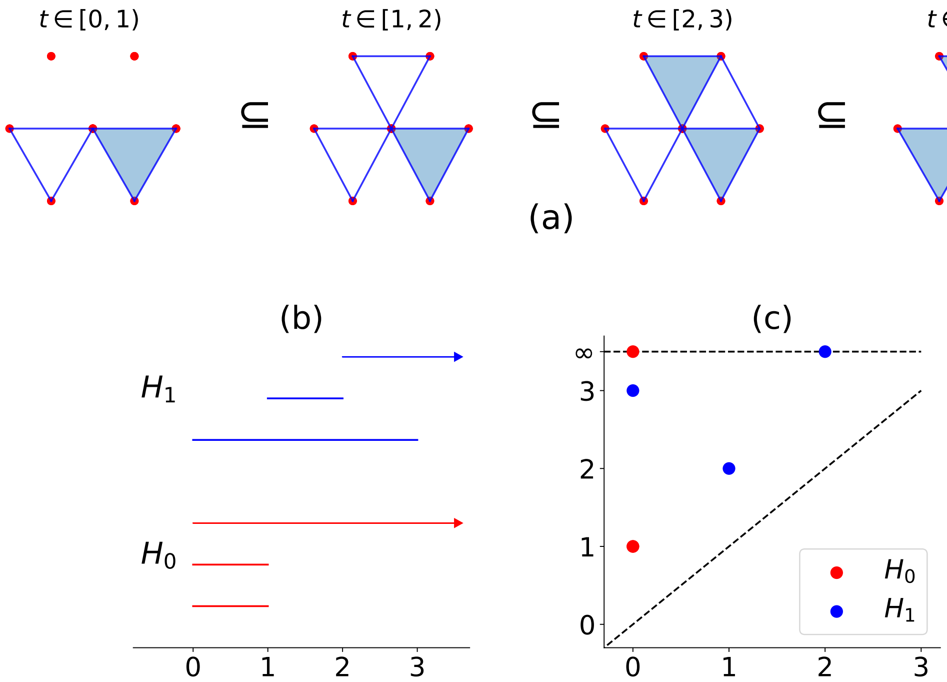

Figure 4: (a) An example of a filtration of simplicial complexes (assume for $t < 0$ we have the empty complex $\emptyset$). (b) The associated H_0 and H_1 persistence barcodes. The arrowheads indicate that the bars carry on to infinity. (c) The associated persistence diagram. Note that the H_0 point at $(0, 1)$ has multiplicity 2.

The interleaving distance between $\mathbb{V}$ and $\mathbb{W}$

$$d_I(\mathbb{V}, \mathbb{W}) = \inf \{ \, \epsilon \geq 0 \mid \exists \text{ an } \epsilon\text{-interleaving between } \mathbb{V} \text{ and } \mathbb{W} \, \}$$

is the smallest ϵ for which there is some ϵ-interleaving between the two.

It is not immediately obvious how to go about computing the interleaving distance, so we also introduce the Wasserstein and bottleneck distances on persistence diagrams which are more readily computable. The bottleneck distance will actually end up being equivalent to the interleaving distance, and this fact turns out to have important consequences for the use of persistence diagrams with real world data.

Definition 2.7 (Wasserstein and Bottleneck Distances)**.** Given persistence diagrams D_1 and D_2, a multiset $\eta \subseteq D_1 \times D_2$ is called a partial matching if it restricts to a bijection $\pi_1\eta \to \pi_2\eta$, where π_1 and π_2 are the projections on to D_1 and D_2 respectively. That is, we can think of it as pairing up some of the points in D_1

with points in D_2, but may leave points in either unmatched. The q-th Wasserstein distance between D_1 and D_2 is

$$d_{W,q}(D_1, D_2) = \inf_{\eta : D_1 \to D_2} \left(\sum_{(x,y)\in\eta} ||x-y||_\infty^q + \sum_{(b,d)\notin \mathrm{dom}(\eta)} \left|\frac{d-b}{2}\right|^q + \sum_{(b,d)\notin \mathrm{im}(\eta)} \left|\frac{d-b}{2}\right|^q \right)^{\frac{1}{q}}$$

where the latter two terms sum the L_∞ distances between the unmatched points and their closest point on the diagonal $\Delta = \{\, (b,b) \mid b \in \mathbb{R} \,\}$. The bottleneck distance

$$d_B(D_1, D_2) = \lim_{q\to\infty} d_{W,q}(D_1, D_2)$$

is calculated by replacing the sum in the above formula by the maximum of any of the L_∞ distances. An example of a partial matching realising the bottleneck distance is shown in Figure 5.

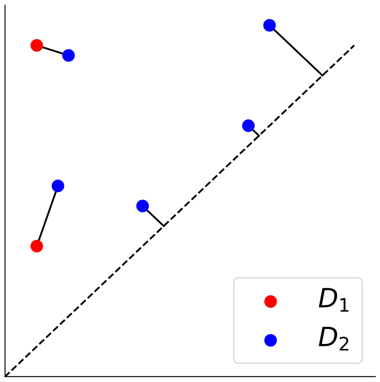

Figure 5: An example of a partial matching that realises the bottleneck distance between two persistence diagrams. Note that we show unmatched points as being matched to the diagonal. Matching the blue and red points in any other way will increase the length of the longest black line segment.

Theorem 2.3 (Isometry Theorem [26, 27]). *Let $\mathbb{V}$ and $\mathbb{W}$ be interval decomposable persistence modules. Then,*

$$d_B(\mathcal{D}(\mathbb{V}), \mathcal{D}(\mathbb{W})) = d_I(\mathbb{V}, \mathbb{W}).$$

The Isometry Theorem not only gives us a practical way to compute the interleaving distance between persistence modules by passing to the persistence diagram, but it also picks up special relevance in relation to the sublevel set filtration. In particular, it is not difficult to see that given two filter functions $f, g : X \to \mathbb{R}$ and $k \in \mathbb{N}$, the persistence modules $H_k(f)$ and $H_k(g)$ are $||f - g||_\infty$-interleaved. If $f(x) \leq t$, then $g(x) = f(x) + (g(x) - f(x)) \leq t + ||f - g||_\infty$. So the maps ϕ_t are simply given by applying H_k to the inclusions

$$f^{-1}(-\infty, t] \hookrightarrow g^{-1}(-\infty, t + ||f - g||_\infty]$$

and similarly for ψ_t. This gives us the following corollary.

Corollary 2.3.1 (Stability). *Let $k \in \mathbb{N}$ and $f, g : X \to \mathbb{R}$ be filter functions on a simplicial or cubical complex X. Then,*

$$d_B(\mathcal{D}(H_k(f)), \mathcal{D}(H_k(g))) \leq ||f - g||_\infty.$$

Thus a small change in our data, in this case represented by a change in the filter function, leads only to a small change in the resulting persistence diagram as measured by the bottleneck distance. Another way to say this is that, as a map between metric spaces, $\mathcal{D} \circ H_k$ is 1-Lipschitz. This is obviously important if our data has any kind of noise, as it guarantees that the persistence diagram will not change dramatically. A fact that was only proved much later is that we also have the same for the q-Wasserstein distances.

Theorem 2.4 (Wasserstein Stability [28]). *Let $k \in \mathbb{N}$ and $f, g : X \to \mathbb{R}$ be filter functions on a simplicial or cubical complex X. Then for all q, we have*

$$d_{W,q}(\mathcal{D}(H_k(f)), \mathcal{D}(H_k(g))) \leq ||f - g||_\infty.$$

Let us summarise the above by giving the typical persistent homology framework. Given some data D, we construct a filtered complex F_D. That is, a $\mathbb{R}$-indexed collection of simplicial or cubical complexes $\{ F_D(s) \}_{s \in \mathbb{R}}$ such that for all $s \leq t \in \mathbb{R}$ we have an inclusion map $F_D(s) \hookrightarrow F_D(t)$. Applying the k-th homology functor to each space and map gives us a persistence module that identifies the k-dimensional holes in each complex and identifies when a hole in one complex is the same as that in another. The interval decomposition theorem then shows that this persistence module can be completely and uniquely described by a collection of intervals which we represent using a persistence diagram consisting of points $(b, d) \in \overline{\mathbb{R}}^2$ indicating a hole that enters the filtered complex at index b and persists through to index d. The stability property of persistent homology tells us that small perturbations to our input filtered complex lead only to small changes in the persistence diagram with respect to a couple of natural metrics.

There are at least two paradigms for using persistent homology to study phase transitions of a given statistical physics model. The first can be called *persistent homology of configuration/data space*, where the topology of the high-dimensional space of model configurations is probed from samples. This approach is based on the *topology hypothesis* for the origin of phase transitions [29, 30]. The idea here is that a thermodynamic phase transition necessarily coincides with a change in the topology of the energy level sets, although such a change does not turn out to be a sufficient condition [31]. In the present work however, we shall make use of a newer paradigm which we call *persistent homology as an observable*. With this approach we take a sampled configuration of a model, construct a filtered complex based on that single configuration, and compute persistence. We can think of this process as a means to reduce the degrees of freedom of the model and produce nonlinear summaries of configurations. Or as a way to focus only on certain topological degrees of freedom. Statistics of these persistence diagrams are then analysed as the system undergoes a phase transition.

2.1.2 Computing Persistent Homology

One of the motivations behind the use of homology was that it is straightforward to compute. It turns out the computation of persistent homology is also straightforward. The following is based on the exposition in [20] and assumes that our field is $\mathbf{F} = \mathbb{Z}_2$ (which it will be throughout this thesis).

Given a filtered complex (simplicial or cubical) recorded as filtration function $f : \Sigma \to \mathbb{R}$ on a set of simplices or cubes, we first define a total order on Σ compatible with f, so that $f(\sigma) < f(\tau)$ implies that $\sigma < \tau$. Let $\sigma_1, \sigma_2, \ldots, \sigma_n$ denote the full collection of simplices or cubes according to this order. We now define a square boundary matrix D of size $n \times n$ so that

$$D_{i,j} = \begin{cases} 1 & \text{if } \sigma_i \text{ is a codimension-1 face of } \sigma_j \\ 0 & \text{otherwise} \end{cases}$$

and for any $n \times n$ $\mathbb{Z}_2$ matrix A we define a partial mapping $\mathrm{low}_A : \{1, \dots, n\} \to \{1, \dots, n\}$ such that

$$\mathrm{low}_A(j) = \begin{cases} \max\{\, i \mid A_{i,j} = 1 \,\} & \text{if column } j \neq 0 \\ \text{undefined} & \text{otherwise.} \end{cases}$$

That is, we can think of $\mathrm{low}_A(j)$ as finding the 1 furthest down the j-th column and returning its row number. What we are interested in is a reduced version of the boundary matrix, where reduced means that low is injective on its domain of definition: no two columns have their furthest down 1 on the same row. The simplest possible algorithm for obtaining a reduced boundary matrix via column operations is shown in Algorithm 1, first introduced in [17].

Algorithm 1 Column algorithm for persistent homology [17]

Input Boundary matrix $D \in M_{2,2}(\mathbb{Z}_2)$
Output Reduced boundary matrix $R \in M_{2,2}(\mathbb{Z}_2)$

```
$R \leftarrow D$
for $j = 1$ to $n$ do
    while $\exists i < j$ with $\mathrm{low}_R(i) = low_R(j)$ do
        $R[:, j] \leftarrow R[:, j] + R[:, i]$                ▷ add column $i$ to column $j$
    end while
end for
```

Once the reduced boundary matrix R is obtained, the barcode can be read as follows. For each $j \in \{1, \dots, n\}$:

- If $\mathrm{low}_R(j) = i$, then there is a homology class that is born from the inclusion of σ_i and which dies with the inclusion of σ_j. We therefore get a bar $[f(\sigma_i), f(\sigma_j))$ in $\mathcal{B}(H_{\dim(\sigma_i)}(f))$.

- Otherwise if there is some k such that $\mathrm{low}_R(k) = j$, then there is a homology class that is born from the inclusion of σ_j and which dies with the inclusion of σ_k. We therefore get a bar $[f(\sigma_j), f(\sigma_k))$ in $\mathcal{B}(H_{\dim(\sigma_j)}(f))$.

- If j is not in the domain or image of low_R at all, then it is unpaired and instead we get an infinite bar $[f(\sigma_j), \infty)$ in $\mathcal{B}(H_{\dim(\sigma_j)}(f))$.

Note that it may be the case that we get bars of zero length from the procedure above and we can simply throw these away. Algorithm 1 has cubic complexity in the number of simplices or cubes in the worst case. However there are several different ways to speed up the computation. In this thesis we use the giotto-tda [32] and Ripser.py [33] software to computer persistent homology. Both of these actually compute persistent cohomology, which turns out to give the same barcode as persistent homology and allows for some optimisations [34, 35], including doing row reduction instead of column reduction. The key theorem is the following.

Theorem 2.5 (Universal coefficient theorem [36, Thm. 3.2])**.** *Given a topological space X and field coefficients $\mathbf{F}$, we have a natural isomorphism*

$$H^k(X) \cong \operatorname{Hom}(H_k(X), \mathbf{F}).$$

Thus the homology and cohomology at each step of the filtration have the same dimension and the naturality means that the induced maps along the filtration have the same rank. This implies the following.

Corollary 2.5.1 (Equivalence of persistent homology and cohomology [34])**.** *Given a filtered complex $X = \{X_t\}_{t \in T}$, we have $\mathcal{B}(H_k(X)) = \mathcal{B}(H^k(X))$.*

2.1.3 Vectorising Persistence Diagrams

Being multisets of points, persistence diagrams do not lend themselves towards immediate use with many typical data analysis and machine learning methods. We do not have an obvious way to compute the mean or variance of a set of diagrams and any machine learning approach that requires an input of fixed size is inapplicable. Methods that rely only on the distance matrix of a set of persistence diagrams theoretically work thanks to the bottleneck and Wasserstein distances, but these become unfeasible to compute when the set of diagrams and the diagrams themselves become large. Instead, one typically uses a vectorisation method, which transforms a set of persistence diagrams into a set of vectors lying in a single vector space. There are many methods and this is an active area of research. A selection of methods include [37–39], but in this thesis we focus on the persistence image [37] as it closely mimics the structure of the persistence diagram and allows one to relate components of the vectorisation back to regions of the diagram.

Definition 2.8 (Persistence Image)**.** Let $g_{a,b} : \mathbb{R}^2 \to \mathbb{R}$ denote a 2D Gaussian of standard deviation σ centered at (a, b):

$$g_{a,b}(x, y) = \frac{1}{2\pi\sigma^2} exp\left[- \frac{(x-a)^2 + (y-b)^2}{2\sigma^2} \right].$$

Fix a non-negative weighting function $w : \mathbb{R}^2 \to \mathbb{R}$ that is zero along the horizontal axis, continuous, and piecewise differentiable. Given a persistence diagram $D = \{(b_i, d_i)\}_{i \in I}$ consisting of a finite number of points with finite birth and death, its persistence surface with respect to w is the function $\rho_D : \mathbb{R}^2 \to \mathbb{R}$ obtained by translating each point $(b, d) \in D$ into birth-persistence coordinates $(b, d - b)$, then placing Gaussians with variance σ^2 on them, weighted by w:

$$\rho_D(x, y) = \sum_{(b,d) \in D} w(b, d - b)\, g_{b,d-b}(x, y).$$

The persistence image $PI(D)$ is obtained by discretizing a rectangular region of the domain of ρ_D into a collection of $n_I \times n_I$ pixels p_i and integrating ρ_D within each:

$$PI^i(D) = \iint_{p_i} \rho_D(x, y) dx dy.$$

In this way we obtain a $(n_I)^2$-dimensional vector representing our persistence diagram. See Figure 6 for an example. So long as we choose the same w, σ and discretization for each diagram, we can compute averages and variances componentwise. As observed in [40], a persistence diagram can be thought of as a discrete measure on $\{(a, b) \in \mathbb{R}^2 \mid a \leq b\}$. If we are sampling data from some distribution and the expected persistence diagram measure has a density with respect to the Lebesgue measure on $\{(a, b) \in \mathbb{R}^2 \mid a \leq b\}$, then the average of the persistence images can be thought of as an estimator for this density, multiplied by an additional weighting of w.

Besides de-emphasising low-persistence points, the weighting by w ensures the stability of the persistence image with respect to the 1-Wasserstein distance.

Theorem 2.6 (Persistence Image Stability [37])**.** *Let D_1, D_2 be persistence diagrams computed with weighting function w, variance σ^2 and the same discretisation. Then for $l \in \{1, 2, \infty\}$,*

$$||PI(D_1) - PI(D_2)||_l \leq \left(\sqrt{5}|\nabla w| + \sqrt{\frac{10}{\pi}} \frac{||w||_\infty}{\sigma} \right) d_{W,1}(D_1, D_2)$$

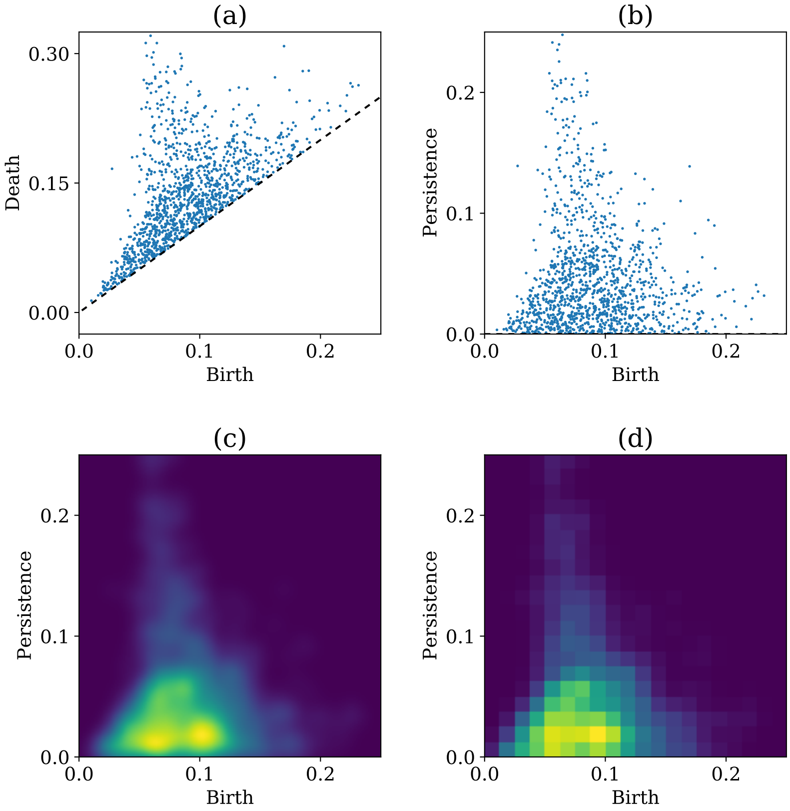

Figure 6: An illustration of how the persistence image is obtained from a persistence diagram (a). It is first transformed into birth-persistence coordinates (b), then the persistence surface (c) is computed before discretisation, yielding the persistence image (d).

where $|\nabla w| = \sup_{(x,y)\in\mathbb{R}^2} ||\nabla w(x,y)||_2$.

Finally we note that, as discussed in [37], machine learning models trained on persistence images are generally insensitive to the resolution and variance parameters n_I and σ.

Another vector representation we will make use of is the Persistence Landscape [39]. First let us define persistent Betti numbers. Recall that Betti numbers record the dimension of homology, $\beta_k(X) = \dim H_k(X)$. Persistent Betti numbers simply record the dimension of homology after a linear map is applied.

Definition 2.9 (Persistent Betti Numbers)**.** Let $\mathbb{V}$ be a persistence module indexed over T consisting of vector spaces V_t and linear maps $v_{s,t} : V_s \to V_t$ for $s \leq t \in T$.

The (s,t)-th persistent Betti number is

$$\beta^{s,t}(\mathbb{V}) = \dim \operatorname{im} v_{s,t}$$

where we will omit $\mathbb{V}$ if it is clear which persistence module we are talking about.

Note that the collection of persistent Betti numbers associated to a persistence module is intimately related to the barcode. The difference $\beta^{s,t} - \beta^{s-\epsilon,t}$ records the number of bars born in the interval $(s-\epsilon, s]$ which live at least as long as t. The difference $\beta^{t-\epsilon,t-\epsilon} - \beta^{t-\epsilon,t}$ tells us the number of bars that die in the interval $(t-\epsilon, t]$. Equipped with this definition we introduce the persistence landscape.

Definition 2.10 (Persistence Landscape [39])**.** Given a persistence module $\mathbb{V}$, its persistence landscape λ is the sequence of functions $\lambda_k(t) : \mathbb{R} \to \overline{\mathbb{R}}$ with

$$\lambda_k(t) = \sup\{m \geq 0 \,|\, \beta^{t-m,t+m} \geq k\}$$

defined for all $k \in \mathbb{N}$. We put the following distances on the space of persistence landscapes:

$$\Lambda_p(\lambda, \lambda') = \big(\sum_{k=0}^{\infty} ||\lambda_k - \lambda'_k||_p\big)^{\frac{1}{p}}.$$

As with persistence images, the persistence landscape is stable with respect to the metrics we put on it and the interleaving or bottleneck distance. In particular, setting $p = \infty$ to obtain the supremum metric, we have the following result.

Theorem 2.7 (Stability of Persistence Landscapes [39])**.** *Let* λ, λ' *be the persistence landscapes associated to persistence modules* $\mathbb{V}, \mathbb{V}'$ *respectively. Then,*

$$\Lambda_\infty(\lambda, \lambda') \leq d_I(\mathbb{V}, \mathbb{V}').$$

However, we do not yet have a finite-dimensional vectorisation. In order to obtain that we simply sample our persistence landscapes at some fixed finite set of sample points.

Definition 2.11 (Discretised Persistence Landscape)**.** Given a persistence landscape λ and fixing a finite point sample $P = \{r_1, \ldots, r_N\} \subset \mathbb{R}$ and some $k_{\max}$, the discretised persistence landscape λ^P is the set of $(k_{\max}+1)$ N-vectors $\lambda^P_k = [\lambda_k(r_1), \ldots, \lambda_k(r_N)]$ for each $0 \leq k \leq k_{\max}$.

So long as we choose the same sample P each time (usually a uniform partition of an interval), we can compute sums and differences of different discretised landscapes.

2.1.4 Representative (Co)cycles

The following is based on the exposition in [41].

Given a filtered complex $\{X_t\}_{t\in T}$, a point (b, d) in the resulting persistence diagram PH_k represents a related set of classes $[c_t] \in H_k(X_t)$ for $b \leq t < d$. One might ask for a representative k-cycle c for the whole interval $[b, d)$. That is,

$$\mathrm{Birth}(c) = \inf\{\, t \mid c \in Z_k(X_t) \,\} = b$$

and

$$\mathrm{Death}(c) = \sup\{\, t \mid [c] \neq [0] \in H_k(X_t) \,\} = d$$

where we let the supremum of an unbounded set be equal to ∞. For example if we detect a vortex in the plane as a 1-dimensional hole, as we will in Chapter 3, we might want a 1-cycle (a closed curve) which bounds that hole, letting us identify the location of the vortex. It turns out that such cycle representatives exist.

Proposition 2.8 (Persistent homology cycle representatives [42])**.** *Every filtration $X = \{X_t\}_{t\in T}$ of simplicial or cubical complexes admits a persistent homology cycle basis $\mathfrak{B}$. That is,*

$$\mathcal{D}(H_k(X)) = \{\, (\mathit{Birth}(c), \mathit{Death}(c)) \mid c \in \mathfrak{B},\ \dim(c) = k \,\}$$

with the appropriate multiplicity, for all k.

In fact, such a basis may be computed using a standard persistent homology algorithm [43] (see the supplementary materials in [41] for more detail). However, there are a couple of issues to address. The first is that this basis is by no means unique and there may be many choices of persistent cycle representative for each point in the persistence diagram. The second is a consequence of the first: cycle representatives are not a stable feature. A small change in the underlying data can cause an arbitrarily large change in cycle representatives. This remains true even if one tries to select a specific representative such as that which minimises volume and then simply measures that volume [44].

Recall also from Section 2.1.2 that the persistent homology software we use computes persistent cohomology instead of homology. While we receive the same barcode or persistence diagram at the end, calculating in this way does not give representative cycles as a byproduct, but rather representative cocycles $c \in C_k^* = \mathrm{Hom}(C_k, \mathbf{F})$. In this thesis we will work exclusively with $\mathbf{F} = \mathbb{Z}_2$, so a cocycle can also be thought of as the set of those k-simplices belonging to any chain that is sent to 1. More precisely, we define the following.

Definition 2.12 (Support of a cochain)**.** Let X be a space. Given a cochain $c \in C^k(X) = \mathrm{Hom}(C_k(X), \mathbf{F})$, its support is the subset

$$\mathrm{Supp}(c) = \sup \left\{ A \subseteq X \;\middle|\; c|_{X \setminus A} = 0 \right\}$$

where the supremum is taken with respect to inclusion.

Intuitively, we ought to think of this subset as representing a cut which kills the corresponding homology class. That is, removing these simplices from the complex would remove the hole. An example is shown in Figure 7.

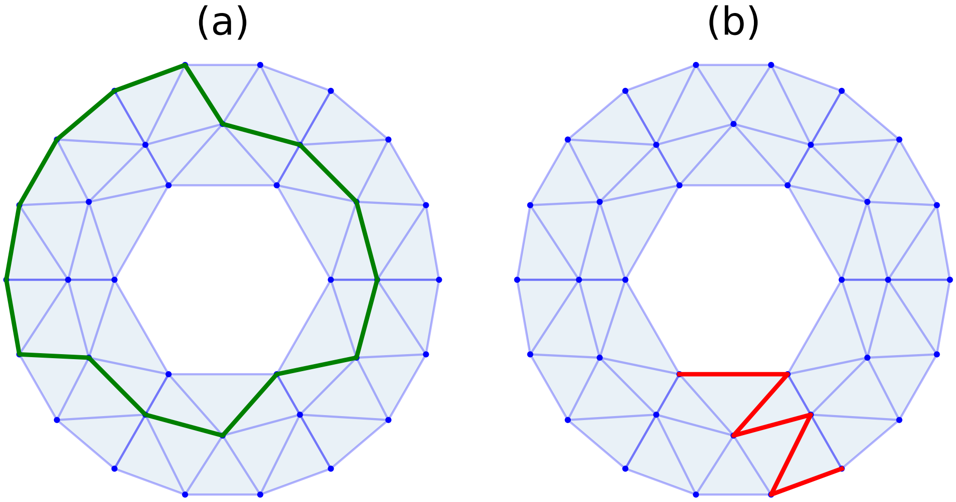

Figure 7: A representative cycle and cocycle for the single non-trivial H_1 class of a simplicial complex X homeomorphic to an annulus. (a) A representative 1-cycle $c \in C_1(X)$ in green such that $[c] \neq [0] \in H_1(X)$. We see that c loops round the hole. (b) A representative 1-cocycle $f \in C^1(X) = \mathrm{Hom}(C_1(X), \mathbb{Z}_2)$ such that $[f] \neq [0] \in H^1(X)$. The red edges indicate the support of f and we see that removing these edges would destroy the hole.

In practise, available software for computing representative cycles or cocycles is sparse. In this thesis we will make use of the Ripser.py [33] software which is able to give representative cocycles for filtered simplicial complexes of a particular form.

2.2 Machine Learning

For our quantitative analysis of phase transitions we will want to assign configurations of models to phases based on their persistence images. In order to do this we will make use of simple machine learning binary classifiers which learn a function $\mathbb{R}^n \to \{0, 1\}$ based on some set of labelled training data. In this case, n will be the number of pixels in our persistence images and the output will indicate which phase the original configuration belongs to. The point is that instead of using more advanced deep-learning classifiers, we introduce only simple models that are either non-parametric or that have only a few parameters to learn, relying on the persistent homology to have selected the relevant features. Then we may more easily interpret the performance of the classifiers in terms of what the persistent homology detects.

2.2.1 Logistic Regression

Following the approach introduced in [45], we will train logistic regression models to map the persistence images obtained from configurations of models onto phases. Logistic regression is a generalised linear model which models a binary dependent variable $y(\mathbf{x}) \in \{0, 1\}$. For input $\mathbf{x} \in \mathbb{R}^N$, a logistic regression model is parameterised by a weight vector $\mathbf{w} = (w_1, \ldots, w_N)^T \in \mathbb{R}^N$ and intercept $b \in \mathbb{R}$. Its output is a logistic function

$$p_{\mathbf{w},b}(\mathbf{x}) = \frac{1}{1 + e^{\mathbf{x}^T\mathbf{w}+b}} \in (0, 1)$$

which can be interpreted as the probability that $y(\mathbf{x}) = 1$, with $1 - p(\mathbf{x})$ giving the probability that $y(\mathbf{x}) = 0$. Given training data $\{(\mathbf{x}_i, y_i)\}$, the weights $\mathbf{w}$ and intercept b are learnt by minimising a cross-entropy loss function

$$J(\mathbf{w}, b) = -\sum_i \left[y_i log(p_{\mathbf{w},b}(\mathbf{x}_i)) + (1 - y_i)log(1 - p_{\mathbf{w},b}(\mathbf{x}_i))\right] + \frac{1}{C}(\mathbf{w}^T\mathbf{w} + b^2).$$

The first term penalises misclassifications with the penalty increasing as the confidence in the incorrect classification increases. The second term implements ℓ_2

regularisation, reducing overfitting by preventing the weights from becoming too large, where C is a hyper-parameter controlling the amount of regularisation.

In our case, $\mathbf{x}$ will be a persistence image, $y(\mathbf{x}) = 0$ will indicate one phase, and $y(\mathbf{x}) = 1$ will indicate another. As in [45] we will train the model using data drawn deep into both phases. However, since we are interested in making a precise estimate of the critical point of a phase transition, we will use data closer to the critical region than in [45]. After successful training, the weights will indicate features in the persistence image characteristic of each phase. Weights $w_j < 0$ will indicate features of one phase, and weights $w_j > 0$ will indicate features of the other. In the intermediate range where there is no training data, the logistic regression model will output an estimated classification $O_{LR} \in \{0, 1\}$ depending on whether p is less than or greater than 0.5. Notice that we clamp the output to 0 or 1 rather than using the direct output of the logistic function. We found that this leads to better finite-size scaling behaviour later on. We may then treat $\langle O_{LR} \rangle$ as a phase indicator. In this thesis we shall be interested in the distribution of O_{LR} at different values of β and different lattice sizes.

2.2.2 k-Nearest Neighbours

We will also make use of k-nearest neighbours (k-NN) classification to map persistence images onto phases. This is a non-parametric model which models a categorical dependent variable $y(\mathbf{x}) \in \mathbb{N}$, where $\mathbf{x} \in \mathbb{R}^N$. The behaviour of the model is determined by the training data $\{(\mathbf{x}_i, y_i)\}$ and a choice of the hyper-parameter $k \in \mathbb{N}$. Given new input $\mathbf{x}$, it finds the k indices $i_1(\mathbf{x}), \ldots, i_k(\mathbf{x})$ that minimise the Euclidean distance $||\mathbf{x} - \mathbf{x}_i||_2$. It then outputs the most common label among the $y_{i_1(\mathbf{x})}, \ldots, y_{i_k(\mathbf{x})}$.

As in the case of logistic regression, $\mathbf{x}$ will be a persistence image, $y(\mathbf{x}) = 0$ will indicate one phase, and $y(\mathbf{x}) = 1$ will indicate the other phase. We will train the model using data drawn in both phases close to the critical region. In the critical region where there is no training data, the k-NN model will output an estimated classification $O_{kNN} \in \{0, 1\}$. We may then treat $\langle O_{kNN} \rangle$ as a phase indicator.

2.3 Critical Exponents and Finite Size Scaling

In this section we assume a basic knowledge of statistical physics and phase transitions. A brief review of these can be found in Appendix A.

A continuous phase transition is marked by diverging correlation length ξ and fluctuations (e.g., the specific heat). One way to characterise the transition is to ask how fast these quantities diverge as we approach the critical parameter β_c. Letting $t = (\beta - \beta_c)/\beta_c$ be a reduced parameter, we find that they diverge with a power law with, for example, $\xi \sim |t|^{-\nu}$ for the correlation length. We call ν the critical exponent of correlation length for the phase transition and there is a whole collection of such exponents associated to each phase transition, with one for each diverging quantity we are looking at. These critical exponents, along with the transition parameter β_c characterise the phase transition. In fact, the same set of values for the critical exponents may characterise several different phase transitions and we say that these transitions lie in the same universality class. For example, we consider a phase transition later in a variant of the 2D XY model which lies in the same universality class as the 2D Ising model. The two models have very different microscopic properties, but it turns out to be the macroscopic properties such as lattice dimension and the $\mathbb{Z}_2$ symmetry in both that determine their critical exponents. A full explanation for this behaviour would require introducing the renormalisation group which is beyond this introduction, but a good reference is [46].

Consider a quantity measured in the continuum limit A_∞ that diverges at the critical point with critical exponent ζ so that $A_\infty(\beta) \sim |t|^{-\zeta}$. Then we have that

$$A_\infty(\beta) \sim |t|^{-\nu\zeta/\nu} \sim \xi^{\zeta/\nu}.$$

But now for a finite system size L close to the critical point we have $L \ll \xi$ since ξ diverges and L takes over as the cutoff, giving us that

$$A_L(\beta) \sim L^{\zeta/\nu},$$

not diverging but displaying a more and more pronounced peak as L increases. This in turn lets us write down the following finite-size scaling ansatz [47]

$$A_L(\beta) = L^{\zeta/\nu}\hat{A}(L^{1/\nu}t) \tag{2}$$

where $\hat{A}$ is a dimensionless function.

In the case of the phase transition in the 2D XY model, known as the Berezinskii-Kosterlitz-Thouless (BKT) transition, we see something slightly different. Here the correlation length diverges much faster $\xi \sim \exp(b|t|^{-\nu})$ where $b > 0$ is some non-universal (i.e., model specific) constant [48]. Now the finite-size scaling ansatz becomes

$$A_L(\beta) \approx L^{\zeta/\nu} \hat{A}(L \exp(-b|t|^{-\nu})) \tag{3}$$

where we have ignored some small logarithmic corrections.

By simulating close to the phase transition on different lattice sizes L we can extract the heights and locations of the different peaks in A_L then fit these to Equation 2 or Equation 3 as appropriate to estimate β_c, ζ and ν. Note that the logarithmic corrections we ignored in the BKT case mean that this method is not typically used for high precision studies, where approaches based on the spin stiffness [49] are more common.

We might expect the persistent homology of a configuration to demonstrate large variations at criticality. We quantify this by looking at the fluctuations in the output O_{LR} and O_{kNN} of the trained logistic regression and k-NN models, measuring the variance

$$\begin{aligned} \chi_{LR}(\beta) &= \langle O_{LR}^2 \rangle_\beta - \langle O_{LR} \rangle_\beta^2 \\ &= \langle O_{LR} \; \rangle_\beta (1 - \langle O_{LR} \rangle_\beta). \end{aligned} \tag{4}$$

Note that the second equation follows since O_{LR} takes values in $\{0, 1\}$. This will display a peak, indicating where the model is least certain about which phase configurations are from, when $\langle O_{LR} \rangle_T$ crosses 0.5. χ_{kNN} is defined similarly. We find evidence that these quantities may also display finite-size scaling behaviour similar to Equations 2 and 3 which we will use to estimate the critical β and the critical exponent of correlation length ν.

If ν is known, we may estimate β_c by fitting the peak locations $\beta_c(L)$ of χ_{LR} and χ_{kNN} obtained from multiple lattice sizes to the ansatz

$$\beta_c(L) - \beta_c(\infty) \propto \frac{1}{L^{1/\nu}}. \tag{5}$$

in the case of a second order transition, or

$$\beta_c(L) - \beta_c(\infty) \propto \frac{1}{\log(L)^{1/\nu}}. \tag{6}$$

for a BKT transition.

To estimate ν (as well as β_c), we can use a curve collapse approach, plotting $y = \chi_{LR}$ or $y = \chi_{kNN}$ for multiple lattice sizes simultaneously against $x = L^{1/\nu} t$ (second order) or $x = L \exp(-bt^{-\nu})$ (BKT). According to Equations 2 and 3, if ν, β_c, and b (in the case of a BKT transition) are correct then the curves for the different lattice sizes should collapse onto one another. Thus finding values of ν and β_c which minimise the distance between the curves using the Nelder-Mead method, as in the procedure described in [50], should give us an estimate of their true values.

2.4 Statistical Techniques

The following are techniques we use to improve and estimate the error in our estimates of critical parameters and exponents.

2.4.1 Histogram Reweighting

Histogram reweighting allows us to express the ensemble average of an observable O at an inverse temperature (or inverse coupling) β' in terms of averages at any other value β according to the equation

$$\langle O \rangle_{\beta'} = \frac{\langle O e^{-(\beta'-\beta)E} \rangle_\beta}{\langle e^{-(\beta'-\beta)E} \rangle_\beta} \tag{7}$$

where E is the energy (or action) of a configuration [51]. However, in practice we can only reweight so far, so that the energy distributions for β and β' have a sizable overlap. To reliably extrapolate to a wider region we can make use of multiple histogram reweighing [52] where we sample at multiple values $\beta_1, \ldots, \beta_R$. Suppose we sample N_i configurations at parameter β_i, then we can iterate the equation

$$e^{-f_\beta} = \sum_{i=1}^{R} \sum_{a=1}^{N_i} \frac{g_i^{-1} e^{-\beta E_i^a}}{\sum_{j=1}^{R} N_j g_j^{-1} e^{-\beta_j E_i^a + f_j}}$$

to estimate the free energies $f_i = f_{\beta_i}$ up to an additive constant, where each g_i is a quantity related to the integrated autocorrelation of the samples in run i. Given the f_i we can estimate

$$\langle O \rangle_{\beta'} = \sum_{i=1}^{R} \sum_{a=1}^{N_i} \frac{O_i^a g_i^{-1} e^{-\beta_k E_i^a + f_{\beta'}}}{\sum_{j=1}^{R} N_j g_j^{-1} e^{-\beta_j E_i^a + f_j}}.$$

2.4.2 Bootstrap Error Estimation

In order to make any reasonable conclusions from the results of our analysis we need to be able to estimate the error in any numerical values obtained. While the error in ensemble averages can be directly estimated from the sample, we also calculate various fits to the data. The way in which error propagates here is not necessarily easy to calculate directly. The idea of bootstrap analysis is to sidestep these concerns by estimating the sampling distribution of a statistic directly. Suppose we obtain N sampled configurations $S = \{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_N\}$ and calculate some numerical statistic $f(S)$ from the data. Given some preset integer N_B, bootstrap analysis proceeds by:

1. resampling S with replacement N_B times to obtain samples $S_1, \ldots, S_{N_B}$ each of size N; then
2. computing $f(S_i)$ for each $i \in \{1, \ldots, N_B\}$.

For large enough N_B, the distribution of the $f(S_i)$ approximates the sampling distribution of f and we can estimate the standard error

$$\sigma_f \approx \sqrt{\frac{1}{N_B} \sum_i \left(f(S_i) - \overline{f(S_j)} \right)^2}.$$

2.5 Previous Work

As mentioned in the introduction to this thesis, there is a growing body of work on using machine learning to analyse phase transitions. An incomplete list of references includes [53–67]. More recently, among other geometric and topological approaches [68–71], there has been an interest in using persistent homology to produce interpretable features which are inherently sensitive to topological objects. These can then be compared in their own right, or fed into a machine learning model [45, 72–76]. A useful review is [77].

Recall that we mentioned two paradigms for using persistent homology to study phase transitions of a given statistical physics model: *persistent homology of configuration/data space* and *persistent homology as an observable*.

The first approach has not seen much activity. In [76], Donato et al. sample configurations at energy levels below, at, and above the critical point of phase transitions in the mean-field XY model and ϕ^4 model. Estimating the geodesic distance between configurations, they compute Vietoris-Rips H_0 and H_1 persistence for each such set of samples and compare the results to see if there has been a topology change across the phase transition, finding some qualitative evidence for this change. A reason this approach has not seen further work, particularly of a more quantitative nature, is that it is very computationally demanding. The configuration spaces involved are very high dimensional – a modestly sized lattice model might have tens of thousands of variables. Therefore, not only is a large number of configurations needed to sufficiently sample the space, but the topology change may be manifested in high dimension so that looking at H_0 and H_1 may not be sufficient. However, including higher dimensions in a Vietoris-Rips complex causes a combinatorial explosion in the number of simplices and therefore computational cost.

The second approach, which we focus on in this thesis, has seen much more work in recent years. Some references include [45, 72–75, 78–80]. Previous works have focused on identifying the different phases in various models in a mostly qualitative manner. While [45] makes some steps towards obtaining quantitative measurements of the multiscale structure of the Ising model at criticality, the work in this thesis constitutes some of the first work on developing a framework for using persistent homology observables to make rigorous numerical estimates of critical temperatures or couplings and critical exponents with quantified error. We will make more specific remarks on some of these previous works in Chapters 3 and 4 as relevant.

3 XY Models

In this chapter we look at the classical 2D XY model as well as two variants. These are spin models defined on a 2-dimensional lattice which exhibit phase transitions driven by topological defects, and therefore serve as a good testbed for the use of topological data analysis methods in statistical physics.

3.1 The XY Model and BKT Transition

The 2-dimensional XY model is defined on an $L \times L$ square lattice Λ by assigning an angle $\theta_i \in S^1$ to each lattice site $i \in \Lambda$. The energy of a given configuration of spins $\boldsymbol{\theta} = \{\theta_i\}_{i \in \Lambda}$ is given by the Hamiltonian

$$H(\boldsymbol{\theta}) = -\sum_{\langle ij \rangle} \cos(\theta_i - \theta_j)$$

where $\langle ij \rangle$ ranges over all pairs of nearest neighbour lattice sites. At low temperatures spins tend to align with their neighbours and at high temperatures they are able to become more and more disordered. We characterise this high temperature behaviour as usual by saying that there is short-range order. That is, the correlation function $C(T, r) = \langle \cos(\theta_0 - \theta_r) \rangle_T$ decays exponentially with distance

$$C(T, r) \sim \exp(-\frac{r}{\xi})$$

with a correlation length ξ depending on temperature. However, unlike the Ising model considered in Appendix A, we do not have long-range order at any temperature. Instead we find that the correlation function decays as a power law

$$C(T, r) \sim (\frac{1}{r})^{\frac{T}{2\pi}}$$

for all temperatures below some temperature T_c. This is called quasi long-range order. There is a qualitative change in the behaviour of the correlation function so we must have a phase transition at critical temperature T_c, but it is unlike most we see in statistical physics. Usually power law decay of correlations only happens at the critical temperature itself, where as now it is happening in the whole interval $T \in (0, T_c]$. The reason for this behaviour was uncovered by Kosterlitz and Thouless [81, 82] for which they received the 2016 Nobel Prize in physics along with Haldane.

Berezinskii had also been able to figure out what drove this phase transition but had not been able to uncover the full picture [83]. We therefore call the phase transition, which fits into neither of the first-order or continuous classifications, a Berezinskii-Kosterlitz-Thouless (BKT) transition. The following rough explanation is inspired by that in [84].

