TOPOLOGICAL METHODS FOR PROBING AND ENGINEERING QUANTUM DYNAMICS

HARVEY FENGRUI CAO

M.Sci, Imperial College London

A THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

CENTRE FOR QUANTUM TECHNOLOGIES

NATIONAL UNIVERSITY OF SINGAPORE

2026

Advisor:

Dimitris G. Angelakis

Declaration

I hereby declare that this thesis is my original work and it has been written by me in its entirety. I have duly acknowledged all the sources of information which have been used in the thesis.

This thesis has also not been submitted for any degree in any university previously.

Harvey Cao

Harvey Fengrui Cao

July 2026

Acknowledgments

This thesis marks the close of one of the most significant chapters of my life, a period in which I have grown considerably, both as a scientist and as a human being. The path that this journey has taken has been anything but linear, marked by countless ebbs and flows, triumphs and setbacks along the way. Looking back, none of this growth would have been possible without the unwavering support of all those who lifted my spirits during the most challenging moments of my PhD, and it is to them that I owe my deepest gratitude.

First of all, I would like to thank my PhD advisor, Dimitris Angelakis, for giving me the opportunity to embark on this journey. From the very beginning, he has been nothing but supportive and understanding, and has given me the confidence and freedom to explore my ideas independently whilst offering guidance whenever I needed it most. I would also like to express my sincere gratitude to Daniel Leykam. He has been a mentor to me throughout my entire PhD, and someone who has been exceptionally generous with his time and expertise. It would not be a stretch to say that every project in this thesis grew out of the countless discussions we had over the whiteboard on the office wall. Watching the thorough, methodical way he approaches every problem has taught me what it means to be a strong and effective researcher.

I also wish to thank all those at CQT who made it a welcoming and inspiring place to work, and for the many lively discussions that made the long hours more memorable, and considerably easier to bear. In particular, I would like to thank Arthur and Yixian, with whom I endured and conquered the most impenetrable of tutorials and problem sheets; Rishabh, who always fascinated me with stories of his adventures over free graduate student club dinners; Beng Yee, who made me realise that my wisdom tooth pain was not something to be ignored; Ben, Supanut, Muhammad and Chee, whose light-hearted chats in the office provided me with lots of comfort throughout the years. My thanks also go to the wonderful staff at CQT, who quietly kept everything running smoothly and were always ready to help, no matter how last-minute the request for poster printing was.

Singapore has also given me many lifelong friends who I want to thank for making this place feel like a home away from home: Nam and Fahad, for being my brothers-in-crime; Lucy, who I could talk to about the importance of Gregg's sausage rolls when no-one else understood; and Siying, for trying her best to get me into bouldering, even if unsuccessful. A special thanks go to Mr. and Mrs. Shen, who treated me like their own son from the moment they picked me up from my quarantine hotel, and Mr. and Mrs. Ramsay King, whose fantastic Christmas celebrations are something I look forward to every single year.

To my parents, grandparents and family: thank you for your unconditional love and support. Wherever you are, in this world or beyond it, thank you for believing in me even when I doubted myself, and I hope that this thesis makes you proud.

Finally, I would like to thank my incredible ~~girlfriend~~ ~~fiancée~~ wife, Aria, for always being my bright burning star. You gave me energy, strength and love in the moments that I needed it the most.

Summary

Topological data analysis (TDA) provides a natural language for characterising the shape of complex, high-dimensional data in a manner that is robust to noise and agnostic to the choice of coordinates. This thesis motivates and examines the use of TDA in extracting and analysing topological features relevant to models from quantum many-body physics. We systematically develop TDA-based pipelines to explore three diverse quantum phenomena, demonstrating the utility and breadth of application that a topological perspective can provide.

First, we use persistent homology to detect signatures of chaos in a nonlinear bosonic cavity. By reconstructing dynamical attractors from experimentally accessible time series and analyzing their topology, this approach distinguishes regular from chaotic regimes even under noisy and sparse data conditions. Next, we introduce an unsupervised learning approach for identifying quantum many-body scar initial states through intrinsic dimension estimation. This reveals a clear separation between low-dimensional scar dynamics and higher-dimensional thermalizing behavior without requiring labeled data or prior knowledge of scar subspaces. Finally, we propose a novel nonlinear Schrödinger model whose nonlinear response depends on the topological structure of the wavefunction intensity profile. We show analytically and numerically that this model supports cusp-like solitons and stabilizes flat-top beams against modulational instability, demonstrating how TDA-based terms inserted as an energy penalty can serve to stabilize novel self-organized states against modulational instability.

Contents

List of Figures

Chapter 1

Introduction and Motivation

1.1 The Shape of Physics

Topology, the mathematical study of properties preserved under continuous deformations, has reshaped the way we study and classify complex systems in modern physics. The key inspiration comes from searching for quantities that measure global properties that belong to the system as a whole, rather than local properties that can be determined by measurements at individual points. The robustness of these topological quantities to perturbations and their inability to vary continuously results in a greater stability of the system's properties than might otherwise be expected, and has enabled the prediction and explanation of many surprisingly robust physical phenomena.

The prominence of topology can be found across many fields, to the point where it has now become ubiquitous in physics. These range from early appearances in classical mechanics, electrodynamics and fluid dynamics [1–4], to more recent applications in quantum physics, condensed matter, optics, and particle physics [5–11]. Perhaps the greatest transformative impact of topology was seen in the study of two-dimensional electronic systems in the 1980s, where the observed quantisation of Hall conductivity was explained using a novel topological phase of matter, the quantum Hall phase [12–14]. This phenomenon is governed by an invariant known as the Chern number, whose topological nature ensures that the quantisation remains extremely precise (accuracy better than one part in a billion) and robust against disorder and sample impurities.

Topology lends itself exceptionally well to quantum physics because while quantum systems suffer inherently from decoherence and disorder, topological features encode information non-locally, offering protection for the physical properties of interest and a way to classify phases and states by global, robust invariants. In many physical settings, such as continuous media or lattices in the thermodynamic limit, topological concepts naturally align with the underlying mathematical framework built upon continuous functions and smooth manifolds. However, there are also cases where the direct application of topological language becomes less straightforward, such as small or discrete systems. Other fields of science face similar challenges with continuity, frequently having to work with sparse data and limited observations in high-dimensional parameter spaces. Nevertheless, despite this very different setting, topological approaches have demonstrated remarkable power and versatility across the sciences [15–17].

1.2 Topology and Machine learning

In parallel with rapid experimental advancements and increasingly vast data, there has been a burgeoning trend in recent years to harness machine learning and other data analysis methods to help with handling the "complexity" that comes with tackling challenging problems which cannot be readily simplified. This complexity can arise from a myriad of factors such as high-dimensionality, nonlinearity, multiscale dynamics, and large numbers of interacting components and mechanisms. Another layer of complexity comes with extracting real-world data from experiment, introducing complications like noise, sparsity and erroneous data.

Neural network models have seen an increasing popularity in tackling such problems, owing to their flexibility across many different types of problems and yielding excellent results given the right training data [18, 19]. However, neural networks are computationally expensive to train, require extensive datasets, and are generally hard to interpret in a meaningful way because the causal relationships that leads the algorithm to their predictions cannot be clearly determined. The opacity of black-box architectures such as these motivate the use of more explainable techniques with clear theoretical foundations, including unsupervised learning methods such as cluster analysis and feature-mapping, which typically leverage domain-specific knowledge and physical intuition behind the problem at hand [20, 21]. Both of these techniques utilize data abstraction to capture the core, relevant structure of the data without significant information loss. They can often fall short, however, at identifying long range order or capturing topological features such as holes and voids.

Topological data analysis (TDA) is a powerful framework that emerged in response to the various problems unresolved by existing unsupervised machine learning methods. It encompasses a suite of computational tools for identifying nonlocal features of discrete high-dimensional datasets, in a fashion that is robust to noise and insensitive to initial parameter and metric choices [22, 23]. In essence, it allows us to systematically define and explore the "shape" of data. Developed in the early 2000s out of rigorous foundations in algebraic topology, TDA was originally applied to problems in life science and data science, where its robustness to noise and ability to capture global features made it suited to analyzing complex, high-dimensional datasets such as biological images, gene expression profiles, and sensor networks [24–27].

In part due to unfamiliarity with its discrete spaces and inaccessibility from lack of pedagogical reviews, TDA was under-utilised by the physics community in its early years. Recently, however, interest in TDA methods has been growing among physicists, following a growing body of accessible literature and successful application across many diverse problems [28, 29]. In condensed matter and materials physics, TDA has been employed to identify and characterise amorphous and glassy phases, providing robust shape descriptors that distinguish structural differences not easily captured by conventional correlation functions [30–32]. In quantum information, TDA-based invariants have been used to classify multiqubit entangled states, offering a basis-independent perspective on entanglement structure [33, 34]. In the study of dynamical and chaotic systems, persistent homology applied to time-series embeddings has revealed qualitative changes in the underlying attractor topology, enabling the detection of regime shifts and topological signatures of phase transitions that are not apparent from standard scalar observables alone [35–37].

In recent years, TDA has also seen significant interest within the domain of quantum algorithms, where quantum linear algebra subroutines such as quantum phase estimation and HHL-type algorithms have been used to compute topological invariants and persistent summaries [38, 39]. Polynomial speedups over classical counterparts have been demonstrated under specific conditions, and proof-of-concept implementations on near-term devices have been reported; however, the practical reach of these methods remains heavily constrained by restrictions on the input data and the structure of the topological complexes, leaving the question of genuine practical quantum advantage in TDA open [40–43]. Nevertheless, certain TDA problems, such as computing numbers of topological features, have been shown to be provably classically intractable, which provides compelling motivation towards continued research in quantum TDA despite the current practical limitations.

1.3 Thesis Motivation and Outline

Feature selection is an important process in the machine learning pipeline, especially for physical applications where data is complex, high-dimensional and riddled with challenges such as erroneous measurement and sparsity. In such cases, the deliberate and careful selection of features using domain-specific expertise becomes essential to identify those that are most relevant and informative to the problem at hand. In addition to improving the performance of models and enhancing the interpretability of results, feature selection has become an important consideration for reducing computational overhead [44]. This is especially prevalent within the context of modern day AI and large language model usage, where heavy energy consumption motivates better feature selection and more efficient architectures which work with information encoded in structured forms such as (hyper)graphs rather than vectors [45, 46].

The central narrative of this thesis is an examination of how quantum systems can encode topological information and to what extent TDA provides a suitable lens for feature selection and feature extraction from these systems in a robust manner. We systematically develop TDA-based methods to explore diverse quantum phenomena, showcasing the breadth and versatility of topological features for novel machine learning applications. The recent emergence of efficient quantum implementations for TDA have attracted considerable interest as a new set of scalable tools, but unfamiliarity with these methods has inhibited widespread adoption by physicists. Our work intends to lay out a pedagogical foundation on topological feature selection techniques to reveal those that best suit problems encountered across quantum physics. In a similar manner to how GPU- or neural network-accelerated simulations of various physical systems each have their own constraints for the specific problem classes they are optimal for, the efficacy of TDA is critically dependent on structural insights to the problem at hand [47–49].

With this objective in mind, in this thesis we pursue three distinct yet complementary lines of research in quantum physics, all of which present complex models that challenge standard analytical approach. The first is a Kerr-nonlinear model whose classical analog exhibits chaotic dynamics; the second is a quantum many-body system which violates typical thermalization behaviour; the third is a nonlinear Schrödinger model hosting localized, self-reinforcing wavepackets. Through taking a topological perspective of the observed phenomena, we provide a new way of understanding complex many-body quantum systems that is beneficial both at the foundational level and in experimental application.

The thesis is structured as follows;

In Chapter 2, we provide a pedagogical introduction to TDA and develop the necessary background on its core methods and techniques used throughout this thesis, with a particular emphasis on persistent homology. This chapter also briefly discusses the unique challenges that quantum many-body systems pose for experimental data analysis. It concludes with a review of key works applying topological features to machine learning applications in physics, establishing the contextual foundation for our contributions.

In Chapter 3, we employ persistent homology to diagnose quantum chaos in a driven Kerr-nonlinear open quantum system, presenting a novel trajectory-based approach that constructs topological data directly from experimentally accessible time series measurements. Our analysis reveals clear topological signatures distinguishing regular from chaotic regimes, even under sparse, noisy measurement conditions. This work establishes persistent homology as a robust, coordinate-free diagnostic that complements spectral methods and is feasible with current quantum optical platforms.

In Chapter 4, we examine quantum many-body scars through an unsupervised learning lens, asking whether scar and thermal initial states can be distinguished directly from dynamical data without prior knowledge of the scar subspace or labelled training sets. Using intrinsic dimension estimation to probe the topology of system dynamics, we identify a sharp distinction between the low-dimensional structure of scar states and high-dimensional thermal bulk. This data-driven approach provides a powerful complement to traditional spectral and entanglement-based diagnostics, highlighting the intrinsic topological signatures of quantum many-body scars.

In Chapter 5, we propose a novel topological nonlinearity in the nonlinear Schrödinger equation, motivated by recent experimental progress in nonlocal nonlinearities and density-dependent gauge fields. The nonlinear response of our model emerges from the topological structure of wavefunction intensity extrema, directly linked to quantities in persistent homology. This work explores its dynamical characteristics, demonstrating numerically and analytically the formation of cusp solitons and stabilisation of flat-top beams against modulational instability.

Chapter 6 summarizes the contributions of this thesis and outlines potential future research directions, including both natural extensions of the present work and more speculative applications.

Chapter 2

Background

2.1 Topological Data Analysis

Topology studies the global quantities of a system that remain unchanged under smooth, continuous transformations such as stretching, bending, or twisting; that is, without tearing, gluing, or passing through itself. Since it throws away intuitive concepts like coordinates and distances, concerning itself instead with properties like continuity and connectedness, it is widely regarded as being one of the most abstract fields in mathematics. This extreme generality is exactly what allows it to apply simultaneously across many domains, such as geometry, analysis, and dynamical systems; it has also become ubiquitous in modern physics.

Deeply rooted in techniques from topology, topological data analysis (TDA) is an emerging field in applied mathematics which provides a general framework to robustly describe the "shape" of data [22, 23]. In contrast to the smooth manifolds and infinite-dimensional vector spaces found in pure topology, TDA handles datasets which are finite, discrete, noisy, and high-dimensional; complex data of this type is universally found across all physical sciences. Through geometric simplification and implicit dimensionality reduction, TDA offers a systematic way to capture higher-order structural relationships within a dataset. This framework does not rely on predefined models or assumptions about functional form; instead, it builds insight from the connectivity of the data itself in an unsupervised manner.

There currently exist a wide variety of tools within the TDA framework, optimised for specific data types, computational constraints, and analytical goals [50, 51]. Nonetheless, they all largely follow the same pipeline which is to take a finite collection of distinct observations, assemble them together based on some notion of proximity, and arrive at a quantitative summary of the topology of the dataset. The primary focus of the following discussion will be on persistent homology, the principal tool of TDA, but there will also be comments made on additional approaches relevant to the present study and prominent within the literature.

Topological preliminaries

At the most fundamental level, we can distinguish between many topological spaces based on a few elementary properties like dimension and connectedness. For example, dimension allows us to distinguish between a (1-dimensional) line and a (2-dimensional) surface; connectedness differentiates between a (connected) circle and two non-intersecting (disconnected) circles. There are fundamental differences in the geometric properties of the two spaces which prevent one from being continuously deformed into the other without illegal actions such as gluing.

In many cases, however, these simple topological properties are not enough to fully encapsulate the difference between two spaces. Consider, for example, a sphere S^2 and a torus T^2 (shown in Fig. 2.1) which can be visualized as the surfaces of a ball and donut respectively. While we intuitively understand that these two shapes are different, both are 2-dimensional surfaces and both are connected, as you can travel anywhere on either surface without jumping gaps. The extra "hole" in the torus carries a different topological property where dimension and connectedness alone fail. This is known as the homology of a topological space, which describes

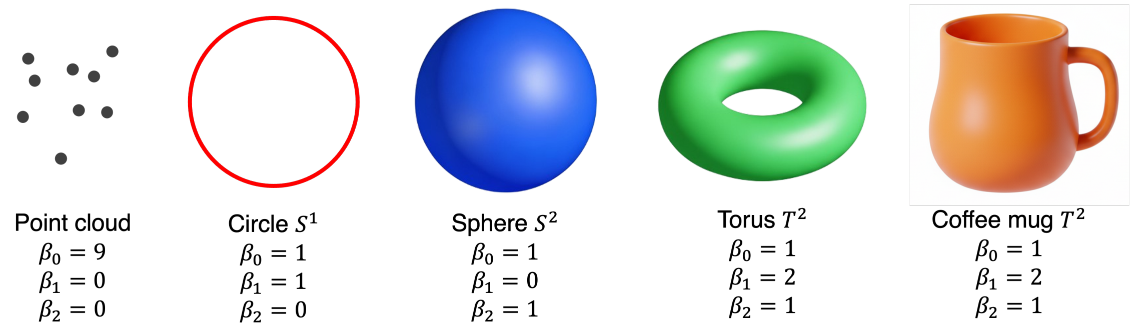

Figure 2.1: Illustration of topological spaces and their corresponding Betti numbers. From left to right: a point cloud of nine disconnected points; a circle with one connected component and one loop; a sphere with one connected component and one enclosed void; a torus with one connected component, two independent loops and one enclosed void; and a coffee mug which is topologically equivalent to the torus.

the topological features in it which cannot be continuously shrunk away or filled, such as holes and voids.

In persistent homology, we are most interested in counting the number of these k-dimensional holes present, where the dimension of a hole is defined to be the dimension of its boundary. For example, the space enclosed by a circle is a 1-dimensional hole since the boundary of the space is a 1-dimensional object (the circle itself). The number of these k-dimensional topological features is given by the kth Betti number B_k:

- The 0th Betti number β_0 counts the number of independent connected components.
- The 1st Betti number β_1 counts the number of non-contractible loops or "holes".
- Higher Betti numbers ($k > 1$) represent higher-dimensional holes. The Betti numbers for all holes with higher dimension than that of the space d is $\beta_{k>d} = 0$

Now, the Betti numbers provide us with a topological invariant that enables the distinction between the sphere and torus in our earlier example; the sphere has Betti numbers $(\beta_0 = 1, \beta_1 = 0, \beta_2 = 1)$ while the torus has Betti numbers $(\beta_0 = 1, \beta_1 = 2, \beta_2 = 1)$. The homology and corresponding Betti numbers, while still not a full topological classification, are sufficient for distinguishing spaces encountered in the majority of applications. Similarly, spaces that appear geometrically and visually distinct can share identical topologies, e.g. the popular equivalence between a coffee mug and a donut (see Fig. 2.1).

In general, computing the Betti numbers for a given topological space is highly non-trivial. Algebraic topology addresses this challenge by approximating spaces with simplicial complexes, which are assembled out of mathematically well-behaved building blocks. In the following sections, we shall see how physical data is used to construct simplicial complexes in TDA which reliably capture the topological features of the underlying manifold from which the data was sampled.

Point clouds

TDA can be applicable to any dataset that reflects discrete samples of an underlying continuous manifold. This normally takes the form of a high-dimensional point cloud $X = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$ embedded within a finite metric space (X, ρ), where $\rho : X \times X \to \mathbb{R}_{\geq 0}$ equips the space with a meaningful similarity measure. This metric can either be the ambient space of the data, such as Euclidean distance in $\mathbb{R}^d$, or one created using domain-specific knowledge. Note that careful selection of the metric often plays a critical role in the resulting topological inferences.

In physics, there is an abundance of opportunity for finding such types of point cloud data. For example, atomic configurations form spatial point clouds in the study of microscopic structures in materials. In fluid dynamics and high-energy physics simulations, the trajectories of particles generate discrete samples in phase space, amenable to topological analysis of vortices, jets, and collision substructures. Similarly, we can find data in more abstract settings that do not possess the same intuitive notion of shape and distance. For example, in the study of multipartite entanglement, point clouds emerge from interpreting individual systems as data points and using bipartite mutual information as a correlation measure.

Simplicial complexes

The next step is to build a continuous topological representation of the data, known as a simplicial complex. This is a collection of simplices, where a k-dimensional simplex is defined to be the convex hull of $k+1$ affinely independent points $v_0, ...v_k$. By an affinely independent set of $k+1$ points, we mean that they do not lie in some hyperplane of dimension less than k, which is equivalent to requiring that the k vectors $v_1 - v_0, ..., v_k - v_0$ are linearly independent. Geometrically, we can think of a k-simplex as a k-dimensional generalization of a triangle; a 0-simplex is a vertex, a 1-simplex is an edge, a 2-simplex is a triangular face, and a 3-simplex is a tetrahedron, as illustrated in Fig. 2.2. These simplices are "glued" together along entire edges and faces such that there are no disjoint connections. The resulting simplicial complex is a finite set containing the entire "glued together" structure, as well as every subset of simplices that it contains.

For any given set of data points X, there are various ways to connect vertices together by edges and faces (and higher-dimensional simplices) which construct a simplicial complex that ideally captures the topological features we are interested in. The particular construction that is most widely used is known as a Vietoris-Rips complex $VR(\varepsilon)$, which includes all simplices σ where the distance between every

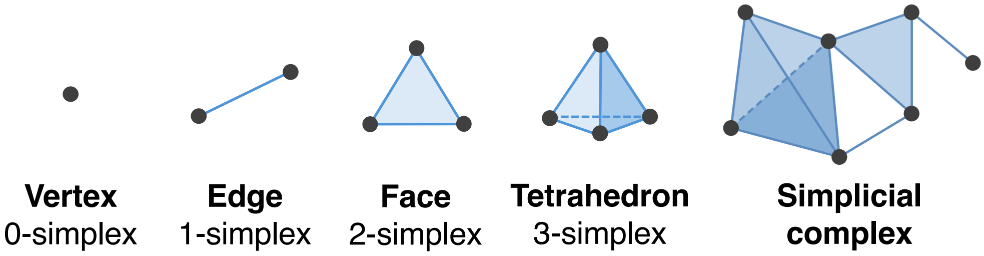

Figure 2.2: Simplices are elementary geometric objects of different dimensionality, which can be combined into more complex structures. A k-simplex can be thought of as a k-dimensional triangle, determined by its $(k+1)$ vertices. From left to right: a vertex (0-simplex), an edge (1-simplex), a filled triangle or face (2-simplex), a tetrahedron (3-simplex), and a set of simplices glued together along certain edges or faces (simplicial complex).

pair of points in σ is at most the cut-off parameter $\varepsilon \geq 0$

$$VR(\varepsilon) = \{\sigma \,|\, d(\sigma) \geq \varepsilon\}, \tag{2.1}$$

where d is the largest distance between any two vertices in the simplex σ.

The Vietoris-Rips complex offers a way to view the data at different length scales ε. At low values of ε, the trivial complex consists purely of isolated vertices; for higher ε, edges and faces emerge as points get “close enough” together; as ε tends towards the diameter of the point cloud, the simplicial complex becomes fully connected.

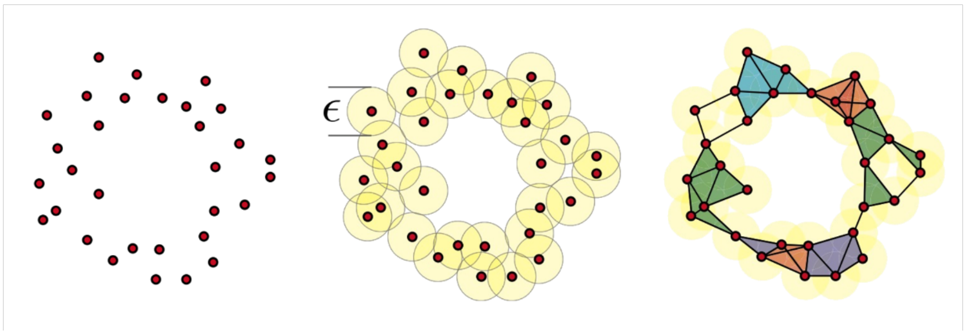

Figure 2.3: Construction of a Vietoris–Rips simplicial complex from a point cloud. (Left) A finite set of points sampled from a circular region. (Centre) Each point is equipped with a closed ball of radius $\epsilon/2$. (Right) Edges and faces are added between any set of points whose balls pairwise overlap, resulting in a Vietoris–Rips complex containing cycles of various scale.

The natural question arises - what choice of ε will result in a simplicial complex that most accurately represents the 'true' topological structure of the underlying data? As illustrated in Fig. 2.3, a single good value of ε may be difficult or even impossible to identify for some data sets. The solution is a relatively straightforward and obvious one, as we shall see in the following section.