Let us consider the continuum limit where a configuration is now a map $\theta : [0, L]^2 \to S^1$ and the Hamiltonian becomes

$$H = \frac{1}{2} \int \nabla(\theta(\vec{r}))^2 d^2\vec{r}.$$

If we simplify the model to allow θ to take values in $\mathbb{R}$ instead of S^1, then field theory calculations show that $C(T, r)$ decays as a power law for all temperatures T. The missing ingredient must have been lost in passing from S^1 to $\mathbb{R}$ and indeed what we lost were vortex configurations. An example of a vortex configuration was discussed in the introduction to this thesis and is characterised by a non-zero vorticity or winding number

$$v = \frac{1}{2\pi} \oint_{\mathcal{C}} \nabla\theta(\vec{r}) d\vec{r}$$

where $\mathcal{C}$ is any curve which encloses the center of the vortex. If $v < 0$ we may also call it an antivortex. The energy of a configuration containing a single rotationally symmetric vortex with $v = \pm 1$ will have an energy like

$$E_v = \pi \ln \frac{L}{a}$$

where a is some short distance cutoff that we can think of as the size of the vortex core. The dependence on L means that for large systems the energy cost will be too high for vortices to form. However if we instead take a configuration with both a vortex ($v = 1$) and an antivortex ($v = -1$), then the energy of this configuration is

$$E_{\text{pair}} = 2\pi \ln \frac{R}{a}$$

where R is the separation between the vortices. Thus for small R, such bound vortex-antivortex pairs may be formed from thermal fluctuations. When bound these still do not lead to exponential decay of the spin-spin correlation. The idea instead is that as the temperature increases past the critical point, it becomes energetically favourable for the vortex-antivortex pairs to unbind, allowing free vortices to proliferate through the system and destroy the quasi long-range order, leaving only the short-range order observed at high temperatures. A rough free energy argument says that the free energy of a single vortex is

$$F = E - TS = \pi \ln \frac{L}{a} - T \ln \frac{L^2}{a^2}$$

where $S = \ln(L^2/a^2)$ is the entropy, estimated from the idea that there are L^2/a^2 possible positions for a vortex with core area a^2 in an $L \times L$ box. Thus at $T_c = \pi/2$ the energy and entropy are balanced so free vortices become favourable. This value of the critical temperature is actually not quite right and in fact it has been numerically estimated to be approximately $T = 0.8929$ [85]. We also note that the critical exponent of correlation length for the BKT transition is $\nu = \frac{1}{2}$.

3.2 Variant Models

Modifying the Hamiltonian of the XY model gives us a couple of different variants which also display BKT transitions. One of these variants also has a second-order phase transition at a lower temperature in addition to the BKT transition.

3.2.1 Constrained XY Model

What we will refer to as the 2-dimensional Constrained XY model was introduced and investigated in [86, 87] where it is called an XY model with a topological lattice action. It is defined similarly to the classical XY model by assigning an angle $\theta_i \in S^1$ to each lattice site $i \in \Lambda$ of an $L \times L$ square lattice Λ. However the Hamiltonian is defined as

$$H(\boldsymbol{\theta}) = \begin{cases} 0 & \text{if } \frac{1}{2\pi}|\theta_i - \theta_j| \leq \delta \text{ for all } \langle i, j \rangle \\ \infty & \text{otherwise.} \end{cases}$$

Therefore all configurations are constrained so that the spins at neighbouring sites cannot differ by more than δ. Since the partition function does not depend on the thermodynamic temperature, we consider the phase portrait as a function of the parameter δ. The model undergoes a BKT transition as δ increases at approximately $\delta = 0.2825$ [87] with $\nu = \frac{1}{2}$. Notice that while $\delta < 0.25$ no (anti)vortices may form. This is interesting as the energy-entropy trade off argument for vortex proliferation no longer works here: the energy is constant for all allowed configurations. Nevertheless, vortices seem to be responsible for the phase transition here too.

3.2.2 Nematic XY Model

There are a variety of generalised XY models with nematic interactions. We will consider the model with Hamiltonian

$$H(\boldsymbol{\theta}) = -\sum_{\langle ij \rangle} \left[\Delta\, cos(\theta_i - \theta_j) + (1-\Delta)\, cos(2(\theta_i - \theta_j))\right]$$

where we will fix $\Delta = 0.15$. The first term is the usual XY interaction, but the second term is a nematic interaction which remains invariant when any individual spin is rotated 180 degrees. We can imagine this as an interaction between the spins considered as headless rods: spins which are parallel contribute less energy, even if they point in opposite directions. The T-Δ phase diagram of this model is explored in [88–90], and we see that at our chosen $\Delta = 0.15$, it undergoes two phase transitions as temperature increases. The first is an Ising-type transition from a magnetic phase to a nematic phase at $T \approx 0.3314$ (as estimated using the magnetic susceptibility) resulting in (anti)vortices (which remain bound into vortex-antivortex pairs) stretching into domain walls with a half-(anti)vortex at each end; across the wall the spins flip by π. See Figure 8 for an example. The second is a BKT transition to a paramagnetic phase at $T \approx 0.7808$ (as estimated using the magnetic susceptibility) driven by the unbinding of these pairs of now-elongated vortices and antivortices.

3.3 Previous Work

While the existing works on the *persistent homology as an observable* paradigm share the same underlying idea, the approaches seen so far have differed significantly, both in how filtrations have been constructed, and how the resulting persistence diagrams have been analysed. Tran, Chen, and Hasegawa investigated phase transitions in the 2D XY model, 1D transverse-field Ising and 1D Bose-Hubbard models [72]. They computed the Vietoris-Rips persistence of point clouds of lattice sites with inter-point distances given by a linear combination of the Euclidean distance in the

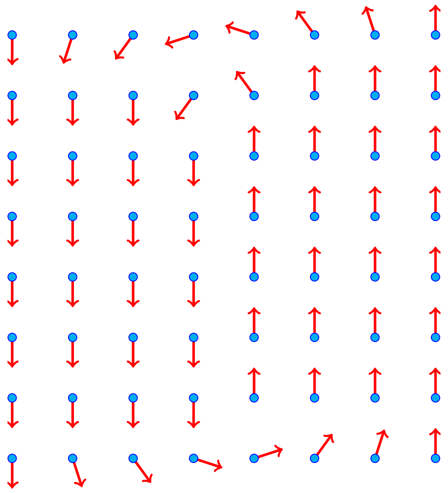

Figure 8: A configuration with a vortex that has stretched out into two half-vortices separated by a domain wall.

lattice and the difference in the spins. They show that clustering configurations based on the the Persistence Fisher kernel [91], persistence entropy and the second moment of persistence of the H_1 diagrams identifies the different phases. They demonstrate that increasing the lattice size produces sharper estimates of the critical temperature. This approach is extended to the XXZ model on a pyrochlore lattice by Olsthoorn, Hellsvik, and Balatsky, approximately separating the six different phases of the model [74].

Cole, Loges, and Shiu apply a different methodology to the previous works. Looking at the 2D Ising, square-ice, XY and fully-frustrated XY models, they introduce general constructions of filtrations for configurations of discrete-valued and circle-valued spin models [45]. In particular, configurations of circle-valued models like the XY model are given a sublevel set filtration of the map $f : \Lambda \to (-\pi, \pi]$ which assigns each site $i \in \Lambda$ in the lattice a parameterisation of its spin $f(i) \in (-\pi, \pi]$. This filtration yields cubical subcomplexes of the lattice. They make use of persistence images [37] to vectorise persistence diagrams, allowing the application of a logistic regression model to separate the phases. They relate some quantitative aspects of the persistence diagrams to the estimation of critical exponents in the case of the Ising model. For discrete models they construct α-complexes on subsets of the lattice sites with the same spin. This is similar to the approach used by Hirakida et al. in [75] who look at the effective Polyakov line model.

Comparing the approach in [72] to that in [45] makes it clear that there is a significant degree of choice in picking the filtration used to compute the persistent homology of a given XY model configuration. We will demonstrate that this choice is an important factor in determining what information about phase transitions one can derive from the persistence. In particular, we find that using two different filtrations is required to detect and analyse the two different phase transitions undergone by the nematic XY model.

3.4 Methods

3.4.1 Simulation

We will make use of the Metropolis algorithm, described in Appendix A, and the Wolff cluster algorithm [92].

Metropolis Algorithm Recall that the Metropolis algorithm produces a sequence of configurations $\{c_t\}_{t\in\mathbb{N}}$ where a configuration c' is proposed as the $t+1$-th configuration c_{t+1} by sampling from a distribution $g(c'|c_t)$ based on the previous configuration c_t (see Appendix A). In the case of the XY model and the variants we study in this chapter, new configurations are proposed by randomly selecting a lattice site $i \in \Lambda$ and angle $\theta' \in S^1$, then proposing the replacement $\theta_i \leftarrow \theta'$. Thus at each step t of the Markov chain we have $g(c'|c_t) = g(c_t|c')$ and the new configuration is accepted or rejected purely based on the resulting change in the value of the Hamiltonian.

Wolff Cluster Algorithm The idea here is that instead of modifying spins one at a time, whole clusters of spins may be flipped across a random line, more quickly exploring the space of configurations while maintaining the conditions of ergodicity and detailed balance required to converge to the correct probability distribution. Once again we set up a Markov chain starting with a random configuration and generating each next step based on the previous. Pseudocode for a single step is given in Algorithm 2.

Algorithm 2 Wolff Cluster Step [92]

Input Configuration $\{\theta_i\}_{i\in\Lambda}$

Output Updated configuration $\{\theta'_i\}_{i\in\Lambda}$

```
for $i \in \Lambda$ do
    $\theta'_i \leftarrow \theta_i$
end for
$S \leftarrow \emptyset$                                   ▷ $S$ is an empty stack of lattice sites
$push(S, \mathrm{UniformRandom}(\Lambda))$                 ▷ Initialise it with a random lattice site
$r \leftarrow \mathrm{UniformRandom}([0, 2\pi))$           ▷ Angle of a random line through the origin
while $S \neq \emptyset$ do
    $i \leftarrow pop(S)$
    for $j \in \Lambda \setminus S$ neighbouring $i$ do
        $P \leftarrow 1 - \exp(\min\{0, -2\beta\cos(\theta_i - r)\cos(\theta_j - r)\})$
        if $\mathrm{UniformRandom}([0, 1]) \leq P$ then
            $push(S, j)$                                   ▷ Add $j$ to the stack with probability $P$
        end if
    end for
    $\theta'_i \leftarrow 2r + \pi/2 - \theta_i$           ▷ Flip $\theta_i$ in the line defined by $r$
end while
```

3.4.2 Filtered Complex

To apply persistent homology we must choose how to define a filtered cubical complex for a given configuration $\boldsymbol{\theta} = \{\theta_i\}$. Our idea is to filter the square tiling of the plane corresponding to the lattice Λ according to the differences in neighbouring spins. For each cube c in this cubical complex, we will specify a filtration index $f(c)$ at which it appears, and then $F_t = f^{-1}(-\infty, t]$ is the subcomplex of the plane consisting of all cubes that have appeared by time t. Denote the smallest angle between spins θ_i and θ_j by $d_{i,j}$. This can also be seen as the length of the shortest arc between θ_i and θ_j on the unit circle.

Definition 3.1 (Angle difference filtration)**.** Taking an $L \times L$ square lattice Λ as a 2-dimensional cubical complex C with 0-cubes $\{i\}$ for $i \in \Lambda$, 1-cubes $\{i, j\}$ for all nearest neighbour pairs of sites i, j in Λ, and 2-cubes $\{i, j, k, l\}$ for all plaquettes with corners i, j, k, l in Λ, the angle difference filtration is the sublevel set filtration defined by $f : C \to \mathbb{R}$ where:

1. We introduce each vertex or 0-cube $i \in \Lambda$ at index
$$f(i) = 0.$$
2. Each edge or 1-cube $\{i, j\}$ is introduced at index
$$f(\{i, j\}) = \frac{1}{2\pi} d_{i,j}.$$
3. Each plaquette or 2-cube $\{i, j, k, l\}$ is introduced at index
$$f(\{i, j, k, l\}) = \max_{a,b \in \{i,j,k,l\}} \frac{1}{2\pi} d_{a,b}.$$

We will also introduce another similar filtration to use with the Nematic XY model in Section 3.2.2. This will instead use a nematic angle difference $d^n_{i,j}$ which denotes the smallest angle between the spins θ_i and θ_j considered as directionless rods. We can think of this as the length of the shortest arc connecting the head of one spin to either the head or tail of the other spin. That is $d^n_{i,j} = \min(d_{i,j}, \pi - d_{i,j})$.

Definition 3.2 (Nematic angle difference filtration)**.** Taking an $L \times L$ square lattice Λ as a 2-dimensional cubical complex C with 0-cubes $\{i\}$ for $i \in \Lambda$, 1-cubes $\{i, j\}$ for all nearest neighbour pairs of sites i, j in Λ, and 2-cubes $\{i, j, k, l\}$ for all plaquettes with corners i, j, k, l in Λ, the nematic angle difference filtration is the sublevel set filtration defined by $f_n : C \to \mathbb{R}$ where:

1. We introduce each vertex or 0-cube $i \in \Lambda$ at index
$$f_n(i) = 0.$$
2. Each edge or 1-cube $\{i, j\}$ is introduced at index
$$f_n(\{i, j\}) = \frac{1}{2\pi} d^n_{ij}.$$
3. Each plaquette or 2-cube $\{i, j, k, l\}$ is introduced at index
$$f_n(\{i, j, k, l\}) = \max_{a,b \in \{i,j,k,l\}} \frac{1}{2\pi} d^n_{ab}.$$

The intuition behind these filtrations originally came from considering the 2D XY model. Regions of the lattice where spins vary slowly will be introduced in the angle difference filtration early, while regions containing rapidly varying spins, such as at the centre of vortices, will enter the filtration later. We should expect then, at least at low temperatures, that each vortex will be manifested as a hole in the filtered lattice which is formed early on in the filtration, and which only gets filled in much later: i.e., a persistent H_1 class. Figure 9 shows an example of this. However we will see that these filtrations can capture other structure such as spin waves, or half-vortices and domain walls when we look at the Nematic XY model. Moreover, compared to the point cloud filtrations used in [72, 74] this class of filtrations has the computational benefit that edges are only introduced between neighbouring lattice sites and only elementary cubes up to dimension 2 are included, greatly speeding up the computation of persistent homology. The persistence diagrams obtained using these filtrations are stable with respect to small perturbations to the spins, in contrast to the sublevel set filtration used in [45].

Proposition 3.1 (Stability of angle-difference filtration)**.** *The maps assigning to each configuration in $(S^1)^\Lambda$ the functions $f, f_n : C \to \mathbb{R}$ are $(1/\pi)$-Lipschitz with respect to the L_∞ metric.*

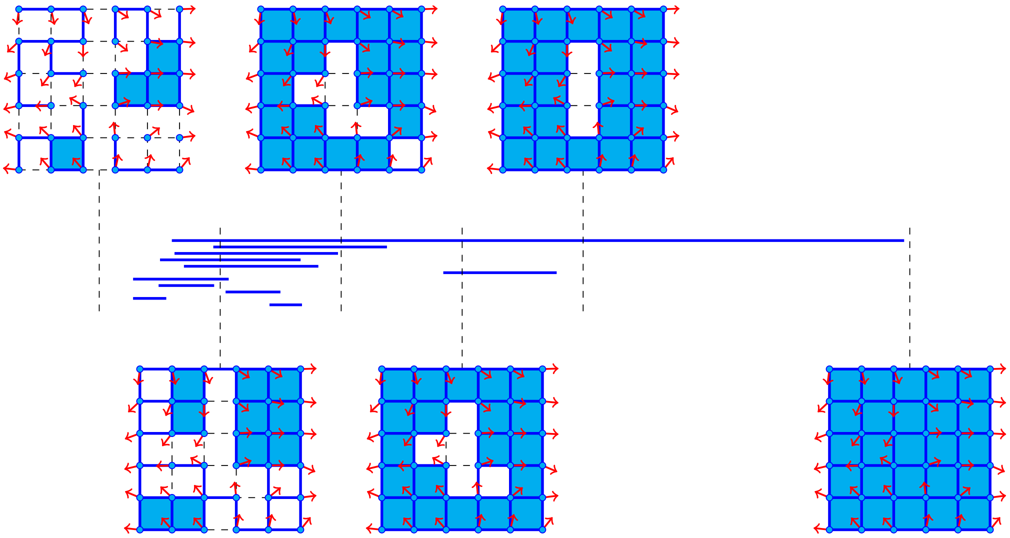

Figure 9: An illustration of the angle difference filtration for a configuration of the XY model with an antivortex. The filtration parameter increases from left to right and the state of the filtration is shown at 6 different stages. On the left-hand side only those neighbouring spins which don't differ too much are connected by edges and plaquettes. As we move towards the right, more and more edges are introduced between more disparate spins. Note the correspondence between the bars and the holes in the filtration. For example, the longest bar corresponds to the hole around the antivortex in the centre of the configuration. This hole is formed early on as the spins far from the centre vary slowly, but survives until the central plaquette is added to the filtration.

Proof. This follows simply from $d_{i,j}$ being a 2-Lipschitz function $((S^1)^2, ||.||_\infty) \to \mathbb{R}$. Suppose we modify each spin of a configuration $\theta'_i = \theta_i + \Delta\theta_i$. Denote the new length of the arc between θ'_i and θ'_j by $d'_{i,j}$. Then for all i, j we have that

$$d_{i,j} - |\Delta\theta_i| - |\Delta\theta_j| \leq d'_{i,j} \leq d_{i,j} + |\Delta\theta_i| + |\Delta\theta_j|,$$

giving us that $||f - f'||_\infty = \frac{1}{2\pi}\max_{i,j}|d_{i,j} - d'_{i,j}| \leq \frac{1}{2\pi}2\max_i|\Delta\theta_i| = \frac{1}{\pi}||\theta - \theta'||_\infty$. Taking the equation $d^n_{i,j} = \min\{d_{i,j}, \pi - d_{i,j}\}$ and working through the cases shows the same for f_n. □

3.5 Detecting Topological Defects

As a first experiment, let us see what the persistent homology of a typical low temperature XY model configuration containing vortices looks like. Figure 10 shows a configuration sampled at a low temperature using the Metropolis algorithm with multiple clearly observed (anti)vortices.

Figure 11 shows the persistence diagram we obtain using the angle difference filtration on the configuration. As expected, we have 4 points on the infinite death line: 1 in H_0, 2 in H_1 and 1 in H_2, recording the homology of the final complex which is homeomorphic to a 2-torus. The 0-dimensional and 2-dimensional persistent homology are not too interesting, but we see some interesting structure in the 1-dimensional persistence. The points split into at least 2 clusters: most lie in a cluster of low-persistence points close to the diagonal, but a few lie in a small cluster of high-persistence points which die close to filtration index 0.5. We immediately hypothesise that this latter cluster corresponds to vortices and antivortices in the configuration. Simply counting the number of vortices and antivortices in the configuration yields a count of 12, but there are 11 points in the cluster in the persistence diagram. Computing some other examples always yields this difference of 1. This turns out to be due to the periodic boundary conditions and a simple topological fact.

Proposition 3.2. *Let T^2 denote the 2-torus, and $x \in T^2$ be a point in the torus. Then,*

$$H_1(T^2 \setminus \{x\}) \cong \mathbb{F}^2 \cong H_1(T^2).$$

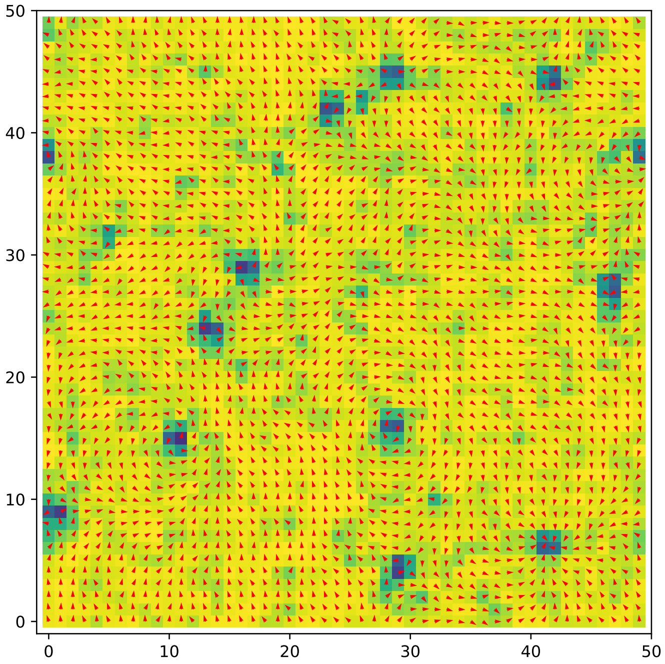

Figure 10: An XY model configuration sampled using the Metropolis algorithm at $T = 0.2$ on a 50×50 lattice. The shading indicates the contribution of each lattice site to the Hamiltonian, helping us identify 12 (anti)vortices (recall that we have periodic boundary conditions).

These are straightforward calculations, but note that we can intuitively see how the first isomorphism holds by retracting the punctured torus $T^2 \setminus \{x\}$ to the bouquet of 2 circles. The point is that if we consider a filtration index where the filtered complex is almost complete except for a single hole centered on a vortex. Then the complex is homeomorphic to a punctured torus and hence its 1-dimensional homology is the same as the torus, so we do not obtain a point in our persistence diagram. We can also think about this difference of 1 as being due to the infinite-death points in H_1. These represent 1-cycles that cross from one side of the configuration to the other, looping around the periodic boundary conditions. These therefore surround all the vortices in the configuration, effectively adding 1 to our count.

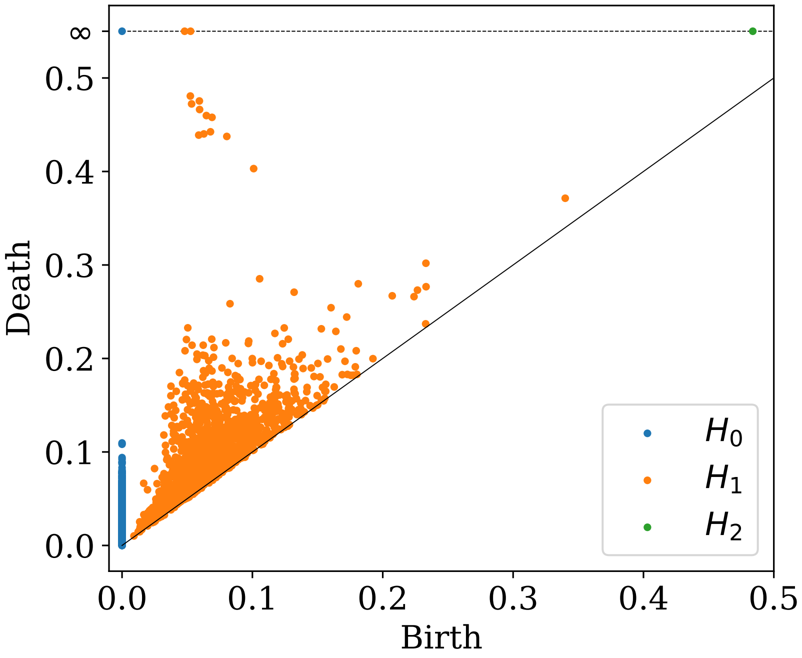

Figure 11: The persistence diagram obtained using the angle difference filtration on the configuration in Figure 10. Note that in the cluster of high persistence, finite-death points in PH_1 there are 11 points.

3.6 Observing the Phase Transitions

In the previous section we looked at the persistence of a single configuration which may or may not produce a persistence diagram representative of its phase. In order to reasonably analyse the persistent homology of different phases we ought to consider the expected persistence. We therefore need to pass to a vectorisation and we choose persistence images since these will allow us to easily relate observations back to features in the persistence diagrams.

Fixing a lattice size of 30×30, the procedure for our analysis of each model at each temperature is as follows:

1. We sample the model using the Wolff cluster algorithm [92] at a range of temperatures spanning the phase transition(s). We perform $50,000$ Wolff cluster flips to properly thermalise the model, and 100 cluster flips between samples to ensure that the autocorrelation is negligible.

2. For each sample, we compute persistence images with 30 × 30 resolution and σ equal to 10% of a pixel.
3. We average the persistence images to estimate the expected value.

XY Model Using the angle difference filtration described in Section 3.4.2 we obtain average persistence images as shown in Figure 12. At low temperatures we see that most points in the persistence diagrams are concentrated in the lower left corner. These come from the presence of spin waves: spins tend to differ more with those in the opposite corner of a plaquette than with their immediate neighbours, producing a short-lived cycle. As the temperature increases we observe that the spin-wave cycles persist longer and longer. At around $T = 0.8$, 0.95, close to the critical point, we begin to see points close to the downwards diagonal $\textit{persistence} = 0.5 - \textit{birth}$, or equivalently $\textit{death} = 0.5$. These represent (anti)vortices: they are born reasonably early, as spins far away from the centre vary slowly, but die much later due to the large difference in spins at the vortex core. In fact, we can check that the sum of the components of the persistence image lying on the diagonal and the two immediate subdiagonals correlates well with the absolute vorticity (the total count of vortices and antivortices) of the configurations as measured by summing the signed angle difference around all plaquettes. For example, computing the Pearson correlation coefficient on 2000 configurations at $T = 1.0$ for $L = 140$ yields a correlation coefficient of $r = 0.70$, $p < 0.001$. At high temperatures we see this concentration of cycles on the diagonal increase and shift rightwards, indicating a disordered phase with many vortices.

Constrained XY Model Using the angle difference filtration described in Section 3.4.2 we obtain average persistence images as shown in Figure 13. We immediately see a resemblance with the persistence images obtained for the XY model in Figure 12 except that we see a cutoff effect at $\textit{birth} = \delta$, since by this point all neighbouring lattice sites must have been connected in the filtration.

Nematic XY Model From Figure 14 we see that the Magnetic-Nematic transition is manifested in the angle difference filtration by the emergence of a cluster

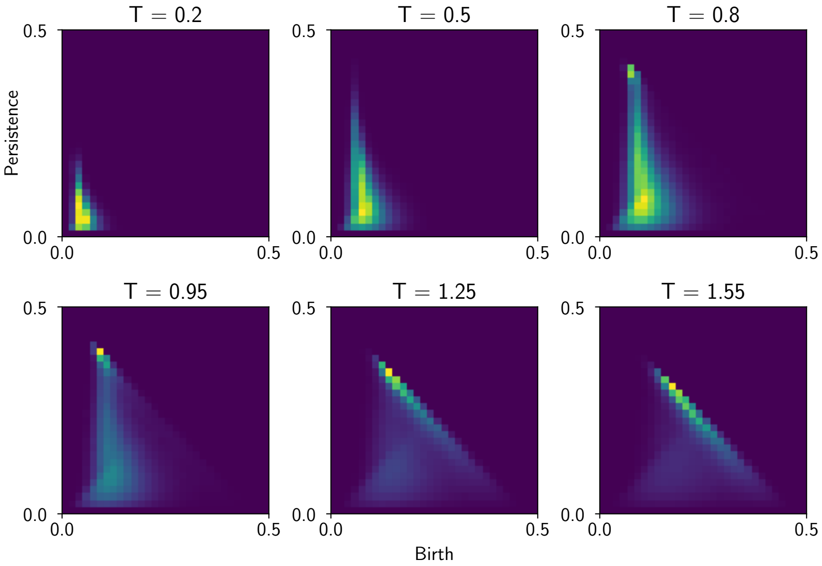

Figure 12: The average H_1 persistence images in birth-persistence coordinates at different temperatures for the XY model with $L = 30$.

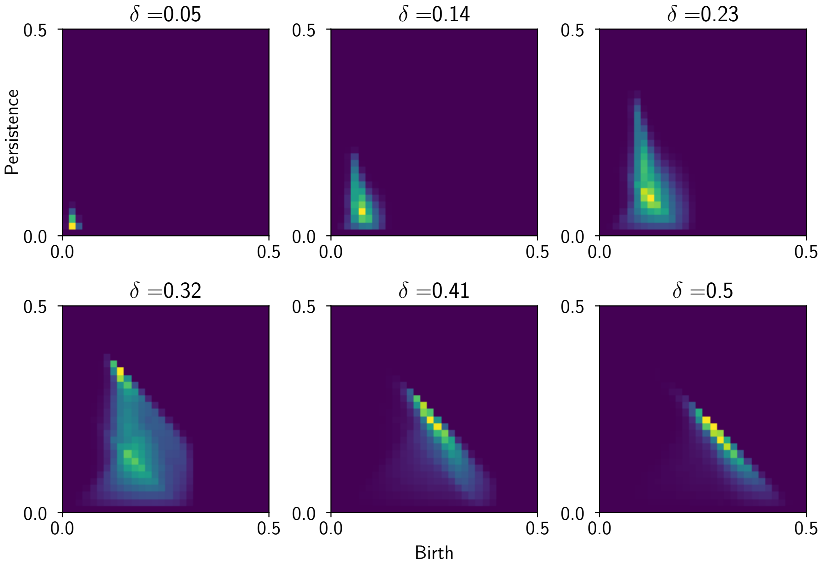

Figure 13: The average H_1 persistence images in birth-persistence coordinates at different temperatures for the Constrained XY model with $L = 30$.

in the bottom right of the persistence image and the rightwards movement of the cluster in the top left. These likely correspond to the appearance of domain walls in configurations. In particular, at a time close to 0.5 in the filtration, the edges which cross domain walls will get added all at once, forming many short-lived cycles. Meanwhile, (anti)vortices get stretched out into strings so that more spins must be connected in the filtration before a hole is formed, generally causing the time at which this happens to increase a little. There is little qualitative difference between the images across the BKT transition however. In Figure 15, showing the average persistence images using the nematic angle difference filtration, we see a familiar picture of the BKT transition which is very similar to that observed in the XY model and Constrained XY model, while the Ising-type transition is not detectable at all. We also looked at a combined angle difference filtration using $\frac{\Delta}{2\pi}d_{i,j} + \frac{1-\Delta}{2\pi}d^n_{i,j}$. This potentially did detect both phase transitions, but the later step of training classification models to identify phases from the persistence images was difficult, so it was not possible to verify this.

3.7 Matching Vortices and Antivortices with Representative Cocycles

Besides simply matching the vortex count and the number of high persistence points in PH_1, we might also attempt to identify each point in PH_1 with a cycle in the cubical complex. Unfortunately, as mentioned in Section 2.1.4, available software tools for this kind of analysis are sparse. Instead we will take a slightly different approach which will turn out to lead to provide some interesting results.

Rather than obtaining representative cycles, we instead obtain representative cocycles. In $\mathbb{Z}_2$ coefficients a 1-cocycle is a function c from the edges of the complex to $\mathbb{Z}_2$, so we can also think of a 1-cocycle as a subset of the edges given by its support $\mathrm{Supp}(c)$. These edges represent a cut of the corresponding cycle, i.e., removing these edges would break the cycle.

We will make use of the Ripser.py software [33] to calculate persistent cohomology and representative cocycles. However, this software works with simplicial complexes and not cubical complexes so we will define a filtered simplicial complex

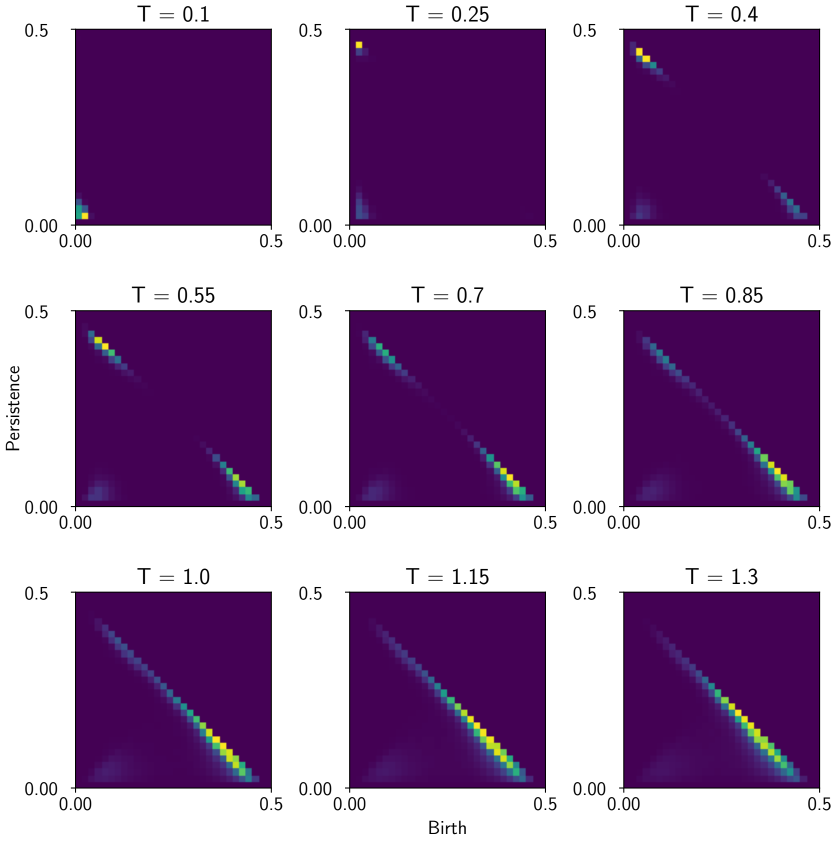

Figure 14: The average H_1 persistence images in birth-persistence coordinates at different temperatures for the Nematic XY model with $L = 30$ using the angle difference filtration.

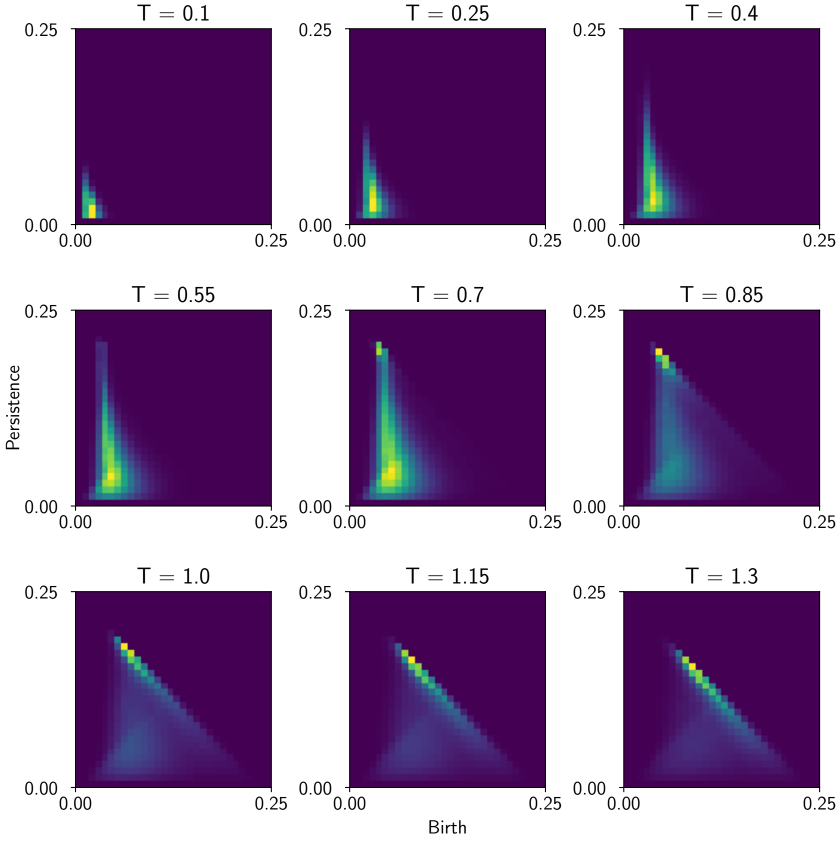

Figure 15: The average H_1 persistence images in birth-persistence coordinates at different temperatures for the Nematic XY model with $L = 30$ using the nematic angle difference filtration.

which yields the same persistent homology as the angle difference filtration.

Definition 3.3 (Simplicial angle difference filtration)**.** Triangulate the lattice so that each plaquette is divided into two triangles by a single diagonal edge (which diagonal is chosen does not matter) to obtain a simplicial complex Σ homeomorphic to the 2-torus. The simplicial angle difference filtration is the sublevel set filtration of $f : \Sigma \to \mathbb{R}$ where:

1. We introduce each vertex (or 0-simplex) $\{i\}$ at index
$$f(\{i\}) = 0.$$
2. If an edge (or 1-simplex) $\{i, j\}$ connects nearest neighbour lattice sites (i.e., is horizontal or vertical), it is introduced at index
$$f(\{i, j\}) = \frac{1}{2\pi} d_{i,j}.$$
Otherwise if it is a diagonal on a plaquette $\square$, it is introduced at index
$$f(\{i, j\}) = \max_{x,y \in \square} \frac{1}{2\pi} d_{x,y}.$$
3. Each triangle (or 2-simplex) $\{i, j, k\}$ lies in some plaquette $\square$ and is introduced at index
$$f(\{i, j, k\}) = \max_{a,b \in \square} \frac{1}{2\pi} d_{a,b}.$$

It is straightforward to see that $\forall t \in \mathbb{R}$, the simplicial angle difference filtered complex at index t is homeomorphic to the (cubical) angle difference filtration at index t. We therefore obtain the same persistence module and thus persistence diagram.

Figure 16 shows a low temperature configuration along with a representative cocycle for one of the high persistence H_1 points. We observe that it forms a path between a vortex and an antivortex.

We might hypothesise that this is always the case for the points in the high persistence cluster, so that the representative cocycles give us a way to match vortices and antivortices. However, in Figure 17 we see that another representative

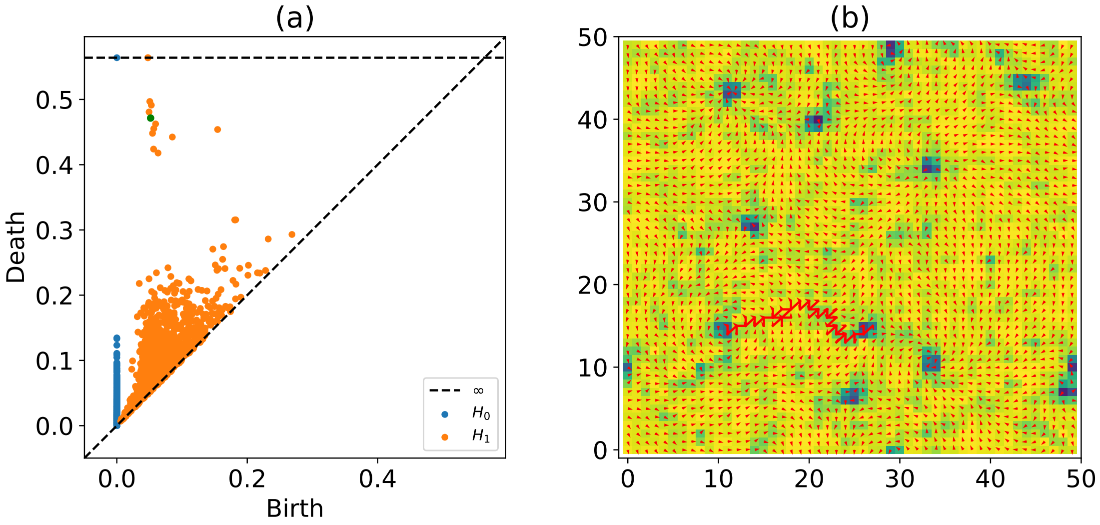

Figure 16: A representative cocycle in an XY model configuration sampled at $T = 0.2$. (a) The persistence diagram obtained using the simplicial angle difference filtration. The point we are looking at is highlighted in green. (b) The configuration coloured by Hamiltonian density with a representative cocycle indicated in red.

cocycle links several vortices. The issue here is that the software is only required to give some representative cocycle c for $[c] \in \ker(\delta^i)/\operatorname{im}(\delta^{i-1})$, among all possible representatives $c + \delta^{i-1}b$ for any $b \in C^*_{i-1}$.

Nevertheless, if the representative cocycles of high persistence points in H_1 yield paths between vortices and antivortices, then the lengths of these paths may signal the phase transition. Given a configuration and its simplicial angle difference persistence, let R_n denote the set of representative cocycles obtained using Ripser.py for the n most persistent finite-death points in H_1. Given a cocycle c, recall that $\operatorname{Supp}(c)$ denotes its support: the set of simplices involved in any chains sent to 1 by c. We define the observable

$$\ell_n = \frac{1}{n} \sum_{c \in R_n} \frac{|\operatorname{Supp}(c)|}{2}$$

which records the average length of the paths for the n most persistent points in H_1. The factor of $1/2$ comes from the fact that we have two edges in the cocycle for each step along the lattice. Note that although Ripser.py makes no guarantees on which cocycle it gives, it is deterministic – given the same data, we get the same

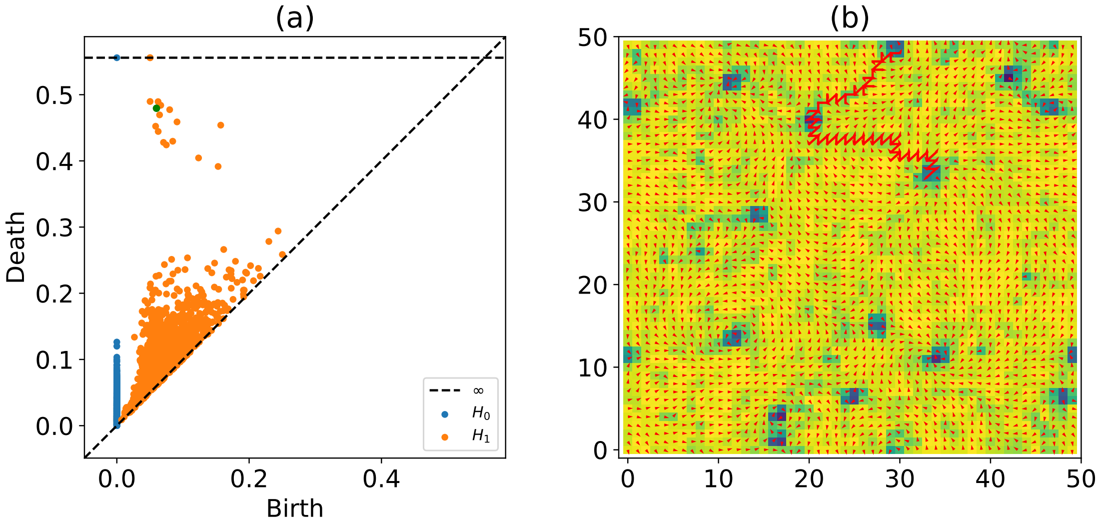

Figure 17: A long representative cocycle in an XY model configuration sampled at $T = 0.2$. (a) The persistence diagram obtained using the simplicial angle difference filtration. The point we are looking at is highlighted in green. (b) The configuration coloured by Hamiltonian density with a representative cocycle indicated in red.

cocycles out – so ℓ_n is well-defined so long as n is not greater than the number of points in the diagram. Figure 18 shows the estimated expectation of ℓ_{15} as a function of temperature for lattice sizes $L \in \{50, 75, 100\}$. For each lattice size and temperature, 500 configurations were generated using the Metropolis algorithm with $200L^2$ steps before the first sample and $10L^2$ steps between each subsequent sample.

We immediately observe that the function peaks close to the critical temperature of the XY model. This agrees with our intuition about the phase transition. As the temperature increases from below, vortices and antivortices begin to separate and spread out, increasing the average cocycle path length. Going past the transition point however, more and more vortices appear in configurations, increasing the vortex density and decreasing the distance between vortices again. The fact that the peaks in the curves appear to grow taller and narrower as the lattice size increases, as well as closer to the transition point, indicates that this observable is sensitive to the phase transition. It may therefore be possible to do a quantitative analysis of the phase transition using this observable, obtaining estimates of the critical temperature and critical exponent of correlation length ν. However, the arbitrary nature

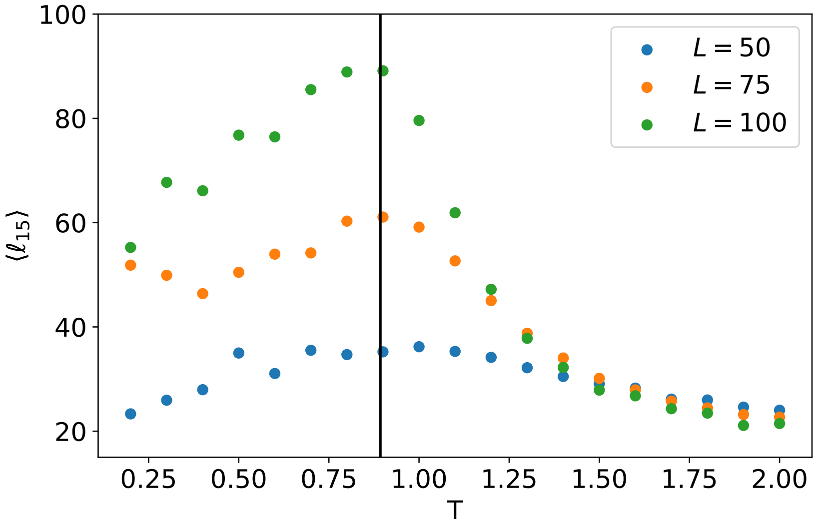

Figure 18: The average lengths of the paths corresponding to the 15 most persistent cocycles as a function of temperature for three different lattice sizes.

of the chosen cocycle representative and the resulting variance in this measurement do not make it ideal for this purpose. Instead we make use of observables based on simple machine learning.

3.8 Quantitative Analysis of Phase Transitions

In this section we apply the the machine learning classifiers and statistical techniques introduced in Chapter 2 to estimate the critical temperatures and critical exponents of correlation length for the different phase transitions.

3.8.1 Numerical Procedures and Error Estimation

The procedure for our analysis of each model at each lattice size is as follows:

1. We sample the model on the given lattice size using the Wolff cluster algorithm [92] at a range of temperatures spanning the phase transition(s). We perform $50,000$ Wolff cluster flips to properly thermalise the model, and 100 cluster flips between samples to ensure that the autocorrelation is negligible.

2. For each sample, we compute persistence images with 30×30 resolution and σ equal to 10% of a pixel.
3. We use persistence images from the low and high temperature phases to train the logistic regression and k-NN models.
4. Using the trained classification models, we assign a predicted phase to each sample from the critical region.
5. Close to the peaks in the variances χ_{LR} and χ_{kNN} of the classifier we apply multiple histogram reweighting to obtain an interpolated curve and a more precise estimate of the location of the peak.

Once we have the interpolated variance curve and peak temperature for each of the lattice sizes, we estimate T_c and ν by fitting the peak temperatures to the appropriate finite-size scaling ansatz (Equations 5 and 6) and optimising the data collapse of the variance curves. For each lattice size we perform two bootstraps: first by resampling the training samples, and second by resampling the samples in the critical region. In each case we resample 500 times, obtaining bootstrap distributions for the estimates of T_c and ν. We estimate the error in these quantities by taking the square root of the sum of the variances of the two bootstrap distributions.

For the constrained XY model we must adjust our methodology slightly since histogram reweighting is not possible. Instead we will sample deltas more densely, then to extract the maximums of χ_{LR} and χ_{kNN} we will fit a parabola to the three highest points.

3.8.2 XY Model

Logistic Regression We trained logistic regression models on samples drawn from the low and high temperature phases given in Table 1 with $10,000$ samples from each temperature. The regularisation hyper-parameter was set to $C = 0.001$. We evaluated the models with $10,000$ samples from each temperature in the critical region.

A plot of the resulting phase indicators is shown in Figure 19 and their variance curves are shown in Figure 20. The plot of the pseudo-critical temperatures against

Region	T
Low T	0.85, 0.85, 0.85
High T	0.91, 0.92, 0.93
Critical	0.88, 0.89, 0.90

Table 1: Values of T sampled at for analysis of the XY model phase transition using logistic regression.

$\log(L)^{-2}$ is shown in Figure 21. We do not observe any significant lattice-size dependence in the pseudo-critical temperatures. They instead seem to be distributed close to $T = 0.89$ which is the midpoint of the training temperatures. A straight line fit yields an extrapolated critical temperature of

$$T_c = 0.8872 \pm 0.0009,$$

well below the expected $T_c = 0.8929$. The curve collapse (Figure 22) procedure gives

$$\begin{aligned} T_c &= 0.8824 \pm 0.0001 \\ \nu &= 0.4968 \pm 0.0055 \\ b &= 0.5098 \pm 0.0068, \end{aligned}$$

not accounting within one standard deviation for the expected values of $T_c = 0.8929$ and $\nu = \frac{1}{2}$.

An advantage of using a generalised linear model like logistic regression, as explored in [45], is that we can easily match the learned weights against the pixels of the persistence images. This allows us to interpret how the classifier distinguishes phases. The weights of the logistic regression model trained on the $L = 140$ XY model data is shown in Figure 23. We see that the low temperature phase is characterised by cycles which are born early and which tend to have low persistence, representing spin waves. The high temperature phase is indicated by cycles with a later birth time and persistence. In particular, the most important region in identifying the high temperature phase is close to $\textit{birth} = 0.1$, $\textit{persistence} = 0.4$ which detects (anti)vortex cycles beginning to change behaviour and move down the diagonal $\textit{persistence} = 0.5 - \textit{birth}$.

k-Nearest Neighbours In the case of the XY model, we found that the k-nearest neighbours classification worked best when trained on a broad range of temperatures. We trained the models on samples drawn from the low and high temperature phases given in Table 2 with 2000 samples from each temperature. The neighbours hyperparameter was set to $k = 30$. We evaluated the models with $10,000$ samples from each of the temperatures in the critical region.

Region	T
Low T	0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50, 0.55, 0.60, 0.65, 0.70, 0.75, 0.80, 0.85
High T	0.95, 1.00, 1.05, 1.10, 1.15, 1.20, 1.25, 1.30, 1.35, 1.40, 1.45, 1.50, 1.55, 1.60
Critical	0.900, 0.905, 0.910, 0.915, 0.920, 0.925, 0.930, 0.935, 0.940, 0.945, 0.950

Table 2: Values of T sampled at for analysis of the XY model phase transition using k-nearest neighbours classification.

A plot of the resulting phase indicators is shown in Figure 24 and their variance curves are shown in Figure 25. The plot of the pseudo-critical temperatures against $\log(L)^{-2}$ is shown in Figure 26. Here we see an asymptotic convergence towards a linear dependence between the pseudo-critical temperatures $T_c(L)$ and $\log(L)^{-2}$. Fitting a straight line to the largest three lattice sizes yields

$$T_c = 0.8935 \pm 0.0043,$$

much closer to the expected $T_c \approx 0.8929$ than the result of the logistic regression approach. The curve collapse (Figure 27) procedure gives

$$\begin{aligned} T_c &= 0.8918 \pm 0.0033 \\ \nu &= 0.4972 \pm 0.0264 \\ b &= 0.5073 \pm 0.0137, \end{aligned}$$

very close to the expected values.

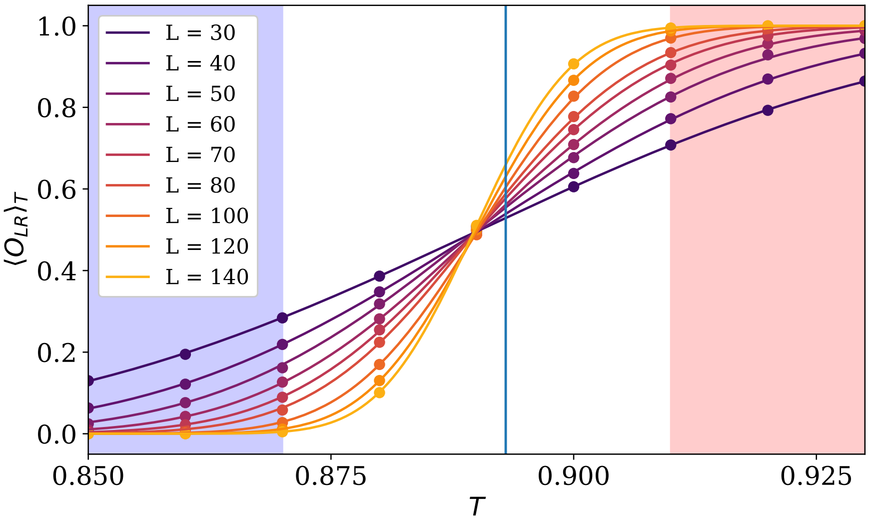

Figure 19: Plot showing $\langle O_{LR}\rangle$ as a function of temperature for each lattice size for the XY model. The shaded regions indicate the temperatures used for the low and high temperature training data. The vertical line shows the location of the expected critical temperature $T_c = 0.8929$.

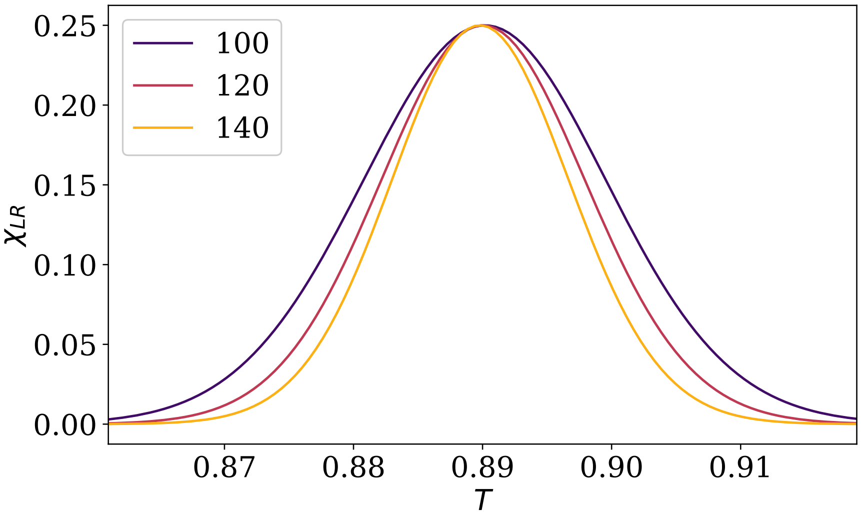

Figure 20: Plot showing χ_{LR} as a function of temperature for the largest three lattice sizes of the XY model. These are what we use to perform the curve collapse procedure.

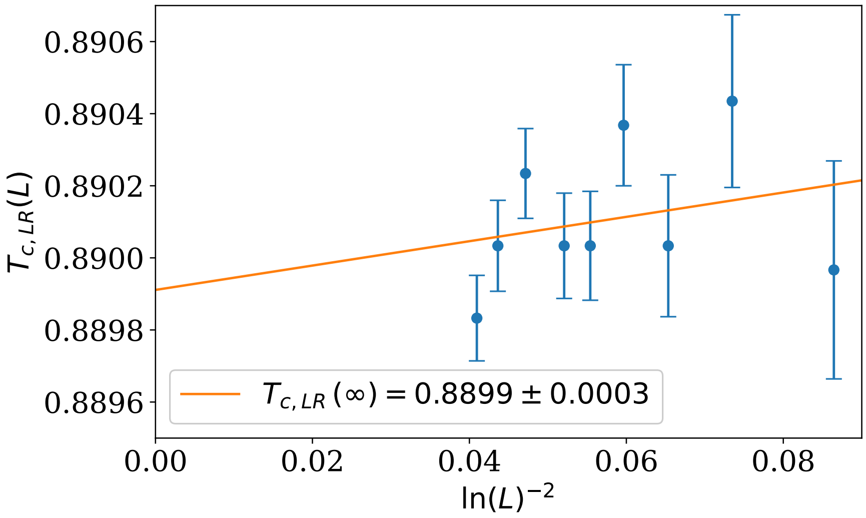

Figure 21: Estimating the critical temperature for the XY model using logistic regression. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{LR}, are fitted to the ansatz in Equation 6. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

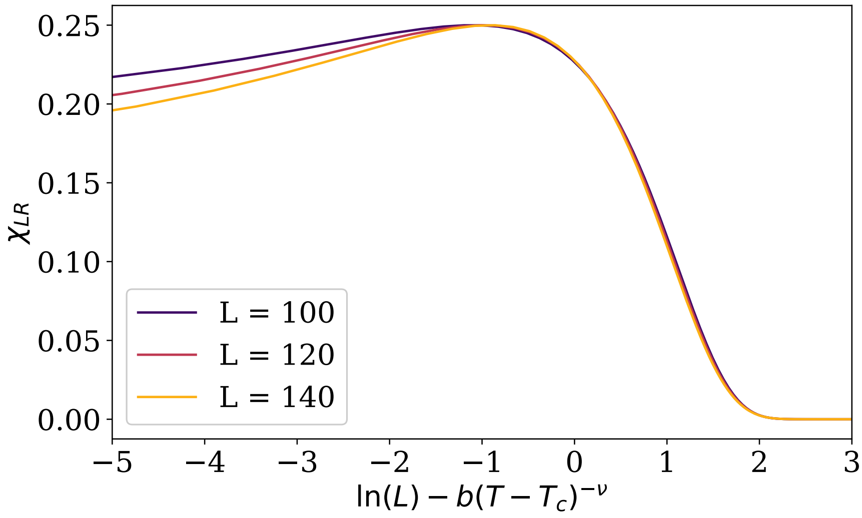

Figure 22: The curve collapse of χ_{LR} for the XY model with $T_c = 0.8824$, $\nu = 0.4968$ and $b = 0.5098$.

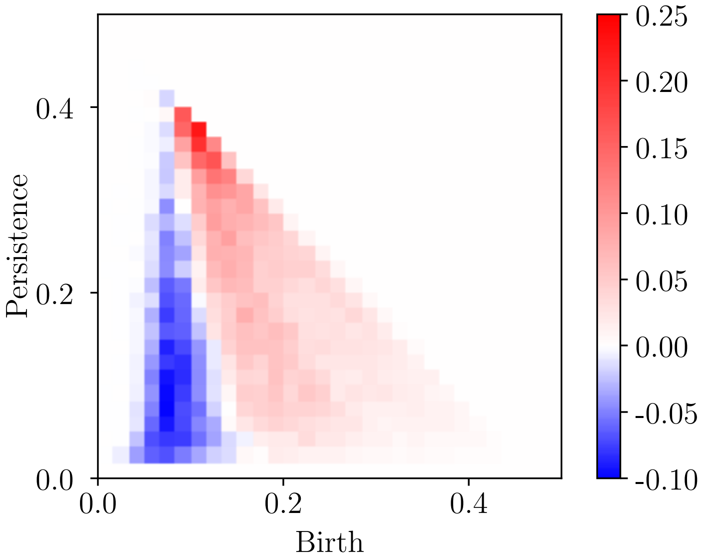

Figure 23: The weights of the logistic regression model trained on the XY model configurations with $L = 140$.

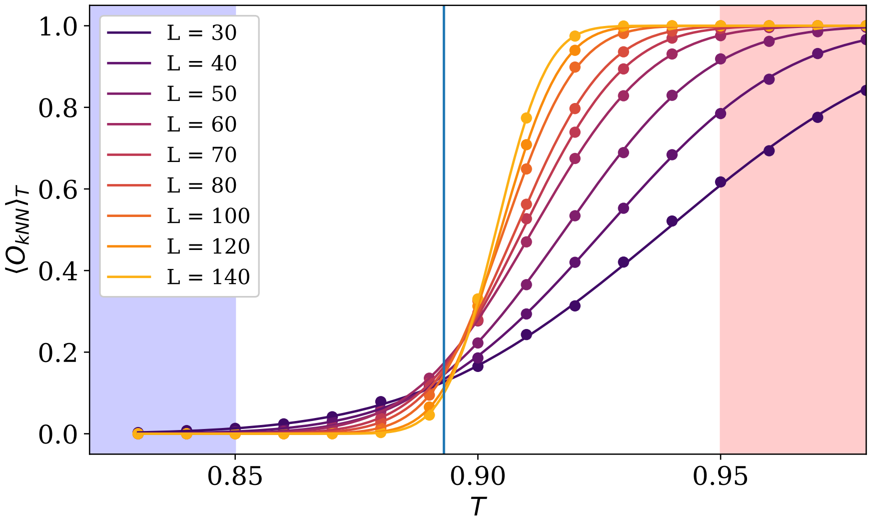

Figure 24: Plot showing $\langle O_{kNN} \rangle$ as a function of temperature for each lattice size for the XY model. The shaded regions indicate the temperatures used for the low and high temperature training data. The vertical line shows the location of the expected critical temperature $T_c = 0.8929$.

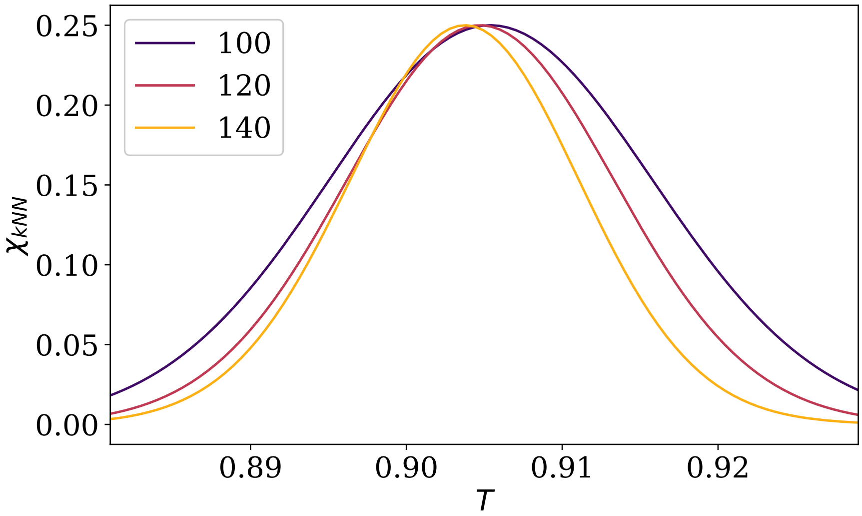

Figure 25: Plot showing χ_{kNN} as a function of temperature for the largest three lattice sizes of the XY model. These are what we use to perform the curve collapse procedure.

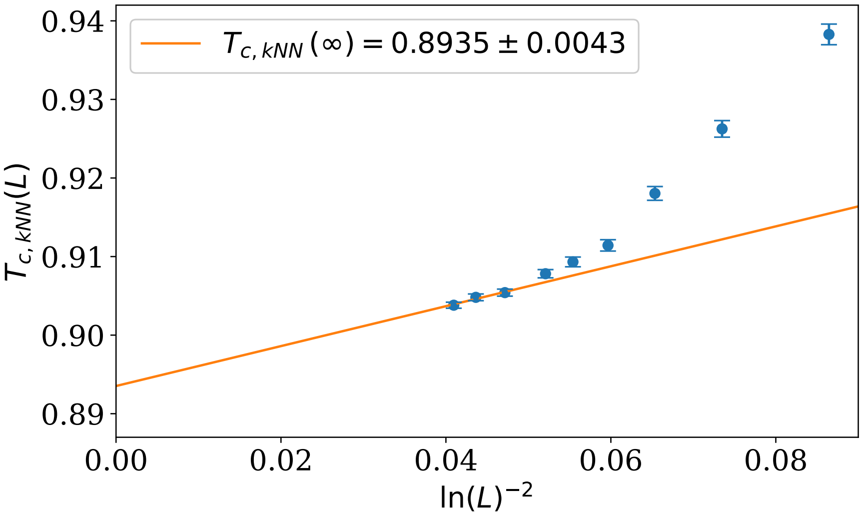

Figure 26: Estimating the critical temperature for the XY model using k-nearest neighbours. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{kNN}, are fitted to the ansatz in Equation 6. We use the largest three lattice sizes in the linear fit. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

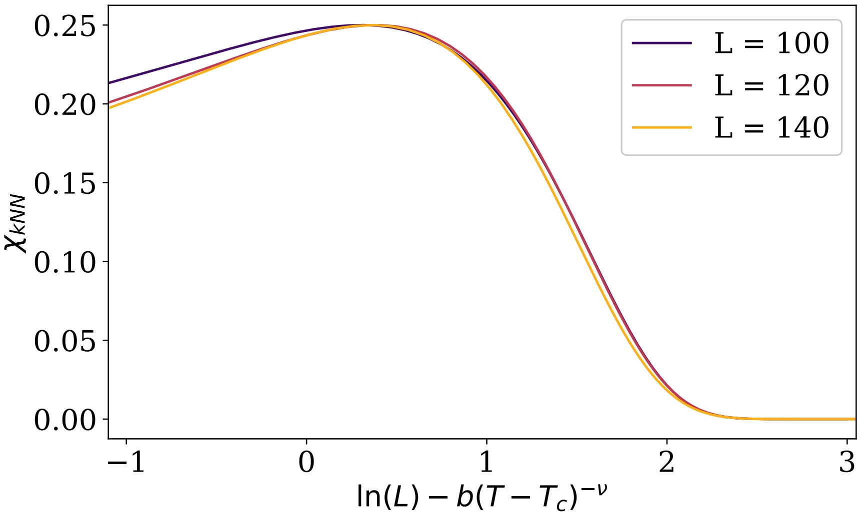

Figure 27: The curve collapse of χ_{kNN} for the XY model with $T_c = 0.8918$, $\nu = 0.4972$ and $b = 0.5073$.

3.8.3 Constrained XY Model

Logistic Regression Analysis We trained logistic regression models on samples drawn from the low delta phase and high delta phase shown in Table 3 with 4000 samples from each. The regularisation hyper-parameter was set to $C = 0.001$. We evaluated the models with 4000 samples from each of δ indicated in the table. A plot of the resulting phase indicators is shown in Figure 28. The plot of the pseudo-critical deltas against $\log(L)^{-2}$ is shown in Figure 29. We do not observe any significant lattice-size dependence in the pseudo-critical deltas. They instead seem to be distributed close to $\delta = 0.283$ which is the midpoint of the training deltas. The curve collapse (Figure 30) procedure gives

Region	δ
Low δ	0.270, 0.272, 0.274, 0.276, 0.278, 0.280
High δ	0.286, 0.288, 0.290, 0.292, 0.294, 0.296
Evaluated	0.270, 0.271, 0.272, 0.273, 0.274, 0.275, 0.276, 0.277, 0.278, 0.279. 0.280, 0.281, 0.282, 0.283, 0.284, 0.285, 0.286, 0.287, 0.288, 0.289, 0.290, 0.291, 0.292, 0.293, 0.294, 0.295, 0.296

Table 3: Values of δ sampled at for analysis of the constrained XY model phase transition using logistic regression.

$$\begin{aligned}\delta_c &= 0.2843 \pm 0.0013\\ \nu &= 0.4999 \pm 0.0189\\ b &= 0.3009 \pm 0.0041,\end{aligned}$$

which accounts for the previously-obtained value of $\delta_c = 0.2825$ within two standard deviations, and supports $\nu = \frac{1}{2}$ within one standard deviation.

The weights of the logistic regression model trained for $L = 140$ are shown in Figure 31. We observe a similarity to the weights learnt for the XY model in Figure 23 although in this case it appears to be more difficult to delineate which regions of the persistence images indicate the two phases.

k-Nearest Neighbours Analysis We trained the k-nearest neighbours models on samples drawn from the low δ phase and high δ phase shown in Table 4 with 4000 samples from each. The neighbours hyper-parameter was set to $k = 30$. We evaluated the models with 4000 samples from each of δ indicated in the table.

Region	δ
Low δ	0.270, 0.272, 0.274, 0.276, 0.278, 0.280
High δ	0.286, 0.288, 0.290, 0.292, 0.294, 0.296
Evaluated	0.270, 0.271, 0.272, 0.273, 0.274, 0.275, 0.276, 0.277, 0.278, 0.279. 0.280, 0.281, 0.282, 0.283, 0.284, 0.285, 0.286, 0.287, 0.288, 0.289, 0.290, 0.291, 0.292, 0.293, 0.294, 0.295, 0.296

Table 4: Values of δ sampled at for analysis of the constrained XY model phase transition using k-nearest neighbours classification.

A plot of the resulting phase indicators is shown in Figure 32. The plot of the pseudo-critical deltas against $\log(L)^{-2}$ is shown in Figure 33. Here we see an asymptotic convergence towards a linear dependence between the pseudo-critical deltas $\delta_c(L)$ and $\log(L)^{-2}$. Fitting a straight line to the largest three lattice sizes yields

$$\delta_c = 0.2821 \pm 0.0014.$$

The curve collapse (Figure 34) procedure gives

$$\begin{aligned} \delta_c &= 0.2818 \pm 0.0017 \\ \nu &= 0.5003 \pm 0.0206 \\ b &= 0.5022 \pm 0.0048, \end{aligned}$$

very close to the expected values.

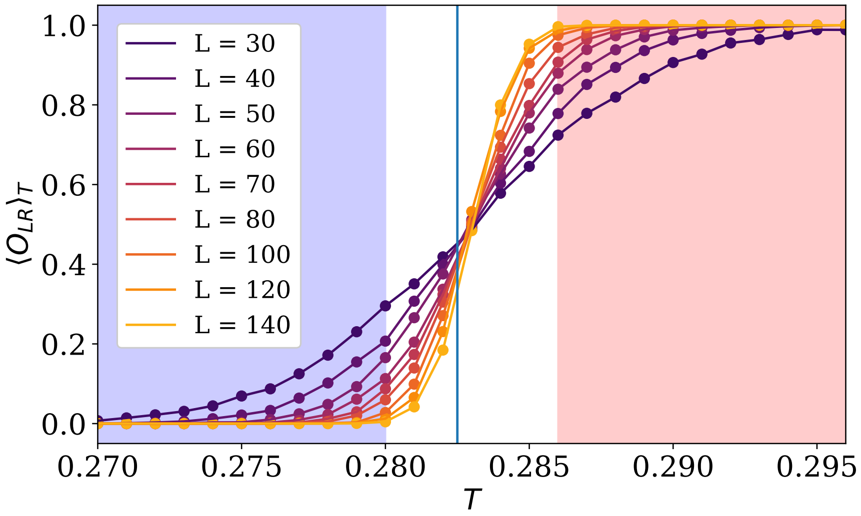

Figure 28: Plot showing $\langle O_{LR} \rangle$ as a function of delta for each lattice size for the Constrained XY model. The shaded regions indicate the deltas used for the low and high delta training data. The vertical line shows the location of the expected critical delta $\delta_c = 0.2825$.

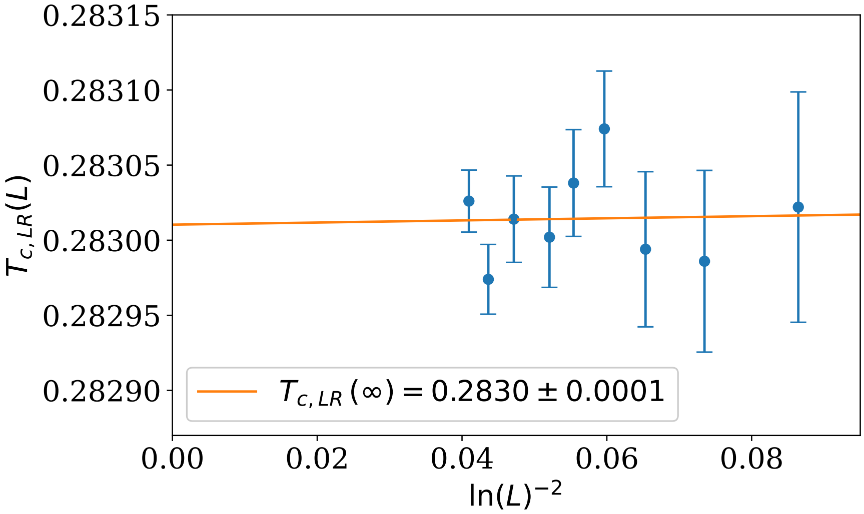

Figure 29: Estimating the critical delta for the Constrained XY model using logistic regression. The pseudo-critical deltas for the different lattice sizes, calculated from finding the peak of χ_{LR}, are fitted to the ansatz in Equation 6. The intercept of the linear fit gives the estimate for $\delta_c(\infty)$. Error bars are estimated by bootstrapping.

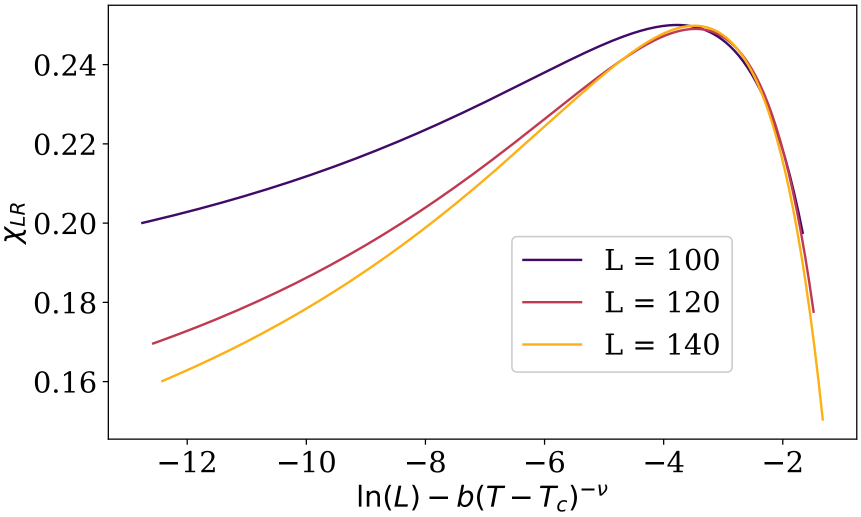

Figure 30: The curve collapse of χ_{LR} for the Constrained XY model with $\delta_c = 0.2843$, $\nu = 0.4999$ and $b = 0.3009$.

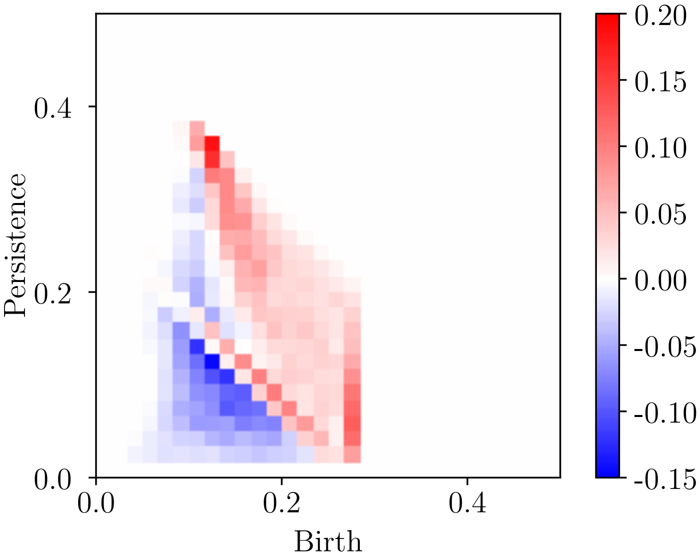

Figure 31: The weights of the logistic regression model trained on configurations over the BKT transition in the Constrained XY model with $L = 140$.

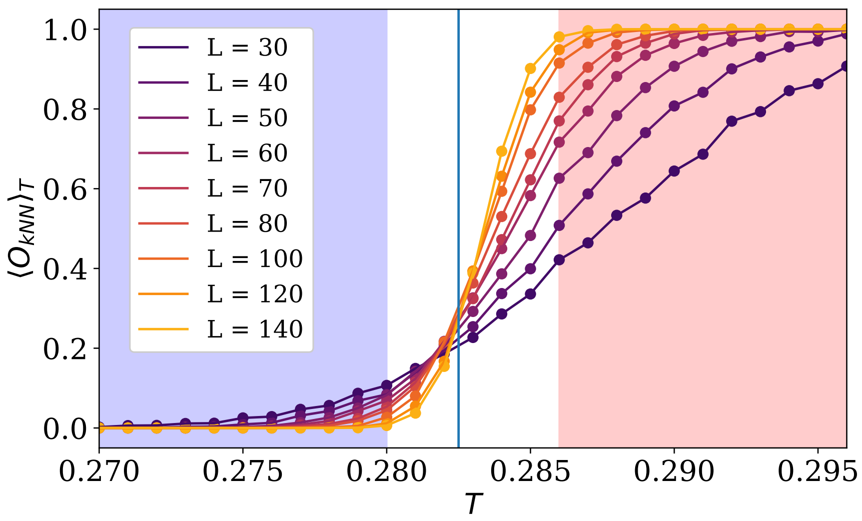

Figure 32: Plot showing $\langle O_{kNN} \rangle$ as a function of delta for each lattice size for the Constrained XY model. The shaded regions indicate the deltas used for the low and high delta training data. The vertical line shows the location of the expected critical delta $\delta_c = 0.2825$.

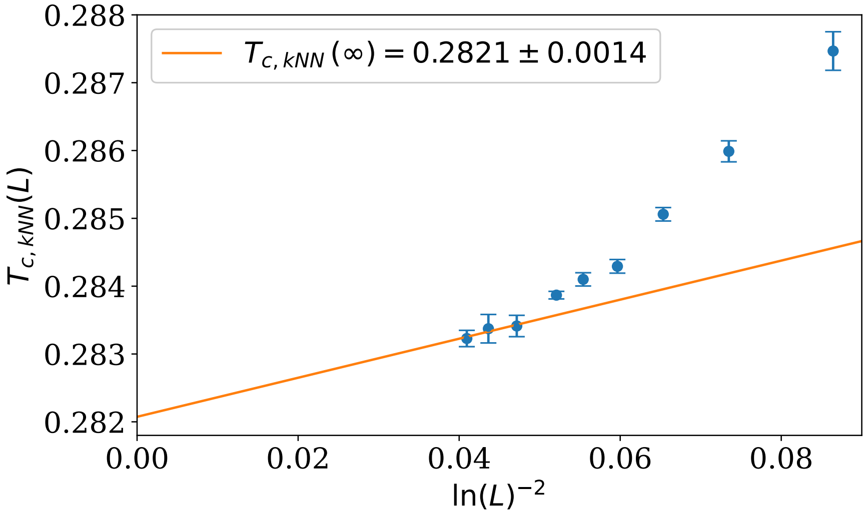

Figure 33: Estimating the critical delta for the Constrained XY model using logistic regression. The pseudo-critical deltas for the different lattice sizes, calculated from finding the peak of χ_{kNN}, are fitted to the ansatz in Equation 6. We fit to the largest three lattice sizes and the intercept of this gives the estimate for $\delta_c(\infty)$. Error bars are estimated by bootstrapping.

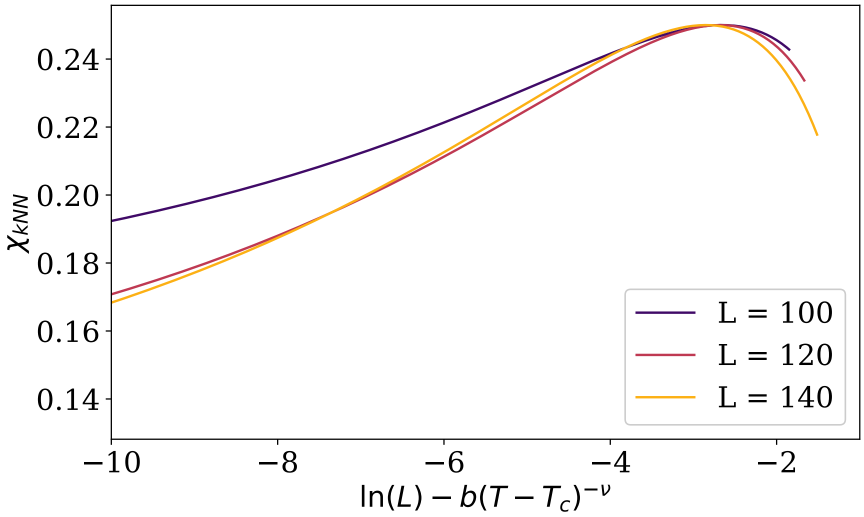

Figure 34: The curve collapse of χ_{kNN} for the Constrained XY model with $\delta_c = 0.2818$, $\nu = 0.5003$ and $b = 0.3022$.

3.8.4 Nematic XY Model

Logistic Regression Analysis of Magnetic-Nematic Transition We trained logistic regression models on samples drawn from the low and high temperature phases given in Table 5 with $10,000$ samples from each temperature. The regularisation hyper-parameter was set to $C = 10^{-6}$. We evaluated the models with $10,000$ samples from each temperature in the critical region.

Region	T
Low T	0.3200, 0.3225
High T	0.3425, 0.3450
Critical	0.33000, 0.33125, 0.33250, 0.33375, 0.33500

Table 5: Values of T sampled at for analysis of the magnetic-nematic transition in the nematic XY model using logistic regression.

A plot of the resulting phase indicators is shown in Figure 35. The plot of the pseudo-critical temperatures against L^{-1} is shown in Figure 36. For the lower lattice sizes $L < 60$, we do not observe any significant lattice-size dependence in the pseudo-critical temperatures. They instead seem to be distributed close to $T = 0.3325$ which is the midpoint of the training temperatures. At the larger lattice sizes $L \geq 60$, a linear dependence on L^{-1} emerges. Fitting a line to the largest four lattice sizes yields an extrapolated critical temperature of

$$T_c = 0.3314 \pm 0.0001.$$

The curve collapse (Figure 37) procedure gives

$$\begin{aligned} T_c &= 0.3315 \pm 0.0001 \\ \nu &= 0.8562 \pm 0.0102. \end{aligned}$$

While these estimates of the critical temperature are good, the expected value of $\nu = 1$ doesn't fall within the error bars estimated with this approach.

The weights of the logistic regression model trained for $L = 140$ are shown in Figure 38. We observe that the classifier learns to detect exactly what we saw in

Figure 14, namely a rightwards shift of the upper left cluster, and the emergence of a cluster in the bottom right, corresponding to domain walls forming in the configurations.

k-Nearest Neighbours Analysis of Magnetic-Nematic Transition We trained the k-nearest neighbours models on samples drawn from the low and high temperature phases given in Table 6 with 4000 samples from each temperature. The neighbours hyper-parameter was set to $k = 30$. We evaluated the models with 4000 samples from each temperature in the critical region.

Region	T
Low T	0.3200, 0.3225
High T	0.3425, 0.3450
Critical	0.33000, 0.33125, 0.33250, 0.33375, 0.33500

Table 6: Values of T sampled at for analysis of the magnetic-nematic transition in the nematic XY model using k-nearest neighbours classification.

A plot of the resulting phase indicators is shown in Figure 39. The plot of the pseudo-critical temperatures against L^{-1} is shown in Figure 40. Here we see that for $L \geq 60$, the pseudo-critical temperatures fit reasonably well on a straight line when plotted against L^{-1}. Fitting a line to the largest four lattice sizes yields

$$T_c = 0.3315 \pm 0.0002.$$

The curve collapse (Figure 41) procedure gives

$$T_c = 0.3316 \pm 0.0002$$
$$\nu = 0.9551 \pm 0.0196,$$

very close to the expected value of $T_c = 0.3314$, but not quite compatible with $\nu = 1$ although better than the logistic regression result.

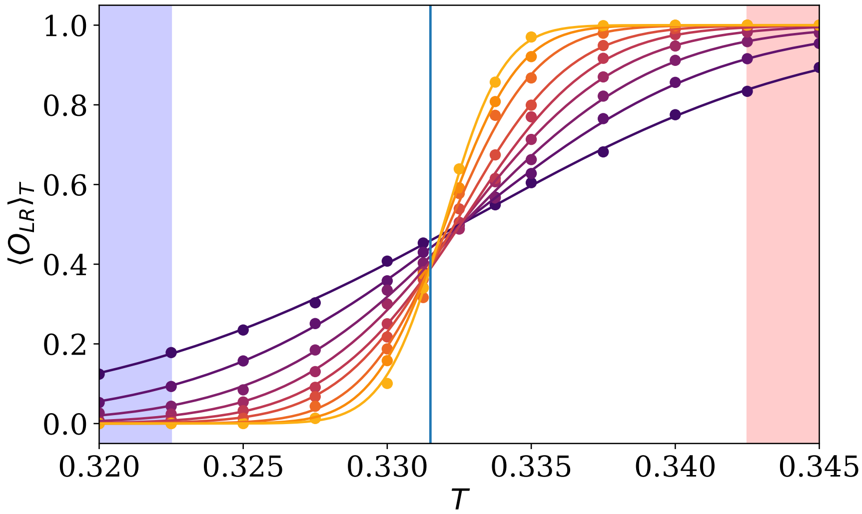

Figure 35: Plot showing $\langle O_{LR} \rangle$ as a function of temperature for each lattice size over the Magnetic-Nematic transition in the Nematic XY model. The shaded regions indicate the temperatures used for the low and high temperature training data. The vertical line shows the location of the expected critical temperature $T_c = 0.3314$.