Persistent homology and filtrations

Persistent homology addresses the challenge of trying to select an optimal length scale ε by instead considering all possible scales (i.e. all possible values of ε). This generates a nested sequence of simplicial complexes known as a filtration

$$VR(\varepsilon_0 = 0) \subset VR(\varepsilon_1) \subset VR(\varepsilon_2) \subset \cdots , \tag{2.2}$$

where $\varepsilon_0 < \varepsilon_1 < \varepsilon_2 < \cdots$ is monotonically increasing. By computing the homology across the filtration of complexes, this allows us to study the topology of the data as a function of the cut-off length scale ε, also often referred to as the filtration parameter. The importance of topological features can then be attributed based on their persistence through the filtration, i.e. the range of scales in which the

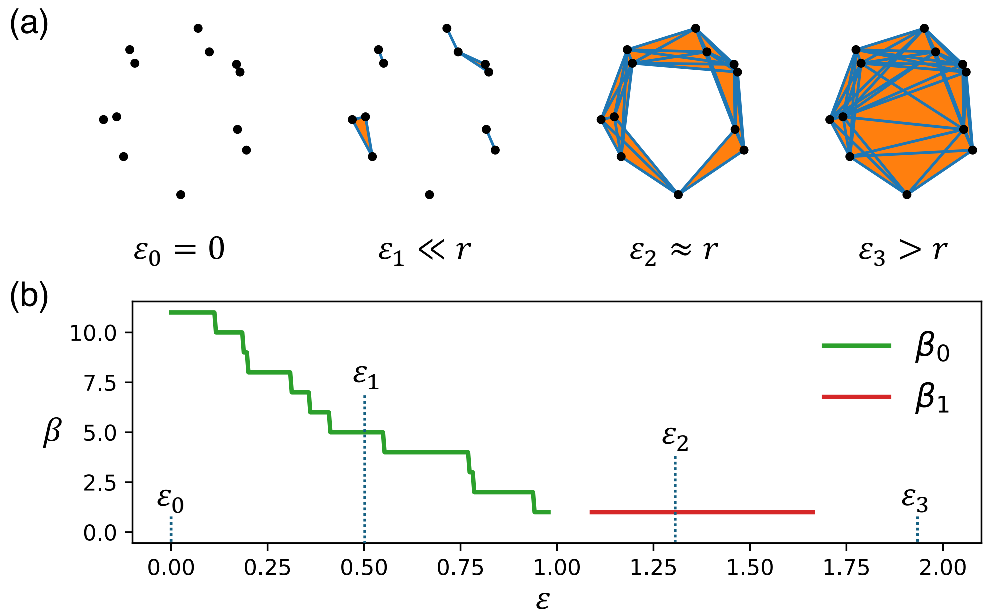

Figure 2.4: (a) Simplicial complexes constructed from a point cloud (noisy samples from a circle with radius $r = 1$) using different cutoff distances ε, where blue lines and orange shaded areas denote edges and faces, respectively. For small cutoff distances all points are disconnected, forming a trivial simplicial complex with no edges ($\varepsilon_0 = 0$). As the cutoff is increased nearby vertices start to become connected by edges and faces ($\varepsilon_1 = 0.5$). Increasing the cutoff further, the simplicial complex has a single connected component hosting a non-trivial cycle ($\varepsilon_2 = 1.3$). For sufficiently large cutoff distances the cycle is destroyed by the addition of faces covering the entire interior of the point cloud ($\varepsilon_3 = 1.9$). (b) Betti numbers β_0 (connected components, green) and β_1 (cycles, red) as a function of filtration scale ε. As ε increases, β_0 decreases as isolated points merge into a single connected component. A one-dimensional cycle is created at $\varepsilon \approx 1$, which remains across the scale of the point cloud, indicating a significant persistent feature. Note that the curve has been omitted for the single cluster at large ε and also for $\beta = 0$.

feature remains present in the simplicial complex. Topological features which persist over a wide range of scales are more robust and should provide a meaningful characterisation of the 'true' topology of the data, while features sensitive to small changes in scale can be attributed to noise.

Consider a noisy point cloud sampled from a circle with radius r, as shown in Fig. 2.4. As we move along the filtration by increasing the filtration parameter ε from 0, initially we see numerous short-lived connected components as they quickly merge together into new connected components. A prominent one-dimensional cycle then emerges around $\varepsilon \approx r$, which approximates the radius of the circle. This cycle persists robustly and unchanging over a wide range of ε, capturing a topological feature which is 'true' to the data.

Topological representations

Persistent diagrams are one stable representation of the scale-dependent topological features that arise from applying persistent homology to a dataset. These are two-dimensional scatter plots where each point (b, d) represents the scale at which a distinct topological feature of dimension k was created (b; birth) and destroyed (d; death) respectively. Since features are always created before they are destroyed, all points lie above the diagonal $b = d$. Points that are further from the diagonal persist over a larger range of scales and are said to have a longer "lifetime" $l = b - d$. These correspond to features that are more characteristic of the global shape of the dataset, whilst points lying close to the diagonal are considered as less important and can often be discarded as noise. At sufficiently large scales, the point cloud for the entire dataset becomes connected to form a single cluster with infinite lifetime. Typically, this is either discarded or plotted on the persistence diagram at a finite d and demarcated using a horizontal dashed line. Fig. 2.5 shows an example point cloud sampled near the unit circle and its corresponding persistence diagram, where the single prominent point far from the diagonal identifies the underlying circular topology. Note that this representation is stable with respect to small variation in the positions of the points used to construct the simplicial complex, but not with respect to the insertion or removal of points. For example, we could imagine that the addition of a point in the center of the point cloud in Fig. 2.5 would result in a drastic change in both the number and persistence of one-dimensional features resulting from the persistence diagram.

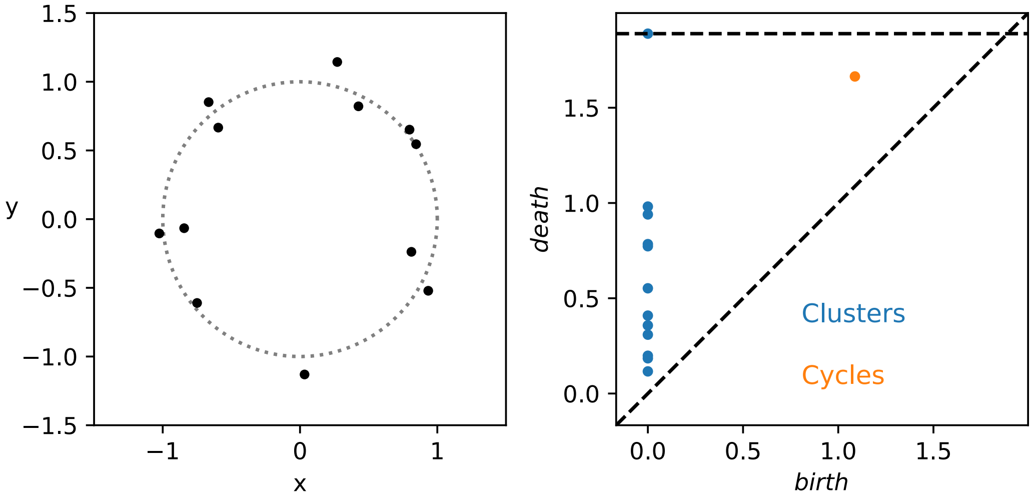

Figure 2.5: Point cloud and persistence diagram for a noisy circle. (Left) Points sampled near the unit circle with added noise. (Right) The corresponding persistence diagram, computed using the Vietoris-Rips complex. Horizontal and vertical axes denote the length scales at which each feature is created (birth) and destroyed (death) respectively. Each point represents a distinct topological feature: the blue points record the connected components that arise from points connecting through the filtration and the single orange point at $(b, d) \approx (1, 1.7)$ encodes the persistent one-dimensional cycle arising from the topology of the circle that the data was sampled from. The single blue point with a horizontal dashed line plotted at the top represents the infinite lifetime cluster that forms when all the points become connected.

Persistence diagrams offer a compact visual summary of the topological features in a dataset and their dependence on scale. Aside from qualitative interpretation, it is often useful to compute statistical quantities that summarize the relevant information contained within the diagram. Typical choices include low-order moments of feature lifetimes

$$\mu_p = \frac{1}{n}\sum_{i=1}^{n}(\ell_i - \bar{\ell})^p, \quad \bar{\ell} = \frac{1}{n}\sum_{i=1}^{n}\ell_i, \tag{2.3}$$

such as the mean ($p = 1$) and variance ($p = 2$), as well as entropy measures that quantify the amount of "disorder" in the lifetimes

$$H = -\sum_{i=1}^{n} p_i \log_2 p_i, \quad p_i = \frac{\ell_i \log_2 \ell_i}{\sum_j \ell_j \log_2 \ell_j}. \tag{2.4}$$

Generally, the persistence diagrams for different datasets will have varying feature counts and differ in their magnitude of scale and degree of noise, so it is not immediately obvious how to compare them directly in a consistent manner. In order to address the need for a rigorous and stable metric, various similarity measures for persistence diagrams were developed. These were accompanied by corresponding stability theorems which ensured that small perturbations in input data resulted in, at most, a similarly small bounded change in the distance between the persistence diagrams of the original data and its perturbed version.

One example is the p-Wasserstein distance, which takes inspiration from optimal transport theory by interpreting diagrams as atomic probability measures on the plane, where each feature (b,d) represents a unit of mass [52]. The distance W_p comes from finding the optimal transport plan γ that minimizes the total cost of 'moving' all off-diagonal features from one diagram $D(f)$ to the other $D(g)$, with unmatched features (i.e. if one diagram has more features) mapped to the diagonal. The distance is defined as

$$W_p(D(f), D(g)) = \left(\inf_{\gamma \in \Gamma} \sum_{(x,y) \in \gamma} \|x - y\|_\infty^p \right)^{1/p} \leq \|f - g\|_\infty, \tag{2.5}$$

where Γ is the set of all partial bijections (one-to-one pairings) between features and the right-most inequality provides the stability bound. For low values of p (e.g. $p = 1$), the distance is sensitive to all features on a comparable level. At higher p, noise-induced outliers become penalized, and in the limit $p \to \infty$ the distance becomes dominated purely by the largest possible deformation of a pair of features, independent of the short-lived features near the diagonal. This distance W_∞ is also popularly referred to as the bottleneck distance [23].

The space of persistent diagrams under these distance metrics do not natively form a Hilbert space, making statistical analysis difficult and preventing the use of familiar tools such as inner products for machine learning pipelines. An alternative approach for characterising persistent diagrams involves vectorisation, which transforms the variable-length information contained in (b, d) pairs into finite-dimensional feature vectors lying in a single vector space. Simple vectorisations can be generated by combining statistical quantities associated with the persistence diagrams, such as the moments or entropy of the feature lifetimes, seen in Eqs. 2.4. In order to minimise the loss of relevant information, more sophisticated, higher-dimensional vectorisations can be obtained through functional representations of the persistence diagrams which are then discretized, such as persistent images and persistent landscapes [53, 54].

From theory to computation

Although the theoretical groundwork required to arrive at the topology of a dataset may appear formidable and involved, the actual computation is surprisingly straightforward and reduces to simple linear algebraic operations. At its core, evaluating the topological invariants (e.g. number of holes or voids) of a simplicial complex amounts to computing the rank and null space of boundary matrices which capture the connectivity between its simplices. For a simplicial complex K, the k-th Betti number can be found as

$$\beta_k(K) = \dim\ker D_k(K) - \operatorname{rank} D_{k+1}(K), \tag{2.6}$$

where D_k is the boundary matrix whose entries are defined as

$$(D_k)_{ij} = \begin{cases} 1 & \text{if simplex } i \text{ is a face of simplex } j, \\ 0 & \text{otherwise.} \end{cases} \tag{2.7}$$

This computation can be readily achieved using matrix reduction techniques, such as Gaussian elimination. This naive approach has worst-case cubic scaling with respect to the number of simplices, but optimized algorithms have been developed which achieve much better (linear) performance [51, 55]

When considering persistent features, this procedure extends naturally from a single simplicial complex to a filtered sequence of complexes by using a boundary matrix whose rows and columns are ordered by filtration value (i.e. at what scale they appear). This can be transformed similarly through matrix reduction into a reduced boundary matrix from which persistent intervals (b, d) can be extracted. Compared with the computation of Betti numbers, this pairing extraction only adds a logarithmic overhead as the full filtered sequence is processed through matrix reduction in one pass [23].

2.2 Quantum many-body systems

Quantum many-body systems pose a unique challenge when it comes to the data that we can extract from them through experiment. Through the development of quantum hardware we have access to an increasing number of degrees of freedom (e.g. particles, spin sites and bosonic modes), each contributing towards the physical Hilbert space. Fig. 2.6 shows representative examples of these systems across different experimental platforms: a Bose-Einstein condensate, a nonlinear optical cavity, and a Rydberg atom spin chain.

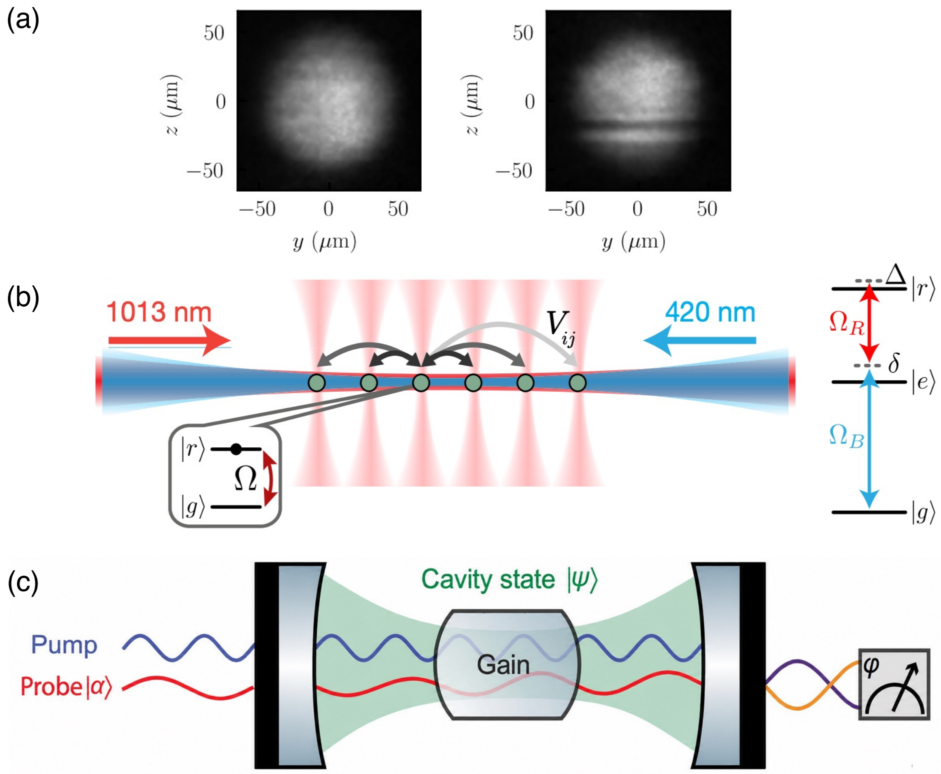

Figure 2.6: Experimental realisations of quantum many-body systems. (a) Absorption images of a Bose-Einstein condensate before (left) and after (right) the formation of a dark soliton, showing a characteristic density notch. Adapted from Fritsch et al. [56]. (b) Schematic of a 51-atom Rydberg quantum simulator, in which atoms are individually trapped in optical tweezers and coupled via van der Waals interactions V_{ij}. A two-photon excitation scheme drives transitions between the ground state $|g\rangle$ and the Rydberg state $|r\rangle$ via an intermediate state $|r\rangle$, using lasers with Rabi frequency Ω_B and Ω_R. Adapted from Bernien et al. [57]. (c) Schematic of a driven nonlinear optical cavity, where a pump and a coherent probe field $|\alpha\rangle$ interact with a gain medium to generate and sustain a non-classical intracavity state $|\psi\rangle$, measured via homodyne detection. Adapted from Choi et al. [58].

For a system of N sites with local Hilbert space dimension d (e.g., $d = 2$ for spin-$\frac{1}{2}$ particles), the full Hilbert space is the tensor product

$$\mathcal{H} = \bigotimes_{i=1}^{N} \mathcal{H}_i, \qquad \dim \mathcal{H} = d^N. \tag{2.8}$$

A generic pure state $|\psi\rangle \in \mathcal{H}$, representing maximal knowledge of the system, is

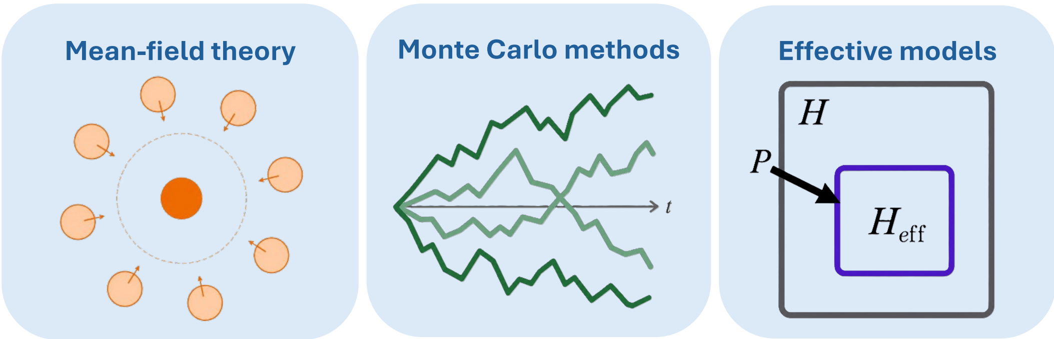

Figure 2.7: Schematic summary of several quantum many-body approximation methods, highlighting mean-field theory, Monte Carlo methods and effective models.

specified by $d^N - 1$ independent complex amplitudes. In practice, when complete knowledge of the system is unavailable due to entanglement with an unobserved environment or classical statistical uncertainty, the state is described by a density operator $\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k|$ satisfying $p_k \geq 0, \sum_k p_k = 1$, $\rho = \rho^\dagger \geq 0$ and $\mathrm{Tr}(\rho) = 1$, which requires $d^{2N} - 1$ independent real parameters.

As N grows, both the Hilbert space dimension and the complexity of the quantum state explode exponentially. However, our ability to probe the state of the system through measurement does not scale accordingly. Measurement access is confined to local projective measurements acting on individual sites or small subsystems. A single projective measurement on a site with d outcomes yields at most $\log_2 d$ bits of information, contributing negligibly to the d^{2N}-dimensional parameter space required for full state tomography. Consequently, with M measurement shots, the resulting empirical data is inevitably sparse in the limit $M \ll d^{2N}$ [59].

To help navigate this complexity, physicists often employ different classes of approximations to describe quantum many-body systems; we briefly discuss three of these here, all of which are relevant to the work covered in this thesis.

Mean-field theory

Mean-field theory is an analytical approach that reduces the intractable many-body problem to an effective single-particle one by replacing all interactions with a self-consistent average field. In second quantization, a generic interacting bosonic system can be described by the Hamiltonian

$$\hat{H} = \int d\mathbf{r}\, \hat{\Psi}^\dagger(\mathbf{r}) \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(\mathbf{r}) \right) \hat{\Psi}(\mathbf{r}) + \frac{1}{2}\int d\mathbf{r}\, d\mathbf{r}'\, \hat{\Psi}^\dagger(\mathbf{r})\hat{\Psi}^\dagger(\mathbf{r}')V(\mathbf{r}-\mathbf{r}')\hat{\Psi}(\mathbf{r}')\hat{\Psi}(\mathbf{r}), \tag{2.9}$$

where $\hat{\Psi}(\mathbf{r}, t)$ is the bosonic field operator satisfying the canonical commutation relations, the first term is the single-particle Hamiltonian and the second is a two-particle interaction part with potential $V(\mathbf{r} - \mathbf{r}')$. The mean-field approximation then decomposes the field into a coherent component and fluctuations,

$$\hat{\Psi}(\mathbf{r}, t) = \langle \hat{\Psi}(\mathbf{r}, t) \rangle + \delta\hat{\Psi}(\mathbf{r}, t), \tag{2.10}$$

neglecting fluctuation terms that are higher order in $\delta\hat{\Psi}$. This amounts to assuming that the many-body state is only weakly correlated and can be well approximated by a mean-field state with no significant long-range entanglement between different modes. Substituting this into the Heisenberg equation of motion,

$$i\hbar\frac{\partial}{\partial t}\hat{\Psi}(\mathbf{r}, t) = [\hat{\Psi}(\mathbf{r}, t), \hat{H}], \tag{2.11}$$

and keeping only leading-order contributions yields a closed evolution equation for the condensate field $\Psi(\mathbf{r}, t) = \langle \hat{\Psi}(\mathbf{r}, t) \rangle$,

$$i\hbar\frac{\partial\Psi(\mathbf{r}, t)}{\partial t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(\mathbf{r}) + \int d\mathbf{r}'\, V(\mathbf{r} - \mathbf{r}')|\Psi(\mathbf{r}', t)|^2 \right) \Psi(\mathbf{r}, t). \tag{2.12}$$

A canonical example relevant to this thesis is the mean-field derivation of the Gross-Pitaevskii equation from the Bose gas with short-range interactions, where in the dilute limit the interaction potential may be approximated by a delta function pseudopotential $V(\mathbf{r}-\mathbf{r}') = g\,\delta(\mathbf{r}-\mathbf{r}')$, where $g = 4\pi\hbar^2 a_s/m$ is a coupling parameter determined by the effective interaction range a_s. This reduces the nonlocal interaction term to a local cubic nonlinearity, giving the mean-field nonlinear Schrodinger equation

$$i\hbar\frac{\partial\Psi}{\partial t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(\mathbf{r}) + g|\Psi|^2\right)\Psi. \tag{2.13}$$

Quantum Monte Carlo methods

Quantum Monte Carlo methods offer a numerical approach by replacing the explicit evaluation of high-dimensional integrals (e.g. expectation values over the full exponentially large many-body configuration space) with stochastic sampling. In these formulations, one constructs a probability distribution over many-body configurations or quantum trajectories and then estimates observables as averages over samples generated by importance sampling or Markov chain Monte Carlo, with the statistical error scaling as $\sim \sigma/\sqrt{M}$ for M samples, independent of the exponentially large Hilbert space dimension. An example relevant to this thesis is the Monte Carlo wavefunction (or quantum jump) method, which approximates the dynamics of an open many-body system described by the Lindblad master equation for the density operator $\rho(t)$,

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[\hat{H},\rho] + \sum_k \hat{L}_k\rho\hat{L}_k^\dagger - \frac{1}{2}\left\{\hat{L}_k^\dagger\hat{L}_k,\rho\right\}, \tag{2.14}$$

where $\hat{H}$ is the many-body Hamiltonian, $\{\hat{L}_k\}$ are Lindblad operators encoding coupling to the environment, and the double commutator term ensures that $\rho(t)$ remains Hermitian, positive semi-definite, and its trace is normalised. Instead of evolving the full density matrix, the equation is unravelled into an ensemble of quantum trajectories, each represented by a pure state $|\psi_n(t)\rangle$ evolving under the non-Hermitian effective Hamiltonian

$$\hat{H}_{\text{eff}} = \hat{H} - \frac{i\hbar}{2}\sum_k \hat{L}_k^\dagger \hat{L}_k, \tag{2.15}$$

and stochastically interrupted by quantum jumps $|\psi_n(t)\rangle \to \hat{L}_k|\psi_n(t)\rangle$ at random times, with jump probability proportional to $\langle\psi_n(t)|\hat{L}_k^\dagger \hat{L}_k|\psi_n(t)\rangle$. The expectation value of an observable $\hat{O}$ is then estimated as

$$\langle\hat{O}(t)\rangle \approx \frac{1}{N_{\text{traj}}}\sum_{n=1}^{N_{\text{traj}}}\langle\psi_n(t)|\hat{O}|\psi_n(t)\rangle, \tag{2.16}$$

which provides a numerically tractable approximation to the full many-body open dynamics by replacing the exponential cost of storing and evolving ρ with an error-controlled procedure whose accuracy is governed by the number of trajectories N_{traj}, with the statistical error scaling as $\sim \sigma/\sqrt{N_{\text{traj}}}$ while remaining exactly equivalent to the Lindblad master equation in the limit $N_{\text{traj}} \to \infty$. Note that tackling fermionic systems using Monte Carlo methods becomes much harder due to the sign problem, which arises from the antisymmetry of fermionic wavefunctions under particle exchange and causes the effective statistical error to grow exponentially.