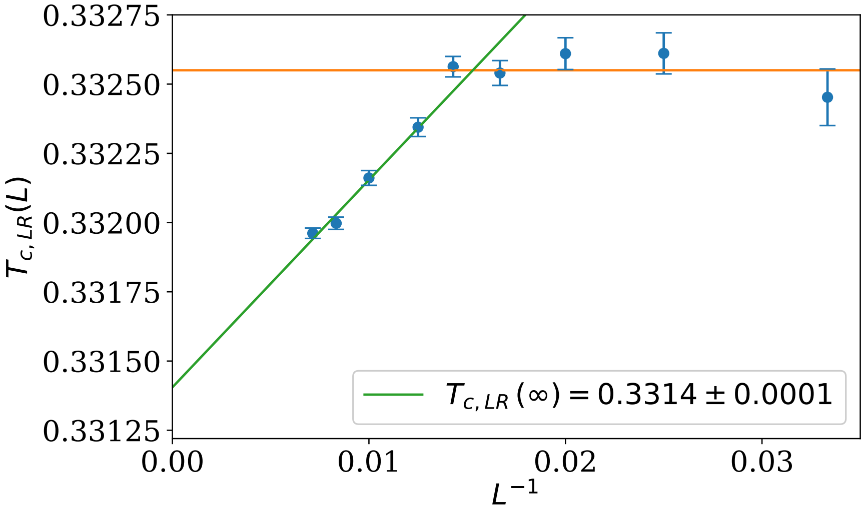

Figure 36: Estimating the critical temperature for the Magnetic-Nematic transition in the Nematic XY model using logistic regression. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{LR}, are fitted to the ansatz in Equation 5. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

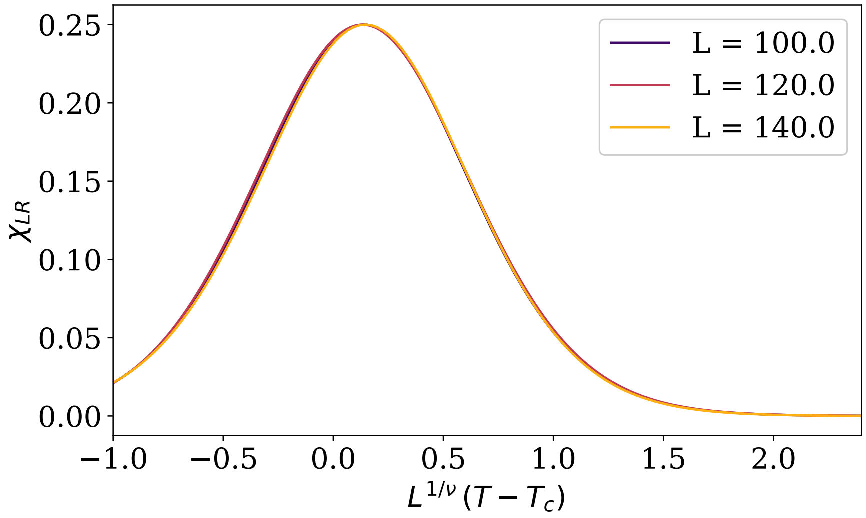

Figure 37: The curve collapse of χ_{LR} for the Magnetic-Nematic transition in the Nematic XY model with $T_c = 0.3315$ and $\nu = 0.8562$.

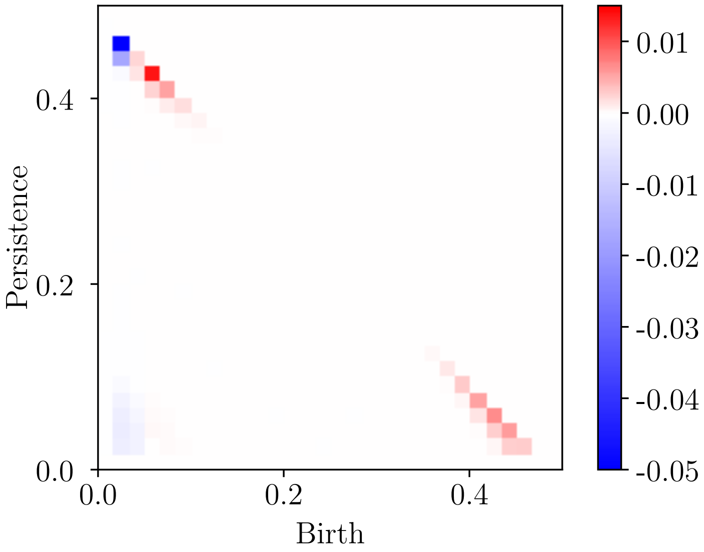

Figure 38: The weights of the logistic regression model trained on configurations over the Magnetic-Nematic transition in the Nematic XY model with $L = 140$.

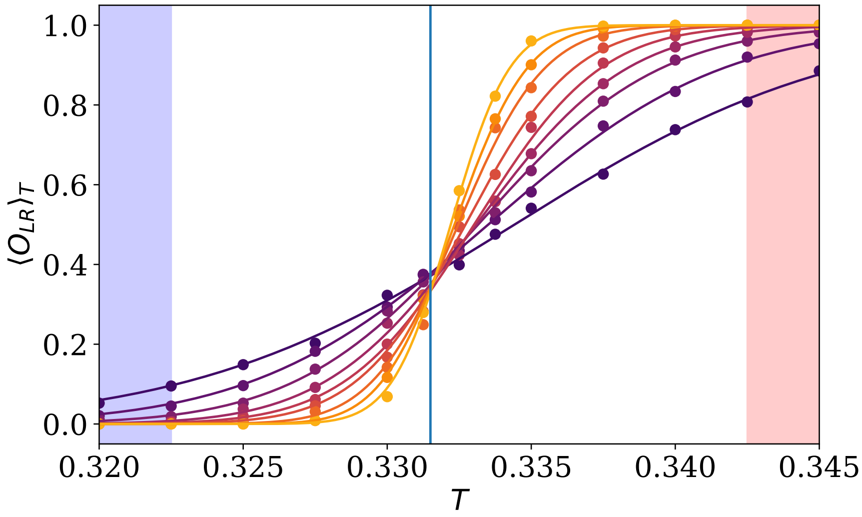

Figure 39: Plot showing $\langle O_{kNN} \rangle$ as a function of temperature for each lattice size over the Magnetic-Nematic transition in the Nematic XY model. The shaded regions indicate the temperatures used for the low and high temperature training data. The vertical line shows the location of the expected critical temperature $T_c = 0.3314$.

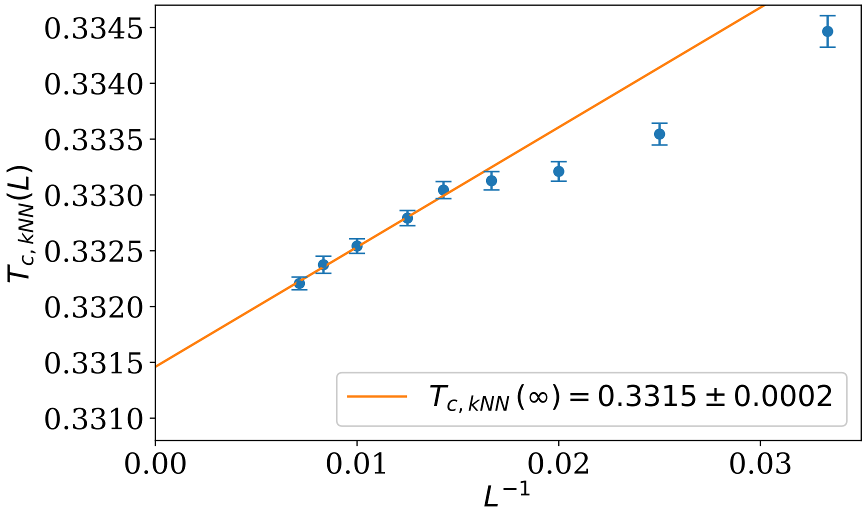

Figure 40: Estimating the critical temperature for the Magnetic-Nematic transition in the Nematic XY model using logistic regression. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{kNN}, are fitted to the ansatz in Equation 5. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

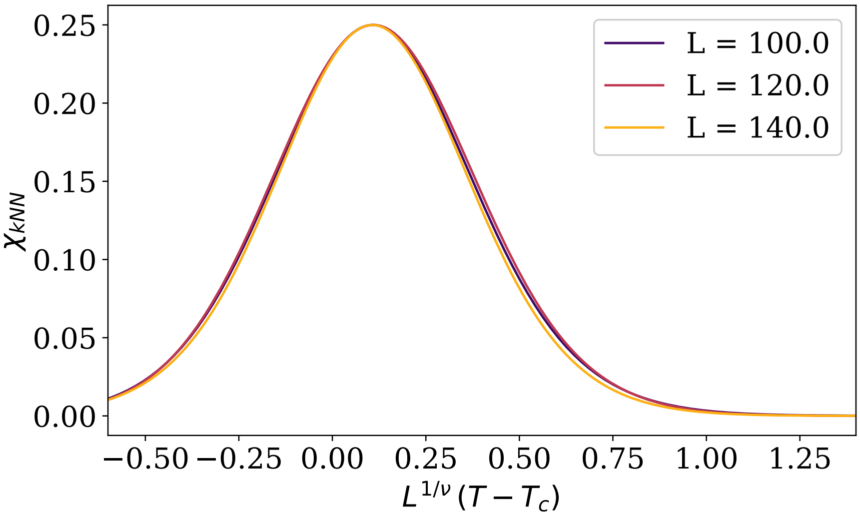

Figure 41: The curve collapse of χ_{kNN} for the Magnetic-Nematic transition in the Nematic XY model with $T_c = 0.3316$ and $\nu = 0.9551$.

3.8.5 Logistic Regression Analysis of Nematic-Paramagnetic Transition

We trained logistic regression models on samples drawn from the low and high temperature phases given in Table 7 with $10,000$ samples from each temperature. The regularisation hyper-parameter was set to $C = 0.001$. We evaluated the models with $10,000$ samples from each temperature in the critical region.

Region	T
Low T	0.74, 0.75, 0.76
High T	0.80, 0.81, 0.82
Critical	0.77, 0.78, 0.79

Table 7: Values of T sampled at for analysis of the nematic-paramagnetic transition in the nematic XY model using logistic regression.

A plot of the resulting phase indicators is shown in Figure 42. The plot of the pseudo-critical temperatures against $\log(L)^{-2}$ is shown in Figure 43. We do not observe any significant lattice-size dependence in the pseudo-critical temperatures. They instead seem to be distributed just above to $T = 0.78$ which is the midpoint of the training temperatures. While a straight line fit to all but the smallest lattice size yields an extrapolated critical temperature of

$$T_c = 0.7804 \pm 0.0002,$$

not too far from the expected $T_c \approx 0.7808$, Figure 43 does not suggest that increasing the statistics would lead to increased accuracy. However, the curve collapse (Figure 44) procedure gives

$$\begin{aligned} T_c &= 0.7803 \pm 0.0025 \\ \nu &= 0.5107 \pm 0.0101 \\ b &= 0.3037 \pm 0.0076, \end{aligned}$$

accounting for the expected value of $T_c = 0.7808$, but giving a potentially questionable result for $\nu = \frac{1}{2}$ which lies just outside one standard deviation.

The weights of the logistic regression model trained for $L = 140$ are shown in Figure 45. We note the similarity to the weights learnt for the XY model in Figure

23 except now the region in the top left represents half-vortices and half-antivortices which change behaviour, shifting down to the right as temperature increases.

3.8.6 k-Nearest Neighbours Analysis of Nematic-Paramagnetic Transition

We trained k-nearest neighbour models on samples drawn from the low and high temperature phases given in Table 8 with 2000 samples from each temperature. The neighbours hyper-parameter was set to $k = 30$. We evaluated the models with $10,000$ samples from each temperature in the critical region.

Region	T
Low T	0.50, 0.55, 0.60, 0.65, 0.70
High T	0.85, 0.90, 0.95, 1.00, 1.05
Critical	0.74, 0.75, 0.76, 0.77, 0.78, 0.79, 0.80, 0.81, 0.82

Table 8: Values of T sampled at for analysis of the nematic-paramagnetic transition in the nematic XY model using k-nearest neighbours classification.

A plot of the resulting phase indicators is shown in Figure 46. The plot of the pseudo-critical temperatures against $\log(L)^{-2}$ is shown in Figure 47. Here we see an asymptotic convergence towards a linear dependence between the pseudo-critical temperatures $T_c(L)$ and $\log(L)^{-2}$. Fitting a straight line to the largest four lattice sizes yields

$$T_c = 0.7766 \pm 0.0034.$$

While this is further from the expected $T_c \approx 0.7808$ than the result of the logistic regression approach, the approach towards the correct finite-size scaling is much clearer. The curve collapse (Figure 48) procedure gives

$$\begin{aligned} T_c &= 0.7757 \pm 0.0064 \\ \nu &= 0.4983 \pm 0.0226 \\ b &= 0.3051 \pm 0.0083, \end{aligned}$$

which is compatible with the expected values of $T_c = 0.7808$ and $\nu = \frac{1}{2}$.

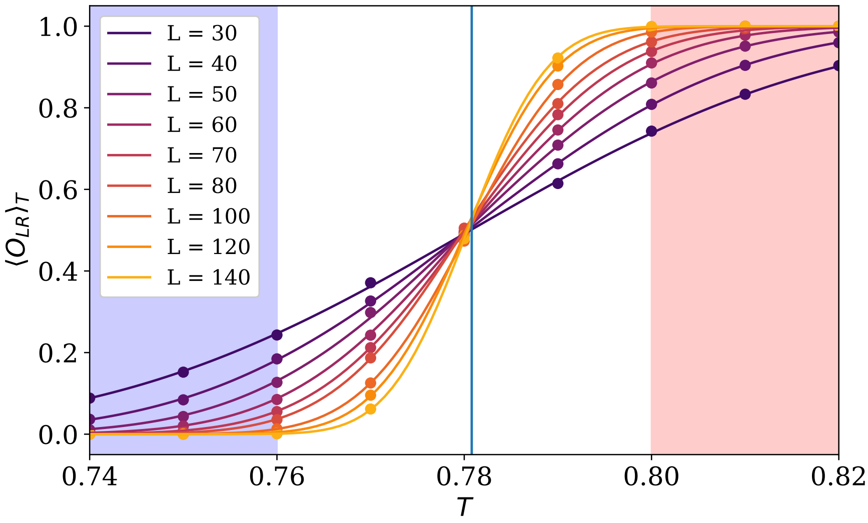

Figure 42: Plot showing $\langle O_{LR} \rangle$ as a function of temperature for each lattice size for the Nematic-Paramagnetic transition in the Nematic XY model. The shaded regions indicate the temperatures used for the low and high temperature training data. The vertical line shows the location of the expected critical temperature $T_c = 0.7808$.

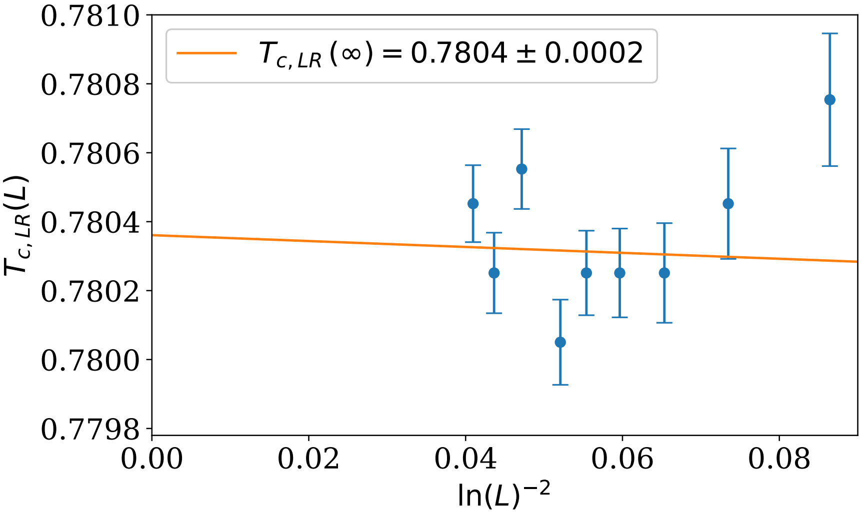

Figure 43: Estimating the critical temperature for the Nematic-Paramagnetic transition in the Nematic XY model using logistic regression. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{LR}, are fitted to the ansatz in Equation 6. We use all the lattice sizes except the smallest in the fit. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

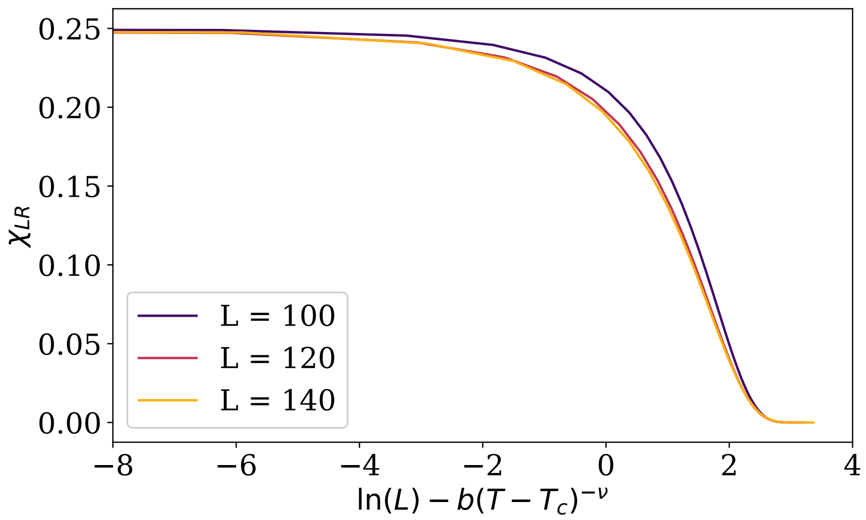

Figure 44: The curve collapse of χ_{LR} for the Nematic-Paramagnetic transition in the Nematic XY model with $T_c = 0.7803$, $\nu = 0.5107$ and $b = 0.3037$.

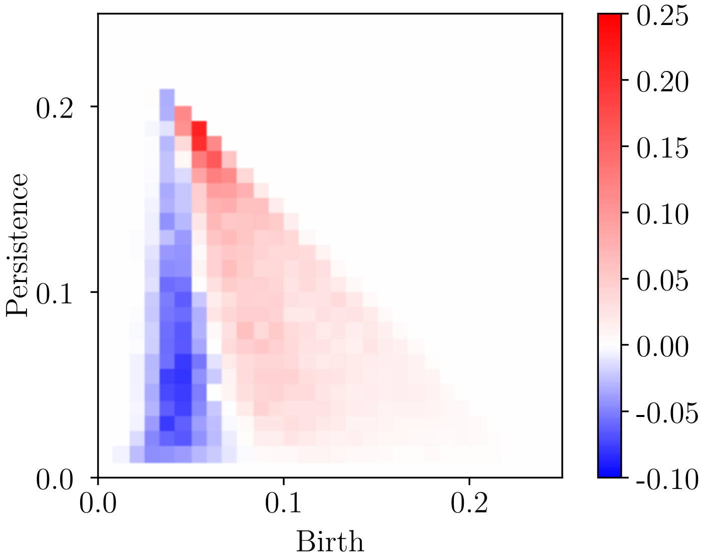

Figure 45: The weights of the logistic regression model trained on configurations over the Nematic-Paramagnetic transition in the Nematic XY model with $L = 140$.

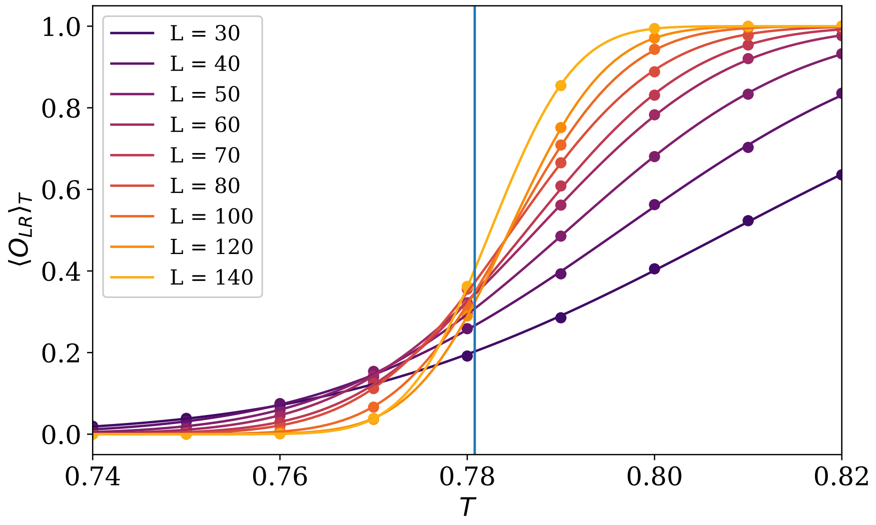

Figure 46: Plot showing $\langle O_{kNN}\rangle$ as a function of temperature for each lattice size for the Nematic-Paramagnetic transition in the Nematic XY model. Note that in this case the training regions lie outside the bounds of the plot. The vertical line shows the location of the expected critical temperature $T_c = 0.7808$.

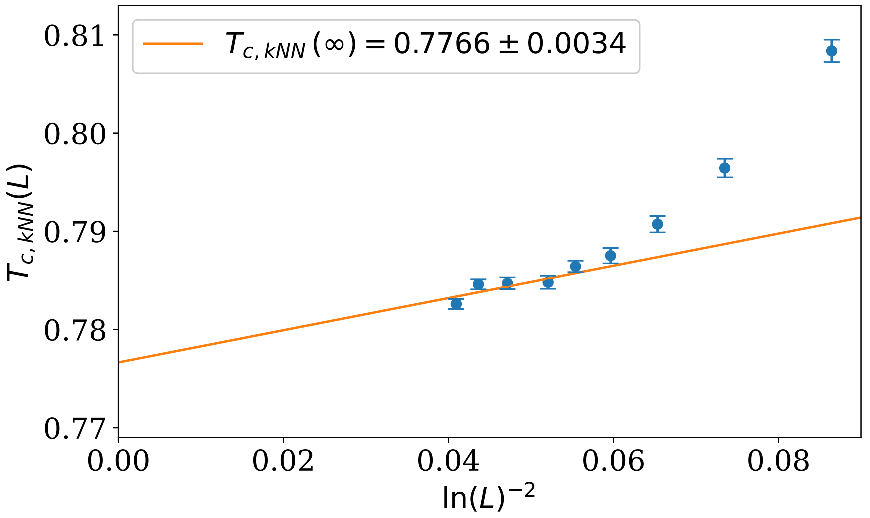

Figure 47: Estimating the critical temperature for the Nematic-Paramagnetic transition in the Nematic XY model using k-nearest neighbours. The pseudo-critical temperatures for the different lattice sizes, calculated from finding the peak of χ_{kNN}, are fitted to the ansatz in Equation 6. We use the largest four lattice sizes for the fit. The intercept gives the estimate for $T_c(\infty)$. Error bars are estimated by bootstrapping.

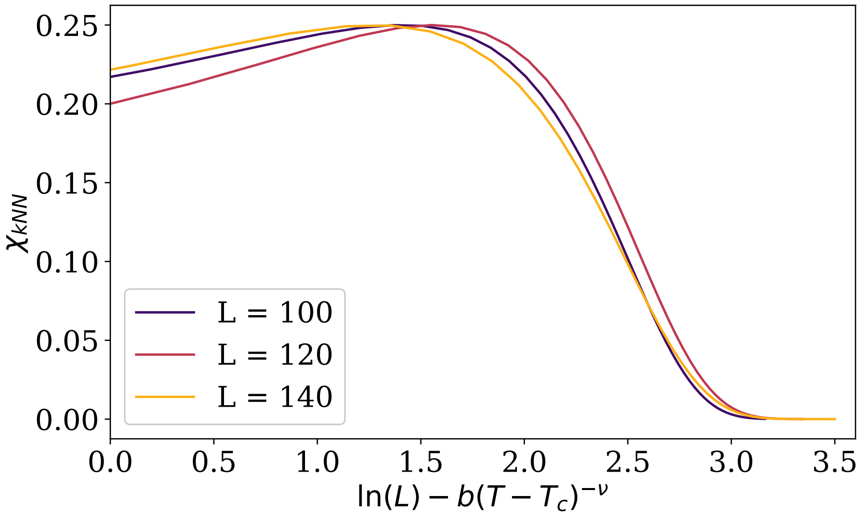

Figure 48: The curve collapse of χ_{kNN} for the Nematic-Paramagnetic transition in the Nematic XY model with $T_c = 0.7757$, $\nu = 0.4983$ and $b = 0.3051$.

3.9 Conclusions and Discussion

We have introduced a way of applying persistent homology to analyse the configurations of lattice spin models and applied this to the 2D XY model with three different Hamiltonians: the standard action, a topological lattice action, and a modified standard action with an additional nematic interaction term. We saw that vortices and antivortices were detected by high persistence points in the H_1 persistence diagram. However, a shortcoming of the filtration considered is that such points may also appear when there is no vortex. Consider the configuration of spins shown in Figure 49. One would not conclude that there is a vortex here, but the angle difference filtered complex would still exhibit a long lasting hole around these spins, yielding a high persistence point in H_1.

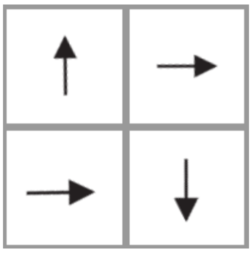

Figure 49: A non-vortex configuration of spins that would result in a high persistence point in persistent H_1 using the angle difference filtration.

We investigated the phase transitions in each model and in each case we were able to successfully identify the transition and estimate its critical temperature and critical exponent of the correlation length by considering the finite-size scaling of observables derived from the persistent homology of configurations. In particular we trained logistic regression and k-nearest neighbours classifiers to identify the phases of the models from persistence images and the critical point was estimated as the temperature at which the variance in the classification reached a maximum. We found that two different filtrations were necessary to detect the two different phase transitions in the nematic XY model. It remains an open question of whether or not the filtrations introduced by other authors for XY models (e.g., in [45, 72]) would detect both simultaneously. However, it certainly leads one to hesitate before proclaiming persistent homology a universal tool for the investigation of phase transitions without prior knowledge. In particular, it is clear now that the choice of filtered complex is important, and good design of these seems to rely on having at least some insight into the mechanism of a phase transition.

We have found that the approach of using logistic regression for classification as introduced in [45], while useful for interpreting which regions of the persistence image indicate the different phases, fails to produce accurate estimates of the critical temperature or exponents in the case of the BKT transitions. Instead it will tend to yield the midpoint between the low and high-temperature training temperatures as the critical temperature. Indeed, using different temperatures for the training causes the estimated critical temperature to shift accordingly. This failure may be because logistic regression is a generalised linear model and the data here is highly non-linear. On the other hand, the non-parametric k-nearest neighbours approach generally produces good results, with a clear asymptotic approach towards the expected finite-size scaling behaviour in all cases.

There are a number of interesting questions and directions for further research:

- The analysis of representative cocycles might be improved by choosing a more specific representative according to some criterion. For example, a cocycle c with minimum length (i.e., minimising $|\operatorname{Supp}(c)|$). This may provide a method to associate vortices with antivortices in the low temperature phase. However the calculation of such a cocycle would involve solving a non-trivial optimisation problem.
- The issue of non-vortex configurations of spins yielding high persistence points in H_1 could potentially be addressed through the use of representative cycles. If a basis of cycles could be found such that each surrounds just one vortex (it may be sufficient to simply take the boundary of the 2-cube which kills the homology class), then a discrete version of the winding number can be calculated around that cycle and used to annotate the corresponding point in the persistence diagram.
- The approach presented could easily be extended to other classical lattice spin models, but it would also be interesting to see if the filtrations presented in Section 3.4.2 could be adapted to the quantum lattice spin models.
- The similarity of the persistence images across the BKT transition in all three models raises the question of the extent to which the persistence is a universal

quantity. This could potentially facilitate a transfer learning approach where classifiers trained on one model can identify phase transitions of the same universality class in another model (see, e.g., [64]).

- It could also be investigated if the use of a vectorisation and a classifier is necessary in the first place. There is a notion of variance for persistence diagrams called Fréchet variance [93] which might show finite-size scaling behaviour directly. However this is computationally expensive to measure.
- Finally, we note that there have been a variety of different filtrations used to compute the persistent homology of configurations of lattice spin models (e.g., [45, 72, 74]). It would be interesting to see how these perform and complement one another on a single data set.

4 SU(2) **Lattice Gauge Theory**

Quantum chromodynamics (QCD) poses several outstanding problems in particle physics, including the mechanism of confinement and the deconfinement phase transition, mass-gap generation, and chiral symmetry breaking [94]. These phenomena are non-perturbative and are therefore typically investigated through the framework of Lattice QCD. Supplementing this traditional approach, there is an emerging body of work exploring the use of machine learning and data analysis tools in extending Monte Carlo analysis of lattice QCD towards generating the insights required to tackle these open problems. References include [95–97] among others. For example, one application of machine learning is to classify phases of the theory based on sampled configurations, learning observables that function as order parameters for the phase transitions undergone by QCD, following the quantitative approach outlined for spin models in [98].

Focusing on confinement, a compelling potential mechanism relies on the presence of topological defects called center vortices in confining gauge configurations [99–101]. Vortex-like configurations have recently been shown to exist in pure gauge theories (e.g., [102–128]) and have recently been observed in lattice simulations of QCD [129]. While order parameters have been constructed for the confinement-deconfinement phase transitions in Yang-Mills theories that are based on the topological symmetry related to the conservation of the number of vortices (see, e.g., [102, 130, 131]), identification of vortices proves to be a more challenging undertaking. In fact, existing methodologies for exposing vortices in gauge theories rely on performing gauge fixing and projection [132]. This procedure suffers from the problem of Gribov ambiguities [133, 134], which, even with careful choices of the gauge fixing condition (for instance, following the prescription of [135]), can be mitigated only in part. Motivated by the physical appeal of a fully gauge-independent description, we investigate the possibility of instead analysing vortices and performing phase classification in a gauge-invariant manner by making use of persistent homology. Rather than considering full QCD, in this chapter we instead develop a methodology for the pure gauge SU(2) lattice gauge theory which also exhibits a deconfinement transition potentially driven by center vortices. In the context of vortex identification, this system represents a simplified toy model of QCD where the quarks have been removed, and the gauge group is simplified from SU(3).

4.1 Model

A configuration of the 4D SU(2) lattice gauge theory is specified by SU(2)-valued variables $U_\mu(x)$ located on each link (x, μ) of an $N_t \times N_s^3$ lattice Λ with periodic boundary conditions, where $\mu \in \{0, 1, 2, 3\}$ describes the direction in which the link emanates from the lattice site $x \in \Lambda$. In practice, $U_\mu(x)$ lies in the fundamental representation of SU(2), taking the form of a 2×2 complex matrix. To simulate at non-zero temperature we ensure $N_t \ll N_s$. The gauge symmetry is generated by gauge transformations $\Omega \in \mathrm{SU}(2)^\Lambda$ sending each $U_\mu(x) \mapsto \Omega(x) U_\mu(x) \Omega^\dagger(x + \hat{\mu})$, where $x + \hat{\mu}$ denotes the lattice site one step in the μ direction from x. It turns out that the observables that are invariant under these transformations are traces of products of the link variables along closed paths C, also known as Wilson loops $W(C)$. The simplest non-trivial example is the Wilson loop around a 1×1 plaquette (x, μ, ν) of the lattice:

$$W_{\mu,\nu}(x) = \frac{1}{2}\,\mathrm{tr}\left[U_\mu(x)\,U_\nu(x + \hat{\mu})\,U_\mu^\dagger(x + \hat{\nu})\,U_\nu^\dagger(x)\right].$$

We use this to define the Wilson action evaluated on a configuration $\mathbf{U} = \{U_\mu(x)\}_{(x,\mu)}$ as

$$S(\mathbf{U}) = -\frac{\beta}{4} \sum_{x,\mu<\nu} W_{\mu,\nu}(x) \tag{8}$$

where β is an inverse gauge coupling parameter. This in turn allows us to define the vacuum expectation value of any given observable $A(\mathbf{U})$ as

$$\langle A \rangle = \frac{\int d\mathbf{U}\, A(\mathbf{U})\, e^{-S(\mathbf{U})}}{\int d\mathbf{U}\, e^{-S(\mathbf{U})}} \tag{9}$$

where $d\mathbf{U} = \prod_{x,\mu} dU_\mu(x)$ is a product of Haar measures over SU(2) for each link variable. In practice we estimate expectations using Monte Carlo methods, where Eq. (9) becomes a simple mean of the observed values.

The model represents a discretised version of the SU(2) Yang-Mills theory after quantisation via functional integrals and Wick rotation into Euclidean space. Letting a denote the lattice spacing as a physical length, the Wilson action approximates the Yang-Mills action as the continuum limit $a \to 0$ is approached with discretisation errors of $\mathcal{O}(a^2)$. Then β is related to the Yang-Mills gauge coupling g by $\beta = 1/4g^2$.

4.2 Confinement and Center Vortices

The model introduced above exhibits two phases – a confined phase at low β and a deconfined phase at high β – and the phase transition between these is known as the deconfinement transition. This transition occurs only in the spatial continuum limit $N_s \to \infty$ and while we care also about the temporal limit $N_t \to \infty$, it's worth noting that a critical β is well defined for each finite N_t and it is these we will be estimating. Confinement in the SU(2) lattice gauge theory can be characterised in a number of ways [94], including:

- **Area law for Wilson loops.** Let $W(C)$ denote the value of a Wilson loop around a closed curve $C = R \times T$, consisting of a rectangle of length R in a space-like direction and T along the temporal direction. We consider the limit of large area $\mathcal{A}(C) = RT$. In the confined phase at low β we have that $\langle W(C)\rangle \propto exp(-\sigma\mathcal{A}(C))$ where σ is known as the string tension. In the deconfined phase at high β we have that $\langle W(C)\rangle$ decays exponentially instead with $\mathcal{P}(C)$, the perimeter of C.
- **Vanishing Polyakov loop.** Define the Polyakov loop at a point x in the lattice as
$$P(x) = \frac{1}{2}\, tr\Big[U_0(x)\, U_0(x+\hat{0})\, \ldots\, U_0(x-\hat{0})\Big]$$
(note this only depends on the spatial coordinates of x). This represents a Wilson loop that has a non-zero winding number around the time direction of the periodic lattice. In the confined phase we have $\langle P(x)\rangle = 0$ whereas in the deconfined phase we have $\langle P(x)\rangle \neq 0$.

Several possible pictures of what drives the deconfinement transition, in both this model and QCD, have been proposed. Here we focus on the center vortex picture [99, 100]. Fix a time slice at time t. Given two closed oriented contours C and C' in that 3-dimensional slice with linking number m, a loop operator $B(C', t)$ can be defined that has the following commutation algebra with the Wilson loop $W(C, t)$:

$$W(C,t)B(C',t) - (-1)^m B(C',t)W(C,t) = 0. \tag{10}$$

This equation defines the 't Hoof algebra [99]. For simplicity, we consider planar non-intersecting curves C and C', for which $m = 0, 1$. The operator $B(C', t)$ is called the 't Hooft loop. When acting on a gauge configuration, $B(C', t)$ creates a magnetic flux with the resulting observable effect of multiplication by -1 of all Wilson loops having support on curves C with linking number 1 with C'. For this reason, the 't Hooft loop is said to be a vortex creation operator. Since the center of the group, which in our case is $Z(\mathrm{SU}(2)) = \{I, -I\} \cong \mathbb{Z}_2$, plays a role in the 't Hooft algebra (as exposed by the factor $(-1)^m$), the vortices created by the 't Hooft loop operator are called center vortices. Fixing the curve C' for all time slices t, we see that a vortex traces out a surface in 4-space, closed by the periodic boundary conditions.

In the limit of weak fields, where the theory is deconfined, all Wilson loops are close to unity. Confining configurations are expected to have Wilson loops that largely deviate from unity. In particular, Wilson loops close to -1 can be obtained from a weak field configuration through the injection of center vortices generated with appropriate insertions of 't Hooft loop operators. Moving from this observation, operationally we can define a center vortex to be a collection of plaquettes in the dual lattice (in the sense of the dual graph) that form a closed surface (with the closedness being a consequence of the Bianchi identities) and that carry a non-trivial charge in $\mathbb{Z}_2$, corresponding to the -1 element. To carry a non-trivial center charge means that any Wilson loop in the lattice that topologically links with this surface is multiplied by that charge.

In the confined phase center vortices are found to form large surfaces, often wrapping round the periodic boundaries, that percolate throughout the lattice [136]. Therefore, given a particular Wilson loop $W(C)$, the number of vortices that link with C is proportional to the enclosed area $\mathcal{A}(C)$, leading to the area law for the suppression of $\langle W(C) \rangle$. In the deconfined phase, the center vortices become smaller and more sparse, ensuring that for sufficiently large Wilson loops, only those vortices close to the curve C have a chance of linking with it, leading to the perimeter law. Similarly we see that only in the confined phase, where vortices may wrap around the periodic boundary conditions of the lattice, is there a chance they may link with a Polyakov loop, suppressing its expectation. For an overview of the evidence supporting the center vortex picture see [137]. In practice, center vortices generated by the system have some finite thickness, so that only larger Wilson loops may fully link with them and obtain the full center charge. Loops that partially link may still obtain a partial charge, some factor lying between I and $-I$ in SU(2).

While the concepts of vorticity and creation of a vortex through the insertion of a 't Hooft loop are well understood in terms of symmetry and boundary conditions in a finite volume (see for instance [99, 102, 130, 131]), quantum fluctuations make vortex identification a much more involved process, with currently used prescriptions not fully validated from first principles. A widely used method to detect and analyse these thick vortices is to transform configurations to the maximal center gauge, where each matrix $U_\mu(x)$ is as close to either I or $-I$ as possible, then project the matrices onto whichever of I or $-I$ is closer. After projection, the Wilson loops of plaquettes take values either 1 or -1 and the latter are identified as projected vortices, or P-vortices. It has been shown that the locations of these correlate with the unprojected thick vortices [138]. However, the projection means that we lose gauge invariance, as well as geometric information such as the thickness of the vortices. In this work we introduce a method designed to look for thin vortices, but we argue that, through the use of persistent homology, the results tell us something about these thick vortices too.

There is no analytical formula for the critical value of β for each value of N_t, but detailed numerical studies have shown good agreement with each other. We will compare our results to those found in [139] which we reproduce in Table 9. We also estimate the critical exponent of correlation length ν which, as a consequence of the Svetitsky–Yaffe conjecture [140], is the same as for the 3D Ising model. We therefore compare our measurement of ν to the high precision estimate $\nu = 0.629971(4)$ from [141].

N_t	β_c
4	2.2986(6)
5	2.37136(54)
6	2.4271(17)

Table 9: Estimates for the critical value of β for the deconfinement phase transition in the SU(2) lattice gauge theory for the values of N_t we consider in this thesis. Reproduced from [139].

4.3 Twisted Boundary Conditions

In order to test the sensitivity of our method to thin vortices we will make use of the trick of imposing twisted boundary conditions [142]. The idea is that we choose some co-closed collection of plaquettes in the lattice, i.e., plaquettes that link with some closed surface in the dual lattice, and negate their contribution to the action. See Figure 50 for an illustration of co-closed collections of plaquettes. In our case we choose the plaquettes

$$T = \{((0,0,y,z),0,1) \mid 0 \leq y,z \leq N_s\}$$

corresponding to a surface wrapping round the latter two spatial dimensions of the lattice. The action with twisted boundary conditions becomes

$$S_T(\mathbf{U}) = -\frac{\beta}{4}\left[\sum_{\substack{x,\mu<\nu \\ (x,\mu,\nu)\notin T}} W_{x,\mu,\nu} - \sum_{\substack{x,\mu<\nu \\ (x,\mu,\nu)\in T}} W_{x,\mu,\nu}\right] \tag{11}$$

which we refer to as the twisted action.

This modification of the action allows the lattice to support an odd number of center vortices wrapping in the yz plane, which is prohibited by the usual periodic boundary conditions of the Wilson action. It's important to note that we are talking about the boundary conditions of the gauge field on the lattice and not the lattice itself. We are not twisting the lattice and forming any kind of Möbius band, rather it is the gauge field which obtains a factor of $-I$ as we loop around the lattice. We can alternatively think of this twisted action as explicitly inserting a thin vortex into the system on the surface defined by T, so that the system is forced to generate a (thick) vortex to cancel it out. We shall denote expectations calculated with respect to this twisted action by $\langle A \rangle_{\text{twist}}$.

Twisted boundary conditions give us an alternative way to characterise confinement and the deconfinement transition. Magnetic and electric flux free energies can be defined in terms of the ratio of partition functions for the twisted and Wilson actions and the behaviour of these can be shown to imply the area law decay for the Wilson loop [143] and therefore confinement.

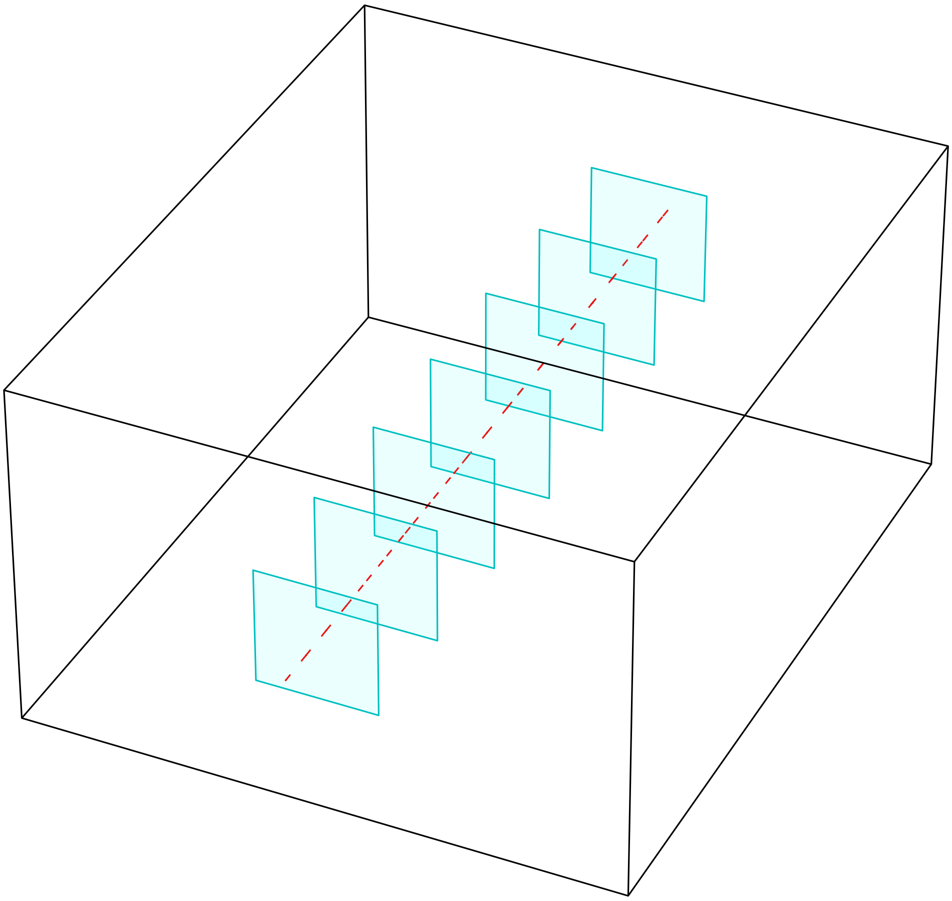

Figure 50: A lower dimensional illustration of a co-closed collection of plaquettes that wraps around the periodic boundary conditions of the lattice. Plaquettes in a 3D lattice link with edges in the dual lattice, so the condition of being co-closed means that the collection of those linking edges forms a closed loop. In this case the loop is closed by the periodic boundary conditions. Going to 4 dimensions, we imagine repeating the co-closed line of plaquettes along the new dimension, forming a co-closed surface of plaquettes.

4.4 Previous Work

We briefly remark on two papers that have used persistent homology to look at QCD and lattice gauge theories. Kashiwa, Hirakida and Kouno noted the sign problem in simulating dense QCD and instead considered an effective model in the form of a modified 3D Potts model where each $\mathbb{Z}_3$ Potts spin corresponds to a Polyakov loop [144], avoiding the need to construct a filtered complex from gauge fields directly. Computing the persistence of pointclouds of lattice sites sharing the same Potts spin, they were able to probe the phase structure of the model using the average and maximum birth-death ratio of points in the resulting persistence diagrams. More recently, Sehayek and Melko investigated the 2D and 3D $\mathbb{Z}_2$ lattice gauge theories [80]. Given a configuration of the model, they place a point at the centre of each spin-down lattice link and compute the Vietoris-Rips filtered complex of the resulting pointcloud. The β_1 Betti curve then provides the number of closed strings of down spins and their sizes. These quantities and the filtered complex are gauge variant, but by averaging the loop count over many configurations they indirectly measure the density of vison defects which produces a clear indicator of the phase transition in the 3D model when plotted as a function of temperature.

4.5 Methods

4.5.1 Simulation

In order to sample the SU(2) lattice gauge theory, we will use the HiRep software [145]. This generates a Markov chain, but rather than using the Metropolis algorithm, it uses the heatbath and over-relaxation algorithms.

Heatbath Algorithm The idea here is to update each link variable $U_\mu(x)$ one by one while holding all the others fixed. In particular, given a current configuration $\mathbf{U} = \{U_\nu(y)\}$ and a particular link (x, μ), consider the probability distribution $P_{(x,\mu),\mathbf{U}}$ on SU(2) such that

$$P_{(x,\mu),\mathbf{U}}(U') \propto \exp\left(\frac{\beta}{4} \sum_{X \text{ staple of } (x,\mu)} \frac{1}{2} \operatorname{tr} X^\dagger U'\right),$$

appropriately normalised and where X ranges over those 6 plaquettes containing the link (x, μ), denoting the product around the other three links. In other words, the sum inside the exponential is the local contribution to the action coming from having the matrix U' on the link (x, μ) and all other links as in $\mathbf{U}$. Then a step of the heatbath algorithm is given in Algorithm 3. The algorithm is named as such because it "consists of successively touching a heat bath to each gauge variable in the system" [146].

Algorithm 3 Heatbath step [146]

Input Configuration $\mathbf{U} = \{U_\mu(x)\}$
Output Updated configuration $\mathbf{U'} = \{U'_\mu(x)\}$

```
for $x \in \Lambda$ do
    for $\mu \in \{0, 1, 2, 3\}$ do
        $U'_\mu(x) \leftarrow \text{Sample}(P_{(x,\mu),\mathbf{U}})$
    end for
end for
```

Over-relaxation This algorithm, introduced in [147], is inspired by over-relaxation approaches for solving linear systems such as those in [148]. The idea is to perform the heatbath algorithm, but to "overcorrect" each link variable and therefore cause the Markov chain to decorrelate faster so that it can be sampled more frequently. Choose some over-relaxation parameter $\omega \in [1, 2]$ and suppose a step of the heatbath algorithm would replace a matrix U at link (x, μ) with a matrix V. Then the rotation from U to V is given by $VU^\dagger$ and we simply increase the rotation in this direction to get a new matrix $U\prime = (VU^\dagger)^\omega U$. In practise, we avoid using fractional powers and use $\omega = 2$, noting that $\omega = 1$ gives the standard heatbath algorithm.

In this thesis all Markov chains generated with HiRep will have 1 heatbath sweep and 4 over-relaxation sweeps between each step of the chain. In general we will sample the chain every 100 such steps (i.e., after 100 heatbath and 400 over-relaxation sweeps) in order to obtain samples with low autocorrelation.

4.5.2 Wilson Flow

The Wilson flow, introduced in [149], is a gradient flow on configurations taking them towards lower action values. The flowed configuration $\boldsymbol{V}(t)$ at Wilson flow time t is defined by the differential equation

$$\frac{\partial V_\mu(t,x)}{\partial t} = -\frac{4}{\beta}\{\partial_{x,\mu} S(\boldsymbol{V}(t))\} V_\mu(t,x)$$

with initial conditions $V_\mu(0,x) = U_\mu(x)$, where $\partial_{x,\mu}$ is a Lie algebra-valued differential operator (see Appendix A in [149] for a precise definition which we shall not need here). We will compute this gradient flow with the HiRep software [145] which uses an adaptive 3rd order Runge-Kutta integrator.

The Wilson flow can be thought of as a way to smooth out short-range singularities from the gauge field which can help with certain measurements such as the topological susceptibility. However, we are primarily interested in it as a means to analyse the stability of center vortices. In order to do this, we will need a reference characteristic Wilson flow time t_0 with which to compare the time that a vortex lives under the flow. Let $C_{\mu,\nu}(x)$ be the average of the product of links around each of the four (μ,ν)-plaquettes around x. Define

$$E_{sym} = \sum_{x,\mu<\nu} tr(C_{\mu,\nu}(x))$$

which is a sort of improved estimator for the action. The reference scale t_0 for SU(3) lattice gauge theory and lattice QCD is defined by the implicit equation $t_0^2 E_{sym}(\boldsymbol{V}(t_0)) = 0.3$ [149]. For SU(2), the authors of [150] suggest scaling the value by the quadratic Casimir, giving the definition

$$t_0^2 E_{sym}(\boldsymbol{V}(t_0)) = 0.3 \times (3/4)^2$$

that we use in this thesis.

As an example, we sample 100 configurations at $\beta = 2.31$ on a 12^4 hypercubic lattice and calculate the Wilson flow, measuring E_{sym} at flow times $t = 0, 0.015, 0.03, \ldots, 1.5$. A plot of the average of E_{sym} as a function of flow time is shown in Figure 51. We use these to calculate t_0 for each configuration, obtaining the distribution shown in Figure 52. This gives us an estimated value of

$$t_0 = 0.99 \pm 0.06$$

for $\beta = 2.31$ which we will also use in Section 4.8.

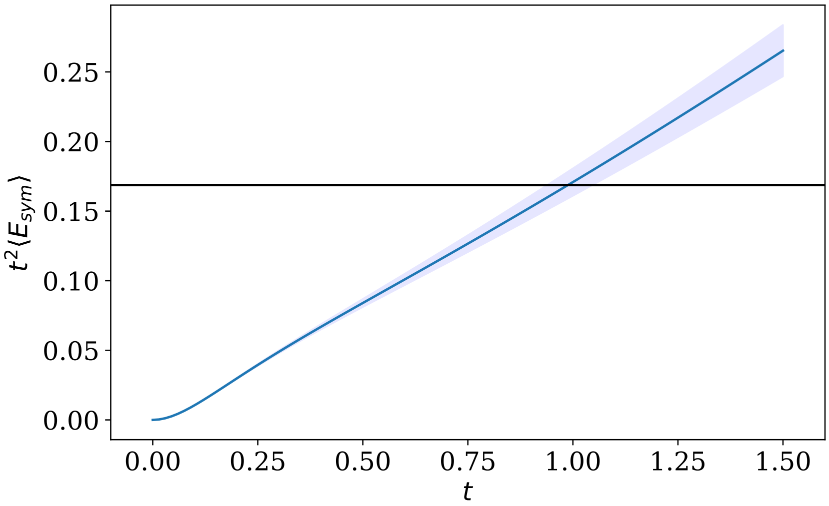

Figure 51: The average of E_{sym} as a function of Wilson flow time on a 12^4 lattice at $\beta = 2.31$. The shaded region indicates plus and minus one standard deviation.

4.5.3 Filtered Complex

To apply persistent homology we must choose how to define a filtered complex for a given configuration $\boldsymbol{U} = \{U_\sigma(x)\}$. We present a filtered complex $F_{\boldsymbol{U}}$ which is constructed based on Wilson loops and which will therefore give gauge-invariant persistence diagrams.

The idea is to explicitly construct a cubical model of vortex surfaces, under the assumption that vortices are thin. Note that since we have periodic boundary conditions, spacetime in this model forms a 4-torus $S^1 \times S^1 \times S^1 \times S^1$. Moreover, both the lattice Λ and its dual Λ^* (as in dual graph) can be considered as providing a way of turning this space into a cubical complex. The lattice Λ gives us a complex X with a vertex for each lattice site and the dual lattice Λ^* gives us a complex Y with a vertex for each 4-cube in X and a 4-cube for each vertex in X. Plaquettes in X are in bijection with those in Y, with a plaquette p of X corresponding to the unique plaquette q in Y that intersects p at a single point. Since vortices live on the dual lattice, we make use of this 1-to-1 correspondence to filter Y according to the value of Wilson loops around plaquettes in X.

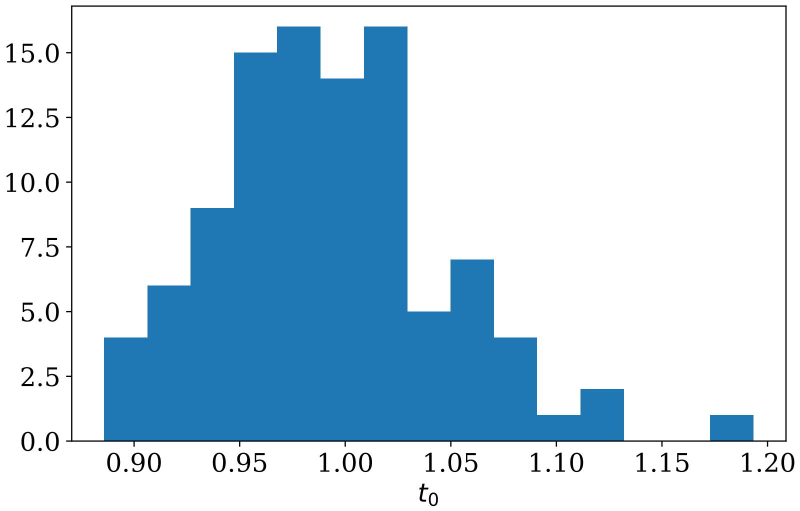

Figure 52: Distribution of t_0 for $\beta = 2.31$.

Denote by $c_A(y) = \prod_{\mu\in A}[y, y + \overrightarrow{\mu}]$ a cube in Y, where $[y, y + \overrightarrow{\mu}]$ is the line segment between lattice site $y \in \Lambda^*$ and $y + \overrightarrow{\mu}$. The cube $c_\emptyset(y)$ is just the point y itself. The dimension d of $c_A(y)$ is $|A|$ and we will refer to it as a d-cube. The boundary $\partial c_A(y)$ of a d-cube is the set of its $d-1$-cube faces. For example $\partial c_{\{\mu\}}(y) = \{c_\emptyset(y), c_\emptyset(y + \overrightarrow{\mu})\}$. With this notation, our observation of the bijection between plaquettes in X and Y becomes that the 2-cube $c_{\{\mu,\nu\}}(y)$ is matched with the 2-cube $c_{\sigma,\tau}(x)$ in X (defined similarly) that is used to define the Wilson loop $W_{\sigma,\tau}(y+\overrightarrow{\mu}+\overrightarrow{\nu})$ where $\{\sigma, \tau\} \cap \{\mu, \nu\} = \emptyset$.

To define the filtered complex we will give a filtration index $f(c_A(y)) \in \mathbb{R}$ for each cube $c_A(y)$ in Y specifying when it appears. Then $F_{\boldsymbol{U}}(s)$ is the subcomplex of Y consisting of all cubes c for which $f(c) \leq s$. That is,

$$F_{\boldsymbol{U}}(s) = f^{-1}(-\infty, s].$$

Since we are attempting to model vortex surfaces, we will initially specify when the 2-cubes are to enter the filtered complex and then introduce the cubes of other dimensions based on these.

Our construction of the function f is the following:

1. We introduce each 2-cube $c_{\{\mu,\nu\}}(y)$ in our filtered complex at index

   $$f(c_{\{\mu,\nu\}}(y)) = W_{\sigma,\tau}(y + \overrightarrow{\mu} + \overrightarrow{\nu})$$

   where $\{\sigma,\tau\} \cap \{\mu,\nu\} = \emptyset$. That is, at an index equal to the value of the Wilson loop around the plaquette in X paired with it by the bijection.

2. Since a 2-cube is not allowed to be included before its constituent 1-cubes and 0-cubes in a cubical complex, we introduce these at the smallest index of all the 2-cubes they are incident to. So

   $$f(c_A(y)) = \min\{\, f(C) \mid c_A(y) \in \partial C \,\}$$

   when $|A| \leq 1$.

3. For the 3-cubes and 4-cubes we follow a clique-like rule where we introduce a cube as soon as all of its boundary cubes are introduced. So

   $$f(c_A(y)) = \max\{\, f(C) \mid C \in \partial c_A(y) \,\}$$

   when $|A| \geq 3$.

Thus for $s < -1$, $F_{\boldsymbol{U}}(s)$ is the empty complex and for $s \geq 1$, $F_{\boldsymbol{U}}(s)$ is the filled in tiling homeomorphic to a 4-torus. Going between these values, the first cubes to enter $F_{\boldsymbol{U}}$ are surfaces made up of plaquettes which link with Wilson loops that are close to -1. The idea therefore is that thin vortex surfaces will enter the filtered complex early. Moreover, since small Wilson loops like those considered here still pick up a partial charge from thick vortices, surfaces representing those thick vortices ought to enter the filtered complex earlier than they otherwise would have. We expect to detect these closed surfaces in persistent H_2 (since we are using $\mathbb{Z}_2$ coefficients, the orientability of the surfaces does not impact this). We may also see other topological features such as the presence of handles or holes in H_1, as well as the transient low-persistence points in persistent H_0 and H_1 that arise as the surface forms near the start of the filtration. An illustration of the connection between Wilson loops and the inclusion of vortices into the filtration is shown in Figure 53. An illustration of how to imagine what the filtered complex is aiming to do is shown in Figure 54.

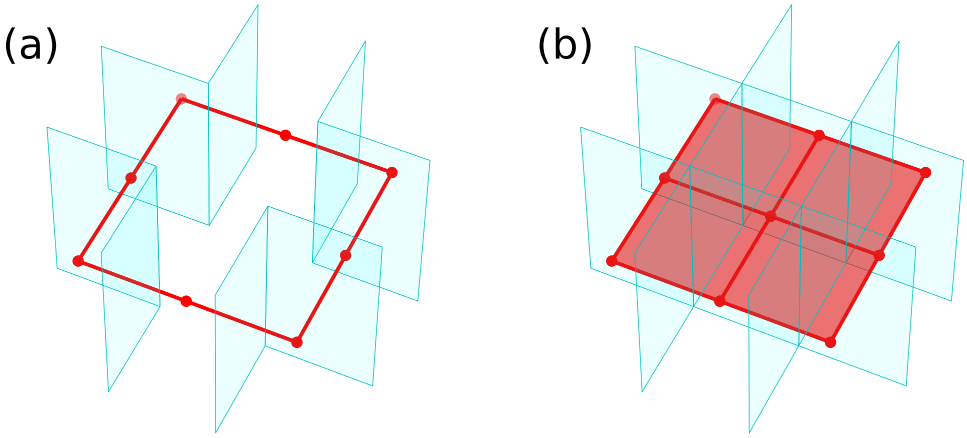

Figure 53: A lower dimension illustration of the idea behind the filtered complex. In 3 dimensions, center vortices form closed 1-dimensional curves that link with 2-dimensional plaquettes. In this setting we would include edges (dark/red) according to the Wilson loop around the plaquette (light/cyan) they link with. (a) Early on in the filtered complex we include edges that link with plaquettes with negative Wilson loop values. In this way we build explicit cubical models of 1-dimensional center vortices which are then detected in PH_1. (b) Later on we eventually fill in the rest of the edges and the 2-cubes between them, destroying the PH_1 features corresponding to the vortices. Moving from 3 dimensions to 4, we are inserting plaquettes instead of edges and instead of a closed curve we obtain a closed surface which we detect with PH_2 (assuming the surface is orientable).

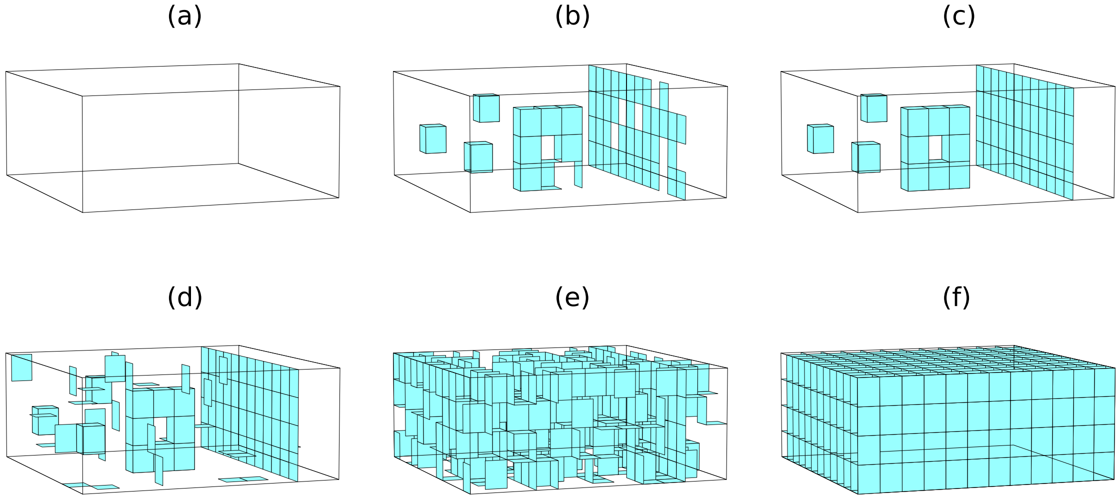

Figure 54: A cartoon of a 3D slice of the filtered complex. (a) At the beginning of the filtration the complex is empty. (b) At low Wilson loop values we begin to form vortex surfaces with transient features detected in H_0 and H_1. (c) The vortex surfaces close, becoming detectable in H_2. (d) At higher Wilson loop values, vortex surfaces persist while other plaquettes begin to be included in the complex. (e) The vortex surfaces may become filled in if all plaquettes inside the surface are included, killing the corresponding H_2 class. (f) Eventually the whole 4-torus cubical complex is filled in.

It is worth noting the difference in approach from our previous work using persistent homology to identify vortices in 2D XY models as in Chapter 3. Vortices there were point defects (located at vertices of the dual lattice), which we aimed to detect by constructing 1-dimensional loops in the original lattice that encircled them. Here, we are modelling the vortex surfaces in the dual lattice directly.

It is straightforward to see that this filtered complex is stable with respect to perturbations of the SU(2) link variables since the Wilson loop $W_{\mu,\nu}(y)$ is a linear map due to the linearity of the trace and is therefore Lipschitz continuous. The stability property of persistent homology [151] therefore ensures that a small perturbation of the link variables only results in a small perturbation of the resulting persistence diagram with respect to the bottleneck distance.

4.6 Detecting Vortices

We first investigate the ability of the persistent homology of our filtered complex to identify an inserted thin vortex, obtained using twisted boundary conditions, as a function of β. For $N_s \in \{12, 16, 20\}$, fixing $N_t = 4$, we generate 200 configurations using the Wilson action (8) and 200 configurations using the twisted action (11) for each $\beta \in \{1.5, 1.6, \ldots 2.9\}$. Configurations are generated using the HiRep software [145] with 1 heatbath step and 4 overelaxation steps for each Monte Carlo step and a sample taken every 100 Monte Carlo steps.

Since the inserted vortex forms a closed surface, we expect to observe it in the PH_2 diagram. Moreover, the surface wraps round the periodic boundary of the lattice in the latter two spatial directions. If it is the first such surface to wrap around those dimensions to enter the filtered complex, then we will observe it as a point in PH_2 with infinite death index since it encloses a two-dimensional hole which remains even in the final complex of the filtered complex (homeomorphic to a 4-torus). Otherwise, it would appear as a point with finite death index. In Figure 55 we compare the persistence diagrams of individual sampled configurations in each phase using the Wilson action and the twisted action. In the confined phase there is no immediate distinction to be made between the persistence diagrams generated using the different actions. We claim that this is because vortices in this phase

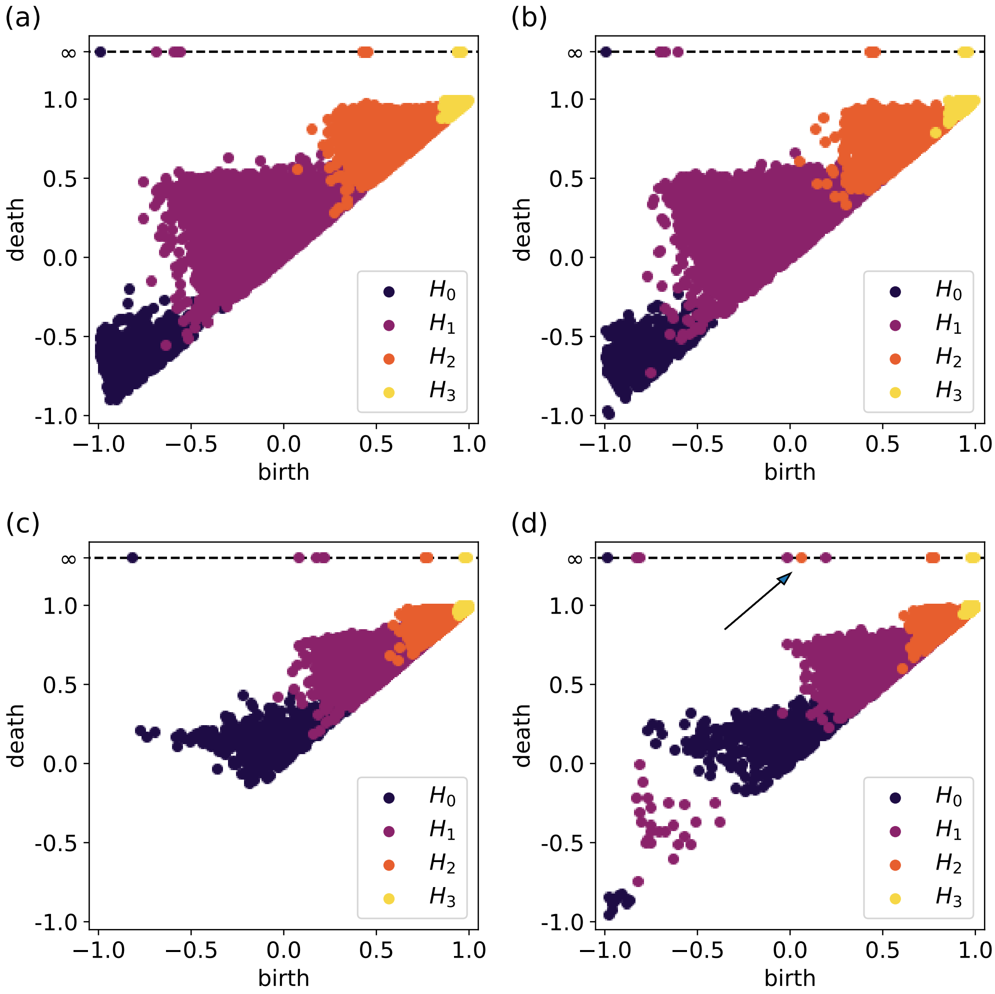

Figure 55: Sample persistence diagrams of individual configurations obtained using the following actions and values of β: (a) Wilson, $\beta = 1.5$ (b) twisted, $\beta = 1.5$ (c) Wilson, $\beta = 2.9$ (d) twisted, $\beta = 2.9$. The arrow in (d) indicates the point $(b, \infty) \in PH_2$ with the smallest birth index b. Note the distance between it and the others.

percolate throughout the system so there are likely to be many vortices that wrap around the periodic boundary conditions of the lattice. Our inserted vortex may then appear as a single point of finite persistence in PH_2, but the persistence of the system is not affected largely. However in the deconfined phase phase there is a clear difference. There is unlikely to be any system-generated vortex surfaces that wrap around the lattice so the inserted vortex becomes the first such surface to enter the filtered complex. We therefore observe that one of the PH_2 points with infinite death time has much lower birth time than the others, allowing us to identify that this point represents our inserted vortex surface. Moreover, we observe a significant change in PH_0 and PH_1 with many low persistence points appearing early on in the filtered complex. These arise as different plaquettes of the inserted vortex enter the filtered complex at different indices, forming transient connected components and holes. Note that these all die by the time we reach the birth index of the point in PH_2.

Following the discussion above, we define an observable based on the persistence diagram of a configuration:

$$m_2 = \min \left\{ b \;\middle|\; (b, \infty) \in PH_2 \right\},$$

the smallest birth index of all points in the H_2 persistence diagram which persist until the end of the filtration (therefore representing generators of the homology of the 4-torus). The expected value of m_2 for different lattice sizes with the Wilson action and twisted action are shown in Figure 56. Note that there is no difference between the expectations estimated using the different actions well into the confined phase (low β), but in the deconfined phase (high β) the curves split apart. As the lattice size increases, the point at which the curves diverge approaches the critical β of the phase transition from below. These observations motivate measuring the difference between the expected values using different actions

$$O_{m_2} = \langle m_2 \rangle - \langle m_2 \rangle_{\text{twist}}$$

as a phase indicator which will be zero in the confined phase and non-zero in the deconfined phase, similar to the definition of an order parameter but without the requirement to detect any symmetry breaking. A finite-size scaling analysis of this quantity yields the curve collapse in Figure 57 with estimates of β_c and ν

$$\beta_c = 2.291 \pm 0.019$$
$$\nu = 0.614 \pm 0.080$$

in agreement with the existing estimate $\beta_c = 2.2986(6)$ in Table 9. Error estimates are obtained by performing 2000 bootstraps. While the error obtained is reasonably large, it should be stressed that these estimates were obtained using only 200 configurations at each value of β and N_s.

Note that we fixed the exponent of -2 for the scaling of O_{m_2} with N_s. Attempting to fit this exponent along with β_c and ν often led to the optimizer returning unrealistic large positive values for the exponent, spoiling the error estimation. The value of -2 was found by hand to give a good fit and we offer a heuristic argument for why. First note that N_s^2 is how the number of plaquettes in the inserted vortex surface scales with N_s, since the surface wraps round the periodic boundary conditions. In the case where it describes the formation of the inserted vortex surface, the value of m_2 is determined by the filtration index of the last plaquette to enter the surface. Now the larger the surface, the more likely it is that there will be at least one plaquette in the surface affected by noise or pierced by another vortex, causing it to enter the filtered complex late and dragging the value of m_2 closer to its average in the Wilson action. Assuming that this likelihood is independent for each plaquette or at least approximately linear in the number of plaquettes, we therefore obtain the quadratic scaling.

Motivated to understand the error bars for the estimates of $\langle m_2 \rangle_{\text{twist}}$ in Figure 56, we look at the distribution of m_2 as measured using the twisted action on a 4×20^3 lattice in Figure 58. We recenter the data using $\langle m_2 \rangle$ in order to compare the effect on O_{m_2} at different values of β.

We see that at lower values of β in the confined phase, the value of m_2 remains close to the average value measured with the untwisted Wilson action. As β increases we observe a bimodal distribution, with some configurations maintaining an m_2 value close to the untwisted average and some joining a lower mode. A likely explanation for this behaviour is that for those configurations in the zero mode, the H_2 generator of the 4-torus responsible for m_2 does not correspond to the inserted vortex in the y-z plane, either because there is another vortex spanning a plane which is more

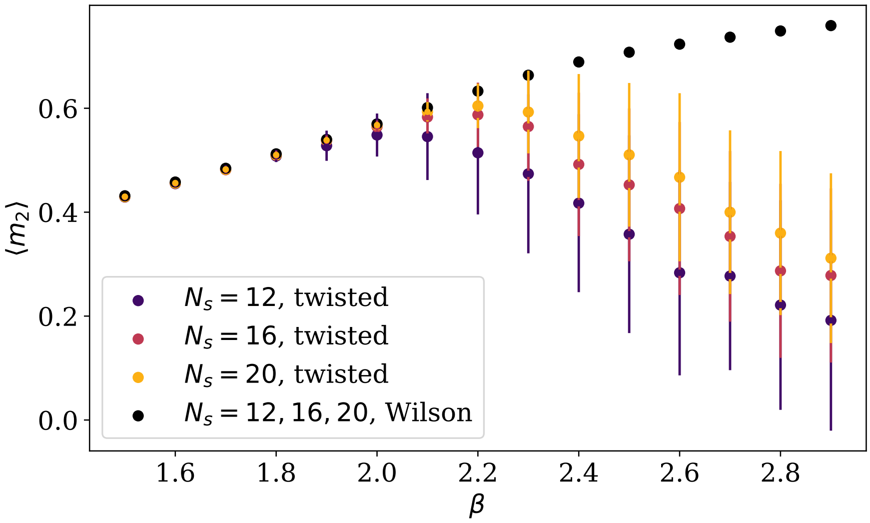

Figure 56: The expected value of the observable m_2 as a function of β plotted for different values of N_s and with the Wilson and twisted actions. The difference between the values for different lattice sizes using the Wilson action is not distinguishable at this scale, so they are plotted as the same points. Their error bars are also small enough so as to not be visible.

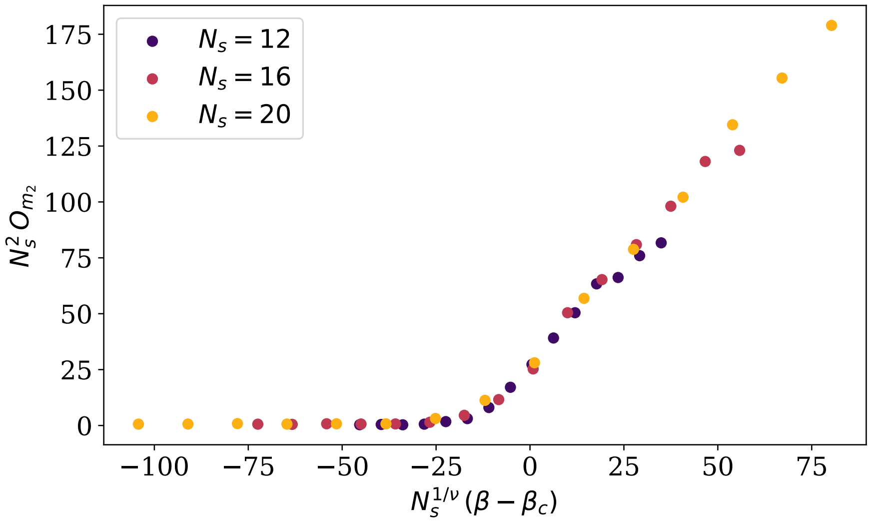

Figure 57: The curve collapse of our phase indicator O_{m_2} using $\beta_c = 2.291$ and $\nu = 0.614$. Error bars are not shown for clarity but are comparable to those in Figure 56.

easily observed or because the formation of a complete vortex surface along the twist in the filtered complex is being impeded, perhaps by intersection with other vortices. By the time we are firmly in the deconfined phase, the majority of configurations lie in the lower mode, signalling that the inserted vortex along the y-z is responsible for the value of m_2.

To verify this picture, we can check if the H_2 generator responsible for m_2 is indeed represented by a yz plane. We do this by recomputing the persistent homology but with the cubical complex only being periodic in the y and z directions and open in the t and x directions. This makes the final complex homeomorphic to the Cartesian product of a 2-torus and $\mathbb{R}^2$, homotopy equivalent to the 2-torus which has a single H_2 generator: the yz plane. Then we can simply check whether or not the birth time of this generator is the same as m_2. Denote by I_{yz} the indicator function for a configuration that is 1 if they match and 0 if they do not. That is, I_{yz} tells us if m_2 is determined by a H_2 generator which spans the yz plane. The mean and variance of this indicator as a function of β are shown in Figure 59.

We see that at low β, the inserted yz vortex sheet is never responsible for the value of m_2. There are likely to be many vortices and those which span planes including the t direction are smaller and likely to be picked up sooner in the persistence. Approaching and passing the transition point, the proportion of configurations for which m_2 describes the birth time of a yz plane increases until close to 1. There are fewer dynamically generated vortices and instead m_2 is determined by the one we inserted via the twisted boundary conditions. We note that peak in the variance of I_{yz} gives us an alternative marker for the phase transition.

4.7 Classification of Vortices

Recall from Section 4.3 that the twisted boundary conditions action S_T depended on some co-closed collection of plaquettes T. For the previous section we fixed T to correspond to a yz plane but we could in principle choose any co-closed collection, yielding inserted thin vortices with any surface we like. Beyond detecting vortices, it seems reasonable to expect that the persistent homology of the filtration we defined will also allow us to distinguish between different surfaces. By choosing different

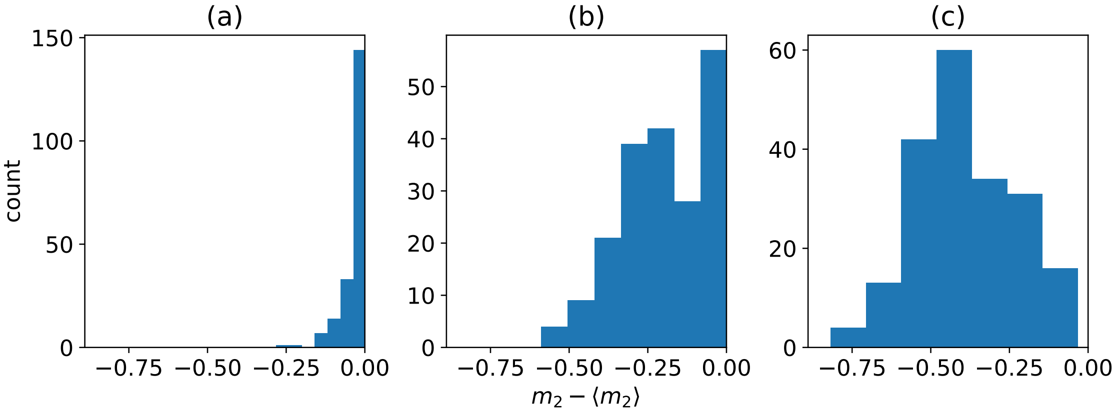

Figure 58: The distribution of $m_2 - \langle m_2 \rangle$ with the twisted action at (a) $\beta = 2.2$, (b) $\beta = 2.5$, and (c) $\beta = 2.8$. Note that $\langle m_2 \rangle$ is the expectation as measured with the Wilson action.

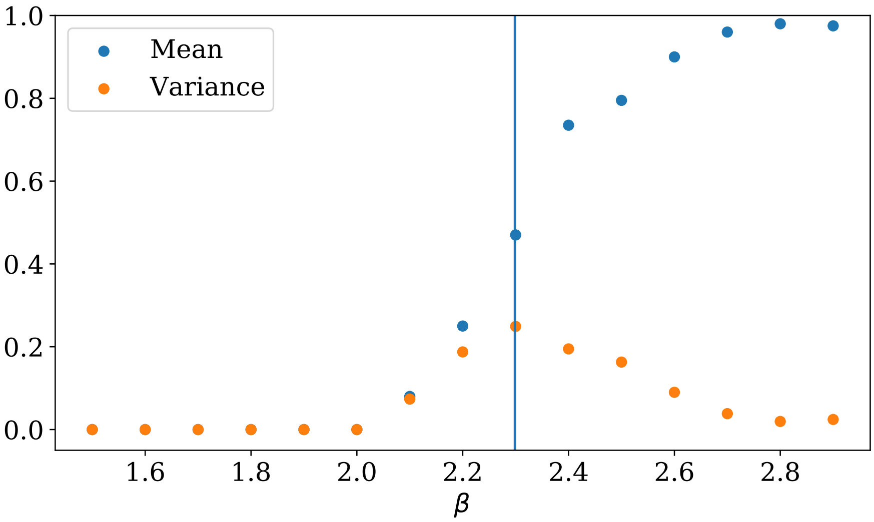

Figure 59: The mean and variance of I_{yz} as a function of β. The vertical line marks the location of β_c.

collections T and sampling using the corresponding action S_T, twisted boundary conditions allow us to test this.

In particular, we consider the following surfaces. Note that the plaquettes given are dual to those making up the surfaces themselves.

Untwisted

$$T = \emptyset,$$

giving the usual Wilson action.

Plane Twist

$$T_{yz} = \{((0,0,y,z),0,1) \mid 0 \leq y, z \leq N_s\},$$

corresponding to a yz plane.

Genus 0 $(2 \times 2 \times 2)$

$$\begin{aligned} T_{g0,2} = \quad & \{((t,x,y,0),2,3) \,|\, 1 \leq t, x \leq 2, y \in \{0,2\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, 1 \leq t, y \leq 2, x \in \{0,2\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, 1 \leq x, y \leq 2, t \in \{0,2\}\}, \end{aligned}$$

corresponding to a $2 \times 2 \times 2$ cube with a corner at $(1,1,1,0)$ as shown in Figure 60.

Genus 0 $(3 \times 3 \times 3)$

$$\begin{aligned} T_{g0,3} = \quad & \{((t,x,y,0),2,3) \,|\, 1 \leq t, x \leq 3, y \in \{0,3\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, 1 \leq t, y \leq 3, x \in \{0,3\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, 1 \leq x, y \leq 3, t \in \{0,3\}\}, \end{aligned}$$

corresponding to a $3 \times 3 \times 3$ cube with a corner at $(1,1,1,0)$ as shown in Figure 61.

2 $\times$ Genus 0 $(2 \times 2 \times 2)$

$$T_{2\times g0,2} = T_{g0,2} \cup \big(T_{g0,2} + (0,8,0,0)\big),$$

corresponding to a pair of $2 \times 2 \times 2$ cubes; one with a corner at $(1,1,1,0)$ and the other with a corner at $(1,9,1,0)$ as shown in Figure 62.

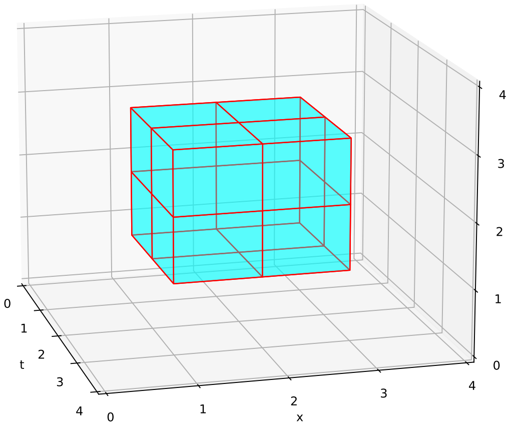

Figure 60: The Genus 0 ($2 \times 2 \times 2$) surface.

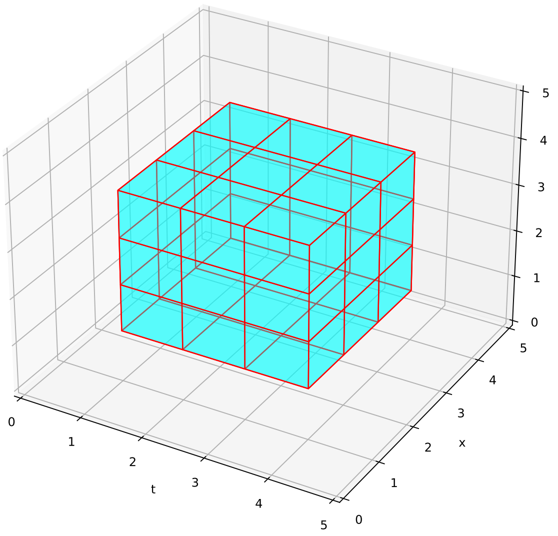

Figure 61: The Genus 0 ($3 \times 3 \times 3$) surface.

Genus 1

$$\begin{aligned} T_{g1} = \quad & \{((1,x,y,0),1,3) \,|\, x \in \{0,3\}, 1 \leq y \leq 3\} \\ & \cup \{((1,x,y,0),2,3) \,|\, 1 \leq x \leq 3, y \in \{0,3\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, t \in \{0,1\}, 1 \leq x,y \leq 3, (x,y) \neq (2,2)\} \\ & \cup \{((1,1,2,0),1,3), ((1,2,2,0),1,3), ((1,2,1,0),2,3), ((1,2,2,0),2,3)\}, \end{aligned}$$

corresponding to a $1 \times 3 \times 3$ torus with a corner at $(1,1,1,0)$ as shown in Figure 63.

Genus 2

$$\begin{aligned} T_{g2} = \quad & \{((1,x,y,0),1,3) \,|\, x \in \{0,5\}, 1 \leq y \leq 3\} \\ & \cup \{((1,x,y,0),2,3) \,|\, 1 \leq x \leq 3, y \in \{0,3\}\} \\ & \cup \{((t,x,y,0),2,3) \,|\, t \in \{0,1\}, 1 \leq x \leq 5, 1 \leq y \leq 3, (x,y) \neq (2,2), (x,y) \neq (4,2)\} \\ & \cup \{((1,1,2,0),1,3), ((1,2,2,0),1,3), ((1,2,1,0),2,3), ((1,2,2,0),2,3)\} \\ & \cup \{((1,3,2,0),1,3), ((1,5,2,0),1,3), ((1,5,1,0),2,3), ((1,5,2,0),2,3)\}, \end{aligned}$$

corresponding to a $1 \times 5 \times 3$ double torus with a corner at $(1,1,1,0)$ as shown in Figure 64.

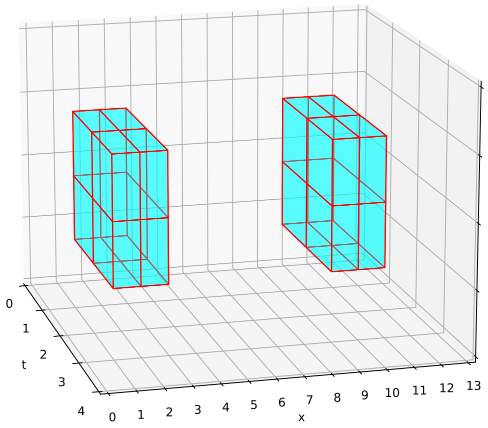

Figure 62: The 2 × Genus 0 (2 × 2 × 2) surface.

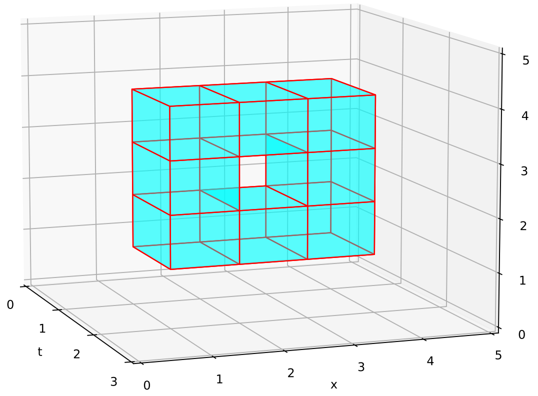

Figure 63: The Genus 1 surface.

Let us fix the lattice to be $4 \times 16 \times 16 \times 16$ and $\beta = 2.31$, just into the deconfined phase. We apply the following procedure.

1. For each of the choices of T, we sample 100 configurations using the action S_T using the HiRep software [145] with 1 heatbath step and 4 overelaxation steps for each Monte Carlo step and a sample taken every 100 Monte Carlo steps.
2. For each sample we compute their PH_0, PH_1, PH_2, PH_3 persistent homology using our filtered complex and compute the corresponding persistence landscapes with 1 layer and a resolution of 100.
3. We train a kNN model ($k = 10$) to predict which action each configuration was generated using.

A 4-fold cross validation produces the confusion matrix shown in Figure 65.