Projection onto effective models

Projection onto effective models represents a third strategy, where the full complexity of a quantum many-body system is circumvented by identifying a physically relevant low-dimensional subspace $\mathcal{H}_{\text{eff}} \subset \mathcal{H}$ (e.g. low-energy subspace, symmetry conditions, or kinematically constrained) and deriving a reduced Hamiltonian H_{eff} that governs the dynamics strictly within this subspace, where the remaining degrees of freedom are projected out. This is formally implemented by a projector $\mathcal{P}$ onto the target subspace, so that the effective Hamiltonian is defined as

$H_{\text{eff}} = \mathcal{P}\, H\, \mathcal{P}$ and the dynamics are restricted to $\mathcal{H}_{\text{eff}} = \mathcal{P}\mathcal{H}$. An illustration of this approach relevant to this thesis is provided by the PXP model, which arises as an effective description of a one-dimensional Rydberg-atom chain governed by the microscopic Hamiltonian

$$H = \sum_i \left[\frac{\Omega}{2} \sigma_i^x + \Delta n_i \right] + \sum_{i<j} V_{ij}\, n_i n_j \tag{2.17}$$

where $\sigma_i^x = |\circ\rangle_i\langle\bullet| + |\bullet\rangle_i\langle\circ|$ is the Pauli x matrix driving transitions between ground states $|\circ\rangle_i$ and Rydberg states $|\bullet\rangle_i$ on site i, $n_i = |\bullet\rangle_i\langle\bullet|$ the Rydberg occupation number, $V_{ij} \propto 1/|i-j|^3$ the van-der-Waals interaction, Ω the laser frequency and Δ the laser detuning. In the strong-blockade limit $V_{i,i+1} \gg \Omega, |\Delta|$, configurations with adjacent Rydberg excitations are energetically suppressed, so the relevant low-energy dynamics is confined to the subspace of configurations in which no two neighbouring sites are simultaneously excited, i.e. $n_i n_{i+1} = 0$ for all i. This constraint is enforced by local projectors $P_i = |\circ\rangle_i\langle\circ|$ onto the atomic ground state at site i, so that the effective projector reads $\mathcal{P} = \prod_i (1 - n_i n_{i+1})$. Within the projected subspace, the bare Rabi term σ_i^x is dressed by its neighboring projectors, yielding the PXP Hamiltonian

$$H_{\text{PXP}} = \Omega \sum_i P_{i-1}\, \sigma_i^x\, P_{i+1} \tag{2.18}$$

which implements a kinetically constrained spin-$\frac{1}{2}$ dynamics where each spin flip at site i is only allowed if both neighboring sites are unexcited, i.e., $\langle\psi|P_{i-1}P_{i+1}|\psi\rangle = 1$.

This constraint reduces the Hilbert space from the full 2^N-dimensional configuration space to a much smaller subspace whose dimension grows roughly as a Fibonacci-like sequence, while still preserving the model's essential features such as non-integrability, weak ergodicity-breaking, and the existence of many-body scarred eigenstates. Nevertheless, projection onto an effective constrained subspace still carries inherent limitations. For the physical Rydberg array approximated by the PXP model, in the strong blockade limit the configurations violating the nearest-neighbour exclusion constraint are suppressed but may still have a small non-zero probability, leading to leakage out of the projected subspace. Additionally, the model neglects longer-range interactions and assumes a perfectly homogeneous chain. Together, assumptions and approximations can lead to errors at long evolution times and in large systems.

Further details on this model are given in Chapter 4.

2.3 Topological feature selection in Physics

In this section, we provide a high-level, general overview of applications of TDA and topological feature selection in physics, highlighting key contributions as well as existing limitations that motivate the work presented in this thesis.

The earliest applications of TDA in physics were typically in settings where the underlying data already has a well-defined shape or graph-like structure that made the construction of simplicial complexes relatively straightforward, such as dynamical systems [60], random clouds of spheres [61], random networks [62], and binary image data [63, 64]. In the case of continuous-time dynamical systems, [60] presented a topological feature-based noise filtering method, taking advantage of the fact that overly-sampled points will trace a single continuous curve in the absence of noise. By analysing the scale-dependent distribution of zeroth Betti numbers B_0 (i.e. number of connected components), they were able to distinguish points belonging to the dynamical trajectory from noise-perturbed points depending on whether they formed connected or isolated clusters. The latter points were filtered out, improving the signal-to-noise ratio and accuracy of estimated Lyapunov exponents. In this work and many others, the use of a single topological feature as a physically motivated feature selector demonstrates that even simple topological invariants can serve as a useful and interpretable diagnostic when the underlying data structure is well-matched to the topological quantity being computed.

Later, persistent homology established itself as a powerful tool for extracting meaningful shape information from raw data, especially in applications where it is necessary to identify structures present at different spatial scales. Initial applications included the study of a broad range of condensed matter and material physics settings such as glassy and crystalline phases [30, 32], amorphous ices [65], granular media [66, 67], lattice spin models and gauge theories [68, 69], and two-dimensional materials [70]. For example, in [68] the authors use persistent homology to identify phases in spin models, including hidden orders that are not captured by conventional symmetry-breaking descriptors. By converting spin configurations into persistence diagrams, where individual sites in the spin system form the point cloud, averaged persistent barcodes across spin configurations yield a characterisation of the phase diagram, with the zeroth Betti number B_0 playing a role analogous to the conventional order parameter of magnetisation density.

This example is a demonstration of how TDA is used to study and discover structure in complex systems for which simple visualisations and explicit shape (such as images, phase space trajectories or material configurations) do not exist. There are various notions of distance that can be applied in contexts where the underlying data lack intuitive shape, and the choice of a suitable distance metric is often crucial for determining the capability of TDA to capture underlying structural changes. For example, in the study of condensed matter systems at finite temperatures, the geodesic distance between different spin configurations and the quantum distance based on the overlap between eigenfunctions are commonly used metrics [35, 71, 72]. However, they may fail to capture emergent low-dimensional structures or nonlocal entanglement features and, for such applications, domain-informed metrics such as mutual information, entanglement entropy and Wasserstein distances [34, 73] become essential.

As exposure to tools in the TDA framework grew within the physics community following successful applications across many diverse fields, several comprehensive reviews emerged to synthesize the rapidly expanding literature and lower the barrier of entry for practitioners. Here we highlight a selection of the most notable reviews that are particularly well-suited towards and primarily targeted at a physics audience. An early review by Carlsson, one of the founders of the field, provides a broad survey of different techniques of TDA and their applications across various areas of science [22]. Another review, by Murugan and Robertson, offers a pedagogical and physicist-friendly introduction to the important techniques of persistent homology and the Mapper algorithm, applying them to the example of an astronomical dataset [29]. Leykam and Angelakis provided a review which focuses on the integration of TDA into physics-targeted machine learning pipelines and discusses performance enhancement through suitable feature selection [28]. A more recent review by the Carlsson team has a similar focus, discussing the appropriate selection of topological tools, representations and machine learning vectorizations for different types of input data [74]. Techniques that go beyond standard persistent homology, such as the use of higher-order Laplacians, are also reviewed and analyzed in this review.

The landscape of TDA applications in physics has evolved rapidly in recent years, exploring a diverse range of physical systems and phenomena. In condensed matter and photonics, Madeleine et al. [75] used persistent homology to classify and quantify structural disorder in plasmonic metasurfaces, demonstrating how topological tools can be used for the fast and accurate design of photonic devices with tailored optical responses or as aid in improving their fabrication methods. In lattice gauge theory, Sale et al. [76] and Spitz et al. [77] used persistent homology observables from monopole currents and center vortices as robust pseudo-order parameters to detect confinement-deconfinement transitions via finite-size scaling. In field theory, Noel and Spitz [78] used persistent homology to characterize field configurations in scalar field theories, identifying topological signatures of symmetry breaking that are not accessible from conventional two-point correlators.

Another particularly active area of research, fueled by the promise of TDA as a novel tool, concerns complexity bounds and the development of quantum algorithms for TDA. Berry et al. [42] analyzed prospects for quantum advantage in estimating Betti numbers, showing that polynomial speedups are achievable under specific conditions but are limited by input model assumptions. Schmidhuber and Lloyd [43] established complementary complexity-theoretic limitations, proving that certain TDA problems are classically intractable, while Crichigno and Kohler [79] showed that clique cohomology is QMA_1-hard. Apers et al. [80] introduced a simple classical algorithm for Betti number estimation that matches some quantum proposals, refining the boundary for quantum advantage. More recently, Gyurik et al. [81] proved that determining whether a topological feature persists across length scales is BQP_1-hard yet contained in BQP, yielding the first provable exponential quantum speedup for a TDA problem. Other algorithms proposed to compute topological invariants on quantum hardware include Betti number estimation from quantum thermal states [82], density of states estimation of combinatorial Laplacians [83], and quantum HodgeRank [84].

Chapter 3

Persistent homology for chaos detection in open quantum systems

TDA has already demonstrated considerable utility in the study of classical dynamical systems: persistent homology has been used to characterise attractors and bifurcations in nonlinear systems, detect phase transitions in time-series data, and classify the topology of invariant manifolds in Hamiltonian mechanics [85–88]. A common theme amongst these applications is that TDA extracts robust topological signatures from trajectory data or state-space point clouds without requiring knowledge of an underlying order parameter, making it well-suited to systems where the relevant structure is global and emergent.

Here we take a first step in extending this approach to the quantum domain. We focus on a nonlinear resonator model supporting a semiclassical limit with known behaviour, which provides a useful starting point for adapting existing topological methods for time series analysis to quantum many-body systems. The material presented in this Chapter is based on the following publication:

- Harvey Cao, Daniel Leykam and Dimitris G. Angelakis, Unravelling quantum chaos using persistent homology, *Physical Review E* **107**, 044204 (2023) [89]

3.1 Introduction

Classical nonlinear systems can exhibit chaotic behaviour if they possess extreme sensitivity to initial conditions and are topologically mixing. The interplay of these two properties leads to long-term unpredictability and seemingly random dynamics, despite being fully deterministic. When dissipative effects play a role, a complex geometrical structure can be observed in the phase space dynamics known as a 'strange attractor'. Strange attractors differ from other phase space attractors in that two points that are initially nearby on the attractor will eventually evolve to become arbitrarily far apart. This divergence is used as a diagnosis of whether a particular system is chaotic, quantified in conventional methods by calculating the maximal Lyapunov exponent (LE), a measure of the average rate of divergence in phase space [90].

When quantum effects are significant, chaotic behaviour at the wavefunction level is precluded due to the linearity of the Schrödinger equation [91]. Instead, the role of chaos in quantum mechanics is studied by identifying signatures of quantum systems which result in a classical counterpart that yields chaotic dynamics. Early efforts in quantifying quantum chaos focused on the spectral statistics of chaotic Hamiltonians, with level spacing statistics corresponding to that of random matrices when the classical limit is chaotic [92–94]. There have also been various quantum generalisations of the classical Lyapunov exponent constructed in terms of quasiclassical phase-space dynamics, the stochastic Schrödinger equation and, more recently, out-of-time-order correlators (OTOCs) [95–98]. Beyond addressing fundamental questions concerning the correspondence principle between quantum and classical limits [99], there is interest in understanding the role of quantum chaos in the thermalization of quantum many-body systems and subsequent implications for quantum information and near-term quantum computation [100–104].

In comparison to the classical case, where the maximal LE can be used unanimously as an indicator of chaos, the toolset for quantifying chaos in open quantum systems remains limited. Furthermore, these methods typically involve quantifiers that are difficult to measure in real experimental settings. For example, calculating the recently proposed OTOC-based LEs requires the ability to precisely reverse the evolution of a system, which is extremely challenging [105, 106].

Recently, there has been growing interest in using TDA methods for the characterization of classical dynamical systems through counting the number of low-dimensional holes of reconstructed attractors [85, 87, 107, 108]. This offers an alternative to methods based on calculating the maximal LE, which can be heavily affected by noise and finite data sizes [109, 110]. The primary objective of this study is to see whether TDA-based methods for chaos detection which have already proven effective in the classical case can be applied to the harder-to-visualise dynamics of open quantum systems, where the attractor in phase space is no longer well-defined by trajectories as in the classical case. As single system trajectories are compared in the classically chaotic case to find exponential divergence, it seems sensible to adopt a quantum approach that does likewise. Thus, we propose an approach based on quantum trajectories that are obtained from unravelling methods. Specifically, we consider the Monte Carlo wavefunction method, which replaces the original master equation with an ensemble of stochastically evolving quantum trajectories, whose average recovers the full density operator evolution [111, 112]. We follow one member of the ensemble and use the time-evolution of suitable observables to approximate a classical trajectory analogue.

We apply this approach to probe the structural stability of cavity-like open quantum systems in order to detect transitions to quantum chaos (that is, regimes where the classical counterpart is chaotic). Namely, we consider a Kerr-nonlinear cavity driven by a periodically modulated external EM field and show that the topological

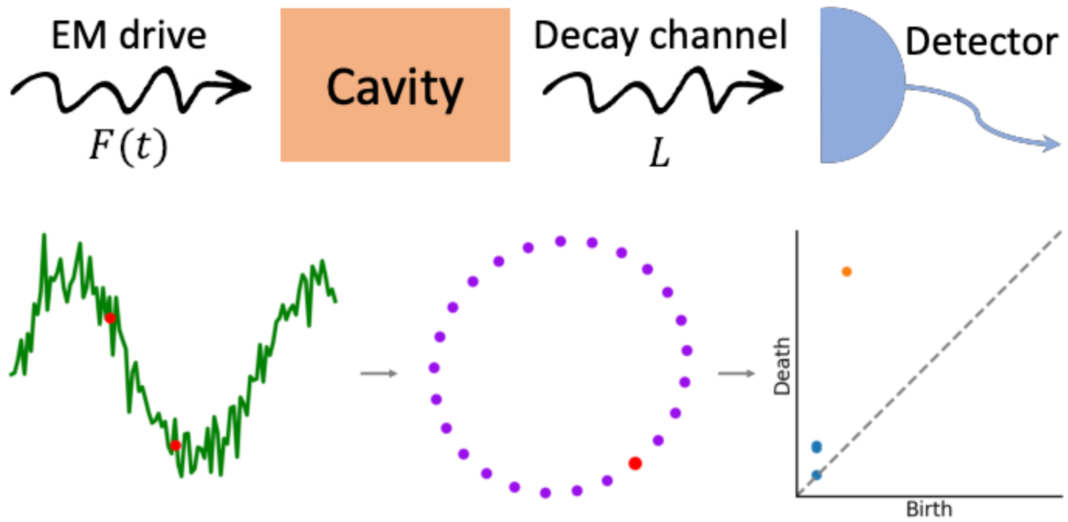

Figure 3.1: Schematic of (a) leaky periodically-driven nonlinear cavity model and (b) TDA pipeline for a noisy periodic time series (left) embedded into two-dimensions using time-delay embedding (middle) with the presence of a single cycle indicated by the corresponding persistence diagram (right) with 0-D (blue) and 1-D (orange) features. To illustrate the embedding of two points (red) in the time series according to Takens' embedding, the corresponding two-dimensional point (red) is indicated on the point cloud.

properties of a single quantum trajectory can be used to distinguish between regular and chaotic parameter regimes of the system. These quantum trajectories can be realised in continuous monitoring schemes, which opens a route for the application of this topological approach in real-life experimental setups [113, 114]. A graphical schematic of the full pipeline is shown in Fig. 3.1. From a broader perspective, this work also bridges the application of TDA to quantum dynamics and we foresee that further research surrounding this connection will be fruitful.

The outline of this chapter is as follows: Section 3.2 provides a brief introduction to attractor reconstruction using time-delay embedding. Section 3.3 establishes the quantum system (and its classical counterpart) considered in the current study and introduces the Monte Carlo wavefunction method. Sections 3.4 and 3.5 present the proposed topological approach for identifying regular and chaotic regimes in classical and quantum systems respectively, demonstrating examples of the full pipeline. Section 3.6 summarises the results obtained using this approach to detect chaotic phases in the parameter space of the system. Section 3.7 makes concluding remarks and suggests future directions.

3.2 State space reconstruction

In general, the attractor of a classical dynamical system may not be directly accessible in an experimental setup because a mathematical description of the system and the dimensionality of its full phase space is unknown. Typically, measurements may only be able to be recorded of one relevant quantity at regular time intervals, resulting in a single time series $x(t)$ that represents a scalar observation function of the full state vector $\mathbf{x}(t)$,

$$x(t) = F(\mathbf{x}(t)). \tag{3.1}$$

Takens' embedding theorem implies that this single time series that has been observed from a potentially unknown dynamical system can be useful in reconstructing the attractor for the full system [115]. Specifically, the observed time series $x(t)$ can be 'embedded' into a d-dimensional vector space in such a way as to preserve the topological properties of the *full* state space $\mathbf{x}(t)$. This *time-delay embedding* involves taking uniformly time-lagged samples of the time series and concatenating them into a single vector,

$$\mathbf{X}(t) = (x(t), x(t-\tau), ..., x(t-(d-1)\tau)), \tag{3.2}$$

where d is the embedding dimension and τ is the embedding time-delay. The embedded points form features, such as loops and voids, that reflect the same topology as the attractor of the dynamical system, guaranteed by Takens' theorem as long as appropriate embedding parameters are chosen [90, 115]. There is no consensus on the most favourable way to choose these and it is largely dependent on the specifics of the data, but popular heuristic methods include using mutual information for estimation of the optimal time-delay and false nearest neighbours for determination of a proper embedding dimension [116, 117].

3.3 Model

For our quantum system of interest, we consider a damped, driven, Kerr-nonlinear cavity which is known to be chaotic in the classical case and whose chaotic regimes have been well studied in the literature [118, 119]. Its unitary dynamics is governed by the following Hamiltonian, written in a frame rotating with the cavity resonance frequency ω,

$$H = \frac{1}{2}\hbar\chi a^{\dagger 2}a^2 + i\hbar F(t)(a^\dagger - a), \tag{3.3}$$

where $a^\dagger$ (a) are photon creation (annihilation) operators respectively, χ is the photon interaction strength and $F(t)$ describes the periodic modulation by an external driving field and time t is normalized by the cavity period $2\pi/\omega$. We consider the same driving as in Ref. [119], where $F(t)$ is a string of rectangular pulses of amplitude A, length $T/2$, and separation $T/2$, such that $F(t) = F(t+T)$. Note that if $F(t)$ is time-independent the system loses the potential for exhibiting chaotic behaviour.

Next, we incorporate damping of the cavity field into the model using the conventional formalism of open quantum systems. The time evolution of the reduced density matrix ρ describing the cavity state is governed by a master equation of the Lindblad form [111],

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[H, \rho] + \mathcal{L}[\rho]. \tag{3.4}$$

The dissipative term $\mathcal{L}[\rho]$ captures the weak coupling of the system to the environment and takes the general form

$$\mathcal{L}[\rho] = \sum_i \left(L_i \rho L_i^\dagger - \frac{1}{2}\{L_i^\dagger L_i, \rho\} \right), \tag{3.5}$$

where $\{A, B\} = AB + BA$, and the jump operators L_i each describe a distinct class of system-environment interaction. Ignoring this term simply reduces Eq. (3.4) to the familiar von-Neumann equation, capturing the unitary dynamics of the system described by the Hamiltonian. Here, we consider a single dissipation channel corresponding to single-photon loss from the leaky cavity, described by the jump operator $L = \sqrt{\gamma}\, a$, where γ is the dissipative coupling constant. We assume that the rate of thermal photon injection from the environment is zero, following Refs. [119, 120], so that Eq. (3.5) reduces to

$$\mathcal{L}[\rho] = \gamma \left(a \rho a^\dagger - \frac{1}{2}\{a^\dagger a, \rho\} \right). \tag{3.6}$$

Though the master Eq. (3.4) can be solved numerically, if the relevant Hilbert space of the quantum system is of dimension N the density operator requires N^2 real numbers to represent. This heavy computational requirement can be reduced using the quantum jump (Monte Carlo wavefunction) method, which constructs quantum trajectories by evolving wavefunctions $|\psi(t)\rangle$ using a pseudo-Hamiltonian of the form

$$H_{MC} = H - \frac{i\hbar}{2} L^\dagger L, \tag{3.7}$$

combined with stochastic jumps that correspond to the environmental action of the jump operator(s). Averaged over an ensemble of many trajectories, this evolution becomes equivalent to the master equation treatment in Eq. (3.4) [112, 121]. Expectation values of observables $\hat{O}$ are similarly calculated by using the expectation values of wavefunctions $\langle\psi(t)|\hat{O}|\psi(t)\rangle$ and taking the ensemble mean over all trajectories. Importantly, this method allows for mimicking the behaviour of individual realizations of the system dynamics and correctly reproduces the density matrix when averaged over many iterations.

In the classical limit of large amplitude coherent states, when the average number of photons in the cavity tends to infinity, following Refs. [118–120] the master equation Eq. (3.4) can be transformed into

$$\frac{d\xi}{dt} = -\frac{1}{2}\gamma\xi + F(t) - i\chi|\xi|^2\xi^*, \tag{3.8}$$

where ξ is a complex dynamical variable whose real and imaginary components represent the position and momentum respectively. This variable can be compared directly with the expectation value of the photon annihilation operator $\langle\hat{a}\rangle_\psi$ in the quantum case. The nonlinear dynamical system described by Eq. (3.8) can be shown to exhibit regions of chaotic behaviour, as visualized in the bifurcation diagram [Fig. 3.2] that shows the stroboscopic parameter dependence. Simple curves correspond to regular dynamics, whilst thicker bands indicate an exploration of the phase space that is characteristic of chaotic behaviour.

3.4 Classical chaos detection

The attractor of a dissipative system describes its long-time evolution in phase space and has distinct topological properties based on whether the system dynamics exhibit regular (non-chaotic) or chaotic behaviour. As discussed in Sec. 3.2, an attractor can be faithfully reconstructed (up to its topology) from the time-series of a single observable through an embedding technique, without knowledge of the full phase space dynamics. The resulting embedded data can then be analysed through TDA methods as persistent homology to characterise the reconstructed attractor and thus the system of interest. This method has been previously proposed as a

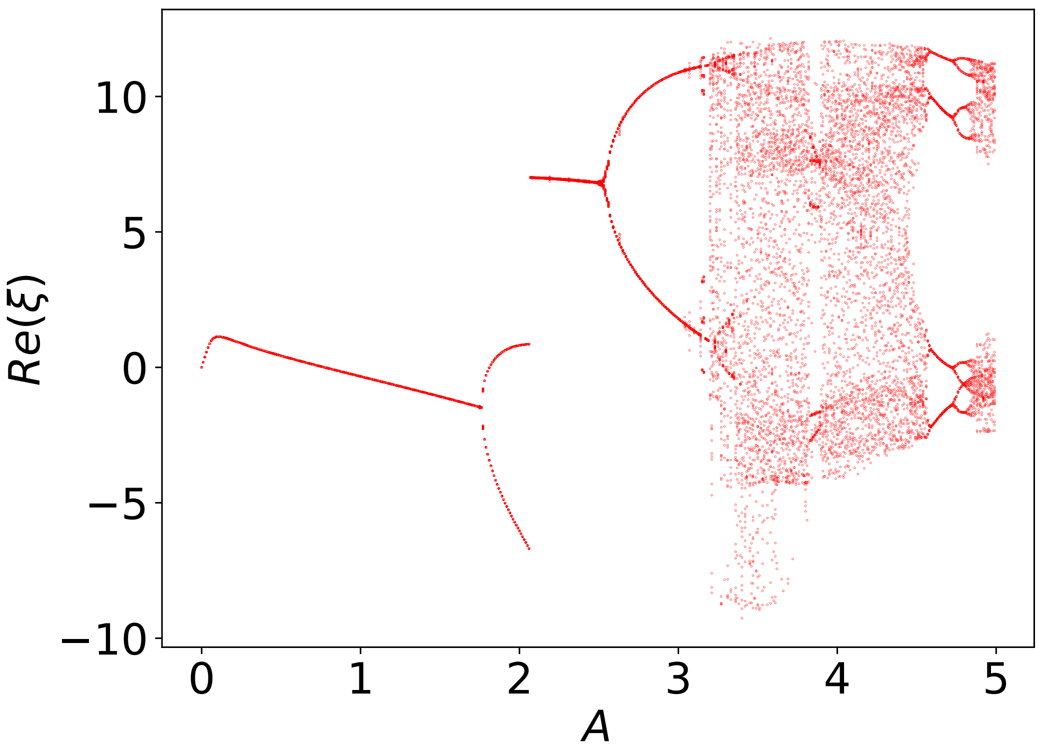

Figure 3.2: Stroboscopic values of Re(ξ) at integer multiples n of the driving period $40 < n < 100$ as a function of the driving amplitude A. This classical bifurcation diagram visually indicates the transitions between regular and chaotic behaviour. Simple curves correspond to regular periodic dynamics and the thicker bands correspond to chaotic trajectories which wander densely in the phase space. The driving period is fixed at $T = 10$.

quantifier for classical chaos, as an alternative to the maximal LE which can be difficult to compute for noisy, finite data [85, 107].