In this representation of the classification accuracy, we want the diagonal to have high values and entries off the diagonal to be small, signalling that the predicted label matches the actual type of the vortex. The diagonal structure we see in Figure 65 therefore tells us that the classification was reasonably successful. However there

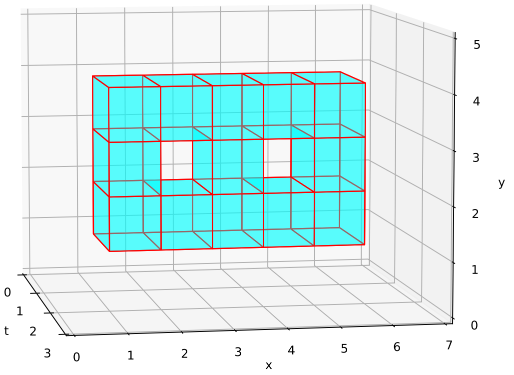

Figure 64: The Genus 2 surface.

is some confusion between certain pairs of vortex types as signaled by the block diagonal structure. In particular, there is confusion between the $2 \times 2 \times 2$ cube vortex surface and the pair of such surfaces, and there is confusion between the genus 1 and genus 2 vortex surfaces. We might hypothesise that this is due to the vortices not entering the filtered complex fully due to noise and the effect of dynamically generated vortices, leaving the surfaces not closed and therefore not appearing in PH_2. Note that these simulations were performed at $\beta = 2.31$, only just into the deconfined phase. Instead, the persistence landscapes may be picking up on the transient 0 and 1-dimensional cycles that appear during the early parts of the filtration. These signatures may look similar in each of the pairs of cases mentioned. One could investigate this hypothesis by computing the classification with only the 0 and 1-dimensional landscapes or by using a generalised linear classifier (such as logistic regression) and analysing the learned weights, but this is future work.

4.8 Stability of Vortices

Having a quantitative way to detect center vortices also allows us to analyse their stability. In the presence of a Higgs field, vortices are minimal action configurations [99] and hence stable under the Wilson flow. However in pure Yang-Mills we are not

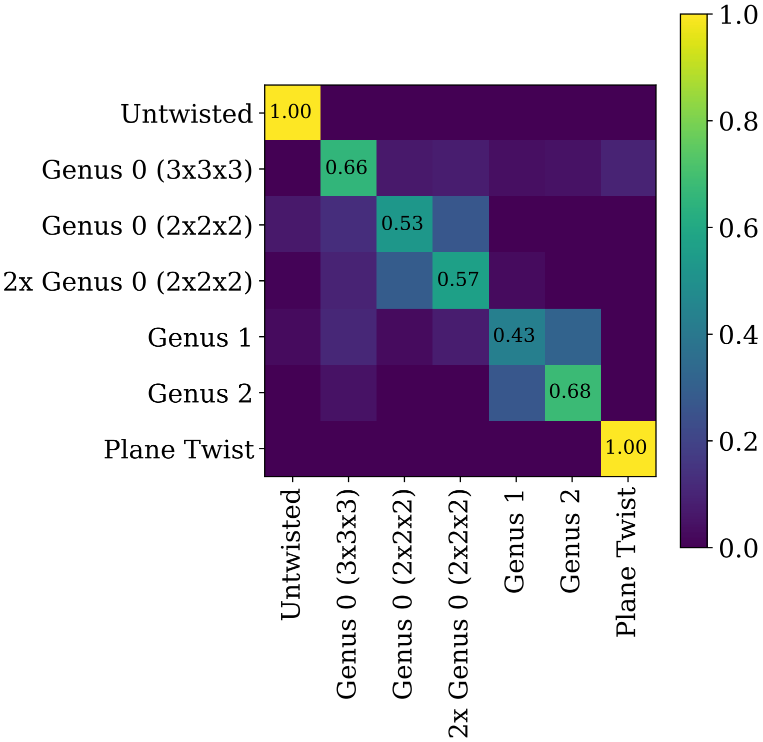

Figure 65: The confusion matrix obtained from cross validating the ability of a kNN classifier to distinguish inserted vortex surfaces from persistence landscapes. The x axis shows the predicted label and the y axis shows the true label.

certain about the stability of vortices under the Wilson flow: i.e., we can ask whether or not the flow will eventually smooth them out. This may also depend on several factors including whether or not we have periodic boundary conditions (in this thesis we do) and whether the vortex surfaces wrap round the periodic boundary. In this section we consider the question of when a vortex remains detectable in persistent homology and use this to identify at which point in the Wilson flow we lose it.

Let us fix the lattice to be $4 \times 16 \times 16 \times 16$ and $\beta = 2.31$, just into the deconfined phase. For three of the twisted actions (Genus 1, Genus 2, and 2 × Genus 0), we apply Wilson flow to the 100 configurations generated in the previous section, saving the flowed configurations at Wilson flow times $t = 0, 0.075, 0.15, \ldots, 1.5$. For each of the configurations we compute their persistent homology using our filtration. We use two different methods to quantify the ability to detect the inserted vortex.

Difference in persistence Since we are working only just above the phase transition and none of the inserted vortex surfaces considered wrap the periodic boundaries of the lattice, assume that the inserted vortex will be responsible for the most persistent finite-death point in PH_2. A naïve approach to quantifying the ability to detect the vortex is to measure the difference in the average of the persistence p_1 of the most persistent such point and the persistence p_2 of the second most persistent. The plot of this quantity as a function of Wilson flow time for each of the vortex surfaces considered is shown in Figure 66.

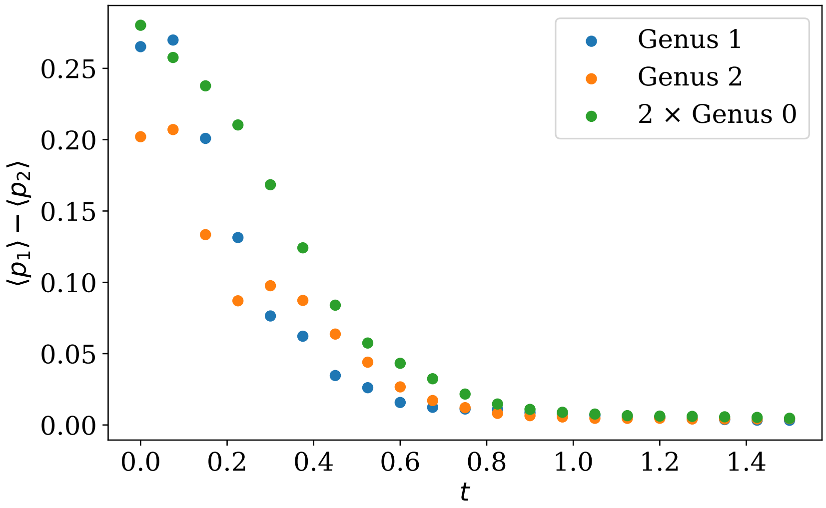

Figure 66: Vortex stability via difference in persistence.

This is by no means a precise measurement, but we do at least observe by eye that the difference becomes negligible close to $t = 1$, approximately around the characteristic flow time of $t_0 = 0.99$ obtained in Section 4.5.2. This would suggest that the vortices are stable for times that correspond to the typical hadronic scale. That is, this analysis does not rule out that vortices may play some role in the dynamics we are interested in.

We can apply a similar analysis for the Plane Twist twisted boundary conditions but where we consider the birth time b_1 of the earliest infinite H_2 point and the birth time b_2 of the second earliest. In particular, note that $b_1 = m_2$ as defined in Section 4.6. Once again we apply Wilson flow to 100 configurations sampled at $\beta = 2.31$, saving the flowed configurations at Wilson flow times $t = 0, 0.075, 0.15, \ldots, 1.5$. A plot of the difference between the average of b_2 and the average of b_1 as a function of t is shown in Figure 67. Note that this difference is effectively the same as for the difference in persistence $p_1 - p_2$ for the finite-death points, but we have avoided trying to make sense of the "equation"

$$p_1 - p_2 = (\infty - b_1) - (\infty - b_2) = b_2 - b_1.$$

The curve we obtain is still therefore to be compared with those in Figure 66 where we see that the Plane Twist vortex is more stable, with the difference only becoming negligible closer to $t = 1.5$.

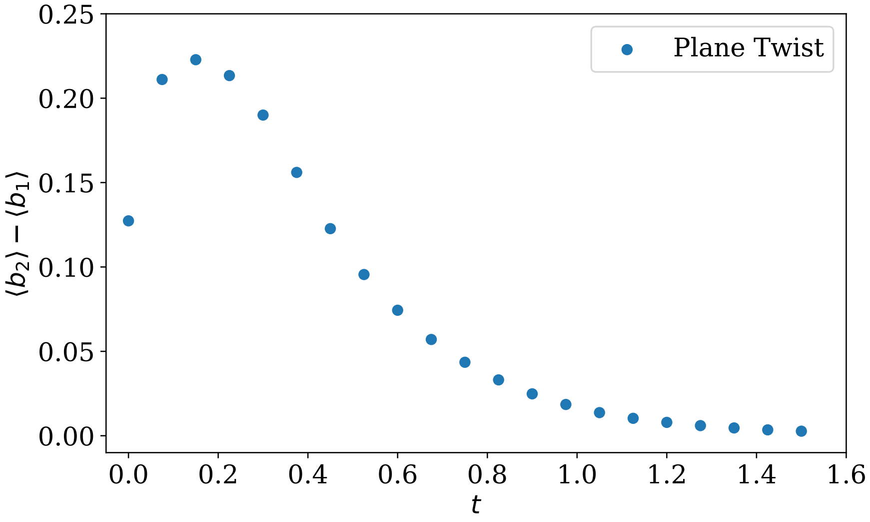

Figure 67: Vortex stability via birth times of infinite H_2 generators.

Overlap via persistence images An alternative approach is inspired by looking at Figure 68. Here we see that the most persistent finite-death point in PH_2 (assumed to be the inserted vortex) is clearly distinct from the other points in PH_2. During Wilson flow, the distance between this point and the others shrinks until it is no longer distinct. The idea is to use persistence images to estimate the density of points in the persistence diagram and quantify when there is significant overlap of the point with the rest. In particular, for each configuration we compute the PH_2 persistence image PI_2 with a resolution of 200×200, σ equal to 2.5% of a pixel and a constant weighting function $f(b, p) = 1$. Note that with this weighting function we are simply approximating the density of points in the persistence diagrams, but

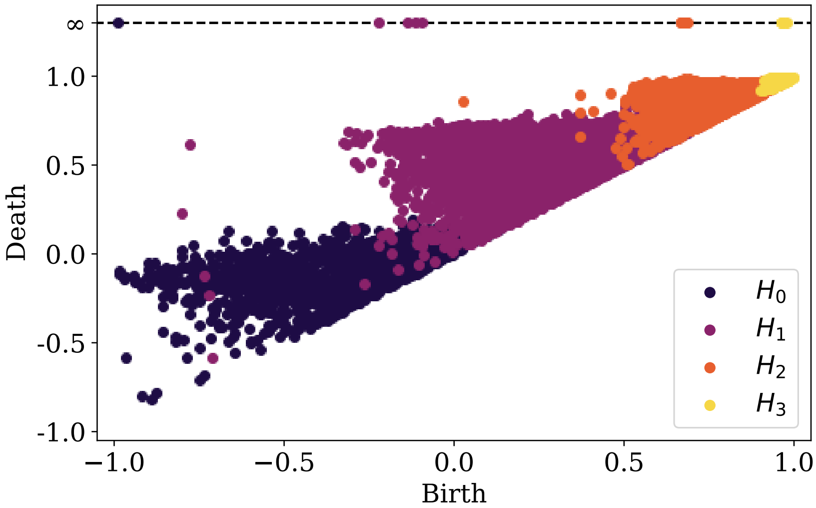

Figure 68: A typical persistence diagram obtained using the Genus 1 twisted boundary conditions. Note the isolated point in PH_2 with maximum persistence among the finite-death points.

lose the stability property of the persistence image since f does not go to zero on the $p = 0$ axis. Define

$$V = PI_2^i \, / \, \langle \sum_j PI_2^j \rangle$$

where i is the pixel containing the most persistent finite-death point in PH_2. The deviation of V from zero gives us an indication of when the vortex point becomes indistinguishable from the background. A plot of the average of V as a function of Wilson flow time for the different vortex surfaces is shown in Figure 69.

This suggests overlap coming in at closer to Wilson flow time $t = 0.6$ which, while less than the characteristic flow time of $t_0 = 0.99$ obtained in Section 4.5.2, is on the same order of magnitude and so also leads us to conclude that vortices are stable long enough to play a role in hadronic physics and they remain visible to the persistent homology of our filtered complex across this scale.

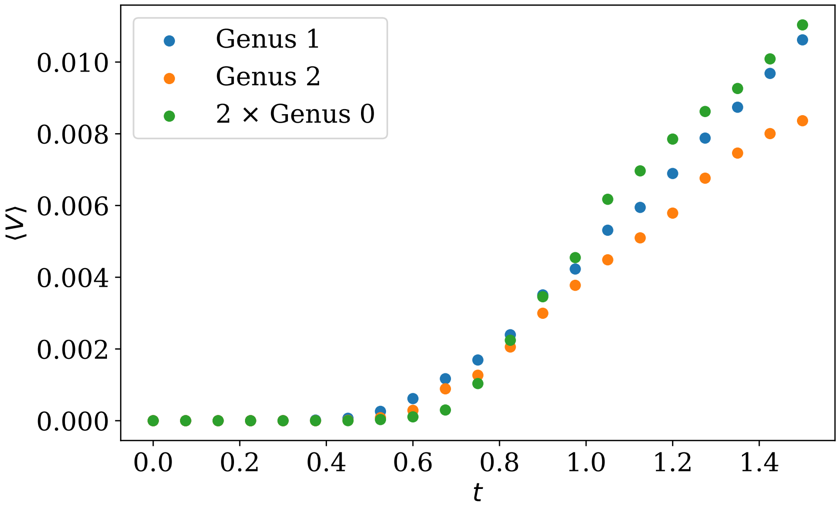

Figure 69: Vortex stability via persistence images.

4.9 Quantitative Analysis of the Deconfinement Transition

We now investigate whether the persistent homology is able to detect the deconfinement phase transition purely from the Wilson action by making use of a simple machine learning framework inspired by that in [1]. In particular, for $N_t \in \{4, 5, 6\}$ we attempt to estimate the critical inverse coupling β_c and the critical exponent of the correlation length ν via a finite-size scaling analysis of the output of a k-nearest neighbours classifier trained on the persistence. In each case we repeat the following procedure for each $N_s \in \{12, 16, 20, 24\}$.

1. Configurations are sampled from a range of values of β (specific values are given in each case below) using the HiRep software [145] with 1 heatbath step and 4 overelaxation steps for each Monte Carlo step and a sample taken every 100 Monte Carlo steps.

2. For each sample we compute their PH_0, PH_1, PH_2 and PH_3 persistent homology using our filtered complex and compute the corresponding persistence images with a resolution of 25×25 and and σ equal to 5% of a pixel. We concatenate the 4 separate images and flatten them into a $4 \times 25 \times 25 = 2500$ dimensional vector.

3. We train a kNN model ($k = 30$) to predict the phase of a configuration based on its concatenated persistence image vector by using vectors from well into the confined and deconfined phases.

4. Using the trained classification model, we define an observable $O_{k\mathrm{NN}}$ which is the predicted phase of a given configuration.

5. Applying multiple histogram reweighting to the variance

$$\begin{aligned}\chi_{k\mathrm{NN}} &= \langle O_{k\mathrm{NN}}^2 \rangle - \langle O_{k\mathrm{NN}} \rangle^2 \\ &= \langle O_{k\mathrm{NN}} \rangle \left(1 - \langle O_{k\mathrm{NN}} \rangle\right)\end{aligned}$$

   (where the second equality follows since $O_{k\mathrm{NN}} \in \{0, 1\}$), we obtain an interpolated curve and a more precise estimate of the location of its peak.

By performing a finite-size scaling analysis of the locations of the peaks obtained for each value of N_s, and a curve collapse of the different reweighted variance curves, we obtain estimates of β_c and ν for the deconfinement phase transition at the given value of N_t.

4.9.1 $N_t = 4$

For lattices of size $4 \times N_s^3$ with $N_s \in \{12, 16, 20, 24\}$, we train a k-nearest neighbours classifier ($k = 30$) on the concatenated PH_0, PH_1, PH_2 and PH_3 persistence images of 200 configurations sampled at each β in the confined and deconfined regions given in Table 10. The classifier is then used to produce a predicted classification $O_{k\mathrm{NN}}$ for 200 configurations sampled from the critical region.

Region	β
Confined	2.2 , 2.21, 2.22, 2.23, 2.24
Deconfined	2.36, 2.37, 2.38, 2.39, 2.4
Critical	2.25, 2.26, 2.27, 2.275, 2.28, 2.285, 2.29, 2.295, 2.298, 2.299, 2.3, 2.301, 2.302, 2.305, 2.31, 2.315, 2.32, 2.325, 2.33, 2.34, 2.35

Table 10: Values of β sampled at for the $N_t = 4$ phase transition.

The resulting estimates of the expectation $\langle O_{k\mathrm{NN}}\rangle(\beta)$ are shown in Figure 70 along with interpolating curves obtained via histogram reweighting. The variance curves $\chi_{k\mathrm{NN}}$ are shown in Figure 71.

From here we proceed with two separate analyses. In the first we estimate β_c from a linear regression assuming the value of the critical exponent ν to be known. In the second we estimate β_c and ν concurrently via a curve collapse procedure. In both cases we perform two separate bootstraps to obtain 500 bootstrap samples from each. One bootstrap is carried out by resampling the configurations for each β used to train the k-nearest neighbours classifier. The other is carried out by resampling the configurations for each β used to estimate $\langle O_{k\mathrm{NN}}\rangle$. Applying the finite-size scaling analysis to both collections of bootstraps yields two separate distributions for β_c and two for ν. The error in these quantities is therefore estimated by combining the standard deviation of the distributions coming from the different bootstrap procedures under the assumption that they are independent.

By defining the pseudo-critical inverse coupling $\beta_c(N_s)$ to be the point at which $\chi_{k\mathrm{NN}}$ peaks, we can plot $\beta_c(N_s)$ against $N_s^{-1/\nu}$ using a previously estimated [141] value of $\nu = 0.629971$. The result is shown in Figure 72. We observe that the points plotted with error bars of 1σ support a straight line fit. The intercept yields $\beta_c = 2.2989 \pm 0.0009$, supporting the previously obtained estimate of $\beta_c = 2.2986(6)$ in Table 9.

To estimate β_c and ν concurrently we employ a numerical curve collapse procedure, plotting $\chi_{k\mathrm{NN}}$ against $N_s^{1/\nu}(\beta - \beta_c)$ and tuning β_c and ν to minimise the distance between the curves using the conjugate gradient method.

The resulting curve collapse is shown in Figure 73 and the obtained estimates of β_c and ν

$$\beta_c = 2.2988 \pm 0.0007$$
$$\nu = 0.634 \pm 0.014$$

are consistent with previous estimates.

To confidently claim that this methodology identifies the phase transition, we also tried using alternative values of β to train the kNN classifier, chosen further away from the transition point and so that the transition point is further from the

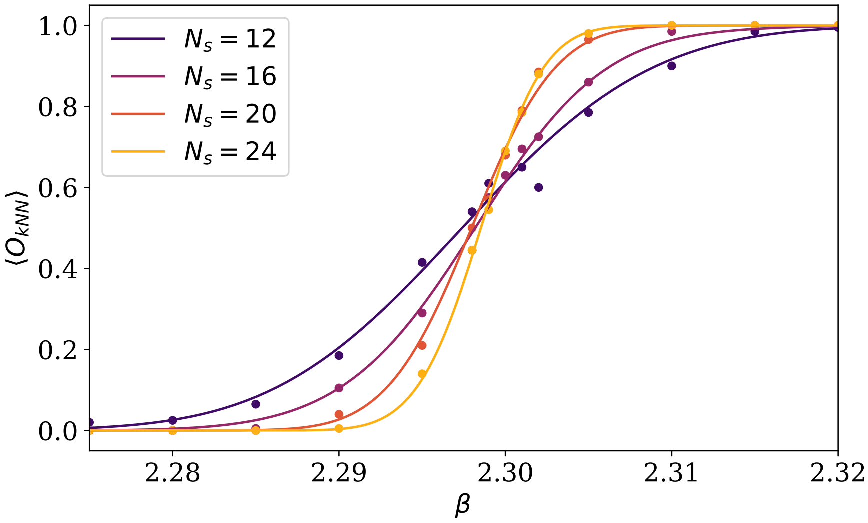

Figure 70: Plot showing our phase indicator $\langle O_{k\mathrm{NN}} \rangle$ as a function of β for $N_t = 4$. The points show the measured expectations and the curve is the output of histogram reweighting these measurements.

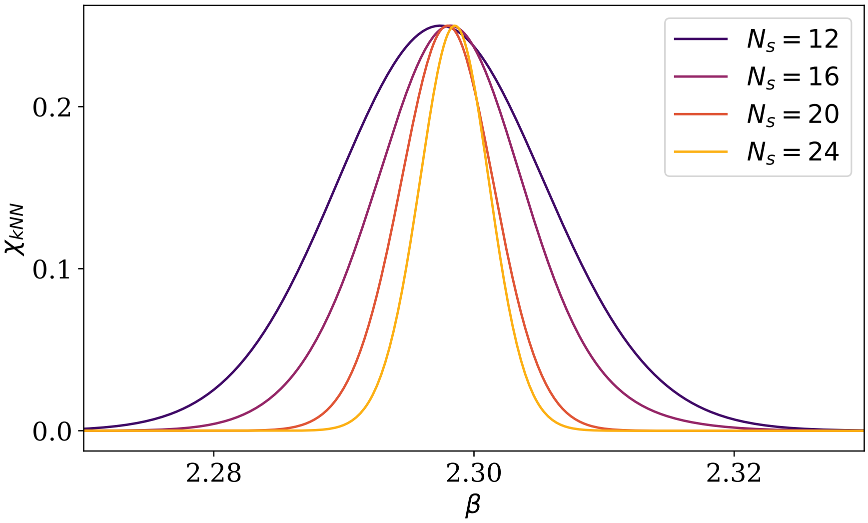

Figure 71: The variance curves $\chi_{k\mathrm{NN}}$ of $O_{k\mathrm{NN}}$ for $N_t = 4$ to which we will apply our curve collapse procedure.

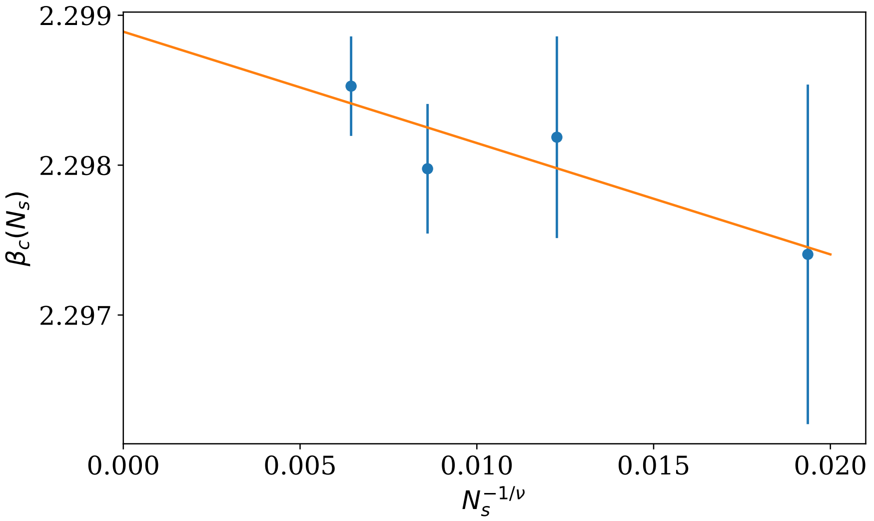

Figure 72: Estimating β_c for $N_t = 4$. The pseudo-critical values of β, obtained from locating the peaks of the variance curves in Figure 71, are fitted to the ansatz (5). Error bars are estimated by bootstrapping.

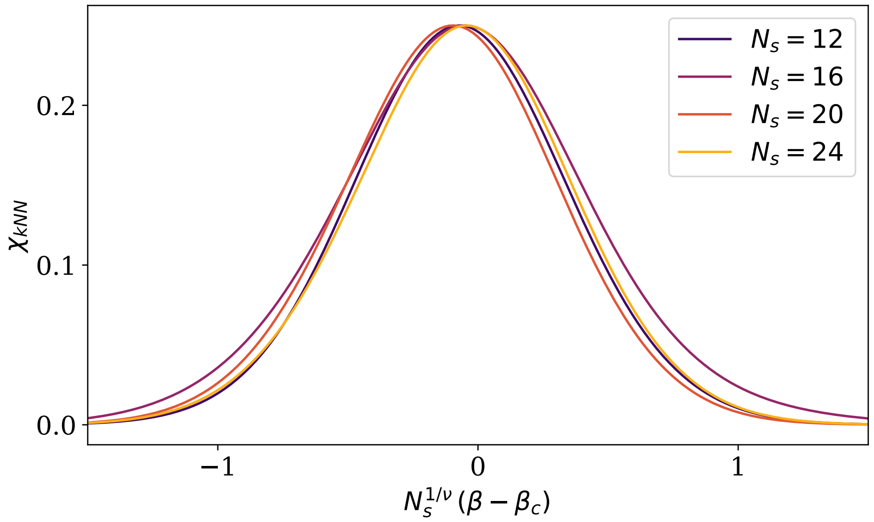

Figure 73: The curve collapse of $\chi_{k\text{NN}}$ for $N_t = 4$ using $\beta_c = 2.2988$ and $\nu = 0.635$.

center point between the highest β in the confined phase and the lowest β in the deconfined phase. The alternative training values are shown in Table 11.

Region	β
Confined	2.21, 2.22, 2.23
Deconfined	2.38, 2.39, 2.4
Critical	2.24, 2.25, 2.26, 2.27, 2.275, 2.28, 2.285, 2.29, 2.295, 2.298, 2.299, 2.3, 2.301, 2.302, 2.305, 2.31, 2.315, 2.32, 2.325, 2.33, 2.34, 2.35, 2.36, 2.37

Table 11: Alternative values of β sampled at for the $N_t = 4$ phase transition to test the sensitivity of the method to the choice of training data.

Using these training values we obtain estimates from the linear fit of

$$\beta_c = 2.2996 \pm 0.0009$$

and the curve collapse of

$$\begin{aligned}\beta_c &= 2.2998 \pm 0.0007\\ \nu &= 0.638 \pm 0.013,\end{aligned}$$

close to our previous ones and still compatible with our reference estimates.

4.9.2 $N_t = 5$

For lattices of size $5 \times N_s^3$ with $N_s \in \{12, 16, 20, 24\}$, we train a k-nearest neighbours classifier ($k = 30$) on the concatenated PH_0, PH_1, PH_2 and PH_3 persistence images of 200 configurations sampled at each β in the confined and deconfined regions given in Table 12. The classifier is then used to produce a predicted classification $O_{k\mathrm{NN}}$ for 200 configurations sampled from the critical region. The resulting estimates of the expectation $\langle O_{k\mathrm{NN}}\rangle(\beta)$ are shown in Figure 74 along with interpolating curves obtained via histogram reweighting.

The plot of the pseudo-critical $\beta_c(N_s)$ against $N_s^{-1/\nu}$ using $\nu = 0.629971$ is shown in Figure 75. Here we fit a straight line to the largest three lattice sizes since this

Region	β
Confined	2.29, 2.3, 2.31, 2.32, 2.33
Deconfined	2.41, 2.42, 2.43, 2.44, 2.45
Critical	2.34, 2.345, 2.35, 2.355, 2.36, 2.365, 2.369, 2.37, 2.371, 2.372, 2.375, 2.38, 2.385, 2.39, 2.395, 2.4

Table 12: Values of β sampled at for the $N_t = 5$ phase transition.

gives a better fit than including $N_s = 12$. The intercept yields $\beta_c = 2.3696 \pm 0.0012$ which is just about compatible with the previously obtained estimate of $\beta_c = 2.37136(54)$ in Table 9.

We also perform the curve collapse on only the highest three lattice sizes and the result is shown in Figure 76. The obtained estimates of β_c and ν

$$\beta_c = 2.3697 \pm 0.0011$$
$$\nu = 0.634 \pm 0.028$$

are consistent with previous estimates.

4.9.3 $N_t = 6$

For lattices of size $6 \times N_s^3$ with $N_s \in \{12, 16, 20, 24\}$, we train a k-nearest neighbours classifier ($k = 30$) on the concatenated PH_0, PH_1, PH_2 and PH_3 persistence images of 200 configurations sampled at each β in the confined and deconfined regions given in Table 13. The classifier is then used to produce a predicted classification $O_{k\mathrm{NN}}$ for 200 configurations sampled from the critical region.

The resulting estimates of the expectation $\langle O_{k\mathrm{NN}} \rangle(\beta)$ are shown in Figure 77 along with interpolating curves obtained via histogram reweighting.

The plot of the pseudo-critical $\beta_c(N_s)$ against $N_s^{-1/\nu}$ using $\nu = 0.629971$ is shown in Figure 78. The intercept of the straight line fit yields $\beta_c = 2.4277 \pm 0.0008$, supporting the previously obtained estimate of $\beta_c = 2.4271(17)$ in Table 9.

The result of the curve collapse is shown in Figure 79. The obtained estimates

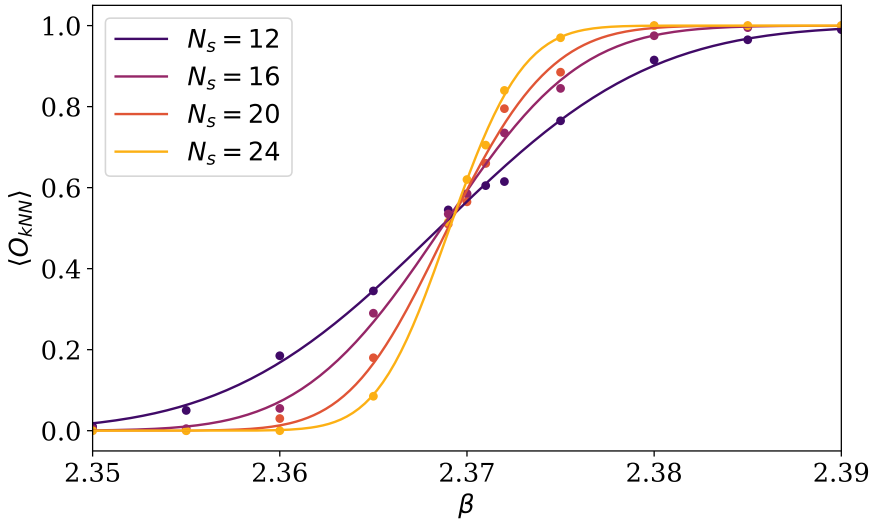

Figure 74: Plot showing our phase indicator $\langle O_{k\mathrm{NN}} \rangle$ as a function of β for $N_t = 5$. The points show the measured expectations and the curve is the output of histogram reweighting these measurements.

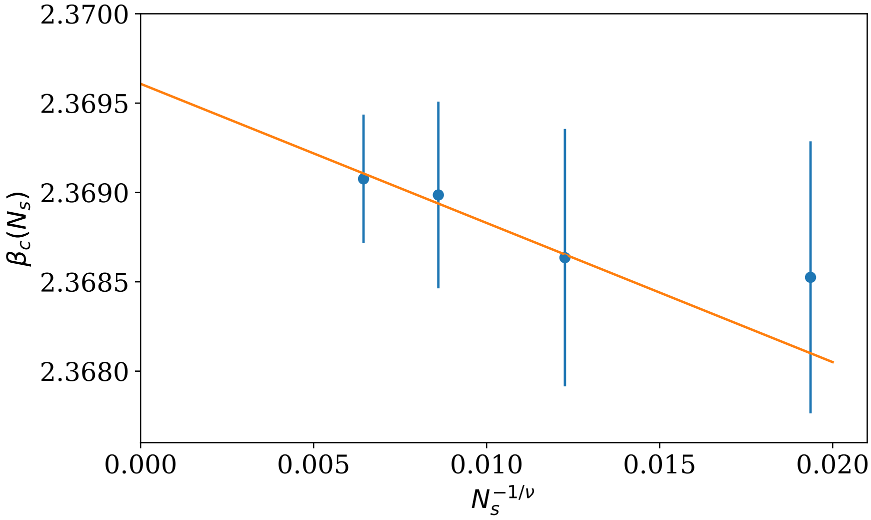

Figure 75: Estimating β_c for $N_t = 5$. The pseudo-critical values of β of the largest three lattice sizes, obtained from locating the peaks of the variance curves, are fitted to the ansatz (5). Error bars are estimated by bootstrapping.

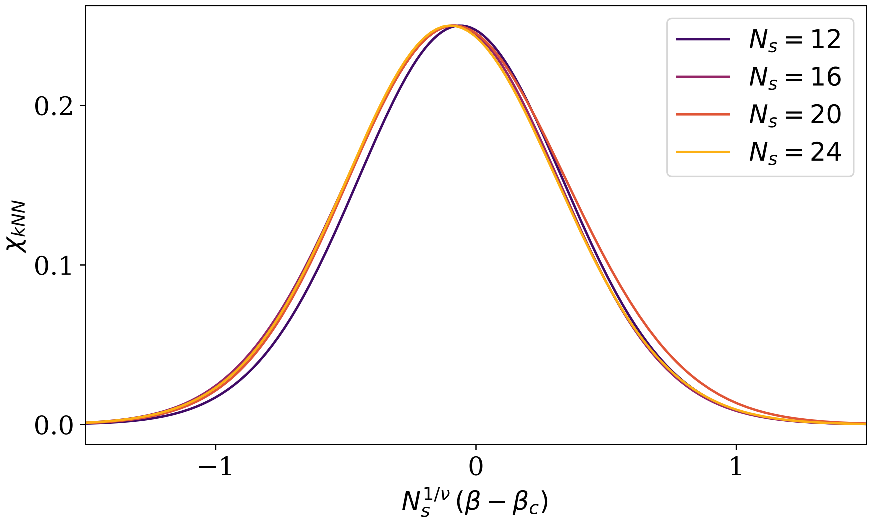

Figure 76: The curve collapse of $\chi_{k\mathrm{NN}}$ for $N_t = 5$ using $\beta_c = 2.3697$ and $\nu = 0.634$.

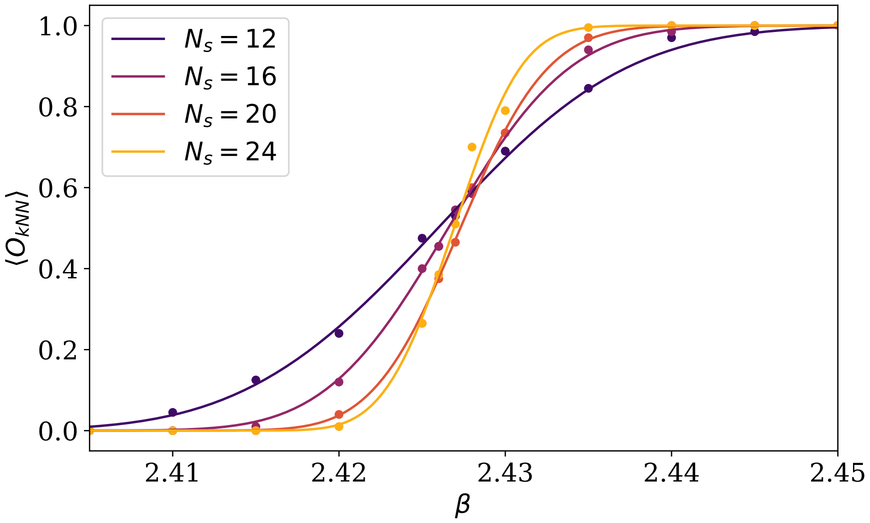

Figure 77: Plot showing our phase indicator $\langle O_{k\mathrm{NN}} \rangle$ as a function of β for $N_t = 6$. The points show the measured expectations and the curve is the output of histogram reweighting these measurements.

Region	β
Confined	2.33, 2.34, 2.35, 2.36, 2.37
Deconfined	2.49, 2.5, 2.51, 2.52, 2.53
Critical	2.38, 2.39, 2.4, 2.405, 2.41, 2.415, 2.42, 2.425, 2.426, 2.427, 2.428, 2.43, 2.435, 2.44, 2.445, 2.45, 2.455, 2.46, 2.47, 2.48

Table 13: Values of β sampled at for the $N_t = 6$ phase transition.

of β_c and ν are

$$\beta_c = 2.4276 \pm 0.0008$$
$$\nu = 0.666 \pm 0.016.$$

The estimate of β_c agrees with the previous estimate, however the previous estimate of $\nu = 0.629971$ we refer to lies just over 2 standard deviations outside of our estimate.

4.10 Discussion

In this chapter we have developed a method to use persistent homology to detect center vortices in configurations of SU(2) lattice gauge theory. We demonstrated that the resulting persistence diagrams are sensitive to vortices in the deconfined phase and investigated some applications of our vortex detection method. Quantitatively, we defined two different persistence-based phase indicators for the deconfinement phase transition — one making use of twisted boundary conditions, and the other using simple machine learning — and successfully used them to estimate the critical β and exponent ν of the transition.

This method was designed to detect and capture the behaviour of vortices, and the efficacy of this was confirmed by being able to detect the use of twisted boundary conditions. Moreover the method was still able to detect the deconfinement phase transition from samples using the original action alone. We argue that this lends support to the center vortex picture of confinement, at least by not ruling it out. For a stronger argument we would need to investigate the relationship between our

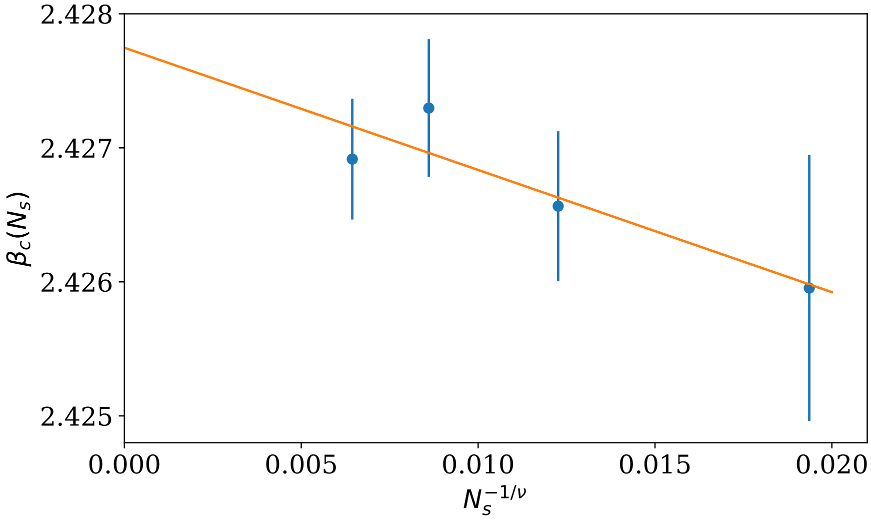

Figure 78: Estimating β_c for $N_t = 6$. The pseudo-critical values of β, obtained from locating the peaks of the variance curves, are fitted to the ansatz (5). Error bars are estimated by bootstrapping.

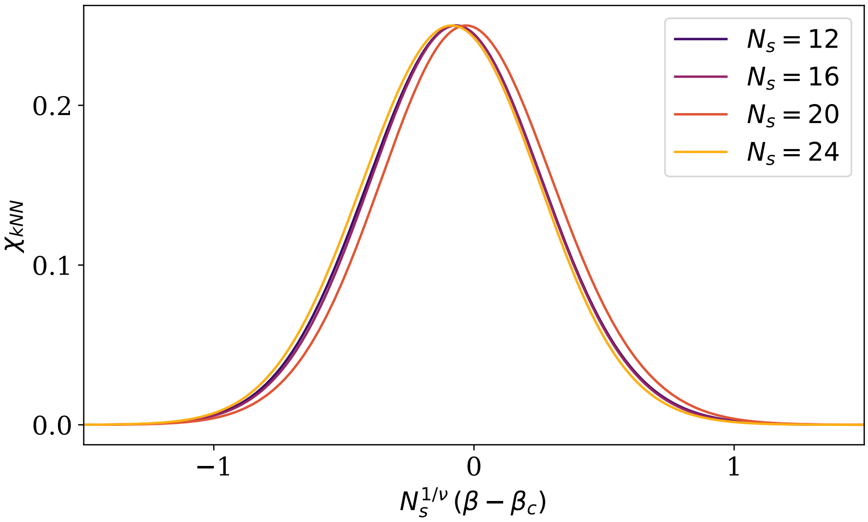

Figure 79: The curve collapse of $\chi_{k\mathrm{NN}}$ for $N_t = 6$ using $\beta_c = 2.4276$ and $\nu = 0.666$.

method and other competing pictures. For example, looking at the monopole picture of confinement (see [152] or [94] for a summary), one could investigate the sensitivity of our method to monopoles. If it were to turn out that the method is sensitive to vortices but not to monopoles, the fact that it captures the phase transition would be evidence for the vortex picture over the monopole picture. However, we are not sure yet whether or not the method is sensitive to monopole configurations and this is work yet to be done, which will require non-trivial adaptations of some of the steps used here. For instance, if the current method is indeed sensitive to monopoles, then one might instead attempt to devise a filtration that exposes monopole-like singularities but not vortices.

Besides a greater degree of interpretability, another advantage of this method over machine learning approaches based on deep learning is that we were able to obtain our results using only a small number of sampled configurations. This is particularly important in view of extending the methodology to full QCD, for which numerical computations for generating gauge configurations near the physical point are very demanding and hence the number of configurations one can use is generally limited.

5 Conclusions

The work described in this thesis lays the groundwork for the use of topological data analysis in producing insights into models from statistical physics and quantum field theory. Beyond using persistent homology to detect topological defects, we have paid close attention to developing methodologies for using it to investigate the quantitative aspects of phase transitions. The idea has been to provide interpretation of the mechanism of the phase transitions by being able to relate them to the behaviour of the topological defects involved.

In particular, we have made the following contributions:

- We introduced a new class of filtered complexes for lattice spin systems, including the angle difference and nematic angle difference filtrations, which allow persistent homology to detect topological defects including point vortices, half vortices, and domain walls.
- We introduced a gauge-invariant construction that produces a filtered complex from a field configuration for SU(2) lattice gauge theory. We found evidence that this detects center vortices by showing that it distinguishes configurations generated using twisted boundary conditions in the deconfined phase. In particular, we show that a phase indicator can be recovered by comparing the H_2 persistence diagrams of configurations generated using twisted boundary conditions and configurations sampled using the usual Wilson action.
- Finding inadequacies with using logistic regression for estimating the critical temperature of XY models, as previously introduced in [45], we introduce a non-parametric method using k-nearest neighbour classification as a tool to estimate the critical point of phase transitions from persistence images. We demonstrated that this yields improved results in the XY model case.
- We demonstrated the applicability of finite-size scaling analysis for persistence-based observables, including the output of k-nearest neighbour classifiers trained on persistence images. In particular, we apply the standard statistical tools of histogram reweighting and bootstrapping to obtain estimates of the critical

point and the critical exponent of the correlation length with quantified error in models from statistical physics and quantum field theory.