Recall from the Section 3.1 that in the study of quantum chaos we are interested in finding signatures of the quantum system in regimes where the corresponding classical system is chaotic. Thus, we start with an analysis of the classical limit described by Eq. (3.8) and use it to determine the parameter regimes under which the system exhibits chaos.

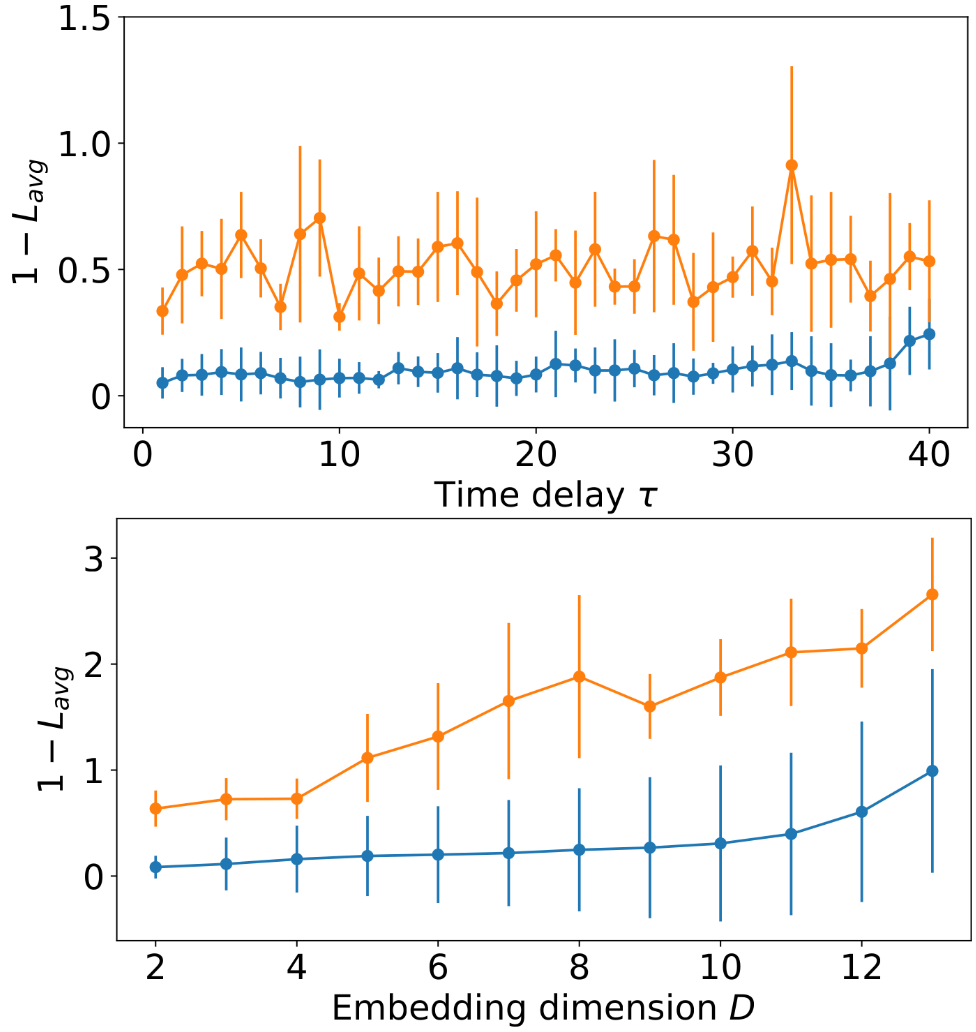

Figure 3.3: Average 1-hole lifetime as a function of (a) embedding time delay ($D = 2$) and (b) embedding dimension ($\tau = 7$) for position expectation value $\langle x \rangle$ of single Monte Carlo trajectory of quantum system [Eq. (3.3)], averaged over 10 points in the regular (blue) and chaotic (orange) phases respectively with corresponding standard deviations.

The embedding and TDA pipeline is demonstrated for the analysis of the classical dynamical system for drive parameter values corresponding to regular and chaotic phases in Fig. 3.4. In this Chapter, for the time-delay embeddings we determine optimal hyperparameters using the heuristic methods discussed in Sec. 3.2. It should be noted that although the specific values of the embedding delay and embedding dimension are not important as long as they satisfy the criteria of Takens' theorem [115], they do have an influence on the features obtained through persistent homology due to the discrete nature of the data. Here, we demonstrate that the results for the quantum case remain robust over a range of hyperparameter values. Figure 3.3(a) shows that for a fixed embedding dimension and varying the time-delay parameter, the topological measure averaged over 10 different chosen points in the regular and chaotic phases respectively can still be easily distinguished. Figure 3.3(b) indicates the same result, but with a fixed time-delay and varying the embedding dimension.

For driving parameters chosen in the regular phase ($A = 1$, $T = 8$), the periodic dynamics obtained by numerically solving Eq. (3.8) and shown in Fig. 3.4(a) lead to a reconstructed attractor that has a simple geometric structure. The embedded point cloud shown in Fig. 3.4(c) exhibits a single cycle. This is quantified by the persistence diagram in Fig. 3.4(e) which identifies the presence of the cycle through the single long-lived one-dimensional feature furthest from the diagonal. The strange attractors that correspond to chaotic dynamics have much more complex structures whose topology cannot be described analytically as in the regular case [90]. Fig. 3.4(b) shows the time series for parameter values chosen in the chaotic phase ($A = 4.5$, $T = 8$), which results in a reconstructed attractor whose points are scattered densely over the phase space in Fig. 3.4(d). Correspondingly, the persistence diagram in Fig. 3.4(f) indicates that at the scale of the system dynamics there are numerous features with a range of lifetimes. During the transition between the regular and chaotic phases, we note that the attractors develop two long-lived features indicative of two cycles. This is highly suggestive of period-doubling behaviour on the route to chaos, also illustrated by the curve bifurcations in Fig. 3.4, as the correlation between the number of cycles observed in phase space trajectories and multi-periodicity of the dynamics has been previously identified [85, 122].

These distinctions between persistent features of different systems can be used as a quantifier for distinguishing regular and chaotic phases. Motivated by the significance of the size and number of cycles outlined in the above discussion, we consider as topological measure the average lifetime of low-dimensional topological features, discussed further in Sec. 3.6.

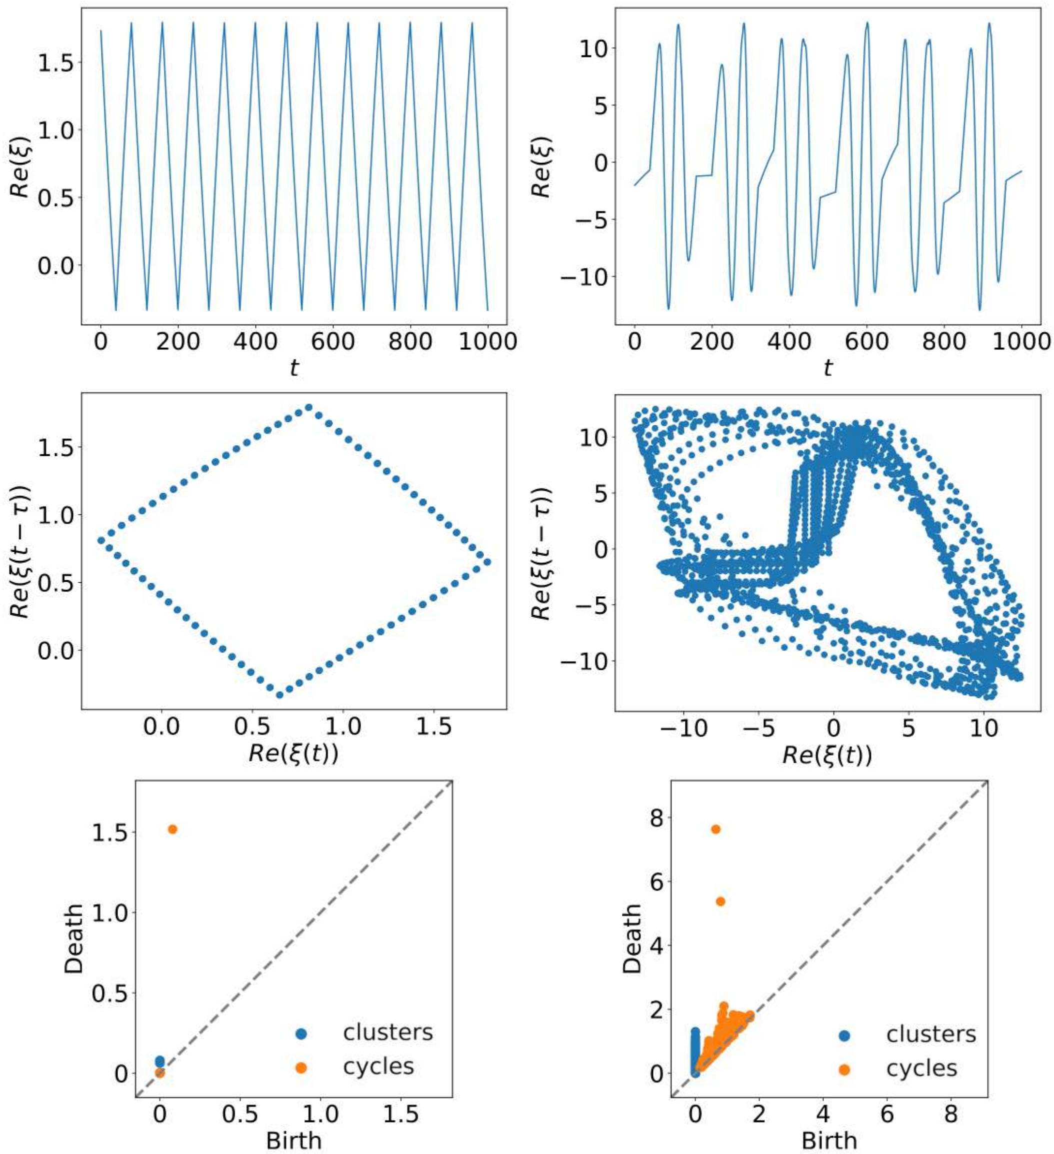

Figure 3.4: Example of TDA pipeline for the classical system [Eq. (3.8)] in regular ($A = 1, T = 8$) (left) and chaotic ($A = 4.5, T = 8$) (right) phases. Time series evolution of real part of dynamical variable (a,b) with corresponding two-dimensional time-delay embedding (c,d), and persistence diagram (e,f) for 0- (blue) and 1-dimensional (orange) features. Embedding time-delay is $\tau = 20$, determined using heuristic methods (see Sec. 3.2).

3.5 Quantum chaos detection

Turning now to the quantum case, we present an approach for detecting chaotic dynamics in open quantum systems using TDA, which draws inspiration from the classical attractor reconstruction method outlined above. The key idea is that instead of the ensemble-based master equation formulation, single quantum trajectories obtained from the quantum jump unravelling can be used as direct analogues to the trajectories in the classical formulation of chaos. Through measurement of a suitable observable, a time series can be obtained whose dynamical properties we can extract through time-delay embedding and the application of TDA as described in Sec. 3.4.

It has been shown previously that as an open quantum system becomes macroscopic, the behaviour of observables $\langle \hat{x} \rangle$ and $\langle \hat{p} \rangle$ gradually resemble that of classical attractors in phase space [123, 124]. However, at small action scales deep in the quantum regime, this classical evolution becomes blurred by decoherence and stochastic dissipative influences of the environment. Our findings suggest that the stochastic evolution of individual states is able to reflect the topology of the underlying motion, robust to the effects of quantum uncertainty. As a consequence, topological analysis of single quantum trajectories provides a quantifier for quantum chaos in open quantum systems. Without requiring the dynamics of the entire density operator, this method is computationally efficient and, more importantly, can be easily applied in experimental settings. For example, single quantum trajectories can be realised through continuous monitoring schemes such as homodyne detection, or through directly measurable observables such as single photon detections (or any other decay/pump channel) corresponding to quantum jumps [113, 114].

We first demonstrate this approach by considering the time evolution of the position expectation value $\langle \hat{x} \rangle$ as our observable to provide a simple comparison with the classical case in Fig. 3.5. The time-evolution of the quantum system described by Eq. (3.3) is simulated using the computational library *QuTiP* [125]. We allow a single quantum trajectory to evolve over $T_{trans} < t < 1000T$, where times before $T_{trans} = 400T$ are discarded to minimise the influence of transient effects. We use Hamiltonian parameters $\chi = 0.008$ and $\gamma = 0.05$, which gives an average photon number 50, and consider a truncated Hilbert space with up to $N = 300$ photons, following Ref. [120]. Using the computed wavefunctions which evolve stochastically under Eq. (3.7), we first obtain the time-evolution for the position expectation value $\langle \hat{x} \rangle$ and apply time-delay embedding as before to obtain a higher-dimensional point cloud that captures an underlying topology of the system dynamics through Takens' theorem. Since the point clouds in Fig. 3.5(c,d) are observed qualitatively to resemble the attractors of the dissipative classical case, we refer to them as 'quantum attractors'. In the classical limit of large mean photon numbers the single photon decays will enter into Eq. (3.8) as a shot noise term that will become negligible in the limit of an infinite number of photons, thus leading to an exact reproduction of the results in Fig. 3.4. For the moderate photon numbers considered here, there is an obvious smearing out of the point clouds due to quantum noise arising from fluctuations in the photon number from stochastic decay events.

Despite the presence of this noise in the quantum regime, the key result to note here is that the topological features of these noisy quantum attractors remain robust. As a consequence, these features can be extracted using persistent homology and used to differentiate dynamical phases as in the classical case discussed in Sec. 3.4. Here, we focus on low-dimensional features which are more efficient to compute than high-dimensional features and capture more intuitively meaningful properties of the topology. We expect, however, that higher dimensional topological features may provide additional useful insights and leave this study to future work.

Given that a true topological description of the dynamics has been captured, we expect the topology extracted by Takens' theorem to be independent of the observable chosen. Thus, we also consider an observable based on the count of quantum jump events, which is easier to obtain experimentally and does not interfere with the intra-cavity dynamics. A time series of binned photon counts is constructed by moving a sliding window across the quantum jump times, as shown in Fig. 3.8(a), which we then process using time-delay embedding in a similar manner. The topological summary results for both observables will be further discussed in Sec. 3.6.

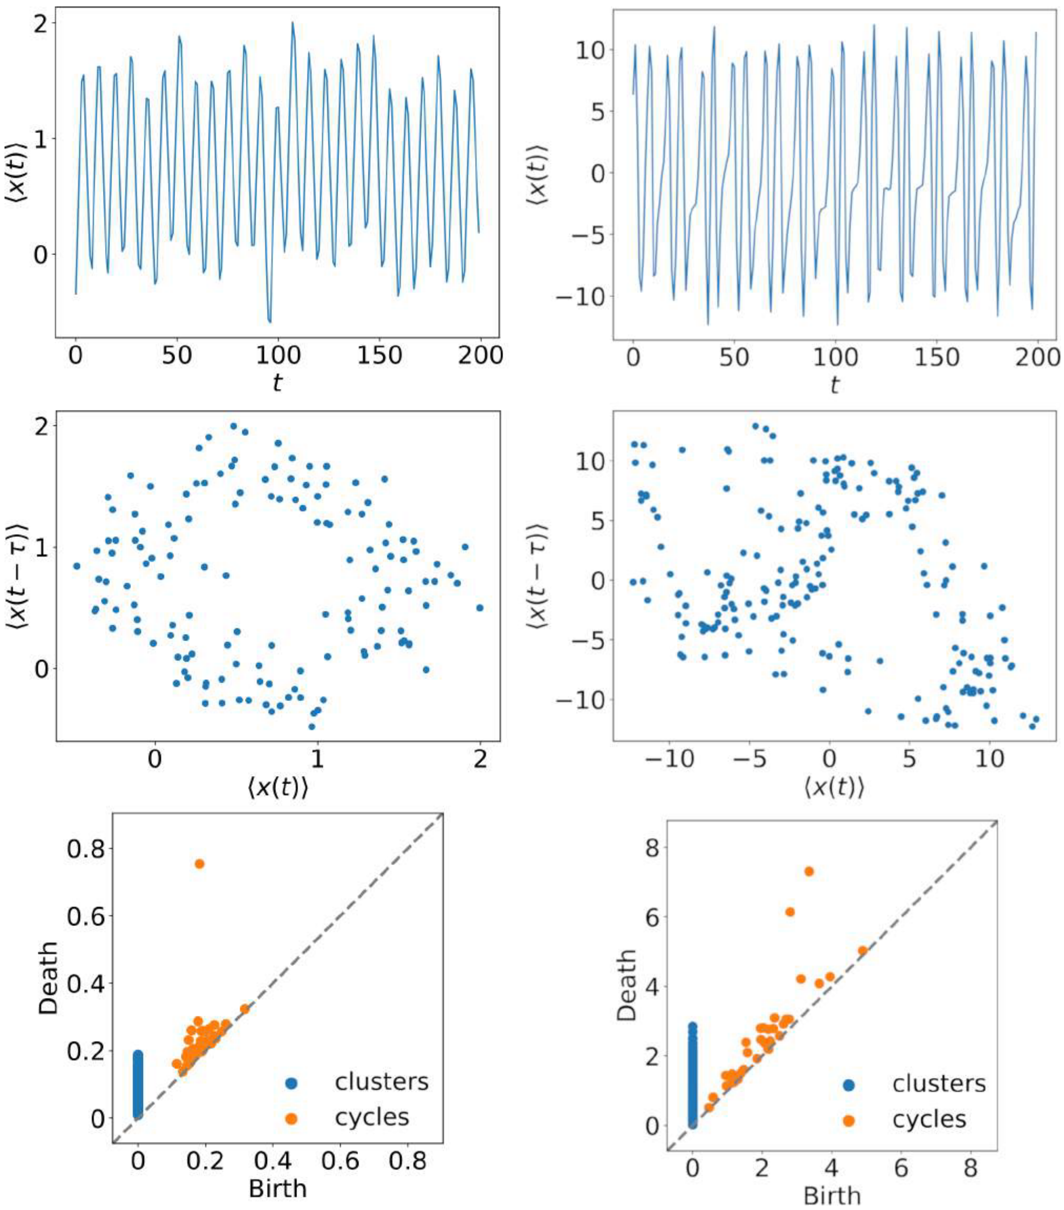

Figure 3.5: Example of TDA pipeline for the quantum system [Eq. (3.3)] in regular ($A = 1, T = 8$) (left) and chaotic ($A = 4.5, T = 8$) (right) phases. Time series evolution of position expectation value $\langle x \rangle$ (a,b) with corresponding two-dimensional time-delay embedding (c,d), and persistence diagram (e,f) for 0- (blue) and 1-dimensional (orange) features. Embedding time-delays are $\tau = 6$ and $\tau = 2$ respectively, determined using heuristic methods (see Sec. 3.2). Note that there are fewer points in the embedded cloud compared to the corresponding classical clouds due to the shorter evolution time simulated.)

3.6 Topological summaries

Various summaries can be used for quantifying the topological information of persistence diagrams, chosen depending on the specific use-case [72, 85, 126]. Here, we consider the average lifetime of low-dimensional features in the persistence diagram, defined as

$$L_{avg}(D) = \frac{1}{|D|} \sum_{h \in D} (d - b), \tag{3.9}$$

where D is the set of all features in the persistence diagram, and b (d) corresponds to the ϵ-value of birth (death) of a particular feature h. This provides a summary similar to that of the Betti number and tells us the average size of topological feature in the reconstructed attractor.

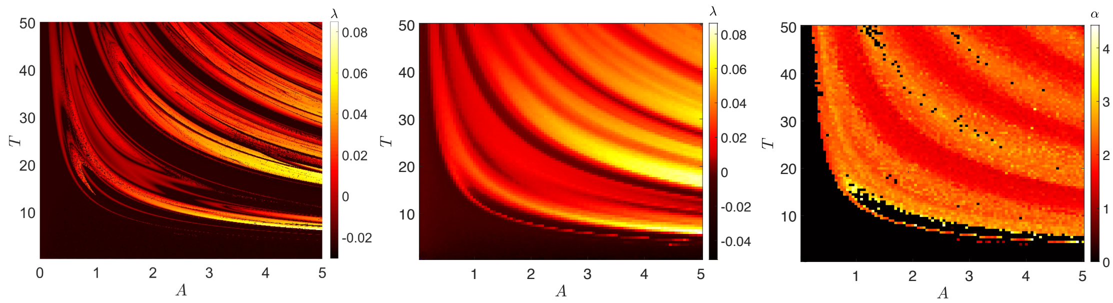

Figure 3.6: (a,b) Largest Lyapunov exponent and (c) power law exponent as a function of amplitude A and period T of the modulations for the (a) classical system and (b,c) quantum system. Adapted from Yusipov et al. [120].

It is known that under different drive parameters A and T, the quantum system of Eq. (3.3) contains regions of regular and chaotic behaviour (see Fig. 3.6) [120, 127]. Here, we consider the driving amplitude between the range of $0 < A < 5$ and driving period between $0 < T < 50$. The average lifetime for 1-dimensional holes of the reconstructed attractors are computed over this (A,T)-parameter space for both the classical and quantum systems, shown in Figs. 3.7 (a) and (b), respectively. The key observation is that the distinct chaotic bands defined by different topological properties of the corresponding quantum attractors are well correlated with the dynamical phases captured by the TDA-based approach applied to the classical system. This suggests that the quantum attractors reconstructed from single quantum trajectories contain topological information about the open quantum system that can be used to distinguish between regular and chaotic dynamics. Note that the blurring of fine structure variations of the classical dynamics within the chaotic phase bands of Fig. 3.7(a) are due to effects arising from the quantum-classical transition [120, 128]. Additionally, we show that this measure is a robust indicator of chaos by demonstrating that the phase distinguishability is independent of free hyperparameters for the time-delay embedding as long as they satisfy the conditions for Takens' theorem (see Section 3.2).

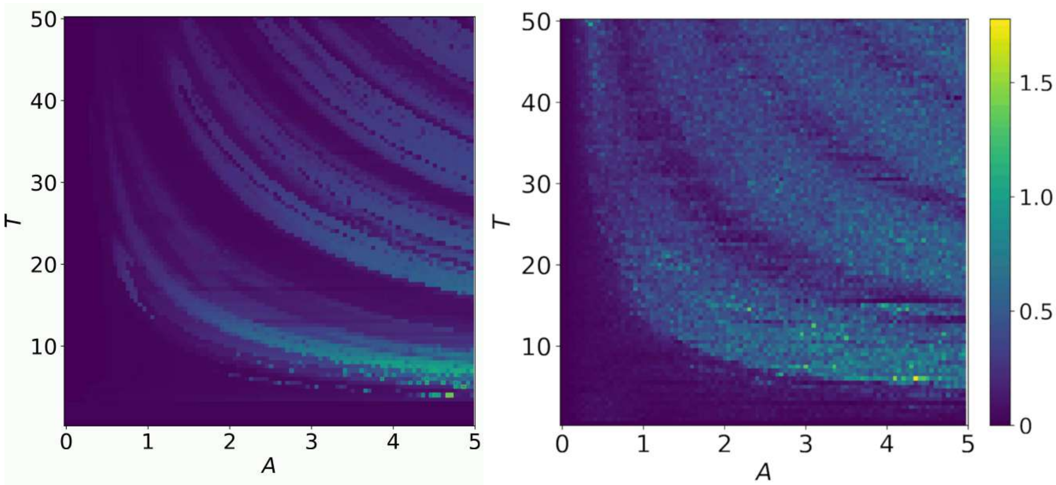

Figure 3.7: Average 1-hole lifetime as a function of driving amplitude A and driving period T for time evolution of (a) real part of dynamical variable for the classical system [Eq. (3.8)] and (b) position expectation value $\langle x \rangle$ of single Monte Carlo trajectory of quantum system [Eq. (3.3)].

Figure 3.8(b) shows the corresponding phase diagram obtained by considering the photon count as the observable and is similarly reminiscent of the classical phase diagram in Fig. 3.7(a). It should be noted that the additional free parameter of the binning interval size has a large influence on these values and the distinction between phases is somewhat less clear. The intervals need to be sufficiently large to filter out zero photon count bins, and consequently the time step-size is larger

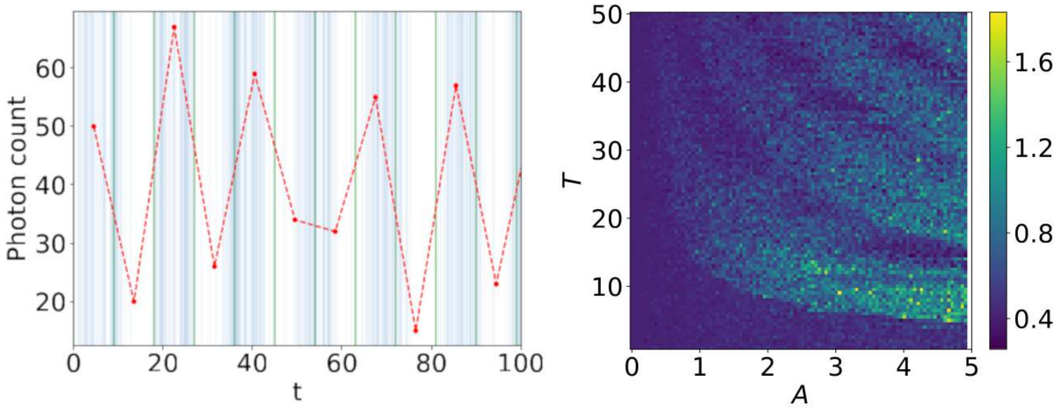

Figure 3.8: (a) Binned photon count time series (red dashed line) for chaotic phase ($A = 4.5, T = 8$) with bin edges (green) and jump event timings (blue). (b) Average 1- hole lifetime as a function of driving amplitude A and driving period T. The binning interval is fixed at 9 and embedding parameters chosen using heuristic methods (see Sec. 3.2).

than for the position operator observable, which results in fewer embedded points for a fixed evolution time. Nonetheless, we emphasize the observation that this simple measurable offers insight into a topological signature of quantum chaos and we leave it to future studies to investigate the optimal observables for probing such topologies of quantum dynamics.

3.7 Conclusion

In this Chapter we introduced a topological approach for the dynamical analysis of open quantum systems, which applies a phase-space embedding method on single quantum trajectories to construct 'quantum attractors' and extract their topological properties using persistent homology. This was inspired by recent applications of topological data analysis to the dynamical analysis of classical systems, underpinned by Takens' theorem which guarantees the preservation of topology under certain transformations of time series data.

We used our quantum pipeline to investigate the dynamical phases of a leaky

Kerr-nonlinear photonic cavity driven by an external field, which can exhibit chaotic behaviour under certain driving parameters. Our results indicate that the single quantum trajectory retains topological information about the entire system dynamics which can be extracted to capture the regular and chaotic phases of the open quantum system. Importantly, these topological measures are found not to be dependent on the observable considered and can be calculated using limited measurement of the system through weak continuous monitoring schemes or other directly measurable observables. Thus, the approach may present an advantage over other methods for quantifying chaos in open quantum systems, which can be experimentally difficult to realize or require measurements of the full wavefunction of set of eigenstates [105, 106, 129–131]. Difficulties arise in part from the global nature of topological properties compared to standard Lyapunov exponent-based techniques, which require many localised measurements to construct a suitable phase space picture.

To the best of our knowledge, this work is the first to apply the framework of topological data analysis for the analysis of chaotic quantum dynamics. We hope that this can open up a new perspective for employing such topological feature selection techniques for studying the dynamics of quantum systems, making use of their inherent robustness to noise and limited measurements, and ability to effectively handle high-dimensional datasets. Another potential application of the approach is as a data reduction step; rather than computing summary statistics of the persistence diagrams, one could feed all the obtained feature lifetimes into a machine learning model such as an artificial neural network that could be trained to distinguish between regular and chaotic dynamics [28, 131].

Chapter 4

Unsupervised learning of quantum many-body scars using intrinsic dimension

In this Chapter, we wanted to see how topological methods could apply to more complicated quantum systems, with quantum many-body scar states chosen as a timely example. In the previous Chapter, reconstruction of the shape of the dynamics assumes access to a sufficiently long time series, whereas decoherence and noise limit the available evolution time or circuit depth in large quantum processors. Moreover, with the huge growth in the Hilbert space size, it is unclear what might be the optimal observable to use for Takens' embedding theorem. A simple observable, such as the average excited state occupation, would not be sensitive to quantum entanglement, while the wrong choice of many-body observable might lead to an extremely weak signal. Accordingly, in this Chapter we needed to device a method that can make the best use of the most easily available measurement data (bit-strings of excited state occupation numbers), allowing us to get to the the same key quantity of interest, the "shape" of the dynamics, in a more measurement-efficient manner.

This Chapter explores the use of intrinsic (effective) dimension as a topological diagnostic tool for detecting and characterising quantum many-body scar states. Unsupervised learning methods have seen growing application in condensed matter and quantum physics, where they have been used to identify phase transitions without labelled data, discover hidden order parameters in strongly correlated systems, and cluster quantum states by their structural properties in an unbiased way [132, 133]. Here we ask whether the distinctive structure in Hilbert space of quantum many-body scar states can be resolved purely from the system dynamics, without prior knowledge of the scar subspace. The material presented in this Chapter is based on the following publication:

- Harvey Cao, Dimitris G. Angelakis, and Daniel Leykam, Unsupervised learning of quantum scars using intrinsic dimension, *Machine Learning: Science and Technology* **5**, 025049 (2024) [134]

4.1 Introduction

Quantum many-body scars (QMBS) are an interesting and recent paradigm describing the weak violation of ergodicity that arises from an intriguing interplay between integrability and chaos in quantum many-body systems. The predictions made by the Eigenstate Thermalization Hypothesis (ETH) no longer remain valid across the full spectrum [135]; a subset of atypical, non-thermal eigenstates exhibit slow relaxation dynamics which retain memory of the system's initial configuration. In particular, these QMBS eigenstates have been found to share high overlap with certain non-equilibrium states which results in persistent oscillations of local observables when the system is quenched from such a state, whereas for generic states the system quickly approaches thermal equilibrium and effectively 'forgets' its initial condition [136].

The set of scar states is exponentially small with respect to the Hilbert space dimension, which raises the question as to whether this phenomenon is too rare to be interesting or experimentally relevant. However, the discovery of anomalously long-lived revivals in a Rydberg blockaded atom array setup showed that not only can QMBS signatures be captured in experiment, but their high overlap with short-range correlated product states makes them surprisingly accessible [57]. Following this initial experimental observation, a lively search for other models hosting scar eigenstates has arisen, as well as analytical construction through direct embedding of scars into the spectrum. This has simultaneously driven work on developing a coherent theory to explain the ergodicity-breaking phenomenon [137–143].

Whilst rapid experimental developments in quantum simulation platforms with low levels of environmental decoherence and high levels of controllability have enabled the exploration of non-equilibrium dynamics of QMBS systems, there are yet to emerge robust analytical methods for identifying systems that belong to non-analytical scarring regimes. The current procedures rely heavily on studying directly measures of return probability or entanglement entropy in order to distinguish them [137, 144–147]. These are not very experimentally accessible, especially as we scale to larger system sizes. Given limited a priori knowledge of the system model, it is important to find a scheme which is not reliant on exponentially many measurements for calculating expectation values or building a full tomographic picture in order to identify the states that exhibit special non-thermal characteristics.

Tools from the field of machine learning have proven to be astonishingly effective at tackling problems that involve the analysis of complex, high-dimensional datasets. In particular, these techniques have found fruitful applications in the context of quantum many-body physics and statistical physics [48, 148]. The outstanding performance of these methods have often been attributed to the idea that this high dimensionality comes about only from the embedding space, while the data itself lies in a manifold with significantly fewer meaningful degrees of freedom. As a

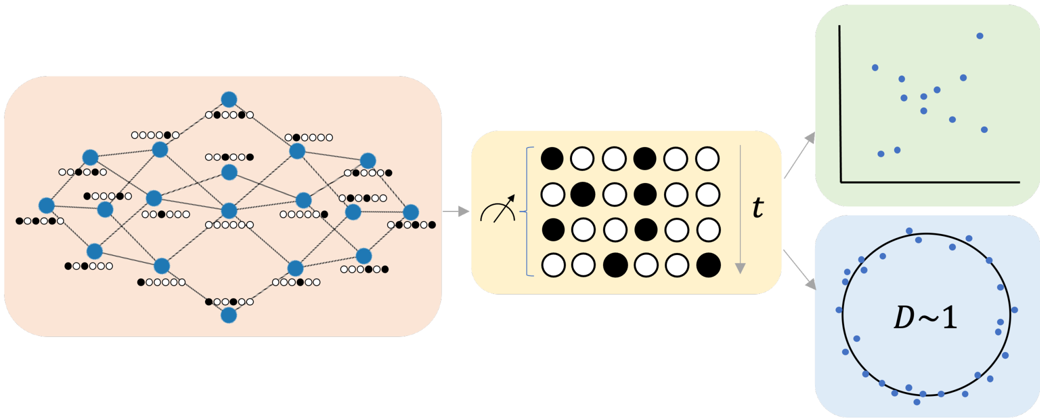

Figure 4.1: Graphical illustration of pipeline for studying quantum scar models using machine learning techniques. Hilbert space graph for PXP model ($L = 6$ atoms) with links between spin configurations where there is an allowed transition (red). The system is sampled P times at N different timesteps (yellow). Multidimensional scaling embeds this data in a lower-dimension while preserving distances in a chosen metric (green) and intrinsic dimension estimation extracts the underlying structure directly from the sampled data (blue).

result, datasets that might be initially hard to describe can have their descriptions condensed into several key features of the latent manifold which allow for easier visualization and effective learning. Within the context of many-body physics, a few recent works have applied dimensionality reduction techniques from machine learning for facilitating ground-state searches and identifying distinct phases of spin models [28, 149–154].

In this Chapter, we consider two ML techniques in the context of manifold learning to study the scar properties of QMBS systems. The first of these is multidimensional scaling (MDS), a nonlinear dimensionality reduction method which allows high dimensional data to be translated into a lower dimension whilst approximately preserving a measure of distance between data points. Observing that there is dynamical structure in the lower-dimensional manifold, we then investigate the property of intrinsic dimension (ID), which describes the effective number of variables required for minimal representation of a dataset. Through applying an estimation technique to extract the ID directly from shot measurements of the system, we find that the manifold for the PXP model can be distinguished from thermal states in an unsupervised manner. In particular, we show that this scheme is robust to limited shots and realistic measurement noise, making it feasible for experimental implementation, e.g. using Rydberg atom quantum simulators [155]. The graphical illustrations in Fig. 4.1 demonstrate the pipeline for our approach, which uses sampled data from the PXP model to both embed in a low-dimensional space and obtain an intrinsic dimension estimate. We note that there has been some recent interest in the detection of quantum scars using neural network approaches [156, 157]. These, however, are not truly unsupervised methods and require partial training of the neural networks using existing knowledge of some scar states, which may not always be accessible. They are also limited by their scalibility to larger system sizes, which further motivates our use of non-parametric ML models.

The outline of this Chapter is as follows: Section 4.2 provides a brief introduction to quantum many-body scarring and the PXP model. Section 4.3 presents a dimensionality reduction scheme for analysing dynamical many-body states using multidimensional scaling and a physics-informed distance measure. Section 4.4 proposes an unsupervised approach for identifying many-body scar and thermalising initial states using an intrinsic dimension estimation method. Section 4.5 discusses the results obtained using these approaches, and Section 4.6 makes concluding remarks and suggests future directions.

4.2 Quantum many-body scar model

In an array of Rydberg atoms, the individual atoms may be viewed as an effective two level system with atoms occupying either a ground or Rydberg excited state, which we denote by $|\circ\rangle$ and $|\bullet\rangle$, respectively. When subject to a driving laser field, the atoms undergo local Rabi oscillations and are free to flip between the two states, as illustrated in Fig. 4.2(a). The Hamiltonian describing a Rydberg chain of length L is given by [57]

$$H_{\text{Rydberg}} = \sum_{i}^{L} \left(\frac{\Omega}{2} X_i - \Delta n_i \right) + \sum_{j<k}^{L} V_{j,k} n_j n_k, \tag{4.1}$$

where the parentheses include a Rabi precession term represented by the Pauli $X_i = |\circ\rangle\langle\bullet| + |\bullet\rangle\langle\circ|$ operator acting on site i driven at frequency Ω and a local density term $n_i = |\bullet_i\rangle\langle\bullet_i|$ with laser detuning Δ. If the Rydberg atoms are assembled close together, atoms in excited states will interact via repulsive van der Waals forces, with an interaction strength $V_{j,k} \propto 1/R^6$, where R is the distance between the atoms j and k. In the limit where the atoms are brought very close together and the strength of nearest neighbour interactions is tuned to be much larger in amplitude than the detuning and Rabi frequency of the laser field ($\Delta = 0, V \gg \Omega$), a regime of the so-called Rydberg blockade can be achieved where simultaneous excitations of neighbouring Rydberg atoms are energetically prohibited [136]. Correspondingly, the system and its Hilbert space becomes kinetically constrained in a way that forbids configurations of the form $|\cdots \bullet \bullet \cdots\rangle$. This model is referred to as the PXP model for the form of its Hamiltonian, expressed as

$$H_{PXP} = \Omega \sum_{i} P_{i-1} X_i P_{i+1}. \tag{4.2}$$

The excitation constraints are enforced by the projectors $P_j = |\circ_j\rangle\langle\circ_j|$, which allows a site j to enter the Rydberg excited state only if both of its neighbours are in the atomic ground state. In this Chapter, we use periodic boundary conditions (PBC) and normalize time t by the Rabi frequency which we set to $\Omega = 1$.

This constrained model was shown in simulations [136] and experiments [57] to exhibit a peculiar absence of thermalization when the system was quenched from

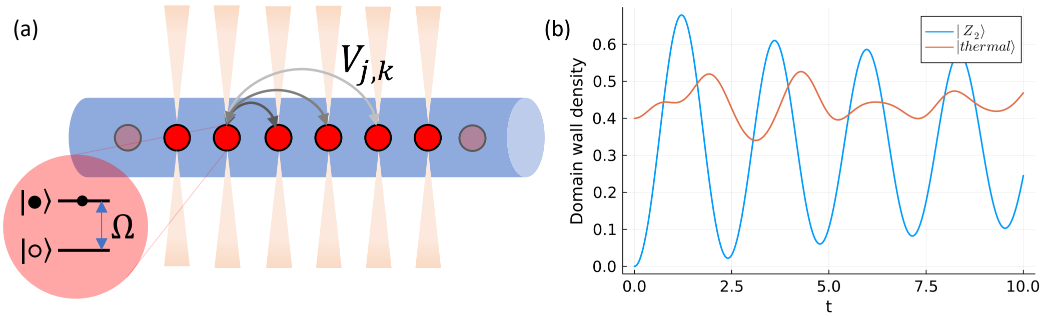

Figure 4.2: PXP system and dynamics: (a) Neutral atoms are trapped in an array using optical tweezers and driven with a laser of Rabi frequency Ω, enabling transitions to a Rydberg state and van der Waals interactions $V_{j,k}$ between the atoms. (b) Persistent revivals in the domain wall density are observed when the PXP system is quenched from the $|Z_2>$ initial state (blue line), but not from other thermalising initial states (orange line). Domain walls occur when two neighbouring atoms are in the same state or a ground state atom is at the edge of the array, and the density is computed by considering all neighbouring atom pairs across the entire chain.

certain special initial states such as the $|Z_2\rangle = |\circ \bullet \cdots \circ \bullet\rangle$ (Néel) state. This was demonstrated through the observation of periodic revival behaviour in local observables such as density of domain walls (see Fig. 4.2(b)), which occur where two neighbouring atoms are in the same state or a ground state atom is at the edge of the array. In contrast, fast relaxation to a thermal value is seen when quenched from other initial states such as $|\mathbf{0}\rangle = |\circ \cdots \circ\rangle$, and remains near that value for future times. The slow relaxation phenomenon of local observables arises directly from the ability of the system to exhibit quantum wavefunction revival when quenched from the special states. The origin of this behaviour can be understood intuitively by examining the dynamics of the PXP model on a Hilbert space graph, as illustrated in Fig. 4.1. The dynamics is sufficiently constrained by the structure of the graph such that the wavefunction finds itself returning to its initial point through serendipitous propagation and reflection when initialized from certain nodes.

Without knowledge of the entire spectrum, there is no universal framework for determining the special initial states that allow these non-thermalization properties to be observed in the system dynamics. Significant effort has been directed to the discovery of novel QMBS models, typically with the scar states embedded directly into the spectrum [158, 159]. However, for kinetically constrained models like the PXP model where the scarring properties are analytically unknown, it is a difficult task to detect where the scar manifold lies. In this Chapter, we employ machine learning techniques to learn properties of the scar manifold and develop an unsupervised scheme for detecting the special initial states of an unknown QMBS model.

4.3 Dimensionality reduction of QMBS

In recent years, the daunting task of navigating the high dimensional Hilbert spaces of quantum many-body systems has been tackled by techniques from machine learning. Their effectiveness in this application stems from the idea that the underlying manifolds of most natural datasets are smaller in dimension than the raw embedded dataset, known as the manifold hypothesis. Thus, dimensionality reduction techniques can be used to convert the unwieldy, high-dimensional data found in quantum many-body systems into a more compact form while still preserving its key features. These include methods such as principal component analysis (PCA) and diffusion maps, which are used for clustering or visualisation by projecting the data to a lower-dimensional space through minimization of an objective error function [150, 152].

Multidimensional scaling (MDS) is another nonlinear optimization-based method for manifold learning that seeks to embed high-dimensional data into a lower-dimensional space. This mapping is found while preserving the pairwise distances between data points as much as possible, where the metric defining the distances can be chosen by the user. If the Euclidean distance metric is used, this method becomes equivalent to PCA, as minimizing linear distance becomes equivalent to maximising linear correlations. MDS is generally useful for visualization purposes and for gaining insights into the underlying structure of data that might not be apparent in its original high-dimensional form [68, 160–162], especially when there is already some intuition as to which measures are appropriate.

Given M objects which lie in some high-dimensional data space and are defined by pairwise distances $d_{i,j}$, MDS finds a mapping into a chosen d-dimensional space while preserving the pairwise distances between the data as much as possible. The output embedded data $x_1, ..., x_M \in R^d$ is such that $||\boldsymbol{x}_i - \boldsymbol{x}_j|| \approx d_{i,j}$ for all $i, j \in 1, ..., M$. The particular embedding is determined through minimisation of the following loss function

$$S_{MDS} = \sum_{i<j} (||\boldsymbol{x}_i - \boldsymbol{x}_j|| - d_{i,j})^2, \tag{4.3}$$

which is typically carried out by numerical optimization techniques such as gradient descent or stress majorization [160, 163]. Note that the specific algorithms used, as well as differences in initialization and presence of noisy data, result in an embedding that is not necessarily unique. However, while the specific configurations of points in the lower-dimensional space may vary, the goal of MDS is to capture the essential relationships of the data, and this method is often used in conjunction with other methods to further explore and validate the data structure.

Generally, considering a set of snapshots of the evolution of a dynamical system, the compressed representation should be isomorphic to the latent underlying intrinsic state of the system. In this Chapter, we employ MDS to reveal potential redundancies in the high-dimensional dynamical full state wavefunction data of the PXP model and demonstrate that the low-dimensional representation obtained via MDS facilitates the identification of scar phases. In order to capture the relationships that we wish to preserve in the embedding, a suitable distance measure between the data must first be chosen.

Here, we consider a probabilistic earth mover's (PEM) distance between two states ψ_A and ψ_B which is inspired by the similarly-named metric from optimal transport theory [162, 164]. If $P_A(i)$ is the probability of obtaining basis state i when the system is in state ψ_A and $P_B(j)$ the probability of obtaining outcome j when the system is in state ψ_B, the PEM is defined as

$$d_{EM}(\psi_A, \psi_B) = \min_{\gamma} \sum_{i,j} \gamma_{ij} d(i,j), \tag{4.4}$$

where $\gamma_{ij} \geq 0$ is a coupling map between the basis states i and j of ψ_A and ψ_B such that $\sum_j \gamma_{ij} = P_A(i)$ and $\sum_i \gamma_{ij} = P_B(j)$. The distance between basis states is given by the Hamming distance $d(i,j)$ and the minimum is taken over all possible couplings $\gamma = (\gamma_{ij})$.

In order to detect scarring through the state dynamics, an appropriate distance measure needs to reflect the structure of the dynamics and capture its constraints. We emphasize that the definition of PEM uses physically relevant information encoded in the metric structure $d(i,j)$ of the underlying observable space. The dynamics generated by the PXP Hamiltonian involves effective single spin flips and so using the PEM metric based on the Hamming distance captures an effective minimal transition time between two different basis states. To capture the dynamics of the system in the MDS embedding, we take the full state wavefunction $\psi_m(t_n)$ at $n = 1, ..., N$ timesteps, quenched from $m = 1, ..., M$ initial states as our input data. A distance matrix for each initial state is obtained by finding the pairwise distances between the states $\psi_m(t_i)$, $\psi_m(t_j)$ under the dynamics-inspired PEM metric. After applying the MDS embedding, the outputs for each evolution from a different initial state are N vectors in a lower-dimensional space. Here, we find that an embedding into only two dimensions is sufficient to uncover the dynamical characteristics of the scar manifold.

Figures 4.3(a) and 4.3(b) show that using the Hilbert-Schmidt distance in MDS does not reveal any distinct structural features that allow the thermal and scar initial states to be distinguished in the 2-D embedded space. By contrast, it is evident that using the PEM metric provides considerably more insight into the dynamical structure of the scarring behaviour compared to this standard distance measure, as illustrated by the plots in Figs. 4.3(c) and 4.3(d). This tighter arrangement for the embedding of the scar initial state dynamics reflects that, within the PEM metric, the states are bound closer together and require fewer degrees of freedom to describe. Similarly dense structures are observed for other special initial states, such as $|Z_3\rangle$.

In order to appreciate how the PEM metric gives us the ability to distringuish between scar and thermal phases, we notice that it hinges on probability amplitudes of bitstrings rather than state fidelities

$$|\psi_n(t)|^2 = \sum_{ij} c_i^* c_j \langle \phi_i | \boldsymbol{n} \rangle \langle \boldsymbol{n} | \phi_j \rangle e^{-i(E_j - E_i)t}, \tag{4.5}$$

where $\boldsymbol{n}$ is a given bitstring measured in the computational basis, ϕ_i are the energy eigenstates, c_i is the amplitude with which the eigenstate ϕ_i is excited, and E_i are the energy eigenvalues.