- By finding representative cocycles for persistent H_1 in the XY model, we directly investigated the unbinding of vortex anti-vortex pairs which drive the BKT phase transition.
- Using the structure of persistence diagrams, we investigated the ability to detect center vortices after Wilson flow as a proxy for vortex stability, concluding that they are stable enough to be involved in the dynamics of interest in SU(2) lattice gauge theory.
- We demonstrated that the persistent homology of our filtered complex construction for SU(2) lattice gauge theory is sensitive to the topology of center vortex surfaces by using persistence landscapes to classify different twisted boundary conditions.

5.1 Future Work

One of the driving motivations of this work is QCD and the approach to analysing the deconfinement transition in lattice QCD. In order to reach this goal, one would first need to extend the approach of Chapter 4 to the SU(3) lattice gauge theory. The center vortex picture here is similar but now, since the center is $Z(\mathrm{SU}(3)) \cong \mathbb{Z}_3$, there are two non-trivial types of center vortex which may interact. This would require the construction of a new filtered complex. Moreover, the deconfinement phase transition in the SU(3) gauge theory is a first-order (or discontinuous) phase transition. It is not *a priori* clear that the quantitative methodology developed in this thesis would be immediately applicable.

Once a methodology for the SU(3) lattice gauge theory is developed, the step to lattice QCD should be straightforward. Indeed lattice QCD can be seen as an extension of the gauge theory where we now consider fermion fields on the lattice sites and additional interaction terms in the action. However, since the center vortex picture concerns only the gauge fields, the gluons, one may immediately apply the SU(3) methodology.

Beyond gauge theories, it would not be unreasonable to expect persistent homology-based analysis to be useful for other classes of systems where topology plays a role such as in certain quantum models. There is already some work in this direction, such as [78, 79, 153]. Moving away from phase transitions, there is potential use for topological data analysis in studying non-equilibrium systems as in [154]. In particular, there is scope to apply the theory of persistence vineyards [43]. The idea here is that for a stable filtered complex, as the underlying data changes continuously over time, so too does the resulting persistence diagram. Points in the diagram therefore trace paths, also known as vines, which allow one to relate the persistence diagrams obtained at different times.

Beyond the extension of the *persistent homology as an observable* approach considered in this thesis, there is potentially interesting work to be done on the *persistent homology of configuration space.* One of the principal drawbacks discussed in [76] is the computational cost of computing persistent homology in degrees higher than 1 for the Vietoris-Rips complexes of large pointclouds. One idea is to more directly focus on finding critical points and Morse indices in a local neighbourhood, as in [155], rather than trying to characterise the topology of the whole sublevel set. It may also be possible to directly compute persistent homology in several degrees higher than [76] by leveraging techniques to parallelise the computation of persistence, as in [156, 157]. Alternatively, a newly emerging idea is quantum computation of not only Betti numbers [158], but also persistent Betti numbers. Such technology will not be available in the immediate future, but may be transformative for applications such as this once it is.

Appendices

A Introduction to Statistical Physics and Phase Transitions

In this thesis we are concerned with the behaviour of models of large physical systems at equilibrium. The natural framework for this is statistical physics where one assigns an energy to each possible state of the system, then considers certain probability distributions over the states based on their energy. By considering statistics computed with respect to these distributions, one can capture the macroscopic properties of the system at hand including drastic phase transitions in the behaviour of the system. We shall briefly review the main ideas in this appendix which is inspired by [46].

A.1 Gibbs Distribution

To describe a model with a single external parameter β we need two things: a space of configurations $\mathcal{C}$ and an energy function (or Hamiltonian) $\mathcal{H} : \mathcal{C} \to \mathbb{R}$. If the space of configurations is countable, then the statistical properties of the system "at equilibrium" and fixed β are described by the partition function

$$Z(\beta) = \sum_{c \in \mathcal{C}} e^{-\beta \mathcal{H}(x)}.$$

If the space of configurations is uncountable, then we instead obtain an integral

$$Z(\beta) = \int_{\mathcal{C}} e^{-\beta \mathcal{H}(x)} dx.$$

Whichever case we find ourselves in, we write a probability distribution function $p_\beta : \mathcal{C} \to \mathbb{R}$ given by

$$p_\beta(x) = \frac{1}{Z(\beta)} e^{-\beta \mathcal{H}(x)}$$

specifying the unique distribution which maximises (information theoretic) entropy for a given value of the average energy $U(\beta) = \int_{\mathcal{C}} p_\beta(x) \mathcal{H}(x) dx$, often called the Boltzmann or Gibbs distribution.

As a justification of why this is the right distribution to look at, in thermodynamics the Gibbs distribution is the one that maximises thermodynamic entropy for a conserved total energy. Such a distribution describes the system at thermodynamic equilibrium with a fixed temperature $T = 1/k_B\beta$, where k_B is Boltzmann's constant. From now on we will ignore k_B or simply pretend it is equal to 1, and call β inverse temperature $1/T$.

We will denote the expectation of functions $A : \mathcal{C} \to \mathbb{R}^n$ with respect to the Gibbs distribution by

$$\langle A \rangle_\beta = \frac{1}{Z(\beta)} \int_{\mathcal{C}} A(x)\, e^{-\beta\mathcal{H}(x)} dx$$

although the expectations of some functions may be calculated from Z in an alternative way, via the free energy

$$\mathcal{F}(\beta) = -\frac{1}{\beta} \ln Z(\beta).$$

For example, a brief calculation shows that the average energy U is a first derivative of the free energy

$$U(\beta) = \langle \mathcal{H} \rangle_\beta = \frac{\partial(\beta\mathcal{F})}{\partial\beta}$$

and the specific heat

$$C(\beta) = -\beta^2 \frac{\partial U}{\partial\beta} = -\beta^2 \frac{\partial^2(\beta\mathcal{F})}{\partial\beta^2},$$

which measures how quickly the average energy responds to changes in temperature, becomes a measure of fluctuation in the energy

$$C(\beta) = \langle \mathcal{H}^2 \rangle_\beta - \langle \mathcal{H} \rangle_\beta^2 = \mathrm{Var}_\beta(\mathcal{H})$$

via a principle called the fluctuation-dissipation theorem.

A.2 An Example: The 2D Ising Model

The configurations of this model are specified by binary variables $\sigma_i \in \{-1, 1\}$ called spins on each site i of a square lattice Λ of size $L \times L$. That is, $\mathcal{C} = \mathbb{Z}_2^{L\times L}$. We will also call a spin $\sigma = 1$ an up spin and $\sigma = -1$ a down spin. The Hamiltonian assigns each configuration $\{\sigma_i\}_{i\in\Lambda}$ an energy based on how well the spins align with their nearest neighbours

$$\mathcal{H}(\{\sigma_i\}_{i\in\Lambda}) = -\sum_{\langle ij\rangle} \sigma_i\sigma_j$$

where $\langle ij\rangle$ runs over all nearest neighbour pairs of sites in the lattice, considering the lattice to have periodic boundary conditions. If many spins are aligned then $\sigma_i\sigma_j$ is usually 1, so the energy is low. If many are misaligned then the product term is usually -1, so the energy is high. The Ising model is usually interpreted as a model of magnetic material where each spin is the magnetic moment for an atom which influences the magnetic moments of neighbouring atoms. There is a known analytic formula for the free energy of the 2D Ising model, as found by Onsager [159]:

$$\frac{1}{L^2}\mathcal{F} = -\frac{\ln(2)}{2} - \ln(\cosh 2\beta) - \frac{1}{2\pi}\int_0^\pi \ln\left(1 + \sqrt{1 - k^2\cos^2\theta}\right)d\theta \tag{12}$$

where $k = 2\sinh 2\beta/\cosh^2 2\beta$. One can therefore also exactly calculate quantities such as the average energy and specific heat by taking the right derivatives as specified above.

A.3 Phase Transitions

In this framework a phase transition is said to occur whenever there is a singularity in a derivative of the free energy. When there is a discontinuity in a first derivative there is said to be a first order transition, and when there is a discontinuity or divergence in a higher derivative there is said to be a continuous phase transition. For example, a first order phase transition might manifest as a discontinuous jump in the average energy U as β is varied. Where as a typical indicator of a continuous phase transition is a divergence in the specific heat C. There is also a third type of transition, often called a topological phase transition, but we leave discussion of that unusual case to the main text in Section 3.

The first thing to note is that by the definition above, phase transitions do not occur on finite lattices Λ. It is straightforward to see that the partition function is a finite sum of analytic functions and hence analytic, so that the free energy is too. Therefore there are no singularities in any of its derivatives. However, this argument can and does fail when we pass to the thermodynamic limit $|\Lambda| \to \infty$. However, in the finite case we often still see the shadow of a phase transition: the divergence of a function might become a pronounced peak that gets taller as $|\Lambda|$ increases. We will gloss over this for now, but exactly how these finite-size effects manifest will be crucial to the work of this thesis and we instead cover this in Section 2.3.

We do not consider any first order phase transitions in this thesis, so let us concentrate on continuous transitions and consider our Ising model example. From Eqn. 12 one finds that the specific heat of the Ising model diverges at temperature

$$T_c = \frac{1}{\beta_c} = \frac{2}{\ln(1+\sqrt{2})}$$

but since having an analytic formula is the exception rather than the rule, it is more instructive to consider how else we might have spotted the phase transition.

Correlation Functions Just looking at typical configurations sampled from a low temperature (high β) and a high temperature (low β), for example those in Figure 80, reveal a stark difference in behaviour. Qualitatively, the spins in Figure 80 (a) seem to be more ordered, forming large regions where all the spins point the same way. Those in Figure 80 (b) are very much disordered; the direction a particular

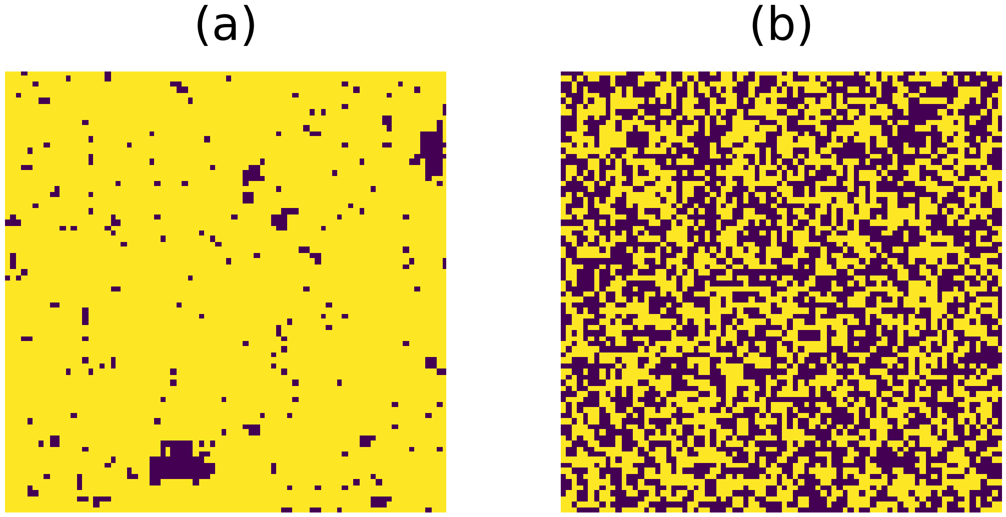

Figure 80: Sample configurations of the Ising model. (a) Sample from low temperature. (b) Sample from high temperature.

spin is pointing in tells us very little about what to expect even just a few lattice sites away. Quantitatively we introduce the following.

Definition A.1 (Spin-Spin Correlation Function)**.** For the Ising model, the spin-spin correlation function for a given value of β and a distance r is

$$C(r,\beta) = \langle \sigma_0 \, \sigma_r \rangle_\beta$$

which measures the correlation between a spin and another at distance r away from it in the lattice. Exactly which spins we look at does not matter due to translation invariance of the system, only the distance between them.

For all values of $\beta \neq \beta_c$ we asymptotically have that

$$C(r,\beta) \sim \frac{1}{r^\tau} \exp\left(-\frac{r}{\xi}\right)$$

where τ is some number depending on the phase and ξ is a characteristic length scale called the correlation length which varies with β.

As we might expect, at high temperatures the correlation function decays to zero and ξ is small. We say that there is only short-range order. Spins are almost independent of each other as it is, but especially as we consider spins further and further apart.

Definition A.2 (Short-Range Order)**.** A spin system is said to display short-range order if the spin-spin correlation function decays exponentially to zero with distance.

At low temperatures the correlation function decays towards a fixed value greater than zero. On average, most of the spins are pointing in the same direction as each other so all spins are correlated. There is said to be long-range order. Perhaps surprisingly, ξ is also small here: a spin flipping to -1 in a sea of $+1$ spins is unlikely to affect many others except its closest neighbours.

Definition A.3 (Long-Range Order)**.** A spin system is said to display long-range order if the spin-spin correlation function decays exponentially towards some $c > 0$ with distance.

Where things get interesting is right at the critical temperature. Here the correlation length ξ diverges, giving us instead a power-law for the correlation function

$$C(r, \beta) \sim \frac{1}{r^{\tau}},$$

decaying slower than the usual exponential decay. For a illustrative cartoon of what is going on here, consider Figure 81. Instead of the large clusters we see at low temperatures and the small clusters and television static we see at high temperatures, here we see clusters over many different scales. The flipping of a single spin will not affect a large cluster immediately, but it may affect a nearby small cluster, which itself may influence a larger cluster, which influences an even large cluster, and so on in a kind of cascade. The multiscale structure of configurations at criticality therefore allows the correlation between spins to carry farther in the lattice. In fact, one can make a more precise statement about scale invariance of the Ising model at criticality, but this is beyond the scope of this introduction. Instead we take away the idea that short and long-range order tell us about the phase of the model we are in, and a diverging correlation length signals the phase transition.

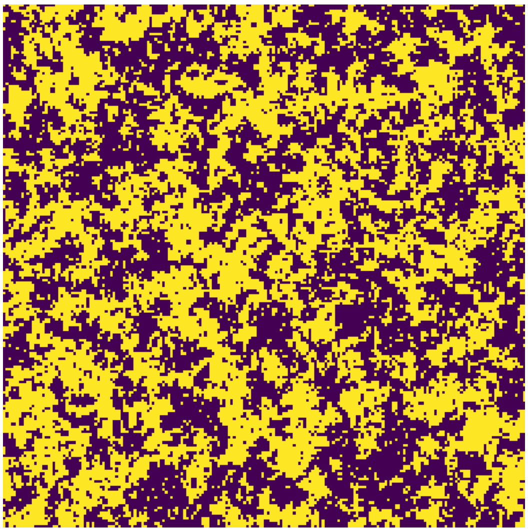

Figure 81: Sample configurations of the Ising model taken close to the critical point.

Order Parameters and Symmetry Breaking Looking at Figure 80 again, a different approach to looking at the phase transition is inspired by noticing that most of the spins in (a) point in the same direction where as in (b) we might estimate that there are as many up spins as down spins. We measure this excess via following observable.

Definition A.4 (Magnetisation)**.** For a lattice spin model, the magnetisation of a configuration $\boldsymbol{\sigma} = \{\sigma_i\}_{i\in\lambda}$ is defined to be

$$M(\boldsymbol{\sigma}) = \frac{1}{|\Lambda|}\sum_{i\in\Lambda}\sigma_i.$$

A cartoon plot of the magnetisation as a function of temperature is shown in Figure 82. Note that is that it is non-zero at low temperatures and zero at high temperatures. It is sensitive to the natural $\mathbb{Z}_2$ symmetry of the Ising model. Flipping all the spins $1 \mapsto -1$, $-1 \mapsto 1$ preserves the Hamiltonian $\mathcal{H}$, but flips the magnetisation M. We call an observable with these two properties an order parameter. How does this let us look at the phase transition? Well at high temperatures we have $M = 0$, so applying the spin flip preserves the magnetisation since $-0 = 0$. But at low temperatures we have $M = m$ for some $m \neq 0$. Applying the spin flip now gives $M = -m \neq m$, no longer preserving the magnetisation. The order parameter therefore allows us to observe the breaking of some $\mathbb{Z}_2$ symmetry. The idea is that this must occur at the phase transition, so we can identify the critical temperature as the point at which the order parameter first goes to zero.

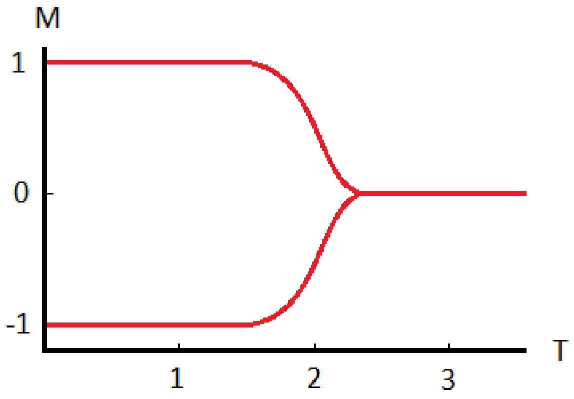

Figure 82: A cartoon of magnetisation in the Ising model.

In this thesis we will often consider observables $O : \mathcal{C} \to \mathbb{R}$ with the property that O is zero in one phase and non-zero in the other but which is not necessarily sensitive to the symmetries of the model in question. Instead of calling these order parameters as is often done, we will prefer to call such quantities phase indicators.

Simulation What if we look at one of the many models which, unlike the 2D Ising model, does not have an exact analytic solution? How can we hope to compute even simple expectations, let alone quantities such as the free energy? The answer is to simply estimate these quantities via Monte Carlo methods. In particular, the standard approach is to set up a Markov Chain which allows us to sample configurations according to the Gibbs distribution. That is, we generate a sequence of configurations $\{c_t\}_{t\in\mathbb{N}}$ whose energies are distributed according to the Gibbs distribution. One way to do this is through the Metropolis Hastings algorithm [160] where the next configuration in the sequence c_{t+1} is obtained from the previous c_t according to some easy to sample distribution $g(c_{t+1}|c_t)$. Algorithm 4 shows how to generate the next step in the Markov chain.

Algorithm 4 Metropolis Step [160]

```
Input Configuration $c_t$
Output Updated configuration $c_{t+1}$
$c' \leftarrow \text{Sample}(g(c'|c_t))$
$A \leftarrow \min\big(1, \exp\big[-\beta(\mathcal{H}(c') - \mathcal{H}(c_t))\big] g(c_t|c')/g(c'|c_t)\big)$
$u \leftarrow \text{Uniform}([0,1])$
if $u \leq A$ then
    $c_{t+1} \leftarrow c'$
else
    $c_{t+1} \leftarrow c_t$
end if
```

The idea behind such algorithms is to generate a Markov chain obeying two conditions. The first, detailed balance, means that the the probability of two consecutive configurations in the chain being A then B is the same as the probability of two consecutive steps being B then A. This is a sufficient condition to ensure that the Markov chain has a stationary distribution. The second, ergodicity, essentially says that the simulation can reach every possible configuration. This ensures that the stationary distribution is unique. The Metropolis algorithm satisfies these two conditions in such a way that the resulting stationary distribution is the Gibbs distribution that we wanted.

One typically begins the chain with either a random configuration or an ordered configuration and then runs several thousand Metropolis steps in order to approach the limiting distribution, or "thermalise". The distribution g is often very simple. For example in the Ising model (or other lattice spin models), one might generate the new configuration by selecting a lattice site at random and flipping the spin at that site (or choosing a random new value for it). Thus the configurations which are close together in the Markov chain can be very highly correlated, so one usually samples the chain every few hundred or thousand steps, or otherwise applies some statistical techniques to deal with this.

Now estimating some expected value $\langle O \rangle_\beta$ is as simple as sampling a large number N of configurations C_i from the Markov chain and calculating

$$\langle O \rangle_\beta \approx \frac{1}{N} \sum_{i=1}^{N} O(C_i).$$

The main reason that Monte Carlo methods are so helpful for statistical physics is that the error in this approximation decreases with $1/\sqrt{N}$ regardless of the dimensionality of the configuration space.

B Introduction to Computational Topology

As a subfield of mathematics, topology covers those properties of topological spaces invariant under continuous maps. This generality is powerful: the ideas and tools of topology can be applied to many weird and wonderful spaces and functions. However we run into an issue if we wish to do computational topology, since computers will only work with more structured and easily-described spaces and maps. As such, we will restrict ourselves to certain classes of topological spaces and functions. In particular we will introduce simplicial and cubical complexes which are constructed out of generalisations of triangles and squares respectively. As it turns out, many compact topological spaces which are sufficiently nice will be homeomorphic to some simplicial and/or cubical complex anyway, so we will obtain the same topological invariants when working with the appropriate simplicial or cubical approximation as with the original space.

Of the different flavours of topology, that which lends itself most towards computation is algebraic topology. Many of the principal tools of algebraic topology rely on the homotopy groups $\pi_k(X)$ which describe the different ways in which we can map a k-sphere into a space X up to continuous deformation. However, the group structure makes practical computation of these difficult, so once again we resort to a more easily computable alternative in the form of homology. The homology vector space $H_k(X)$ simply counts the number of k-dimensional holes in X via its dimension in a way that we shall soon make more precise.

B.1 Simplicial Complexes

Introductory courses in algebraic topology often start by introducing simplicial complexes, precisely because they are easy to construct and compute invariants with. In that scenario the intent is to perform computations by hand, but this ease of computation extends to computers too. A simplicial complex is an abstract construction, but it comes equipped with a geometric realisation which is a topological space constructed by gluing together triangles of different dimensions.

Definition B.1 (Abstract simplicial complex)**.** An abstract simplicial complex $K = (V, \Sigma)$ consists of a set V called the set of vertices of the complex, along with a set $\Sigma \subseteq \mathcal{P}(V)$ called the set of simplices, obeying the following two conditions:

1. For each $v \in V$, we have $\{v\} \in \Sigma$ (singletons are simplices).
2. If $\tau \in \Sigma$ and $\sigma \subseteq \tau$, then $\sigma \in \Sigma$ (the set of simplices is downwards closed).

We may simply refer to the abstract simplicial complex as Σ, where V is understood to be given by the singletons in Σ. The dimension of a simplex σ is $k = |\sigma| - 1$ and we call it a k-simplex.

In order to obtain the geometric realisation we first need the following definition.

Definition B.2 (Standard n-simplex)**.** The standard n-simplex is the subset

$$\Delta_n = \{\, (a_0, \dots, a_n) \,|\, \sum_i a_i = 1 \text{ and } a_i \geq 0 \,\forall i \,\} \subset \mathbb{R}^{n+1}.$$

The standard 0-simplex is a point, the 1-simplex is a line segment, the 2-simplex is a triangle, the 3-simplex is a tetrahedron and so on.

Definition B.3 (Geometric realisation)**.** Given an abstract simplicial complex $K = (V, \Sigma)$, its geometric realisation $|K|$ is the topological space obtained from the following procedure.

1. For each $\sigma \in \Sigma$ take a copy of the standard $(|\sigma| - 1)$-simplex and call it Δ_σ. Label its vertices with the elements of σ.
2. Whenever $\tau \subset \sigma \in \Sigma$, identify Δ_τ with a subset of Δ_σ via the face inclusion that sends the elements of τ to the corresponding elements of σ.

Thus the geometric realisation is a quotient space $|K| = \bigsqcup_{\sigma \in \Sigma} \Delta_{|\sigma|-1} \big/ \sim$ where $\sim$ is the equivalence relation described in Step 2.

Now let us look at how to define the homology of a simplicial complex. We first need to algebraise the notion of a collection of simplices. In the following we will fix some ordering on the vertex set.

Definition B.4 (Simplicial chains)**.** Fix some field $\mathbf{F}$ and some $k \in \mathbb{N}$. Given a simplicial complex Σ, the space of simplicial k-chains is the $\mathbf{F}$-vector space

$$C_k = \Big\{ \sum_i a_i \sigma_i \,|\, a_i \in \mathbf{F},\, \sigma_i \in \Sigma,\, |\sigma_i| = k + 1 \;\forall i \Big\},$$

consisting of finite formal sums of k-simplices with coefficients in $\mathbf{F}$. Given a chain $c \in C_k$, we say it has dimension $\dim(c) = k$.

What we are now interested in is how to obtain the boundary of a simplex or collection of simplices. Consider a 2-simplex $\{i, j, k\}$ which (by the geometric realisation) we can think of as a triangle with vertices i, j and k. Its boundary ought therefore to be made up of three edges or 1-simplices $\{i, j\}$, $\{i, k\}$ and $\{j, k\}$:

all the simplices obtained by removing one vertex from the original. Rather than just giving the sum of these simplices, we give the alternating sum which will help us keep track of orientations.

Definition B.5 (Simplicial boundary map)**.** Given a k-simplex $\sigma = \{v_0, \dots, v_k\}$ (where the labels v_j respect the ordering on the vertex set), its boundary is the $(k-1)$-chain

$$\partial\sigma = \sum_{j=0}^{k} (-1)^j \{v_0, \dots, v_{j-1}, v_{j+1}, \dots, v_k\}.$$

This extends linearly to the boundary map $\partial_k : C_k \to C_{k-1}$ sending

$$\sum_i a_i \sigma_i \mapsto \sum_i a_i \partial\sigma_i.$$

The alternating sum ensures that "interior" simplices are cancelled out. For example, consider the sum of two 1-simplices $\{i, j\} + \{j, k\}$ corresponding to two line segments joined end-to-end. The boundary would be

$$(\{i\} - \{j\}) + (\{j\} - \{k\}) = \{i\} - \{k\},$$

containing only the outer endpoints. The interior simplex $\{j\}$ was cancelled out.

In order to find holes we are interested in collections of simplices, algebraised as simplicial chains, which are closed. For example a 1-dimensional hole is enclosed by a closed curve and a 2-dimensional hole should be enclosed by a closed surface. Another way to say this is that the simplicial chain has no boundary.

Definition B.6 (Simplicial cycles)**.** Given a simplicial complex, its k-cycles are given by the subspace of C_k:

$$Z_k = \ker \partial_k.$$

However, not all cycles describe holes. Take our triangle $\{i, j, k\}$. The boundary $\{j, k\} - \{i, k\} + \{i, j\}$ is a 1-cycle: taking its boundary causes all the resulting 0-chains to cancel out. But we certainly would not say that there is a hole in the triangle. The issue here is that this cycle bounds a higher dimensional simplex which "fills" the hole.

Definition B.7 (Simplicial boundaries)**.** Given a simplicial complex, its k-boundaries are given by the subspace of C_k:

$$B_k = \operatorname{im} \partial_{k+1}.$$

It is easy to verify using the formula for ∂_k that $B_k \subseteq Z_k$ for all k, which is equivalent to saying that $\partial_k \circ \partial_{k+1} = 0$. This finally allows us to define homology by looking at cycles that are not filled in by higher dimensional simplices, or indeed that simply differ by the addition or subtraction of the boundaries of higher dimensional simplices.

Definition B.8 (Simplicial homology)**.** Given a simplicial complex K, its k-th simplicial homology is the quotient $\mathbf{F}$-vector space

$$H_k(K) = Z_k / B_k.$$

We therefore refer to homology classes $[c] \in H_k$ (or $c + B_k$ in coset notation), where $c \in Z_k$ is a representative cycle. For example, the dimensionality of H_0 records the connected components of the simplicial complex. Take the complex consisting of a single line segment $\{i, j\}$. Both $\{i\}$ and $\{j\}$ are 0-cycles, but we clearly do not have two connected components. This is reflected in H_0 by the fact that these two cycles differ by the boundary $\partial_1 \{i, j\} = \{i\} - \{j\}$. Therefore both $\{i\}$ and $\{j\}$ are representatives of the same homology class.

To rephrase the above as a purely algebraic picture, given a simplicial complex, we have a sequence of vector spaces and linear maps

$$\ldots \to C_3 \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

describing how to go from a collection of simplices to its boundary, with the property that $\partial_k \circ \partial_{k+1} = 0$. Such a sequence is called a chain complex and given any chain complex we can define its k-th homology as the vector space $H_k = \ker \partial_k / \operatorname{im} \partial_{k+1}$. Homology therefore maps simplicial complexes to vector spaces. We can ask if this mapping is functorial: if homology also sends simplicial maps $X \to Y$ between simplicial complexes to linear maps $H_k(X) \to H_k(Y)$ between vector spaces. But first we need to know what the appropriate maps between simplicial complexes are.

Definition B.9 (Simplicial map)**.** Given two simplicial complexes $K = (V_K, \Sigma_K)$ and $L = (V_L, \Sigma_L)$. A simplicial map $f : K \to L$ is described by a map $f_V : V_K \to V_L$ such that $\forall \sigma \in \Sigma_K$ we have that $f(\sigma) = \{f(v) \,|\, v \in \sigma\} \in \Sigma_L$.

Now we can extend a simplicial map f to give a map between the spaces of k-chains of K and L, by defining $f_k : C_k(K) \to C_k(L)$ as the map

$$\sum_i a_i \sigma_i \mapsto \sum_{\substack{i \\ \dim f(\sigma_i) = k}} a_i f(\sigma_i).$$

The key result now is that this map descends to the quotient.

Proposition B.1 (Functoriality of homology)**.** *Given simplicial complexes K and L and a simplicial map $f : K \to L$, the induced linear map f_k on k-chains descends to a map $f_k : H_k(K) \to H_k(L)$ on k-th homology given by*

$$[c] \mapsto [f(c)].$$

The identity map $id : K \to K$ induces the identity map on homology. Moreover, if $g : L \to M$ is another simplicial map, then

$$(g \circ f)_k = g_k \circ f_k : H_k(K) \to H_k(M).$$

The fact that there is an induced map on homology which respects composition is referred to by saying that homology is a functor from the category of simplicial complexes with simplicial maps to the category of vector spaces with linear maps. This will be important for the definition of persistent homology in Section 2.1.1.

Finally, we should also define cohomology since it comes up briefly in Sections 2.1.2 and 2.1.4. The idea is to take our sequence of chain spaces and boundary maps

$$\ldots \to C_3 \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

and simply dualise everything. That is, we replace C_k with $C_k^* = \mathrm{Hom}(C_k, \mathbf{F})$ and the boundary maps $\partial_k : C_k \to C_{k-1}$ with coboundary maps $\delta^k : C_k^* \to C_{k+1}^*$ in the opposite direction giving

$$\ldots \leftarrow C_3^* \xleftarrow{\delta_2} C_2^* \xleftarrow{\delta_1} C_1^* \xleftarrow{\delta_0} C_0^* \leftarrow 0,$$

where δ_k is defined on a cochain $b \in C_k^*$ as

$$\delta_k(b)(c) = b(\partial_{k+1} c)$$

for all chains $c \in C_{k+1}$.

Definition B.10 (Simplicial cohomology)**.** Given a simplicial complex K, its k-th simplicial cohomology is the quotient $\mathbf{F}$-vector space

$$H^k(K) = \ker(\delta_k)/\operatorname{im}(\delta_{k-1}).$$

Cohomology is also a functor from simplicial complexes to vector spaces, but with the arrows flipped (it is a contravariant functor).

Proposition B.2 (Functoriality of cohomology)**.** *Given simplicial complexes K and L and a simplicial map $f : K \to L$, there is an induced linear map $f^k : H^k(L) \to H^k(K)$ on k-th cohomology given by*

$$[b] \mapsto [b \circ f_k]$$

where $f_k : C_k(K) \to C_k(L)$ is the induced map on chains. The identity map id : $K \to K$ induces the identity map on cohomology. Moreover, if $g : L \to M$ is another simplicial map, then

$$(g \circ f)^k = g^k \circ f^k : H^k(M) \to H^k(K).$$

B.2 Cubical Complexes

In most of this thesis it will make more sense to work with a slightly different type of complex. If we instead build our spaces out of generalisations of squares and cubes rather than triangles, we obtain cubical complexes.

The story here closely mirrors that for simplicial complexes. We can make the following definitions, following [161].

Definition B.11 (Abstract cubical complexes)**.** An abstract cubical complex $K = (V, \Sigma)$ consists of a vertex set V along with a set $\Sigma \subseteq \mathcal{P}(V)$ called the cubes of the complex, obeying:

1. For each $v \in V$, we have $\{v\} \in \Sigma$.
2. For every $\sigma \in \Sigma$ the elements of σ can be represented as the vertices of a finite dimensional cube, where the faces contained in σ are exactly the vertex sets of the faces of this cube.
3. If $\sigma, \tau \in \Sigma$, then $\sigma \cap \tau = \emptyset$ or $\sigma \cap \tau \in \Sigma$.

A map $f : V_K \to V_L$ between the vertex sets of two cubical complexes is a cubical map if:

1. For every $\sigma \in \Sigma_K$, $f(\sigma) \subseteq \tau$ for some $\tau \in \Sigma_L$.
2. f takes adjacent vertices in K (connected by a 1-cube) to either adjacent vertices in L, or the same vertex.

We could then simply copy and paste the rest of the definitions and results for simplicial complexes with only minor modifications to get to the notion of homology. However, it may be more instructive to approach the definition in a slightly different way. Rather than starting with this abstract data structure, one can define cubical complexes concretely as subsets of Euclidean space from the start, as in [162].

Definition B.12 (Elementary intervals and cubes)**.** An elementary interval is an interval of the form $[i, i+1] \subset \mathbb{R}$ (non-degenerate) or $[i, i] = \{n\}$ (degenerate) for some choice of $i \in \mathbb{Z}$. An elementary cube is a finite product of elementary intervals $Q = I_1 \times \ldots \times I_n \subset \mathbb{R}^n$, where n is some fixed embedding dimension. Its dimension $\dim Q$ is the number of non-degenerate intervals in the product.

Definition B.13 (Concrete cubical complex)**.** A cubical complex C is a subset of $\mathbb{R}^n$ that is a union of elementary cubes which may be of varying dimension. n is known as the embedding dimension. The dimension of the cubical complex $\dim C$ is the maximum dimension of any elementary cube contained in C.

As before, in order to define homology we need to inject some algebra. Fix a field $\mathbf{F}$.

Definition B.14 (Cubical chains)**.** Given $k \in \mathbb{N}$, the space of cubical k-chains is the $\mathbf{F}$-vector space

$$C_k = \{\sum \alpha_i Q_i \mid Q_i \subseteq C,\, \dim Q_i = k,\, \alpha_i \in \mathbf{F}\},$$

consisting of finite formal sums of elementary cubes.

Definition B.15 (Cubical boundary map)**.** The boundary of a non-degenerate elementary interval is given by the formal sum $\partial[i, i+1] = [i+1, i+1] - [i, i]$. For a degenerate elementary interval the boundary is zero. The boundary of an elementary cube $Q = (I_1 \times \ldots \times I_n)$ is a formal sum

$$\partial Q = \sum_{j=1}^{n} (-1)^{\sum_{i=1}^{j-1} \dim Q_i} (I_1 \times \ldots \times \partial I_j \times \ldots \times I_n) \tag{13}$$

where we consider $\times$ as distributing over the formal summation. We can see that for $\dim Q \geq 1$ we have $\dim \partial Q = \dim Q - 1$. Therefore we can extend ∂ to linear maps $\partial_k : C_k \to C_{k-1}$ via the mapping $\sum \alpha_i Q_i \mapsto \sum \alpha_i (\partial Q_i)$.

Once again, we are interested in the closed chains and chains that are obtained in the image of the boundary map.

Definition B.16 (Cubical cycles and boundaries)**.** The space of k-cycles is the subspace

$$Z_k = \{\, c \in C_k \,|\, \partial_k c = 0 \,\} \subseteq C_k$$

and the space of k-boundaries is the subspace

$$B_k = \{\, \partial c \,|\, c \in C_{k+1} \,\} \subseteq C_k.$$

Now a key observation is that, since $\partial\partial I = 0$ for any elementary interval I, we also have that $\partial_k \circ \partial_{k+1} = 0$ for all $k \in \mathbb{N}$, so that $B_k \subseteq Z_k$. See Figure 83 for an example.

The sequence of linear maps

$$\ldots \to C_3 \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

therefore forms a chain complex.

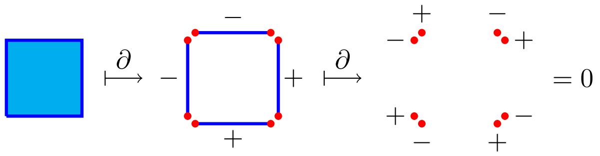

Figure 83: Example of how the boundary operator ∂ acts on a simple cubical complex consisting of a single 2-dimensional cube. Note how sum in equation (13) being alternating ensures that $\partial\partial = 0$.

Definition B.17 (Cubical homology)**.** The k^{th} cubical homology of C is defined to be the quotient $\mathbf{F}$-vector space

$$H_k(C) = Z_k/B_k.$$

We therefore have collections of cubes forming closed loops, surfaces, etc. which are not simply the boundaries of higher-dimensional cubes: i.e. holes.

For the definition of persistent homology in Section 2.1.1 to work for cubical complexes, we need to know that cubical homology is functorial: given a suitable definition of a cubical map $f : C \to D$ between cubical complexes, there is an induced map $f_k : H_k(C) \to H_k(D)$ for each $k \in \mathbb{N}$ which respects composition. Unfortunately this is when the concrete approach to cubical complexes let us down. It is certainly possible to define cubical maps and obtain an induced map on homology as is done in [162], but the story is much more complicated. Instead we shall simply state the following.

Proposition B.3 (Functoriality of cubical homology)**.** *Given a cubical map $f : C \to D$, there is an induced map on cubical homology $f_k : H_k(C) \to H_k(D)$ for all $k \in \mathbb{N}$. Moreover, given another cubical map $g : D \to E$, then*

$$(g \circ f)_k = g_k \circ f_k : H_k(C) \to H_k(E).$$

In particular for our purposes, given a subset $C \subseteq D$ of a cubical complex D, the inclusion map $C \hookrightarrow D$ is a cubical map and thus induces a map on homology.

Cohomology is obtained in exactly the same way as with simplicial complexes: by dualising the chain complex of chain spaces and boundary maps, then computing the quotient of coboundary maps.

References

[1] Nicholas Sale, Jeffrey Giansiracusa, and Biagio Lucini. “Quantitative analysis of phase transitions in two-dimensional XY models using persistent homology”. In: *Phys. Rev. E* 105 (2 Feb. 2022), p. 024121. DOI: `10.1103/PhysRevE.105.024121`.

[2] Nicholas Sale, Biagio Lucini, and Jeffrey Giansiracusa. *Probing center vortices and deconfinement in* SU(2) *lattice gauge theory with persistent homology.* 2022. DOI: `10.48550/ARXIV.2207.13392`. arXiv: `2207.13392 [hep-lat]`.

[3] N. D. Mermin. “The topological theory of defects in ordered media”. In: *Rev. Mod. Phys.* 51 (3 July 1979), pp. 591–648. DOI: `10.1103/RevModPhys.51.591`.

[4] Charles Nash and Siddhartha Sen. *Topology and Geometry for Physicists.* Academic Press, 1988. ISBN: 0125140819.

[5] Raoul Bott. “An application of the Morse theory to the topology of Lie-groups”. In: *Bulletin de la Société Mathématique de France* 84 (1956), pp. 251–281. URL: `http://www.numdam.org/article/BSMF_1956__84__251_0.pdf`.

[6] Elchanan Solomon, Alexander Wagner, and Paul Bendich. *From Geometry to Topology: Inverse Theorems for Distributed Persistence.* 2021. DOI: `10.48550/ARXIV.2101.12288`. arXiv: `2101.12288 [math.AT]`.

[7] Vin De Silva, Robert Ghrist, et al. “Homological sensor networks”. In: *Notices of the American mathematical society* 54.1 (2007), pp. 10–17. URL: `https://www.ams.org/notices/200701/fea-ghrist.pdf`.

[8] Chunyuan Li, Maks Ovsjanikov, and Frederic Chazal. “Persistence-Based Structural Recognition”. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition.* 2014, pp. 2003–2010. DOI: `10.1109/CVPR.2014.257`.

[9] Violeta Kovacev-Nikolic, Peter Bubenik, Dragan Nikolić, and Giseon Heo. “Using persistent homology and dynamical distances to analyze protein binding”. In: *Statistical Applications in Genetics and Molecular Biology* 15.1 (2016), pp. 19–38. DOI: `doi:10.1515/sagmb-2015-0057`.

[10] Adrien Poulenard, Primoz Skraba, and Maks Ovsjanikov. "Topological Function Optimization for Continuous Shape Matching". In: *Computer Graphics Forum* 37 (2018). DOI: `10.1111/cgf.13487`.

[11] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl. "Deep Learning with Topological Signatures". In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. Curran Associates, Inc., 2017. URL: `https://proceedings.neurips.cc/paper/2017/file/883e881bb4d22a7add958f2d6b052c9f-Paper.pdf`.

[12] Mathieu Carrière and Raúl Rabadán. "Topological Data Analysis of Single-Cell Hi-C Contact Maps". In: *Topological Data Analysis*. Ed. by Nils A. Baas, Gunnar E. Carlsson, Gereon Quick, Markus Szymik, and Marius Thaule. Cham: Springer International Publishing, 2020, pp. 147–162. ISBN: 978-3-030-43408-3. DOI: `10.1007/978-3-030-43408-3_6`.

[13] Anuraag Bukkuri, Noemi Andor, and Isabel K. Darcy. "Applications of Topological Data Analysis in Oncology". In: *Frontiers in Artificial Intelligence* 4 (2021). ISSN: 2624-8212. DOI: `10.3389/frai.2021.659037`.

[14] Tananun Songdechakraiwut and Moo K. Chung. "Dynamic Topological Data Analysis for Functional Brain Signals". In: *2020 IEEE 17th International Symposium on Biomedical Imaging Workshops (ISBI Workshops)*. 2020, pp. 1–4. DOI: `10.1109/ISBIWorkshops50223.2020.9153431`.

[15] Allison Godwin, Aaron Robert Hamilton Thielmeyer, Jacqueline Ann Rohde, Dina Verdín, Brianna Shani Benedict, Rachel Ann Baker, and Jacqueline Doyle. "Using topological data analysis in social science research: Unpacking decisions and opportunities for a new method". English (US). In: *ASEE Annual Conference and Exposition, Conference Proceedings* (June 2019). ISSN: 2153-5965. DOI: `10.18260/1-2--33522`.

[16] Mickaël Buchet, Yasuaki Hiraoka, and Ippei Obayashi. "Persistent Homology and Materials Informatics". In: *Nanoinformatics*. Ed. by Isao Tanaka. Singapore: Springer Singapore, 2018, pp. 75–95. ISBN: 978-981-10-7617-6. DOI: `10.1007/978-981-10-7617-6_5`.

[17] H. Edelsbrunner, D. Letscher, and A. Zomorodian. "Topological Persistence and Simplification". In: *Discrete & Computational Geometry* 28 (2002), pp. 511–533. DOI: `10.1007/s00454-002-2885-2`.

[18] G. Carlsson. "Topology and data". In: *Bulletin of the American Mathematical Society* 46 (2009), pp. 255–308. DOI: `10.1090/S0273-0979-09-01249-X`.

[19] R. Ghrist. "Barcodes: The persistent topology of data". In: *Bulletin of the American Mathematical Society* 45 (2007), pp. 61–75. DOI: `10.1090/S0273-0979-07-01191-3`.