This measure is predominantly dependent on two quantities: the projection of the initial state on the computational basis and the distribution of expansion coefficients in the energy eigenbasis. For thermal initial states, the individual initial state

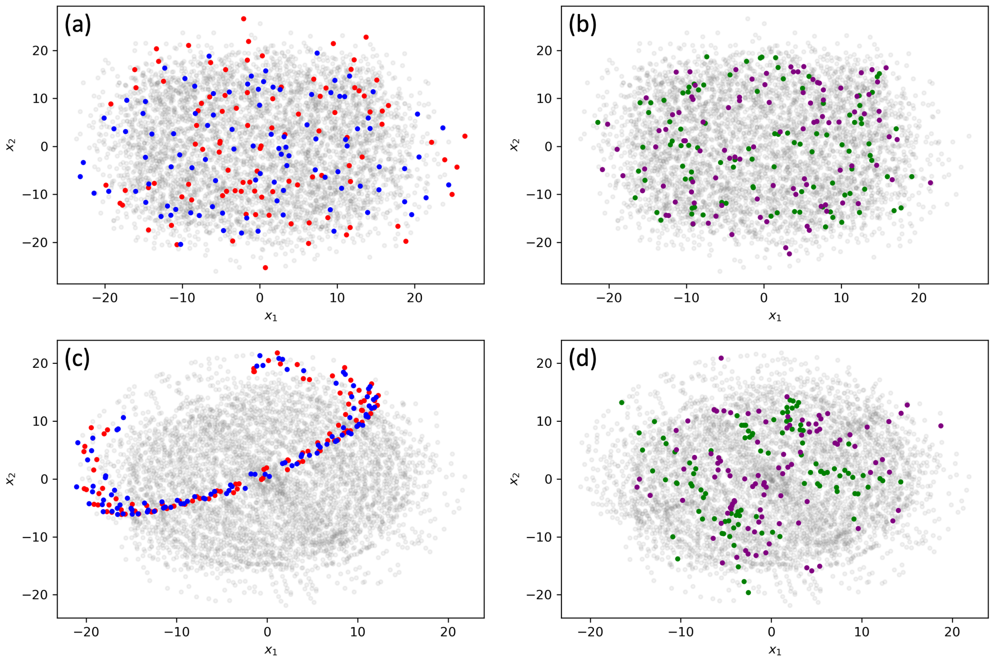

Figure 4.3: Multidimensional scaling 2D embedding of PXP dynamical state ($L = 10$) using (a,b) Hilbert-Schmidt distance and (c,d) PEM distance, quenched from all 123 possible initial spin configurations (grey). (a,c) Embedded points obtained from evolution trajectories of two different scar initial states, highlighted in red and blue. (b,d) Embedded points obtained from evolution trajectories of two different thermal initial states, highlighted in green and purple. Note that a separate MDS embedding is obtained from the evolution trajectory of each initial state, and the projection axes (x_1, x_2) have been aligned to illustrate the similarities between the two metrics of the collective embedded data.

amplitudes projected onto the ith eigenstate become vanishingly small during the system evolution due to spreading across many states. Furthermore, we find that the expansion coefficients in the energy eigenbasis are distributed randomly (see Fig. 4.4(b)), which leads to dephasing over terms in the distance measure. Scarring initial states, on the other hand, do not spread significantly and maintain periodic proximity to eigenstates. They also have a discrete distribution of expansion coefficients (see Fig. 4.4(b)) caused by towers of scars in the energy spectrum [165]. The combination of these effects enables the phases to be distinguished by the distance measure which exhibits periodic revival behaviour for scar dynamics,

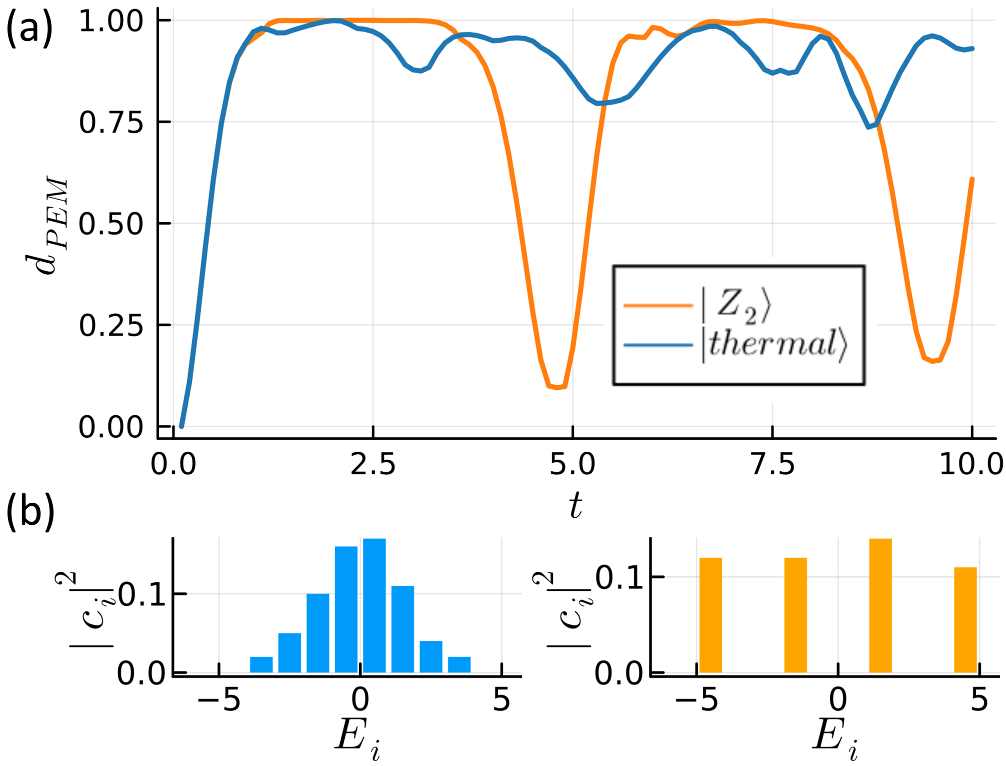

Figure 4.4: Plots showing (a) probabilistic earth mover's metric between initial and evolved state for scar (orange) and thermal (blue) initial states and (b) corresponding distributions of expansion coefficients in energy eigenbasis, sampled over 100 evolutions.

as shown in Fig. 4.4(a). Through MDS, which aims to preserve the inter-point distances, this periodic shrinking of the measure is manifested in the tight structure of the resulting embedded points for the scar dynamics, as seen by the red and blue points in Fig. 4.3(c).

4.4 Intrinsic dimension of scar dynamics

In a typical experiment, access to measures based on the full wavefunction data is limited because exponentially many measurements are required to define a complete set of state parameters [166]. The following section proposes a potential avenue for overcoming the experimental limitations of the scheme discussed in

Section 4.3, by targeting the dimensionality reduction process itself rather than the embedded result. This allows for capturing the desired structural features of the system without a complete description of the dynamics. Specifically, we target a property known as intrinsic dimension (ID), which describes the number of variables required for minimal representation of a dataset without significant information loss [167, 168]. In general for natural datasets, the ID is expected to be much lower than the dimensionality of the data due to the presence of structural correlations between variables. This concept can be visualized by imagining a set of Cartesian data points extracted from a circle, as visualised in Fig. 4.1. Although two input coordinates are required to describe the data, they are collected from a manifold with ID = 1 which causes the points to be strongly correlated.

Whilst this quantity is often considered implicitly in dimensionality reduction techniques to determine to what extent a dataset can be compressed into a lower-dimension, information about the ID can also be used directly to characterize a system in an unsupervised manner. Recently, ID estimation was used to study phase transitions in spin lattices and dynamics of BECs, showing that ID can be seen as an order parameter which signals a transition between different structures in configuration space [154, 169–171]. For dimensionality reduction methods, the dimension of the identified subspace is often viewed as an estimation of the ID. However, these methods generally come with various conditions which make the estimation of ID a nontrivial task. For example, PCA fails to be effective if the lower dimensional manifold is curved or nonlinear. As a result, various methods have been developed which provide a direct estimate for the ID from the dataset as a whole, without any projection steps that may affect the structural relationships within the data [172, 173].

When dealing with discrete spaces, an ID estimator will try to find the dimension of a hypercubic lattice which locally approximates the structure of the original data. The method used in this Chapter considers the statistics of points within lattice volumes surrounding each data point and makes use of certain relationships in polytope theory to infer the dimension of the entire region, which approximates the ID of the original discrete dataset [174]. The user has the freedom of deciding the scale at which the ID is computed by selecting a length scale parameter for the lattice. In general, it is unknown what scale may be appropriate for a given dataset, so the ID is first computed over a reasonable range of scales. A suitable scale is then decided upon by seeing at which value the ID plateaus, indicating the scale at which the global ID dominates. More details are given in Section 4.4.

Since we would like to improve the experimental accessibility of the method, instead of the full state probability distribution we take as data projective measurements of the state at different evolution times. Specifically, we sample N bitstrings from each wavefunction $\psi_m(t_n)$ at $n = 1, ..., P$ timesteps. This set of bitstrings for each initial state quench m are then respectively input into the ID estimation method with the Hamming distance as a metric. As the subset of initial states leading to scar dynamics is exponentially small, we expect to distinguish these as outliers to the intrinsic dimension of dynamics from thermal initial states. Therefore, given sufficient data points, we can use a simple statistical summary to identify the outlying ID points, allowing an unsupervised clustering of the initial states into scar and thermal initial states.

Intrinsic dimension estimation for discrete spaces

Here we outline the method used in this Chapter for ID estimation of discrete data from [162]. Assuming a set of independently generated data points with density ρ, the probability of observing n points in a given region A follows a Poisson distribution. Now, consider a specific data point i which lies within two regions A and B, where one contains the other, such that $i \in A \subset B$. It can be shown that the conditional probability of having n points in A given that there are k points in B follows a binomial distribution whose success probability p is a ratio of the volumes of the two regions. Remarkably, this probability is independent of the density ρ, which allows the estimator to be used even when the density varies across large distance scales.

Now, we can write a likelihood function of the observations n_i as

$$L(n_i|k_i, p_i) = \prod_{i=1}^{N} \binom{k_i}{n_i} p_i^{n_i} (1 - p_i)^{k_i - n_i}, \tag{4.6}$$

where p_i and k_i are parameters which may possibly be point-dependent. Taking our space to be a lattice with the L_1 metric as a natural choice, the volume of a given region $V(A)$ is simply the number of lattice points contained in A. According to Ehrhart theory of polytopes, the number of lattice points within distance t in dimension d from a given point amounts to

$$V(t, d) = \binom{d+t}{d} {}_2F_1(-d, -t, -d - t, -1), \tag{4.7}$$

where ${}_2F_1(a, b, c, z)$ is the ordinary hypergeometric function. For a given t, this is a polynomial in d of order t, which makes the ratio of volumes in the likelihood function a ratio of polynomials in d. Given a dataset, the choice of two distances t_1 and t_2 which define the scale at which the data is probed fixes the likelihood function in Eq. (4.6). This can be maximised with respect to the dimension d to infer an estimate for the ID of the data manifold, given by the root of the following equation which can be obtained via standard optimization techniques

$$\frac{V(t_1, d)}{V(t_2, d)} - \frac{\langle n \rangle}{\langle k \rangle} = 0, \tag{4.8}$$

where the mean values over n and k are computed over all data points [174]. Note that, the observation resolution determined by choices of t_1 and t_2 has considerable impact on the estimated ID if the scale is comparable to the noise level and curvature of the data. Thus, the ID is first estimated over a range of scales and then decided upon by seeing at which value it plateaus, indicating the scale at which the global ID dominates. This multi-scale approach bears particular similarity to persistent homology, which also considers a range of scales to find the most meaningful signal.

4.5 Results

Here, we use boxplots to provide a statistical summary of the resulting ID estimates, shown in Fig. 4.5. An emphasis is placed on outlying data (fliers) shown in red points, which allows a distinction to be seen between the thermal majority and scarring initial states. Figure 4.5(a) shows the estimated values from samples of each initial configuration in the large measurement limit. The two initial states which have outlying sample ID correspond respectively to the $|Z_2\rangle$ and $|Z_2'\rangle = |\bullet \circ \cdots \bullet \circ\rangle$ states, which are the special initial states that result in scar dynamics for the PXP model. This is an unsurprising result as we are effectively considering the full probability distributions which we noticed in Section 4.3, when the scar state dynamics are embedded in a lower-dimension, are restricted in the embedding space as compared to the thermal states. The estimated ID value of $\sim$1.2 can be interpreted as the data having approximately an ID of $\sim$1 but containing noise which inflates the dimension slightly.

The main band of points spread around ID ≈ 2 correspond to the thermal majority of initial states. This value of the ID similarly agrees with the results from Fig. 4.3(d), which suggests that when the wavefunction data is projected using MDS, the thermal initial states result in an embedding which extends across a two-dimensional space. These dimensions arise from the time evolution of the dynamical system and the deviation away from the scar manifold. The variation can again be attributed the presence of noise, although we note that, the $|Z_3\rangle$ and other periodic charge-density states can be found lying at the bottom of the band. These states have been found to exhibit some characteristics of scarring behaviour such as periodic quantum revival, although with much higher rates of decay [136].

Here, we investigate the robustness of this ID distinguishability between scar and thermal initial states under experimentally realistic limitations. This includes reducing the number of samples made at each timestep and the total number of timesteps, as well as introducing sampling errors. These are readout errors caused by erroneous measurement operations which lead to sites being incorrectly identified as being in the ground or excited state. We consider error rates up to typical experimental error rates of around 1% [175].

Figures 4.5(b) and 4.5(c) illustrate that with limited measurement times P or limited number of measurements N respectively, the range and IQR are increased but the ID remains sufficiently distinct between thermal and scar dynamics. We note that when both P and N are reduced significantly, the failure of the data to sufficiently represent the dynamics causes this distinction to become blurred out. In a similar fashion, Fig. 4.5(d) shows the ID estimation boxplot when measurement errors are applied to the shot samples. Up to an error rate of 5%, the ID estimation can still be used to identify the scar initial states. Note that the introduction of noise shifts the ID estimation of all states upwards due to the addition of bitstrings normally forbidden by the constraints of the PXP model. However, the magnitude of this shift is similar for the thermal and scar states, so their distinguishability is unaffected.

In order to show that our method is generally applicable to other quantum many-

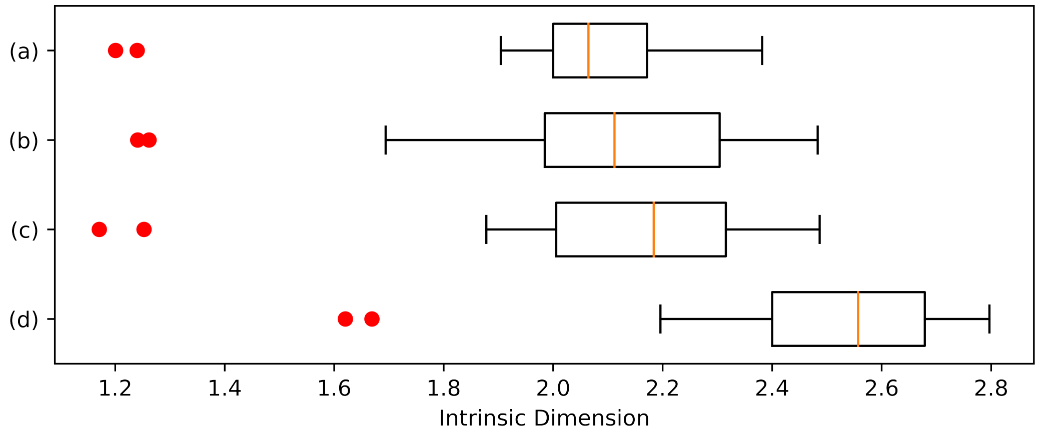

Figure 4.5: Boxplots showing the quartile summaries of the intrinsic dimension estimation for varying shots per time step N, time steps P and measurement errors ϵ: (a) $N = 500, P = 20$, (b) $N = 50, P = 10$, (c) $N = 10, P = 50$, (d) $N = 50, P = 10, \epsilon = 0.05$. The orange vertical line corresponds to the median of the data and the red points correspond to fliers lying outside of the inter-quartile range.

body scar systems, we also demonstrate the ID pipeline on a one-dimensional spin-1 XY model which was recently shown to host scar states [137]. The model Hamiltonian is given as

$$H = J \sum_{\langle ij \rangle} (S_i^x S_j^x + S_i^y S_j^y) + h \sum_i S_i^z + D \sum_i (S_i^z)^2, \tag{4.9}$$

where $S_i^\alpha (\alpha = x, y, z)$ are spin-1 operators acting on sites i of the chain, $\langle ij \rangle$ represents nearest neighbours, and the parameters J, h and D respectively are the coupling constant, magnetic field strength and single-site anisotropy. While there is a no longer a hard dynamical constraint (as was the case for the PXP model), the evolution still involves changes to the z component of the spin, suggesting the Hamming distance should should still capture an effective transition time between different spin configurations measured in the S_z basis.

Under quenching from certain initial states, this model exhibits nonthermal behaviour. The following “nematic Néel" initial state in Eq. (4.10) was shown to exhibit periodic revivals in the many-body fidelity, while generic initial states decay rapidly [137]

$$|Z_{Nem.}\rangle = \bigotimes_i \frac{|m_i = +1\rangle - e^{ir_i\pi}|m_i = -1\rangle}{\sqrt{2}}, \tag{4.10}$$

where $m_i = -1, 0, +1$ are the eigenvalues of S_i^z and r_i are the coordinates of the i-th spin. In Fig. 4.6, we demonstrate that the ID for dynamics from the nematic Néel initial state is distinctly separated from the ID of generic thermal initial states, allowing the scarring nature of the nematic Néel state to be inferred directly from the ID. Note that the ID values are higher than those for the PXP model due to the increased number of degrees of freedom and size of Hilbert space, but the separation between the thermal and scar sample IDs remain similar.

4.6 Discussion and Outlook

Quantum many-body scar models exhibit a deviation from the expected thermal behaviour of many-body systems, resulting in a robustness to decoherence that sees potential applications in quantum control and quantum metrology [155, 176, 177]. Without a priori knowledge of the model and access to its full spectrum or entanglement scaling, it is difficult to determine the initial states for system quenching which allow the properties of scarring to be observed in the dynamics. Thus, we employ techniques from machine learning, and in particular dimensionality reduction, to learn these scar properties directly from experimentally accessible data. A low-dimensional embedding of the full wavefunction evolution using multidimensional scaling reveals structural distinctions between thermal and scar states and is suggestive of dynamics in a lower-dimensional manifold than the state space. We aim to learn structural properties of this manifold directly from the intrinsic dimension of the system dynamics, which we estimate using a recent

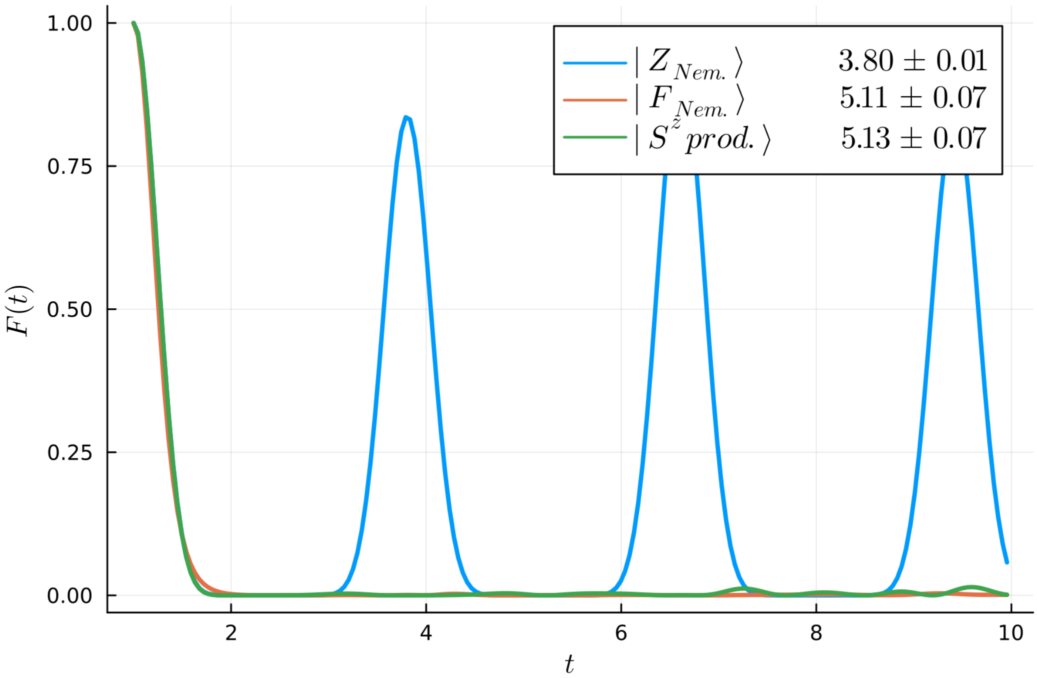

Figure 4.6: Many-body fidelity $F(t) = |\langle\psi(0)|\psi(t)\rangle|^2$ for various initial states for the spin-1 XY model (L=8, J=1, h=1, D=0.1) with PBCs. Persistent revivals are observed when quenched from the nematic Néel $|Z_{Nem.}\rangle$ state, but not from other thermalising initial states (we consider the nematic ferro $|F_{Nem.}\rangle$ state and $|S^z\, prod.\rangle$ state as in Fig. 3 of [137]). (Inset) The intrinsic dimension estimation values (N=100, P=200) for the three initial states discussed above. Two random initial states were also tested, giving ID estimates of 5.11 ± 0.07 and 5.10 ± 0.07 respectively.

method designed for discrete spaces. As such, this method can be applied in an unsupervised fashion to shot data sampled from the system at different evolution times when quenched from different initial product state configurations. The results indicate that the intrinsic dimension for scarring and thermalizing states can be used to distinguish between these phases. This method is also shown to be robust to limited samples and measurement times, as well as typical experimental measurement errors.

The discovery of many-body scarring arose as a result of experimental advancements in cold atom setups using optical tweezers and now there have already been a variety of scarring behaviours observed in other experimental platforms, such as superconducting qubit processors and bosonic quantum gases [178, 179]. We envisage that the approach introduced in this Chapter provides a general and experimentally accessible method of detecting scarring states, requiring only a (polynomially) small number of measurements in the computational basis and being robust to generic measurement errors. As such, given a set of initial states and quench dynamical samples, our method allows efficient and unsupervised clustering of the initial states into thermal and scar classes. In future studies, this connection between the intrinsic dimension and dynamical signatures could be extended to finding suitable initial states over an unknown landscape, for example using gradient-based optimization techniques.

Our work provides a novel perspective on the analysis of QMBS from the perspective of dynamical manifold intrinsic dimension, which may see potential theoretical links to network and complexity theory in future studies [180–183]. We also note that the structures observed in the MDS embedding of scar dynamics are reminiscent of those seen in the phase-space of single-body quantum scar eigenstates, which are similarly embedded in a lower dimensional manifold defined by the classical periodic orbits. This may be suggestive of a deeper connection related to near-integrability or dynamical symmetries [184–186].

Chapter 5

Topology-inspired nonlinear response functions for the Schrödinger equation

In Chapters 3 and 4, we have applied computational topology to understand the dynamics of different quantum systems, providing a new lens for studying and identifying (known) dynamical phases. In this Chapter, we turn this idea around and ask a fundamentally different question: can computational topology be used not merely as an analytical framework, but as a means of engineering and exploring new physics through defining new kinds of nonlinear or many-body dynamics?

Recent experimental advances have enabled the realisation of nonlocal nonlinearities and density-dependent gauge fields in photonic lattices and Bose–Einstein condensates [187, 188]. Inspired by these advances, we introduce a persistent-homology-inspired nonlinearity in the nonlinear Schrödinger equation, whose nonlinear response is governed by the topological structure of wavefunction intensity extrema rather than conventional local Kerr effects, and demonstrate that this construction gives rise to interesting and experimentally-observable dynamical phenomena. The material presented in this Chapter is based on the following publication:

- Harvey Cao and Daniel Leykam, Cusp solitons mediated by a topological nonlinearity, *Chaos* **36**, 033104 (2026) [189]

5.1 Introduction

The nonlinear Schrödinger equation (NLSE) appears throughout physics to model systems where wavelike self-interactions occur. When this nonlinearity is balanced by dispersion, it gives rise to interesting phenomena such as modulational instability and the formation of stable, non-dispersive topological structures known as solitons. These self-stabilizing solutions are widely studied for both their mathematical properties as well as their considerable potential applications in information and communications technology [190–194].

Such behaviour has been found to be universal and solitons have been identified in a wide range of physical systems: within the field of ultracold atomic physics, the NLSE describes Bose-Einstein condensates (BECs) in the limit of a dilute gas of weakly-interacting bosons; in optical systems, the NLSE describes light propagation through nonlinear media, resulting in behaviour such as self-focusing beams; solitons also feature in the study of plasmas under certain conditions [192, 195–197].

Since the initial investigations of the NLSE, a wide range of studies have explored its diverse generalizations, revealing a variety of nonlinear phenomena and an understanding of the criteria for soliton stability [198, 199]. In particular, extensions to incorporate higher-order and nonlocal nonlinearities have demonstrated that solitons can be stabilized against dynamical collapse under appropriate conditions [200–203]. Crucially, these nonlocal nonlinear responses are intrinsically linked to the local geometry and curvature of the wavefunction intensity, which enables new mechanisms for controlling nonlinear excitations that are not achievable with purely local nonlinearities.

Further theoretical work on topologically-motivated nonlinearities has demonstrated the formation of topological solitons and opened new perspectives for quantum simulations [204–206]. In particular, it has been shown that density-dependent gauge fields in the discrete NLSE can promote the formation of stable domain walls between low and high density regions - topological defects - whose dynamical response to the gauge field is observed to be drastically different from the bare atoms [207–209]. Parallel experimental advances have shown how these exotic forms of nonlinearities can be realized using schemes such as Floquet modulation based periodic modulation of the lattice potential and nonlinearity strength [187,188,210].

Motivated by these recent works on nonlinear gauge fields, in this Chapter we introduce and analyze a nonlocal nonlinear Schrödinger model inspired by persistent homology, which is an unsupervised machine learning method for identifying significant topological features within a dataset. The considered nonlocal nonlinearity is proportional to the sign of the gradient of the intensity. After a suitable transformation, this nonlocal nonlinearity corresponds to an energy penalty or reduction for each local extremum of the intensity. Similar nonlinearities may arise in optics via light-induced gradient forces acting on small particles [211–213]. We investigate the implications of this peculiar purely nonlocal nonlinearity, which leads to the formation of cusp solitons with different topological features depending on the sign and strength of the nonlinearity.