[20] Nina Otter, Mason Porter, Ulrike Tillmann, Peter Grindrod, and Heather Harrington. "A roadmap for the computation of persistent homology". In: *EPJ Data Science* 6 (June 2015). DOI: `10.1140/epjds/s13688-017-0109-5`.

[21] Herbert Edelsbrunner and John Harer. "Persistent homology - a survey". In: *Discrete & Computational Geometry - DCG* 453 (Jan. 2008). DOI: `10.1090/conm/453/08802`.

[22] Gunnar Carlsson. *Persistent Homology and Applied Homotopy Theory*. 2020. DOI: `10.48550/ARXIV.2004.00738`. arXiv: `2004.00738 [math.AT]`.

[23] Steve Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*. Vol. 209. Mathematical Surveys and Monographs. American Mathematical Society, 2015. ISBN: 978-1-4704-3443-4.

[24] William Crawley-Boevey. *Decomposition of pointwise finite-dimensional persistence modules*. 2012. DOI: `10.48550/ARXIV.1210.0819`. arXiv: `1210.0819 [math.RT]`.

[25] Gorô Azumaya. "Corrections and Supplementaries to My Paper concerning Krull-Remak-Schmidt's Theorem". In: *Nagoya Mathematical Journal* 1 (1950), pp. 117–124. DOI: `10.1017/S002776300002290X`.

[26] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The Structure and Stability of Persistence Modules*. Springer International Publishing, 2016. DOI: `10.1007/978-3-319-42545-0`.

[27] Michael Lesnick. “The Theory of the Interleaving Distance on Multidimensional Persistence Modules”. In: *Foundations of Computational Mathematics* 15.3 (June 2015), pp. 613–650. ISSN: 1615-3383. DOI: `10.1007/s10208-015-9255-y`.

[28] Primoz Skraba and Katharine Turner. *Wasserstein Stability for Persistence Diagrams*. June 2020. arXiv: `2006.16824 [math.AT]`.

[29] Michael Kastner. “Phase transitions and configuration space topology”. In: *Rev. Mod. Phys.* 80 (1 Jan. 2008), pp. 167–187. DOI: `10.1103/RevModPhys.80.167`.

[30] Lando Caiani, Lapo Casetti, Cecilia Clementi, and Marco Pettini. “Geometry of Dynamics, Lyapunov Exponents, and Phase Transitions”. In: *Phys. Rev. Lett.* 79 (22 Dec. 1997), pp. 4361–4364. DOI: `10.1103/PhysRevLett.79.4361`.

[31] Roberto Franzosi and Marco Pettini. “Theorem on the Origin of Phase Transitions”. In: *Phys. Rev. Lett.* 92 (6 Feb. 2004), p. 060601. DOI: `10.1103/PhysRevLett.92.060601`.

[32] Guillaume Tauzin, Umberto Lupo, Lewis Tunstall, Julian Burella Pérez, Matteo Caorsi, Anibal Medina-Mardones, Alberto Dassatti, and Kathryn Hess. *giotto-tda: A Topological Data Analysis Toolkit for Machine Learning and Data Exploration*. 2020. DOI: `10.48550/ARXIV.2004.02551`. arXiv: `2004.02551 [cs.LG]`.

[33] Christopher Tralie, Nathaniel Saul, and Rann Bar-On. “Ripser.py: A Lean Persistent Homology Library for Python”. In: *The Journal of Open Source Software* 3.29 (Sept. 2018), p. 925. DOI: `10.21105/joss.00925`.

[34] Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. “Dualities in persistent (co)homology”. In: *Inverse Problems* 27 (2011), p. 124003. DOI: `10.1088/0266-5611/27/12/124003`.

[35] Ulrich Bauer. “Ripser: efficient computation of Vietoris-Rips persistence barcodes”. In: *Journal of Applied and Computational Topology* (2021). DOI: `10.1007/s41468-021-00071-5`.

[36] Allen Hatcher. *Algebraic topology*. Cambridge: Cambridge Univ. Press, 2000. URL: `https://cds.cern.ch/record/478079`.

[37] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. "Persistence Images: A Stable Vector Representation of Persistent Homology". In: *J. Mach. Learn. Res.* 18.1 (Jan. 2017), pp. 218–252. ISSN: 1532-4435. URL: `http://jmlr.org/papers/v18/16-337.html`.

[38] Bartosz Zieliński, Michał Lipiński, Mateusz Juda, Matthias Zeppelzauer, and Paweł Dłotko. "Persistence Bag-of-Words for Topological Data Analysis". In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19.* International Joint Conferences on Artificial Intelligence Organization, July 2019, pp. 4489–4495. DOI: `10.24963/ijcai.2019/624`.

[39] Peter Bubenik. "Statistical Topological Data Analysis Using Persistence Landscapes". In: *J. Mach. Learn. Res.* 16.1 (Jan. 2015), pp. 77–102. URL: `https://dl.acm.org/doi/10.5555/2789272.2789275`.

[40] Vincent Divol and Frédéric Chazal. "The density of expected persistence diagrams and its kernel based estimation". In: *JoCG* 10.2 (2019), pp. 127–153. DOI: `10.20382/jocg.v10i2a7`.

[41] Lu Li, Connor Thompson, Gregory Henselman-Petrusek, Chad Giusti, and Lori Ziegelmeier. "Minimal Cycle Representatives in Persistent Homology Using Linear Programming: An Empirical Study With User's Guide". In: *Frontiers Artif. Intell.* 4 (2021), p. 681117. DOI: `10.3389/frai.2021.681117`.

[42] Afra Zomorodian and Gunnar E. Carlsson. "Computing Persistent Homology". In: *Discrete & Computational Geometry* 33 (2004), pp. 249–274. DOI: `10.1007/s00454-004-1146-y`.

[43] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. "Vines and Vineyards by Updating Persistence in Linear Time". In: *Proceedings of the Twenty-Second Annual Symposium on Computational Geometry.* SCG '06. Sedona, Arizona, USA: Association for Computing Machinery, 2006, pp. 119–126. ISBN: 1595933409. DOI: `10.1145/1137856.1137877`.

[44] Tamal K. Dey, Tao Hou, and Sayan Mandal. "Persistent 1-Cycles: Definition, Computation, and Its Application". In: *Computational Topology in Image Context*. Ed. by Rebeca Marfil, Mariletty Calderón, Fernando Díaz del Río, Pedro Real, and Antonio Bandera. Cham: Springer International Publishing, 2019, pp. 123–136. ISBN: 978-3-030-10828-1. DOI: `10.1007/978-3-030-10828-1_10`.

[45] Alex Cole, Gregory J. Loges, and Gary Shiu. "Quantitative and interpretable order parameters for phase transitions from persistent homology". In: *Phys. Rev. B* 104 (10 Sept. 2021), p. 104426. DOI: `10.1103/PhysRevB.104.104426`.

[46] J. M. Yeomans. *Statistical mechanics of phase transitions.* Oxford science publications. Oxford: Clarendon Press, 1992. ISBN: 0198517300.

[47] Michael E. Fisher and Michael N. Barber. "Scaling Theory for Finite-Size Effects in the Critical Region". In: *Phys. Rev. Lett.* 28 (23 June 1972), pp. 1516–1519. DOI: `10.1103/PhysRevLett.28.1516`.

[48] J M Kosterlitz. "The critical properties of the two-dimensional xy model". In: *Journal of Physics C: Solid State Physics* 7.6 (Mar. 1974), pp. 1046–1060. DOI: `10.1088/0022-3719/7/6/005`.

[49] Michael E. Fisher, Michael N. Barber, and David Jasnow. "Helicity Modulus, Superfluidity, and Scaling in Isotropic Systems". In: *Phys. Rev. A* 8 (2 Aug. 1973), pp. 1111–1124. DOI: `10.1103/PhysRevA.8.1111`.

[50] S. M. Bhattacharjee and F. Seno. "A measure of data collapse for scaling". In: *Journal of Physics A* 34 (2001), pp. 6375–6380. DOI: `10.1088/0305-4470/34/33/302`.

[51] Alan M. Ferrenberg and Robert H. Swendsen. "New Monte Carlo technique for studying phase transitions". In: *Phys. Rev. Lett.* 61 (23 Dec. 1988), pp. 2635–2638. DOI: `10.1103/PhysRevLett.61.2635`.

[52] Alan M. Ferrenberg and Robert H. Swendsen. "Optimized Monte Carlo data analysis". In: *Phys. Rev. Lett.* 63 (12 Sept. 1989), pp. 1195–1198. DOI: `10.1103/PhysRevLett.63.1195`.

[53] Askery Canabarro, Felipe Fernandes Fanchini, André Luiz Malvezzi, Rodrigo Pereira, and Rafael Chaves. "Unveiling phase transitions with machine learning". In: *Phys. Rev. B* 100 (4 July 2019), p. 045129. DOI: `10.1103/PhysRevB.100.045129`.

[54] Jonas Greitemann, Ke Liu, and Lode Pollet. "Probing hidden spin order with interpretable machine learning". In: *Phys. Rev. B* 99 (6 Feb. 2019), p. 060404. DOI: `10.1103/PhysRevB.99.060404`.

[55] Sebastian J. Wetzel and Manuel Scherzer. "Machine learning of explicit order parameters: From the Ising model to SU(2) lattice gauge theory". In: *Phys. Rev. B* 96 (18 Nov. 2017), p. 184410. DOI: `10.1103/PhysRevB.96.184410`.

[56] E. V. Nieuwenburg, Y. Liu, and S. Huber. "Learning phase transitions by confusion". In: *Nature Physics* 13 (2017), pp. 435–439. DOI: `10.1038/nphys4037`.

[57] Tanaka Akinori and Tomiya Akio. "Detection of Phase Transition via Convolutional Neural Networks". In: *Journal of the Physical Society of Japan* 86 (2017), p. 063001. DOI: `10.7566/JPSJ.86.063001`.

[58] Lei Wang. "Discovering phase transitions with unsupervised learning". In: *Phys. Rev. B* 94 (19 Nov. 2016), p. 195105. DOI: `10.1103/PhysRevB.94.195105`.

[59] Juan Carrasquilla and Roger G. Melko. "Machine learning phases of matter". In: *Nature Physics* 13.5 (May 2017), pp. 431–434. DOI: `10.1038/nphys4035`.

[60] Cinzia Giannetti, Biagio Lucini, and Davide Vadacchino. "Machine Learning as a universal tool for quantitative investigations of phase transitions". In: *Nuclear Physics B* 944 (2019), p. 114639. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/j.nuclphysb.2019.114639`.

[61] Dimitrios Bachtis, Gert Aarts, Francesco Di Renzo, and Biagio Lucini. "Inverse Renormalization Group in Quantum Field Theory". In: *Phys. Rev. Lett.* 128 (8 Feb. 2022), p. 081603. DOI: `10.1103/PhysRevLett.128.081603`.

[62] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini. “Quantum field-theoretic machine learning”. In: *Phys. Rev. D* 103.7 (2021), p. 074510. DOI: `10.1103/PhysRevD.103.074510`. arXiv: `2102.09449 [hep-lat]`.

[63] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini. “Adding machine learning within Hamiltonians: Renormalization group transformations, symmetry breaking and restoration”. In: *Phys. Rev. Res.* 3.1 (2021), p. 013134. DOI: `10.1103/PhysRevResearch.3.013134`. arXiv: `2010.00054 [hep-lat]`.

[64] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini. “Mapping distinct phase transitions to a neural network”. In: *Phys. Rev. E* 102.5 (2020), p. 053306. DOI: `10.1103/PhysRevE.102.053306`. arXiv: `2007.00355 [cond-mat.stat-m`

[65] Dimitrios Bachtis, Gert Aarts, and Biagio Lucini. “Extending machine learning classification capabilities with histogram reweighting”. In: *Phys. Rev. E* 102.3 (2020), p. 033303. DOI: `10.1103/PhysRevE.102.033303`. arXiv: `2004.14341 [cond-mat.stat-mech]`.

[66] W. Rzadkowski, N. Defenu, S. Chiacchiera, A. Trombettoni, and G. Bighin. “Detecting composite orders in layered models via machine learning”. In: *New J. Phys.* 22.9 (2020), p. 093026. DOI: `10.1088/1367-2630/abae44`. arXiv: `1907.05417 [cond-mat.dis-nn]`.

[67] Constantia Alexandrou, Andreas Athenodorou, Charalambos Chrysostomou, and Srijit Paul. “The critical temperature of the 2D-Ising model through Deep Learning Autoencoders”. In: *Eur. Phys. J. B* 93.12 (2020), p. 226. DOI: `10.1140/epjb/e2020-100506-5`. arXiv: `1903.03506 [cond-mat.stat-mech`

[68] Ankita Chakrabarti, S. R. Hassan, and R. Shankar. “Intrinsic and extrinsic geometries of correlated many-body states”. In: *Phys. Rev. B* 99 (8 Feb. 2019), p. 085138. DOI: `10.1103/PhysRevB.99.085138`.

[69] Fernando A. N. Santos, Ernesto P. Raposo, Maurício D. Coutinho-Filho, Mauro Copelli, Cornelis J. Stam, and Linda Douw. “Topological phase transitions in functional brain networks”. In: *Phys. Rev. E* 100 (3 Sept. 2019), p. 032414. DOI: `10.1103/PhysRevE.100.032414`.

[70] T. Mendes-Santos, X. Turkeshi, M. Dalmonte, and Alex Rodriguez. "Unsupervised Learning Universal Critical Behavior via the Intrinsic Dimension". In: *Phys. Rev. X* 11 (1 Feb. 2021), p. 011040. DOI: `10.1103/PhysRevX.11.011040`.

[71] Joaquin F. Rodriguez-Nieva and Mathias S. Scheurer. "Identifying topological order through unsupervised machine learning". In: *Nature Physics* 15.8 (2019), pp. 790–795. ISSN: 1745-2481. DOI: `10.1038/s41567-019-0512-x`.

[72] Quoc Hoan Tran, Mark Chen, and Yoshihiko Hasegawa. "Topological persistence machine of phase transitions". In: *Phys. Rev. E* 103 (5 May 2021), p. 052127. DOI: `10.1103/PhysRevE.103.052127`.

[73] Leo Speidel, Heather A. Harrington, S. Jonathan Chapman, and Mason A. Porter. "Topological data analysis of continuum percolation with disks". In: *Phys. Rev. E* 98 (1 July 2018), p. 012318. DOI: `10.1103/PhysRevE.98.012318`.

[74] Bart Olsthoorn, Johan Hellsvik, and Alexander V. Balatsky. "Finding hidden order in spin models with persistent homology". In: *Phys. Rev. Research* 2 (4 Dec. 2020), p. 043308. DOI: `10.1103/PhysRevResearch.2.043308`.

[75] Takehiro Hirakida, K. Kashiwa, Junpei Sugano, J. Takahashi, H. Kouno, and M. Yahiro. *Persistent homology analysis of deconfinement transition in effective Polyakov-line model.* 2018. DOI: `10.48550/ARXIV.1810.07635`. arXiv: `1810.07635 [hep-lat]`.

[76] Irene Donato, Matteo Gori, Marco Pettini, Giovanni Petri, Sarah De Nigris, Roberto Franzosi, and Francesco Vaccarino. "Persistent homology analysis of phase transitions". In: *Phys. Rev. E* 93 (5 May 2016), p. 052138. DOI: `10.1103/PhysRevE.93.052138`.

[77] Daniel Leykam and Dimitris G. Angelakis. *Topological data analysis and machine learning.* 2022. DOI: `10.48550/ARXIV.2206.15075`. arXiv: `2206.15075 [cond-mat.mes-hall]`.

[78] Andrea Tirelli and Natanael C. Costa. "Learning quantum phase transitions through topological data analysis". In: *Phys. Rev. B* 104 (23 Dec. 2021), p. 235146. DOI: `10.1103/PhysRevB.104.235146`.

[79] Yu He, Shiqi Xia, Dimitris G. Angelakis, Daohong Song, Zhigang Chen, and Daniel Leykam. *Persistent homology analysis of a generalized Aubry-André-Harper model.* 2022. DOI: `10.48550/ARXIV.2204.13276`. arXiv: `2204.13276 [cond-mat.dis-nn]`.

[80] Dan Sehayek and Roger G. Melko. "Persistent homology of $\mathbb{Z}_2$ gauge theories". In: *Phys. Rev. B* 106 (8 Aug. 2022), p. 085111. DOI: `10.1103/PhysRevB.106.085111`.

[81] J M Kosterlitz and D J Thouless. "Long range order and metastability in two dimensional solids and superfluids. (Application of dislocation theory)". In: *Journal of Physics C: Solid State Physics* 5.11 (June 1972), pp. L124–L126. DOI: `10.1088/0022-3719/5/11/002`.

[82] J M Kosterlitz and D J Thouless. "Ordering, metastability and phase transitions in two-dimensional systems". In: *Journal of Physics C: Solid State Physics* 6.7 (Apr. 1973), pp. 1181–1203. DOI: `10.1088/0022-3719/6/7/010`.

[83] V. L. Berezinsky. "Destruction of long range order in one-dimensional and two-dimensional systems having a continuous symmetry group. I. Classical systems". In: *Sov. Phys. JETP* 32 (1971), pp. 493–500. URL: `https://inspirehep.net/files/f55503250f690969aedfda4ceaf9b4f9`.

[84] Class for Physics of the Royal Swedish Academy of Sciences. *Topological phase transitions and topological phases of matter, Scientific Background on the Nobel Prize in Physics 2016.* Oct. 2016. URL: `https://www.nobelprize.org/uploads/2018/06/advanced-physicsprize2016.pdf`.

[85] M. Hasenbusch. "The two-dimensional XY model at the transition temperature : a high-precision Monte Carlo study". In: *Journal of Physics A* 38 (2005), pp. 5869–5883. DOI: `10.1088/0305-4470/38/26/003`.

[86] W. Bietenholz, U. Gerber, M. Pepe, and U. Wiese. "Topological lattice actions". In: *Journal of High Energy Physics* 2010 (2010), pp. 1–36. DOI: `10.1007/JHEP12(2010)020`.

[87] W. Bietenholz, M. Bögli, F. Niedermayer, M. Pepe, F. Rej'on-Barrera, and U.-J. Wiese. "Topological lattice actions for the 2d XY model". In: *Journal of*

High Energy Physics 2013 (2013), pp. 1–23. DOI: `10.1007/JHEP03(2013)141`.

[88] Pablo Serna, J. T. Chalker, and Paul Fendley. “Deconfinement transitions in a generalised XY model”. In: *J. Phys. A* 50.42 (2017), p. 424003. DOI: `10.1088/1751-8121/aa89a1`. arXiv: `1706.01475 [cond-mat.stat-mech]`.

[89] Duong Xuan Nui, Le Tuan, Nguyen Duc Trung Kien, Pham Thanh Huy, Hung T. Dang, and Dao Xuan Viet. “Correlation length in a generalized two-dimensional XY model”. In: *Phys. Rev. B* 98 (14 Oct. 2018), p. 144421. DOI: `10.1103/PhysRevB.98.144421`.

[90] Gabriel A. Canova, Yan Levin, and Jeferson J. Arenzon. “Competing nematic interactions in a generalized XY model in two and three dimensions”. In: *Phys. Rev. E* 94 (3 Sept. 2016), p. 032140. DOI: `10.1103/PhysRevE.94.032140`.

[91] Tam Le and Makoto Yamada. “Persistence Fisher Kernel: A Riemannian Manifold Kernel for Persistence Diagrams”. In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Inc., 2018, pp. 10007–10018. URL: `http://papers.nips.cc/paper/8205-persistence-fisher-kernel-a-riemannian-manifold-kernel-for-persistence-diagrams.pdf`.

[92] Ulli Wolff. “Collective Monte Carlo Updating for Spin Systems”. In: *Phys. Rev. Lett.* 62 (4 Jan. 1989), pp. 361–364. DOI: `10.1103/PhysRevLett.62.361`.

[93] Katharine Turner, Yuriy Mileyko, S. Mukherjee, and J. Harer. “Fréchet Means for Distributions of Persistence Diagrams”. In: *Discrete & Computational Geometry* 52 (2014), pp. 44–70. DOI: `10.1007/s00454-014-9604-7`.

[94] J. Greensite. *An Introduction to the Confinement Problem*. Lecture Notes in Physics. Springer Berlin Heidelberg, 2011. ISBN: 9783642143823.

[95] Andrea Palermo, Lucio Anderlini, Maria Paola Lombardo, Andrey Kotov, and Anton Trunin. *Machine learning approaches to the QCD transition*. 2021. DOI: `10.48550/ARXIV.2111.05216`. arXiv: `2111.05216 [hep-lat]`.

[96] Sebastian J. Wetzel and Manuel Scherzer. “Machine learning of explicit order parameters: From the Ising model to SU(2) lattice gauge theory”. In: *Phys. Rev. B* 96 (18 Nov. 2017), p. 184410. DOI: `10.1103/PhysRevB.96.184410`.

[97] D. L. Boyda, M. N. Chernodub, N. V. Gerasimeniuk, V. A. Goy, S. D. Liubimov, and A. V. Molochkov. “Finding the deconfinement temperature in lattice Yang-Mills theories from outside the scaling window with machine learning”. In: *Phys. Rev. D* 103 (1 Jan. 2021), p. 014509. DOI: `10.1103/PhysRevD.103.014509`.

[98] Cinzia Giannetti, Biagio Lucini, and Davide Vadacchino. “Machine Learning as a universal tool for quantitative investigations of phase transitions”. In: *Nucl. Phys. B* 944 (2019), p. 114639. DOI: `10.1016/j.nuclphysb.2019.114639`. arXiv: `1812.06726 [cond-mat.stat-mech]`.

[99] G. ’t Hooft. “On the phase transition towards permanent quark confinement”. In: *Nuclear Physics B* 138.1 (1978), pp. 1–25. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(78)90153-0`.

[100] John M. Cornwall. “Quark confinement and vortices in massive gauge-invariant QCD”. In: *Nuclear Physics B* 157.3 (1979), pp. 392–412. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(79)90111-1`.

[101] Gerard ’t Hooft. “A Property of Electric and Magnetic Flux in Nonabelian Gauge Theories”. In: *Nucl. Phys. B* 153 (1979), pp. 141–160. DOI: `10.1016/0550-3213(79)90595-9`.

[102] Philippe de Forcrand and Massimo D’Elia. “On the relevance of center vortices to QCD”. In: *Phys. Rev. Lett.* 82 (1999), pp. 4582–4585. DOI: `10.1103/PhysRevLett.82.4582`. arXiv: `hep-lat/9901020`.

[103] L. Del Debbio, Manfried Faber, J. Greensite, and S. Olejnik. “Center dominance and Z(2) vortices in SU(2) lattice gauge theory”. In: *Phys. Rev. D* 55 (1997), pp. 2298–2306. DOI: `10.1103/PhysRevD.55.2298`. arXiv: `hep-lat/9610005`.

[104] Manfried Faber, J. Greensite, and S. Olejnik. “Casimir scaling from center vortices: Towards an understanding of the adjoint string tension”. In: *Phys.*

Rev. D 57 (1998), pp. 2603–2609. DOI: `10.1103/PhysRevD.57.2603`. arXiv: `hep-lat/9710039`.

[105] L. Del Debbio, Manfried Faber, J. Giedt, J. Greensite, and S. Olejnik. “Detection of center vortices in the lattice Yang-Mills vacuum”. In: *Phys. Rev. D* 58 (1998), p. 094501. DOI: `10.1103/PhysRevD.58.094501`. arXiv: `hep-lat/9801027`.

[106] R. Bertle, Manfried Faber, J. Greensite, and S. Olejnik. “The Structure of projected center vortices in lattice gauge theory”. In: *JHEP* 03 (1999), p. 019. DOI: `10.1088/1126-6708/1999/03/019`. arXiv: `hep-lat/9903023`.

[107] Manfried Faber, J. Greensite, S. Olejnik, and Daiske Yamada. “The Vortex finding property of maximal center (and other) gauges”. In: *JHEP* 12 (1999), p. 012. DOI: `10.1088/1126-6708/1999/12/012`. arXiv: `hep-lat/9910033`.

[108] M. Engelhardt, K. Langfeld, H. Reinhardt, and O. Tennert. “Deconfinement in SU(2) Yang-Mills theory as a center vortex percolation transition”. In: *Phys. Rev. D* 61 (2000), p. 054504. DOI: `10.1103/PhysRevD.61.054504`. arXiv: `hep-lat/9904004`.

[109] M. Engelhardt and H. Reinhardt. “Center projection vortices in continuum Yang-Mills theory”. In: *Nucl. Phys. B* 567 (2000), p. 249. DOI: `10.1016/S0550-3213(99)00727-0`. arXiv: `hep-th/9907139`.

[110] Michael Engelhardt. “Center vortex model for the infrared sector of Yang-Mills theory: Topological susceptibility”. In: *Nucl. Phys. B* 585 (2000), p. 614. DOI: `10.1016/S0550-3213(00)00350-3`. arXiv: `hep-lat/0004013`.

[111] Roman Bertle, Manfried Faber, Jeff Greensite, and Stefan Olejnik. “P vortices, gauge copies, and lattice size”. In: *JHEP* 10 (2000), p. 007. DOI: `10.1088/1126-6708/2000/10/007`. arXiv: `hep-lat/0007043`.

[112] K. Langfeld, H. Reinhardt, and J. Gattnar. “Gluon propagators and quark confinement”. In: *Nucl. Phys. B* 621 (2002), pp. 131–156. DOI: `10.1016/S0550-3213(01)00574-0`. arXiv: `hep-ph/0107141`.

[113] J. Greensite. “The Confinement problem in lattice gauge theory”. In: *Prog. Part. Nucl. Phys.* 51 (2003), p. 1. DOI: `10.1016/S0146-6410(03)90012-3`. arXiv: `hep-lat/0301023`.

[114] Falk Bruckmann and Michael Engelhardt. "Writhe of center vortices and topological charge: An Explicit example". In: *Phys. Rev. D* 68 (2003), p. 105011. DOI: `10.1103/PhysRevD.68.105011`. arXiv: `hep-th/0307219`.

[115] M. Engelhardt, M. Quandt, and H. Reinhardt. "Center vortex model for the infrared sector of SU(3) Yang-Mills theory: Confinement and deconfinement". In: *Nucl. Phys. B* 685 (2004), pp. 227–248. DOI: `10.1016/j.nuclphysb.2004.02.036`. arXiv: `hep-lat/0311029`.

[116] P. Yu. Boyko, V. G. Bornyakov, E. -M. Ilgenfritz, A. V. Kovalenko, B. V. Martemyanov, M. Muller-Preussker, M. I. Polikarpov, and A. I. Veselov. "Once more on the interrelation between Abelian monopoles and P-vortices in SU(2) LGT". In: *Nucl. Phys. B* 756 (2006), pp. 71–85. DOI: `10.1016/j.nuclphysb.2006.08.025`. arXiv: `hep-lat/0607003`.

[117] Ernst-Michael Ilgenfritz, Karl Koller, Yoshiaki Koma, Gerrit Schierholz, Thomas Streuer, Volker Weinberg, and Markus Quandt. "Localization of overlap modes and topological charge, vortices and monopoles in SU(3) LGT". In: *PoS* LATTICE2007 (2007). Ed. by Gunnar Bali, Vladimir Braun, Christof Gattringer, Meinulf Gockeler, Andreas Schafer, Peter Weisz, and Tilo Wettig, p. 311. DOI: `10.22323/1.042.0311`. arXiv: `0710.2607 [hep-lat]`.

[118] V. G. Bornyakov, E. -M. Ilgenfritz, B. V. Martemyanov, S. M. Morozov, M. Muller-Preussker, and A. I. Veselov. "Interrelation between monopoles, vortices, topological charge and chiral symmetry breaking: Analysis using overlap fermions for SU(2)". In: *Phys. Rev. D* 77 (2008), p. 074507. DOI: `10.1103/PhysRevD.77.074507`. arXiv: `0708.3335 [hep-lat]`.

[119] Alan O'Cais, Waseem Kamleh, Kurt Langfeld, Ben Lasscock, Derek Leinweber, Peter Moran, Andre Sternbeck, and Lorenz von Smekal. "Preconditioning Maximal Center Gauge with Stout Link Smearing in SU(3)". In: *Phys. Rev. D* 82 (2010), p. 114512. DOI: `10.1103/PhysRevD.82.114512`. arXiv: `0807.0264 [hep-lat]`.

[120] Michael Engelhardt. "Center vortex model for the infrared sector of SU(3) Yang-Mills theory: Topological susceptibility". In: *Phys. Rev. D* 83 (2011), p. 025015. DOI: `10.1103/PhysRevD.83.025015`. arXiv: `1008.4953 [hep-lat]`.

[121] Patrick O. Bowman, Kurt Langfeld, Derek B. Leinweber, Andre Sternbeck, Lorenz von Smekal, and Anthony G. Williams. “Role of center vortices in chiral symmetry breaking in SU(3) gauge theory”. In: *Phys. Rev. D* 84 (2011), p. 034501. DOI: `10.1103/PhysRevD.84.034501`. arXiv: `1010.4624 [hep-lat]`.

[122] Elyse-Ann O’Malley, Waseem Kamleh, Derek Leinweber, and Peter Moran. “SU(3) centre vortices underpin confinement and dynamical chiral symmetry breaking”. In: *Phys. Rev. D* 86 (2012), p. 054503. DOI: `10.1103/PhysRevD.86.054503`. arXiv: `1112.2490 [hep-lat]`.

[123] Daniel Trewartha, Waseem Kamleh, and Derek Leinweber. “Connection between center vortices and instantons through gauge-field smoothing”. In: *Phys. Rev. D* 92.7 (2015), p. 074507. DOI: `10.1103/PhysRevD.92.074507`. arXiv: `1509.05518 [hep-lat]`.

[124] Daniel Trewartha, Waseem Kamleh, and Derek Leinweber. “Evidence that centre vortices underpin dynamical chiral symmetry breaking in SU(3) gauge theory”. In: *Phys. Lett. B* 747 (2015), pp. 373–377. DOI: `10.1016/j.physletb.2015.06.025`. arXiv: `1502.06753 [hep-lat]`.

[125] Jeff Greensite. “Confinement from Center Vortices: A review of old and new results”. In: *EPJ Web Conf.* 137 (2017). Ed. by Y. Foka, N. Brambilla, and V. Kovalenko, p. 01009. DOI: `10.1051/epjconf/201713701009`. arXiv: `1610.06221 [hep-lat]`.

[126] Daniel Trewartha, Waseem Kamleh, and Derek B. Leinweber. “Centre vortex removal restores chiral symmetry”. In: *J. Phys. G* 44.12 (2017), p. 125002. DOI: `10.1088/1361-6471/aa9443`. arXiv: `1708.06789 [hep-lat]`.

[127] James C. Biddle, Waseem Kamleh, and Derek B. Leinweber. “Gluon propagator on a center-vortex background”. In: *Phys. Rev. D* 98.9 (2018), p. 094504. DOI: `10.1103/PhysRevD.98.094504`. arXiv: `1806.04305 [hep-lat]`.

[128] Felix Spengler, Markus Quandt, and Hugo Reinhardt. “Branching of Center Vortices in SU(3) Lattice Gauge Theory”. In: *Phys. Rev. D* 98.9 (2018), p. 094508. DOI: `10.1103/PhysRevD.98.094508`. arXiv: `1810.04072 [hep-th]`.

[129] James Biddle, Waseem Kamleh, and Derek Leinweber. *Static quark potential from centre vortices in the presence of dynamical fermions*. June 2022. DOI: `10.48550/ARXIV.2206.00844`. arXiv: `2206.00844 [hep-lat]`.

[130] L. Del Debbio, A. Di Giacomo, and B Lucini. "Vortices, monopoles and confinement". In: *Nucl. Phys. B* 594 (2001), pp. 287–300. DOI: `10.1016/S0550-3213(00)00651-9`. arXiv: `hep-lat/0006028`.

[131] L. Del Debbio, A. Di Giacomo, and B. Lucini. "Monopoles, vortices and confinement in SU(3) gauge theory". In: *Phys. Lett. B* 500 (2001), pp. 326–329. DOI: `10.1016/S0370-2693(01)00091-0`. arXiv: `hep-lat/0011048`.

[132] Rudolf Golubich and Manfried Faber. "A Possible Resolution to Troubles of SU(2) Center Vortex Detection in Smooth Lattice Configurations". In: *Universe* 7.5 (2021). ISSN: 2218-1997. DOI: `10.3390/universe7050122`.

[133] John D. Stack and William W. Tucker. "The Gribov ambiguity for maximal abelian and center gauges in SU(2) lattice gauge theory". In: *Nuclear Physics B - Proceedings Supplements* 94.1 (2001). Proceedings of the XVIIIth International Symposium on Lattice Field Theory, pp. 529–531. ISSN: 0920-5632. DOI: `https://doi.org/10.1016/S0920-5632(01)00903-3`.

[134] V.N. Gribov. "Quantization of non-Abelian gauge theories". In: *Nuclear Physics B* 139.1 (1978), pp. 1–19. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(78)90175-X`.

[135] P. de Forcrand and M. Pepe. "Center vortices and monopoles without lattice Gribov copies". In: *Nucl. Phys. B* 598 (2001), pp. 557–577. DOI: `10.1016/S0550-3213(01)00009-8`. arXiv: `hep-lat/0008016`.

[136] M. Engelhardt, K. Langfeld, H. Reinhardt, and O. Tennert. "Deconfinement in SU(2) Yang-Mills theory as a center vortex percolation transition". In: *Phys. Rev. D* 61 (5 Feb. 2000), p. 054504. DOI: `10.1103/PhysRevD.61.054504`.

[137] Greensite, Jeff. "Confinement from Center Vortices: A review of old and new results". In: *EPJ Web Conf.* 137 (2017), p. 01009. DOI: `10.1051/epjconf/201713701009`.

[138] Manfried Faber, Jeff Greensite, and Stefan Olejník. "Direct Laplacian Center Gauge". In: *Journal of High Energy Physics* 2001.11 (Nov. 2001), pp. 053–053. DOI: `10.1088/1126-6708/2001/11/053`.

[139] Biagio Lucini, Michael Teper, and Urs Wenger. "The high temperature phase transition in SU(N) gauge theories". In: *Journal of High Energy Physics* 2004 (July 2003). DOI: `10.1088/1126-6708/2004/01/061`.

[140] Benjamin Svetitsky and Laurence G. Yaffe. "Critical behavior at finite-temperature confinement transitions". In: *Nuclear Physics B* 210.4 (1982), pp. 423–447. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(82)90172-9`.

[141] Filip Kos, David Poland, David Simmons-Duffin, and Alessandro Vichi. "Precision Islands in the Ising and $O(N)$ Models". In: *Journal of High Energy Physics* 2016 (Mar. 2016). DOI: `10.1007/JHEP08(2016)036`.

[142] G. 't Hooft. "A property of electric and magnetic flux in non-Abelian gauge theories". In: *Nuclear Physics B* 153 (1979), pp. 141–160. ISSN: 0550-3213. DOI: `https://doi.org/10.1016/0550-3213(79)90595-9`.

[143] E. T. Tomboulis and L. G. Yaffe. "Finite temperature SU(2) lattice gauge theory". In: *Communications in Mathematical Physics* 100.3 (Sept. 1985), pp. 313–341. ISSN: 1432-0916. DOI: `10.1007/BF01206134`.

[144] Kouji Kashiwa, Takehiro Hirakida, and Hiroaki Kouno. *Persistent homology analysis for dense QCD effective model with heavy quarks*. 2021. DOI: `10.48550/ARXIV.2103.12554`. arXiv: `2103.12554 [hep-lat]`.

[145] Luigi Del Debbio, Agostino Patella, and Claudio Pica. "Higher representations on the lattice: Numerical simulations, SU(2) with adjoint fermions". In: *Phys. Rev. D* 81 (9 May 2010), p. 094503. DOI: `10.1103/PhysRevD.81.094503`.

[146] Michael Creutz. "Monte Carlo study of quantized SU(2) gauge theory". In: *Phys. Rev. D* 21 (8 Apr. 1980), pp. 2308–2315. DOI: `10.1103/PhysRevD.21.2308`.

[147] Stephen L. Adler. "Overrelaxation algorithms for lattice field theories". In: *Phys. Rev. D* 37 (2 Jan. 1988), pp. 458–471. DOI: `10.1103/PhysRevD.37.458`.

[148] David Young. “Iterative Methods for Solving Partial Difference Equations of Elliptic Type”. In: *Transactions of the American Mathematical Society* 76.1 (1954), pp. 92–111. ISSN: 00029947. URL: `http://www.jstor.org/stable/1990745` (visited on 08/04/2022).

[149] Martin Lüscher. “Properties and uses of the Wilson flow in lattice QCD”. In: *Journal of High Energy Physics* 2010 (2010), pp. 1–18. DOI: `10.1007/JHEP08(2010)071`.

[150] Ed Bennett, Jack Holligan, Deog Ki Hong, Ho Hsiao, Jong-Wan Lee, C. J. David Lin, Biagio Lucini, Michele Mesiti, Maurizio Piai, and Davide Vadacchino. “Progress in $Sp(2N)$ lattice gauge theories”. In: *PoS* LATTICE2021 (2022), p. 308. DOI: `10.22323/1.396.0308`. arXiv: `2111.14544 [hep-lat]`.

[151] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. “Stability of Persistence Diagrams”. In: vol. 37. Jan. 2005, pp. 263–271. DOI: `10.1007/s00454-006-1276-5`.

[152] Gerard ’t Hooft and Falk Bruckmann. *Monopoles, Instantons and Confinement*. 2000. DOI: `10.48550/ARXIV.HEP-TH/0010225`. arXiv: `0010225 [hep-th]`.

[153] Bart Olsthoorn and Alexander V. Balatsky. *Persistent homology of quantum entanglement*. 2021. DOI: `10.48550/ARXIV.2110.10214`. arXiv: `2110.10214 [cond-mat.stat-mech]`.

[154] Daniel Spitz, Jürgen Berges, Markus Oberthaler, and Anna Wienhard. “Finding self-similar behavior in quantum many-body dynamics via persistent homology”. In: *SciPost Physics* 11.3 (Sept. 2021). DOI: `10.21468/scipostphys.11.3.060`.

[155] Ka Man Yim and Vidit Nanda. *Topological Inference of the Conley Index*. 2022. DOI: `10.48550/ARXIV.2206.10198`. arXiv: `2206.10198 [math.DS]`.

[156] Alvaro Torras Casas. *PerMaViss: Persistence Mayer Vietoris spectral sequence*. Version v0.2. Aug. 2021. DOI: `10.5281/zenodo.5266475`.

[157] Alvaro Torras Casas. *Distributing Persistent Homology via Spectral Sequences*. 2019. DOI: `10.48550/ARXIV.1907.05228`. arXiv: `1907.05228 [math.AT]`.

[158] Casper Gyurik, Chris Cade, and Vedran Dunjko. *Towards quantum advantage via topological data analysis*. 2020. DOI: `10.48550/ARXIV.2005.02607`. arXiv: `2005.02607 [quant-ph]`.

[159] Lars Onsager. "Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition". In: *Phys. Rev.* 65 (3-4 Feb. 1944), pp. 117–149. DOI: `10.1103/PhysRev.65.117`.

[160] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. "Equation of State Calculations by Fast Computing Machines". In: *The Journal of Chemical Physics* 21.6 (June 1953), pp. 1087–1092. DOI: `10.1063/1.1699114`.

[161] Frédéric Meunier. "Polytopal complexes: maps, chain complexes and necklaces". In: *Electronic Notes in Discrete Mathematics* 31 (2008). The International Conference on Topological and Geometric Graph Theory, pp. 183–188. ISSN: 1571-0653. DOI: `https://doi.org/10.1016/j.endm.2008.06.037`.

[162] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational Homology*. Applied Mathematical Sciences. Springer New York, 2004. ISBN: 9780387408538. URL: `https://books.google.nl/books?id=AShKtpi3GecC`.

[163] M. E. J. Newman and G. T. Barkema. *Monte Carlo methods in statistical physics*. Oxford: Clarendon Press, 1999.

[164] Jonas Greitemann, Ke Liu, and Lode Pollet. "Probing hidden spin order with interpretable machine learning". In: *Phys. Rev. B* 99 (6 Feb. 2019), p. 060404. DOI: `10.1103/PhysRevB.99.060404`.

[165] Peter Bubenik and J. Scott. "Categorification of Persistent Homology". In: *Discrete & Computational Geometry* 51 (2014), pp. 600–627. DOI: `10.1007/s00454-014-9573-x`.

[166] Ulli Wolff. "Collective Monte Carlo Updating for Spin Systems". In: *Phys. Rev. Lett.* 62 (1989), p. 361. DOI: `10.1103/PhysRevLett.62.361`.

[167] B. Efron. "Bootstrap Methods: Another Look at the Jackknife". In: *Ann. Statist.* 7.1 (Jan. 1979), pp. 1–26. DOI: `10.1214/aos/1176344552`.

[168] Sacha Friedli and Yvan Velenik. *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, 2017. ISBN: 978-1-107-18482-4. DOI: `10.1017/9781316882603`.

[169] R. Peierls and M. Born. "On Ising's model of ferromagnetism". In: *Proceedings of the Cambridge Philosophical Society* 32.3 (Jan. 1936), p. 477. DOI: `10.1017/S0305004100019174`.

[170] James C. Biddle, Waseem Kamleh, and Derek B. Leinweber. "Visualization of center vortex structure". In: *Phys. Rev. D* 102 (3 Aug. 2020), p. 034504. DOI: `10.1103/PhysRevD.102.034504`.

[171] Herbert Edelsbrunner and John Harer. *Computational Topology - an Introduction.* American Mathematical Society, 2010, pp. I–XII, 1–241. ISBN: 978-0-8218-4925-5.

[172] L. Del Debbio, M. Faber, J. Giedt, J. Greensite, and Š. Olejník. "Detection of center vortices in the lattice Yang-Mills vacuum". In: *Phys. Rev. D* 58 (9 Sept. 1998), p. 094501. DOI: `10.1103/PhysRevD.58.094501`.

[173] Robert Ghrist. "Barcodes: The persistent topology of data". In: *BULLETIN (New Series) OF THE AMERICAN MATHEMATICAL SOCIETY* 45 (Feb. 2008). DOI: `10.1090/S0273-0979-07-01191-3`.

[174] Nina Otter, Mason Porter, Ulrike Tillmann, Peter Grindrod, and Heather Harrington. "A roadmap for the computation of persistent homology". In: *EPJ Data Science* 6 (June 2015). DOI: `10.1140/epjds/s13688-017-0109-5`.

[175] Michael R. Shirts and John D. Chodera. "Statistically optimal analysis of samples from multiple equilibrium states." In: *The Journal of chemical physics* 129 12 (2008), p. 124105. DOI: `10.1063/1.2978177`.

[176] Philippe de Forcrand, Massimo D'Elia, and Michele Pepe. "A Study of the 't Hooft loop in SU(2) Yang-Mills theory". In: *Phys. Rev. Lett.* 86 (2001), p. 1438. DOI: `10.1103/PhysRevLett.86.1438`. arXiv: `hep-lat/0007034`.