The outline of this Chapter is as follows: Section 5.2 provides an introduction to our NLSE model, placing it in the context of other nonlocal nonlinear Schrödinger equations which have been studied in the literature, and outlines how our model may be realized using photonic feedforward lattices. Section 5.3 draws a connection to sublevel set filtration persistent homology. Sections 5.4 and 5.5 detail the analysis of the model and its plane wave linear stability. In Section 5.6, we outline the propagation results from three types of initial beam profile, discuss the potential implications, and analyse the ground states of the model. Section 5.7 makes concluding remarks and suggests future research directions.

5.2 Model

We consider the one-dimensional nonlocal NLSE. Nonlocal NLSEs arise in various contexts, particularly nonlinear optics, and can be written generically in dimensionless units as

$$i\partial_t \psi = -\partial_x^2 \psi + N[|\psi|^2]\psi, \tag{5.1}$$

where the wave field $\psi(t,x)$ simultaneously undergoes diffraction (described by the ∂_x^2 term) and interaction with an nonlinear potential $N[|\psi|^2]$ dependent on the field intensity profile $|\psi(x)|^2$. One commonly-employed model for nonlocal nonlinearity is

$$N[I] = g \int dx' R(x'-x)|\psi(x')|^2, \tag{5.2}$$

where the nonlocal response function R, which is assumed to be real, symmetric and normalized $\int_{-\infty}^{\infty} R(x)dx = 1$, describes how the potential shift depends on the intensity profile and the parameter g determines the strength and sign of the nonlinearity. For example, in the case of thermal nonlinearities, heat diffusion will tend to smear out the induced potential to a finite size determined by the diffusion constant of the medium [214, 215].

When the response is short-ranged, it can be approximated by a δ function, $R(x) = g\delta(x)$, which yields the standard local NLSE,

$$i\partial_t \psi = -\partial_x^2 \psi + g|\psi|^2\psi. \tag{5.3}$$

The simplest nonlocal correction to this equation is to assume R is non-singular but strongly localized, such that the width of the nonlocal response is much less than the width of the beam ψ. In this case one can expand the nonlocal response function as

$$N[|\psi|^2] \approx g(1 + \gamma\partial_x^2)|\psi|^2, \tag{5.4}$$

where the parameter $\gamma = \frac{1}{2}\int dx\, x^2 R(x)$ characterizes the strength of the nonlocality [192] and is independent of the fine details of the nonlocal response function.

While this response has been often neglected in the weakly nonlocal limit ($\gamma \ll 1$), studies show that such nonlocality can have significant impact on narrow optical beams in self-focusing media [216, 217]. There is an additional shift to the effective potential, proportional to the curvature of intensity, which affects the width and power thresholds of bright and dark solitons [190, 191]. In the extreme limit where the local part of the nonlinear response is negligible, arising in plasmas under certain conditions, cusp solitons with a characteristic exponential intensity profile can emerge [196].

Let us now consider a purely nonlocal response where the nonlinear potential peaks at changes in the sign of the intensity gradient, that is

$$N[|\psi|^2] = \alpha\partial_x \mathrm{sgn}(\partial_x|\psi|^2) \approx \alpha\partial_x \left[\tanh\left(\frac{\partial_x|\psi|^2}{w}\right)\right], \tag{5.5}$$

where w is a regularization parameter. The limit $w \to 0$ yields the sign function response, while $w \to \infty$ reproduces the weakly nonlocal limit of Eq. (5.4) with $\gamma = \alpha/w$. For uniform intensity profiles the nonlinearity vanishes. Non-uniform fields will be subjected to a nonlinear potential peaked at the local minima and maxima of $|\psi|^2$.

Substituting Eq. (5.5) into Eq. (5.1) gives

$$i\partial_t\psi = -\partial_x^2\psi + \alpha\partial_x\left[\tanh\left(\frac{\partial_x|\psi|^2}{w}\right)\right]\psi. \tag{5.6}$$

Physically, nonlocal nonlinear Schrödinger equations similar Eq. (5.6) may arise for nonlinearities mediated by small particles suspended in a fluid and accelerated by optical gradient forces, leading to peaks in the wavefunction intensity at local extrema of the intensity where $\partial_x|\psi|^2 = 0$ [211–213]. The regularization parameter w then accounts for a finite spread of the particles about the extrema due to repulsion or other effects. In the following, we will consider the limit $w \to 0$ for our analytical results and use a small finite w in numerical simulations to avoid numerical instabilities.

Realization using photonic feedforward lattices

Here we discuss a potential realization of Eq. (5.6) making use of feed-forward optoelectronic nonlinearities in coupled fiber loops, which are emerging as a powerful platform for nonlinear photonics and quantum simulation [218, 219]. In this platform, discrete time delays n of pulses circulating within each loop plays the role of the position x, while the number of round trips through the loops m plays the role of "time". We consider the setup shown in Fig. 5.1, which is similar to the one used in Ref. [220], where light propagates through two coupled loops with slightly different lengths, described by optical field amplitudes u_n^m, v_n^m. Feed-forward nonlinearities have been realized by outcoupling a small fraction of the light within each loop, using photodiodes to obtain electrical signals proportional to $|u_n^m|^2$ and $|v_n^m|^2$. Phase modulators then convert these signals into intensity-dependent phase shifts on the circulating light.

Using this scheme, a nonlinearity dependent on the intensity gradient can be approximated by using the difference between the photodiode signals. This signal can then be fed into circuits designed to approximate a tanh-like response, which are of independent interest in the context of analog hardware for neuromorphic photonics [221–223]. This setup is described by the evolution equations

$$u_n^{m+1} = \left[(\cos\theta)u_{n+1}^m + i(\sin\theta)v_{n+1}^m\right]\exp(i[\beta_{n+1}^m - \beta_{n-1}^m]), \tag{5.7}$$

$$v_n^{m+1} = \left[(\cos\theta)v_{n-1}^m + i(\sin\theta)u_{n-1}^m\right]\exp(i[\beta_{n+1}^m - \beta_{n-1}^m]), \tag{5.8}$$

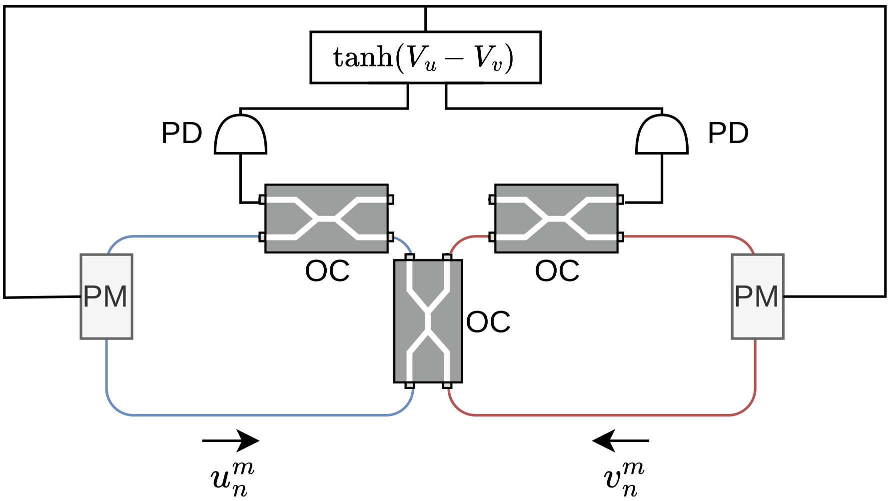

Figure 5.1: Schematic of two coupled fiber loops with feedforward nonlinearity. Circulating optical fields in each loop (u_n^m, v_n^m) are coupled using optical couplers (OCs). A small fraction of each field is diverted to photodiodes (PDs). The transformed difference in their output voltages drives phase modulators (PMs) which impart field-dependent phase shifts realizing the topological nonlinearity.

where θ is the coupling parameter between the loops and $\beta_n^m = \frac{\alpha}{2}\tanh[(|u_n^m|^2 - |v_n^m|^2)/w]$ is the nonlinear phase shift created using the feedforward nonlinearity.

To show how this setup can reproduce our model equation (5.6), we take the continuum limit. We set $\theta = \frac{\pi}{2} - J$, where $J \ll 1$, retaining terms only up to linear in J as follows:

$$u_n^{m+1} \approx \left[J u_{n+1}^m + i v_{n+1}^m\right] e^{i\beta_n^m}, \tag{5.9}$$

$$v_n^{m+1} \approx \left[J v_{n-1}^m + i u_{n-1}^m\right] e^{i\beta_n^m}, \tag{5.10}$$

The field after two round trips is

$$
\begin{aligned}
u_n^{m+2} &\approx \left[J u_{n+1}^{m+1} + i v_{n+1}^{m+1}\right] e^{i\beta_n^{m+1}}, \\
&= \left[iJ(v_{n+1}^{m} + v_n^m) - u_n^m\right] e^{i(\beta_n^{m+1}+\beta_n^m)} \\
v_n^{m+2} &\approx \left[J v_{n-1}^{m+1} + i u_{n-1}^{m+1}\right] e^{i\beta_n^{m+1}}, \\
&= \left[iJ(u_{n-1}^{m} + u_n^m) - v_n^m\right] e^{i(\beta_n^{m+1}+\beta_n^m)}
\end{aligned}
$$

Assuming the feedforward phase shift α is also small, allowing us to neglect terms of order $O(\alpha J)$, we obtain

$$
\begin{aligned}
i(u_n^{m+2} + u_n^m) &= -J(v_{n+1}^m + v_n^m) - (\beta_n^{m+1} + \beta_n^m)u_n^m, \\
i(v_n^{m+2} + v_n^m) &= -J(u_{n-1}^m + u_n^m) - (\beta_n^{m+1} + \beta_n^m)v_n^m,
\end{aligned}
$$

As discussed in Ref. [220], the amplitudes u_n^m and v_n^m are thus analogous to the wavefunction amplitude at odd and even spatial grid points, making $(|u_n^m|^2 - |v_n^m|^2)$ proportional to the intensity gradient in the continuum limit of this model. Assuming the field intensity is slowly varying with m and n, the left hand side of these equations will approximate $i\partial_t$, and the terms proportional to J will approximate ∂_x^2, thus reproducing Eq. (5.6).

5.3 Topological perspective

The conserved energy associated with the topological nonlinearity Eq. (5.5) can be found as

$$
\begin{aligned}
E[\psi, \psi^*] &= \int_{-\infty}^{\infty} |\partial_x \psi|^2 dx - \int_{-\infty}^{\infty} \alpha \,\mathrm{sgn}(\partial_x|\psi|^2)(\partial_x|\psi|^2) dx, \\
&\equiv E_L - E_{NL},
\end{aligned}
\tag{5.11}
$$

which we have written as an integral over the entire field profile. Here, we show that it is possible to view the nonlinearity from a topological perspective, in which the nonlinear energy shift E_{NL} depends only on the local maxima and minima of intensity and not on the precise details of the intensity distribution.

To do so, we assume a localized (normalizable) field profile with $\psi \to 0$ as $|x| \to \infty$ and evaluate E_{NL} by splitting the integral into intervals where $\partial_x |\psi|^2$ is strictly increasing or decreasing, corresponding to $\mathrm{sgn}(\partial_x |\psi|^2) = \pm 1$. Evaluating the E_{NL} integral from the first minimum (intensity $b_1 = 0$ at $x = -\infty$) to the first local maximum, we obtain

$$\int_{-\infty}^{x_1} \partial_x |\psi|^2 dx = |\psi(x_1)|^2 \equiv d_1 - b_1, \tag{5.12}$$

since $\mathrm{sgn}(\partial_x |\psi|^2) = 1$ in this interval. Continuing the integral up to the next local minimum,

$$\begin{aligned}\int_{x_1}^{x_2} \mathrm{sgn}(\partial_x |\psi|^2) \partial_x |\psi|^2 dx &= -\int_{x_1}^{x_2} \partial_x |\psi|^2 dx, \\ &= d_1 - |\psi(x_2)|^2, \\ &\equiv d_1 - b_2, \end{aligned} \tag{5.13}$$

followed by the next local maximum,

$$\begin{aligned}\int_{x_2}^{x_3} \mathrm{sgn}(\partial_x |\psi|^2) \partial_x |\psi|^2 dx &= \int_{x_2}^{x_3} \partial_x |\psi|^2 dx, \\ &= |\psi(x_3)|^2 - b_2, \\ &\equiv d_2 - b_2, \end{aligned} \tag{5.14}$$

and so on. Thus, we can write the nonlinear contribution to the conserved energy as a sum of the intensities at the local maxima (d_i), minus the intensities at the local minima (b_i),

$$E_{NL} = 2\sum_i (d_i - b_i). \tag{5.15}$$

Remarkably, the same quantity emerges as a topological summary statistic when applying persistent homology to quantify the shape of the intensity profile $|\psi(x)|^2$. Specifically, persistent homology considers the sublevel set of points X where a real-valued function $f(x)$ is below a certain threshold λ

$$X = f^{-1}(-\infty, \lambda] \tag{5.16}$$

By varying λ, a sequence of nested topological spaces known as a filtration is generated. Observing the change in topological features throughout this sequence provides a quantitative topological summary of the whole intensity profile. This provides an indication as to which topological features (such as connected components, cycles, and higher-dimensional holes) are more or less significant based on how long they persist through the filtration [224]. We note that the birth and death of connected components correspond to local maxima and local minima of the function, respectively.

The homology of a sublevel set filtration only changes when critical points of the function are swept across, such as minima, maxima, and saddle points, as shown in Fig. 5.2. For quantum physics, persistent homology has been particularly useful for studying properties of energy landscapes and quantum states, offering insights into phase transitions, entanglement, and the behaviour of systems under external perturbations [28, 29, 34, 35, 73, 89].

In persistent homology, it is common to summarise the persistence of topological features (in this case, the local maxima and minima) by computing a norm of the

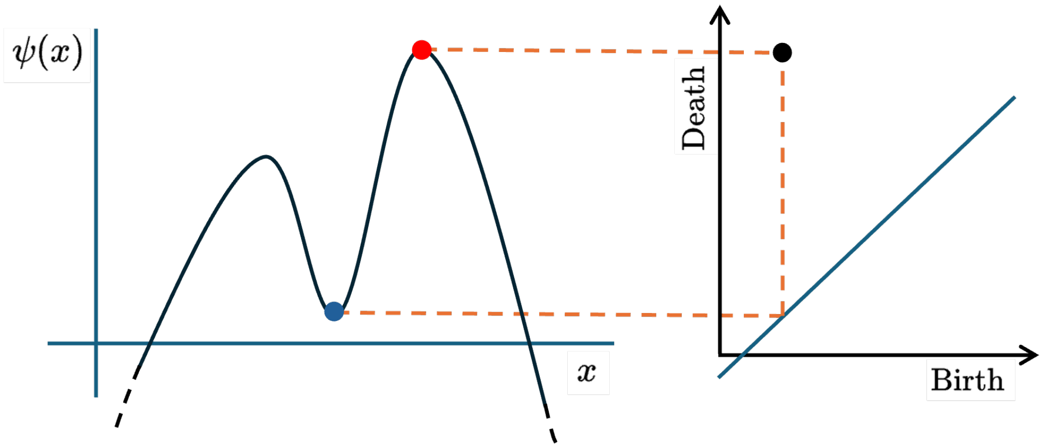

Figure 5.2: Illustration of sublevel set filtration persistent homology from a 1D profile. One example persistent feature (black dot) is shown, which is born at the local minimum (blue dot) and dies at the subsequent local maximum (red dot).

feature lifetimes, e.g. [29, 35, 225–228]

$$P_p = (\sum_i (d_i - b_i)^p)^{\frac{1}{p}}, \tag{5.17}$$

where the summation is made over all features and the choice of p affects the sensitivity to the persistence of features. The case $p = 1$, corresponding to the sum of feature lifetimes $l_i = d_i - b_i$, is precisely the nonlinear energy term appearing in Eq. (5.11). Thus, our model introduces a novel line of research which incorporates for the first time quantities deriving from persistent homology directly into quantum dynamics. Note that this quantity is related to the average feature lifetime considered in Chapter 3.

Because the total energy is conserved, the dynamics will exhibit a trade-off between the kinetic and nonlinear parts of the energy. For example, when $\alpha > 0$, increasing the height of a peak can become energetically favourable despite the kinetic energy penalty it imposes (due to stronger localization) because local maxima induce attractive potentials. Conversely, $\alpha < 0$ induces repulsive (attractive) potentials at local maxima (minima), tending to smooth out sharp features, as we will see in the next Section.

5.4 Energy functional

In field theory we can obtain an evolution equation from the energy functional by taking the variational derivative with respect to the complex conjugate of the wavefunction

$$i\frac{d\psi}{dt} = \frac{\partial E}{\partial \psi^*} - \nabla \cdot \frac{\partial E}{\partial(\nabla\psi^*)}. \tag{5.18}$$

Computing the corresponding variational derivative for the terms in Eq.(5.11), the first term E_L leaves us with the standard diffraction part (∂_x^2) of the Schrödinger equation. The term resulting from the proposed nonlinear energy penalty E_{NL} can be found as

$$\frac{\delta E_{NL}}{\delta \psi^*} = \alpha[(\partial_x\psi) - \partial_x\psi][\delta(\partial_x|\psi|^2)(\partial_x|\psi|^2) + \mathrm{sgn}(\partial_x|\psi|^2)] \tag{5.19}$$

$$\equiv (\partial_x A_x)\psi \tag{5.20}$$

where

$$A_x = \alpha[\delta(\partial_x|\psi|^2)(\partial_x|\psi|^2) + \mathrm{sgn}(\partial_x|\psi|^2)]. \tag{5.21}$$

Noting that $x\delta(x) = 0$, this simplifies to

$$A_x = \alpha\,\mathrm{sgn}(\partial_x|\psi|^2), \tag{5.22}$$

and thus we recover the step-like response in Eq. (5.5).

5.5 Modulational instability

To perform a linear stability analysis we consider the regularized nonlinear response in Eq. (5.6). The model admits a constant amplitude background solution $\psi_0 = A$, which may be taken as real without loss of generality. To study the modulational instability, we peturb this solution as $\psi_0 \to \psi_0 + p(x,t)$, where the perturbation is assumed to be small, $|p(x,t)| \ll A$. The intensity under this perturbation becomes $|\psi_0 + p|^2 = A^2 + A(p + p^*)$. Substitution into Eq. (5.6), assuming $p \ll w/A$, such that we can approximate $\tanh(\frac{A\partial_x(p+p^*)}{w}) \approx \frac{A\partial_x(p+p^*)}{w}$, and neglecting terms of order p^2 yields an equation for the evolution of the perturbation amplitude,

$$i\partial_t p = -\partial_x^2 p + \frac{A^2\alpha}{w}\partial_x^2(p + p^*). \tag{5.23}$$

The goal is to determine the time-dependence of the Fourier components of p, specifically whether they remain bounded in time or grow, which distinguishes linearly stable and unstable regimes [229, 230]. This can be done by using a trial function of the form

$$p(x,t) = c_1 e^{i(kx-\omega t)} + c_2^* e^{-i(kx-\omega^* t)}, \tag{5.24}$$

where k and ω are the real wavenumber and complex angular frequency of the perturbation, and c_1 and c_2^* are constants.

After substitution into Eq. (5.23) and collecting coefficients of $e^{\pm i(kx-\omega t)}$, the linearized dynamics reduces to a linear system coupling c_1 and c_2. This can be written in matrix form as

$$\begin{pmatrix} (1-\alpha A^2/w)k^2 - \omega & -\alpha A^2 k^2/w \\ -\alpha A^2 k^2/w & (1-\alpha A^2/w)k^2 + \omega \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = 0, \tag{5.25}$$

The condition for nontrivial solutions is that the determinant of this matrix vanish, which yields the dispersion relation

$$\omega(k) = \pm k^2 \sqrt{1 - 2\alpha A^2/w}. \tag{5.26}$$

Thus, $\alpha > 0$ results in instability for a sufficiently large wave amplitude, whereas the uniform background remains linearly stable for all $\alpha < 0$. Thus, we can associate $\alpha > 0$ and $\alpha < 0$ with being analogous to focusing and defocusing nonlinearities, respectively.

5.6 Dynamics

In order to understand the propagation characteristics of our model, we use the split-step method to simulate beam propagation from a few different initial profiles. The evolution of the intensity profiles in both spatial and Fourier space will reveal how dispersion and nonlinearity interact to shape the beam dynamics, including features such as self-focusing or defocusing, soliton formation, and the development of spatial or spectral instabilities, depending on the sign and strength of the nonlinearity. Here, we consider three types of initial profile - a Gaussian beam, a dark soliton-type beam, and a flat-top beam.

For all simulations, the spatial domain is discretized over $N = 1024$ grid points, the system size is $L = 400$, ensuring negligible boundary effects, and the time domain step-size is $\Delta t = 0.01$ which is small enough to ensure numerical stability. The regularization parameter in Eq. (5.6) is taken to be $w = 10^{-5}$. All initial states are normalized such that $\int |\psi(x)|^2 dx = 1$.

Numerical Methods

In general, it is difficult to find analytical solutions to nonlocal NLSEs, so we use numerical techniques to explore its dynamics and stationary states. The general solution to the NLSE can be written as

$$\psi(t+dt) = e^{-iHdt}\psi(t), \tag{5.27}$$

where H is the Hamiltonian operator. This can be divided into a kinetic energy component ∂_x^2, which is diagonal in Fourier space, and an effective potential energy component, which is diagonal in real space. Splitting the full evolution operator into N steps corresponding to time intervals $\Delta t = t/N$, the evolution operator for a single time step is

$$e^{-iH\Delta t} = e^{-iV\Delta t/2}e^{i\partial_x^2\Delta t}e^{iV\Delta t/2}e^{O(\Delta t^3)}. \tag{5.28}$$

Numerical time evolution of this type is known as the symmetrized split-step method, which takes advantage of the separable structure of the NLSE by alternating between linear dispersion in Fourier space and nonlinear evolution in real space. In this way, computational expenses are minimized compared to the direct calculation using finite difference methods. Provided that the step size is sufficient small, the error is of order dt^3 for a time step of size dt and energy will be conserved to a good approximation [231].

For beam propagation on a discrete lattice, numerical regularization is essential in order to address numerical instabilities and control spurious high-frequency components, particularly for long time evolutions. We address this by making the substitution $k^2 \to 2J\cos k$, where J corresponds to the hopping strength of an optical lattice realization. In the long wavelength limit, this reduces to the

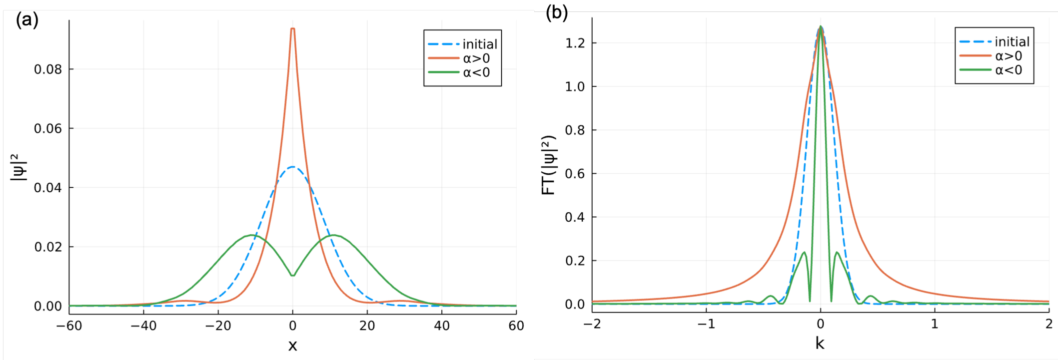

Figure 5.3: Beam propagation of an initial normalized Gaussian profile (blue dotted line) under the topological nonlinearity model in (a) real and (b) Fourier space. For $\alpha > 0$ (red line), the beam undergoes a sharpening effect, whereas for $\alpha < 0$ (green line) it forms a strong dip, splitting the beam into two. Parameters: $T = 500$, $\sigma = 12$, $|\alpha| = 0.1$.

continuum model $2J \cos k \approx 2J - Jk^2$ if $J = 1$. Secondly, we smooth the sign nonlinearity term which may be numerically problematic due to the discontinuity by taking $\text{sgn}(\partial_x |\psi|^2) \to \tanh(\partial_x |\psi|^2 / w)$, where w is a cutoff length scale. The smoothing, as well as the use of a finer mesh than the 'width' of the discontinuity, allows for the gradient of the term to be obtained using finite difference methods. In the limit of many sites and broad wavepackets compared to w, or equivalently the large w limit, this will reduce to the ideal continuum model, which is a purely nonlocal NLSE with a nonlinear term proportional to $\frac{1}{w}\partial_x^2 |\psi|^2$.

We use imaginary time evolution (ITE) to approximate the ground state of our NLSE model. This is a widely used numerical technique which converges a system to its ground state, as long as the system has some initial support on the ground state. By transforming the real-time evolution $t \to i\tau$, the time-dependent NLSE becomes a diffusion-like equation that exponentially suppresses higher-energy components of an initial state and thereby projecting it toward the ground state [232, 233].

Gaussian beam

Self-focusing regime

The Gaussian beam serves as a canonical initial condition for probing the interplay between dispersion and nonlinearity in both the conventional NLSE and our persistent homology-inspired model. We consider the following initial state

$$\psi_G(x) = \frac{1}{\sqrt{\mathcal{N}_G}} \, e^{-\frac{x^2}{2\sigma^2}}, \tag{5.29}$$

where σ is the beam width and the factor $\mathcal{N}_G = \sigma\sqrt{\pi}$ ensures normalization of the initial state. In the standard cubic NLSE, the evolution of a Gaussian profile is well understood: for focusing nonlinearity, self-interactions counteract diffraction, leading to self-focusing and, above a critical power threshold, the formation of bright solitons - spatially localized, non-dispersive wave packets [234, 235].

In our model, for $\alpha > 0$, the nonlinearity acts as an effective attractive potential at local maxima in the intensity profile. Thus an initial Gaussian beam rapidly self-focuses, forming a highly localized structure with a pronounced cusp at its center and slow attenuation towards infinity (see Fig. 5.3(a)). The corresponding Fourier spectrum in Fig. 5.3(b) is Lorentzian in form, which would be consistent with the $\exp(-w|x|)$ real space profile, since the Fourier transform of an exponential decay produces a Lorentzian lineshape. This spectral signature contrasts with the Gaussian Fourier spectrum of the initial condition, highlighting the nonlinearity-induced reshaping. These kinds of Fourier spectra are relevant for optical lattice realizations with time-of-flight imaging, which allows measuring the intensity profile in k-space directly [236].

Defocusing regime

In contrast, the defocusing regime of the conventional NLSE enhances the diffraction of Gaussian beams and can result in the formation of dark solitons. These are usually finite intensity dips on a uniform background, accompanied with a phase shift across the dip.

For our model, we notice that for $\alpha < 0$, where the potential is now repulsive at local maxima. Thus, the initial intensity maximum at the beam centre is transformed to a local minimum, with two peaks appearing on either side. Notably, the topology of the sublevel set filtration is altered under this regime since the number of connected components (0-D topological features) is increased; the initial single peak beam is transformed to a double-peaked beam.

Dark soliton state

Dark solitons are fundamental nonlinear excitations of the defocusing NLSE, characterized by localized intensity dips embedded in a continuous nonzero background field. Given the structural similarity between the self-focusing soliton solutions of the conventional cubic NLSE and those of our model, we investigate whether Gaussian-shaped dark solitons exhibit analogous behavior under the influence of the topological nonlinearity. We consider a standard approximation of the tanh-like dark soliton state [216], whose intensity profile takes the form of a Gaussian dip against a uniform background

$$\psi_{DS}(x) = \frac{1}{\sqrt{\mathcal{N}_{DS}}} \left(1 - A\, e^{-\frac{x^2}{2\sigma^2}} \right) , \tag{5.30}$$

where A is the dip amplitude, σ its width, and the factor $\mathcal{N}_{DS} = L + \sqrt{\pi}\sigma A(A - 2\sqrt{2})$ is the normalization factor.

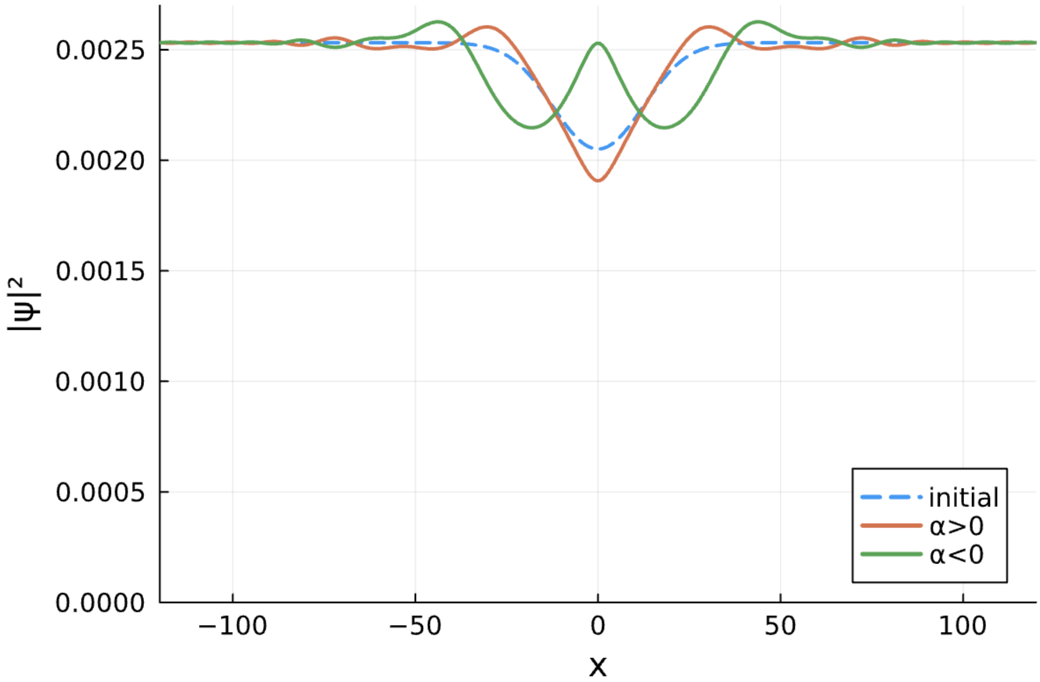

Figure 5.4: Beam propagation of an initial normalized Gaussian-shaped dark soliton (blue dotted line) under the topological nonlinearity model in real space. The behaviour is qualitatively similar as for the Gaussian initial profile in Fig. 5.3, whereby the peak is either sharpened (red line) or inverted (green line) depending on the sign of α. Parameters: $T = 500$, $A = 0.1$, $\sigma = 12$, $|\alpha| = 0.01$.

The numerical simulations reveal that the evolution of dark solitons in our model qualitatively parallels that of Gaussian beams (see Fig. 5.4): for one $\alpha > 0$, corresponding to effective self-focusing of Gaussian beams, the local curvature of the intensity profile is enhanced, reinforcing the dip structure. In other words, the local minimum induces a repulsive potential which acts to reinforce the minimum. Remarkably, the same nonlinearity sign supports bright cusp solitons and dark soliton-like states on a uniform background. Conversely, for the opposite sign, the local minimum in intensity inverts into a local maximum.

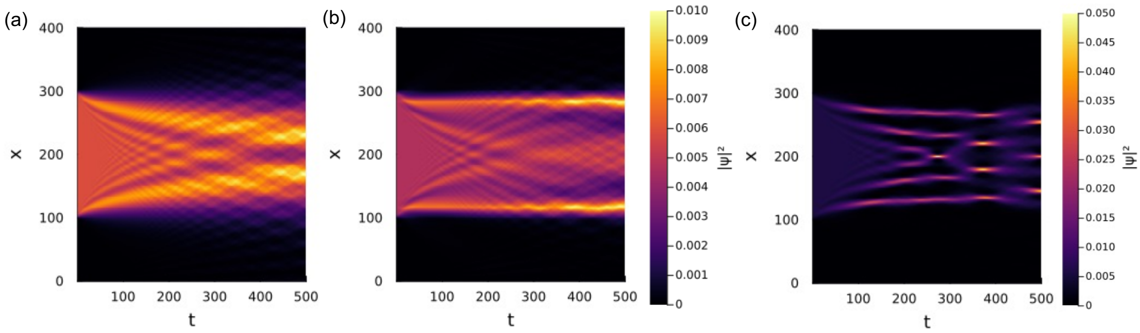

Figure 5.5: Beam propagation of an initial normalized flat-top profile under the linear (left), topological nonlinear (middle) and focusing cubic nonlinear (right) Schrödinger models in real space. Notably, the propagation under the topological nonlinear model differs from the linear and cubic nonlinear cases as it maintains the sharp edges of the flat-top profile. Under the unstable conventional model, the exponential growth of small modulations leads to high-intensity localized peaks and the breakup of the plateau. Parameters: $l = 100$, $p = 20$, $|\alpha| = 0.1$, $g = -5$. Note that the magnitude of the cubic nonlinearity term g is balanced through dimensionality scaling such that the term is noticeable at the soliton scale, and all three initial states are the same.

Flat-top beam

In the conventional NLSE, the stability of a uniform solution is determined by the sign of the nonlinearity. In the defocusing regime, the uniform state is stable; in the focusing regime, however, the uniform solution is subject to modulational instability (MI), whereby small perturbations are exponentially amplified, eventually leading to the formation of localized structures such as solitons or breathers [237].

In our model, the uniform state remains a trivial solution, since the absence of local extrema in $|\psi|^2$ means the topological energy penalty vanishes. To probe the stability of this solution further, we consider a wide flat-top beam with sharp transitions at the edges,

$$\psi_{FT}(x) = \frac{1}{\sqrt{\mathcal{N}_{FT}}} e^{-(x/l)^{2p}}, \tag{5.31}$$

where the parameters l and p set the width and edge steepness, and $\mathcal{N}_{FT} \approx 2l$ is the normalization factor. This type of super-Gaussian configuration is often used to investigate boundary-induced instabilities and soliton formation and is also experimentally relevant as flat-top beams serve as finite-domain approximations for infinite uniform states [238, 239]. Under the conventional NLSE, such a beam typically evolves as follows: for defocusing nonlinearity, the flat top gradually disperses and approaches a Gaussian-like profile as the plateau narrows and the edges smooth out; for focusing nonlinearity, MI leads to the rapid formation of localized spikes and the eventual breakup of the plateau (see Fig. 5.5(c)). In the linear case where $g = 0$, the beam will also lose its flat-top shape as it propagates due to strong diffraction at the edges, evolving towards a smoother, bell-shaped distribution (see Fig. 5.5(a)).

The topological nonlinearity leads to qualitatively distinct evolution dynamics. The sharp edges which define the width of the flat-top beam are enhanced, and we observe the emergence of stable, localized features (peaks or dips) just inside the beam boundaries, depending on the sign of α ($\alpha > 0$ shown in Fig. 5.5(b)). These edge features remain stationary and spatially robust throughout the simulation, indicating a suppression of MI-driven breakup and a stabilization of the plateau against both long- and short-wavelength perturbations. This behavior can be understood as a consequence of the topological energy term, which penalizes the creation of new local extrema and thus energetically disfavors the amplification of small fluctuations that would otherwise seed MI. Instead, the dynamics within the beam in Fig. 5.5(b) resembles the linear diffraction in Fig. 5.5(a), at least for moderate times $t \leq 200$.

The stabilization of flat-top beams by a topological nonlinearity has several notable implications. First, it suggests a new mechanism for controlling the stability of nonlinear systems and suppressing MI using geometric and topological constraints:

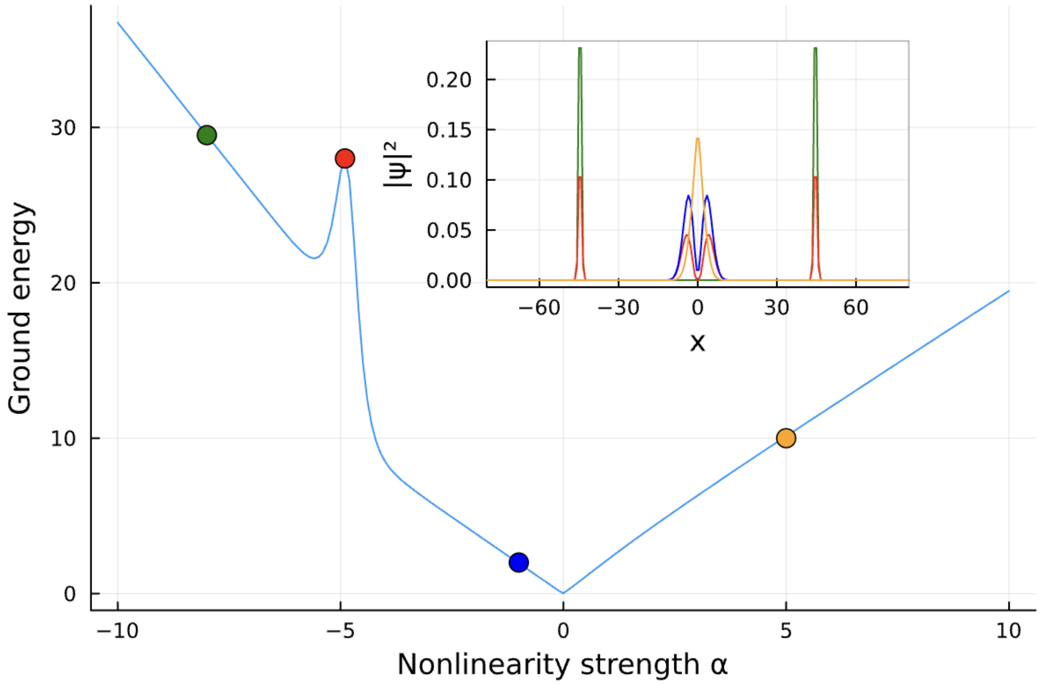

Figure 5.6: Relationship between ground state energy and the nonlinearity strength α for the topological NLSE (blue curve). The coloured points identify distinct ground state phases, the intensity profiles for which are shown in the inset figure with corresponding colours.

The nonlinearity concentrates at the edges, inducing self-localizing potential barriers, whereas nonlinear effects within the beam itself remain weak due to the near-uniform intensity profile. Secondly, the persistent edge features may find application in edge detection for optical systems or in machine learning, where the natural enhancement and stabilization of edge features could serve as a physics-inspired preprocessing step for identifying boundaries in structured data.

Ground states

To corroborate the beam propagation simulation results, we numerically obtain the ground states of our model as a function of the nonlinearity strength α. Starting from a normalized Gaussian trial wavefunction, we propagate the system in imaginary time τ using a symmetrized split-step ITE method. Similar to the bright solitons in the conventional focusing NLSE, we observe sharply localized cusp-like solitons for $\alpha > 0$, where the topological term energetically favours the formation of local extrema. Conversely, for $\alpha < 0$, the energy landscape favors the proliferation of intensity minima, leading to bifurcated or multi-peak structures, reminiscent of the delocalized or dark soliton regimes in the defocusing NLSE.

Analytically, we can obtain some qualitative insights to the behavior of our model through seeking a self-consistent ansatz. Assuming an exponential profile $\psi(x) = A\exp(-\kappa|x|)$, where $1/\kappa$ is the localization length, we try to find a self-consistent solution to the nonlinear wave equation. The nonlinear potential $\alpha\partial_x\delta(\partial_x|\psi|^2)$ generates a δ function potential at $x = 0$ for the exponential profile. In the regions where $|x| > 0$, the nonlinear potential is zero, and we find from substituting into the wave equation and choosing A such that the wavefunction is bounded and continuous, $E = -\frac{\kappa^2}{2} < 0$. By integrating the Schrödinger equation through the δ function discontinuity at $x = 0$, we find a relationship between the amplitude and localization strength, suggesting a self-consistent solution of the form $\psi(x) = \sqrt{\alpha}e^{-\alpha|x|}$. These cusp solitons have an energy $E < 0$, indicating that they are favoured over the (linear) $k = 0$ plane wave ground state. Similar cusp solitons also exist as exact solutions of the nonlinear Schrödinger equation with a purely nonlocal nonlinearity [192, 196].

5.7 Conclusion

In this Chapter, we have introduced a novel nonlinear Schrödinger-type model whose nonlinearity is explicitly dependent on the local minima and maxima of the field intensity. Whilst the form of soliton solutions is similar to those observed in the conventional NLSE, our model captures nontrivial geometric effects absent in amplitude-based formulations. Our analysis shows that the inclusion of a topological energy penalty, derived from the persistent homology of the intensity profile, fundamentally alters both the stability and the dynamical evolution of localized wave packets. Specifically, we find that the sign-dependent nonlinearity leads to the emergence of stable cusp-like soliton solutions, modifies the modulational stability of flat-top beams, and can promote the formation of multi-peak structures through topological bifurcations. These effects are robust to numerical regularization and persist across a range of initial conditions and feasible parameter regimes.

In the dimensionless NLSE Eq. (5.1), x is measured in units of some characteristic length scale a, time is measured in units of $\tau = 2m_{\text{eff}}a^2/\hbar$, where m_{eff} is the particle effective mass. The Bose-Einstein condensate with a synthetic gauge field experiment of Ref. [209] used ^{1}33Cs atoms with $m_{\text{eff}} = 1.8 \times 10^{-25}$ kg; taking $a = 1\mu$m yields $\tau \approx 0.5$ms. Therefore, the results in Figs. 5.3 – 5.5 correspond to an evolution time of about 300 ms, within the measured 1/e lifetime of 700 ms under periodic driving. On the other hand, experiments on optical force-induced nonlinearities and self-focusing [213] have used a green laser with a wavelength $\lambda = 532$ nm, with $m_{\text{eff}} = 2\pi\hbar/\lambda$, leading to $\tau \approx 0.02$ mm, meaning our simulations correspond to feasible propagation lengths of about 1 cm.

Since the nonlinear energy depends only on the intensity local minima and maxima, it is invariant under arbitrary transformations that do not change the intensity at these points. For example, the (super)Gaussian beams $\psi(x) = \exp[-(x^2/l^2)^p]$ and cusp beams $\psi(x) = \exp(-|x|/l)$ have the same nonlinear energy, independent of their width l. This is in contrast to conventional local or nonlocal nonlinearities, which lack this kind of invariance. Another way of viewing this is that the nonlinearity only acts at isolated points (the critical points). Thus, this class of models is intermediate between linear dynamics (no nonlinearity anywhere) and regular nonlinear wave dynamics (nonlinearity acting everywhere the intensity is not negligible).

By bridging concepts from computational topology and nonlinear physics, our work opens new conceptual and practical avenues for the controlled manipulation of solitons, the stabilization of nonlinear modes, and the realization of topologically protected states in photonic and optical systems. Future research may extend these ideas to higher dimensions, explore the interplay with disorder and external potentials, and seek experimental realizations in optical lattices or Bose-Einstein condensates. Our findings suggest that topological nonlinearities could serve as a versatile tool for both fundamental studies and technological applications in nonlinear quantum systems.

Chapter 6

Closing remarks and Outlook

The core objective of this thesis is understanding to what extent we are able to capture the "shape" of the dynamics in quantum systems in a way that is physically significant and experimentally accessible. To this end, we have systematically explored the applicability of computational topology and the use of topological features for characterising the dynamics of various complex quantum systems. We have developed novel TDA-based methods that provide a new lens for studying and identifying dynamical phases, as well as TDA-inspired approaches for the engineering and exploring of new physics. We ultimately conclude that applying a computational topological perspective to quantum many-body physics serves to be an effective and versatile framework for probing underlying geometric structures, provided that sufficient physical insight is given in the construction of the topological representation - that is, in how the raw (accessible) data of quantum dynamics are encoded into topological objects in a meaningful way. A primary strength of TDA-based methods compared to conventional machine learning approaches lies in the intrinsic qualities inherited from applied topology, enabling a robustness to noise, efficiency in data requirement, and geometric interpretability that are well-suited to many use cases in physics. As with the adoption of all powerful yet unfamiliar tools, the challenge lies in identifying where best to direct them, and we believe that this thesis provides several significant steps in that direction.

6.1 Summary of Results

In Chapter 3, we began by establishing a persistent homology pipeline for diagnosing chaos in open quantum systems, adapting existing topological methods for classical time series analysis. As a demonstrative first example, we considered a nonlinear leaky resonator model which supports known behaviour in the semiclassical limit and showed that a phase-space embedding of quantum trajectories yields dynamical attractors whose topological properties reliably distinguish regular from chaotic dynamical phases. An important finding was the robustness of these topological signatures to measurement backaction, with consistent results found across different quantum observables such as position expectation values and photon counts. This suggested that the extracted features reflect an intrinsic property of the system rather than artefacts of the measurement scheme.

Subsequently, in Chapter 4, we extended the topological lens to the more complicated setting of quantum many-body scar systems. We introduced a topology-based framework for distinguishing initial states that belong to scarring and thermalising manifolds based on the topological dimension of the dynamics, demonstrating this on a kinetically-constrained Rydberg atom chain. While our approach in Chapter 3 assumed access to a sufficiently long time series in order for reliable reconstruction of the shape of the dynamics, decoherence and noise limit the available evolution time in large quantum processors. Accordingly, our method in Chapter 4 makes optimal use of the most easily available measurement data (bit-strings of excited state occupation numbers), allowing us to extract the same key quantity of interest, the "shape" of the dynamics (effective dimension), in a more measurement-efficient manner.

In Chapter 5, we posed a fundamentally different question: rather than employing TDA as an analytical tool, we asked whether topological quantities could be used to define new physical models. We answered this affirmatively by introducing a persistent-homology-inspired nonlinearity in the nonlinear Schrödinger equation, governed by the topological structure of wavefunction intensity extrema. This model was shown to support stable cusp-like soliton solutions, modified modulational instability, and admits topological bifuractions into multi-peak structures - effects absent in conventional local or nonlocal nonlinear models, and within reach of current BEC and optical experiments.

6.2 Future Prospects

On the theory side, there are various obvious directions to take to develop our work further. The topological feature selection in Chapters 3 and 4 have been empirical demonstrations that these kinds of methods can work in distinguishing different kinds of quantum many-body dynamics. A natural next step is to put these numerical observations on a firmer theoretical footing. For example, this can be achieved by analyzing other toy models to understand when and how these methods might offer better performance in terms of robustness to noise or number of measurements required compared to standard approaches. For the model introduced in Chapter 5, various generalizations can be explored. The most immediate is the extension to higher spatial dimensions, where the persistent homology of the intensity profile becomes much richer, admitting saddle points and higher-dimensional topological features. Another direction is to explore the interplay of the topological nonlinearity with external potentials and disorder.

- In Chapter 3 we considered a system whose classical-quantum correspondence is well established, but we envision that the ideas developed can be used to study the behaviour of topological properties for quantum chaotic systems without a classical counterpart such as the driven Rabi model [240] or systems in which measurement-induced backaction can modify the boundary between regular and chaotic regimes [241]. Another interesting class of systems where topological tools may be useful are those where energy level statistics with the Brody distribution, where phase space is segmented into separate regions of regular and chaotic dynamics [242]. Since energy levels also constitute 1-dimensional data, it may be possible to apply similar time-series analysis methods to distinguish between energy intervals with regular and chaotic statistics. This could be cast as a classification, clustering, or anomaly detection problem; for each of these tasks there are well-established applications of TDA to classical datasets.

Since the TDA-based approach can be applied to both the classical and quantum regimes, another possible future route of investigation would be to see whether these topological descriptions of quantum systems can offer any insight into critical points during the quantum-classical transition. For instance, the TDA method of zig-zag persistence may provide a convenient way of tracking the persistence of specific features through a filtration of quantum attractors that correspond to changing size/action scales of the dynamical system [243, 244].

Another direction for which this TDA approach may aid in uncovering new physics is understanding quantum phase transitions in relation to the time dependence of entanglement structures, following recent work using persistent homology to extract experimentally-intractable entanglement measures of quantum many-body systems [104, 130, 245].

- Our work in Chapter 4 provides a novel perspective on the analysis of QMBS from the perspective of dynamical manifold intrinsic dimension, which has potential theoretical links to network and complexity theory in future studies

[180–183]. We also note that the structures observed in the MDS embedding of scar dynamics are reminiscent of those seen in the phase-space of single-body quantum scar eigenstates, which are similarly embedded in a lower dimensional manifold defined by the classical periodic orbits. This may be suggestive of a deeper connection related to near-integrability or dynamical symmetries [184–186].

Several recent works [246, 247] are similarly based on exploiting the fact that scar states tend to lie on low-dimensional manifolds, using this in order to seek a compressed representation of the scar states using matrix product states or quantum convolutional neural networks. They find that if compression is possible then the states can be identified as many-body scars, whereas thermal states have a more complicated structure that cannot be efficiently compressed. Our work demonstrates that this idea of compressed representations can be effectively exploited even without having to directly measure the entanglement properties of the scar states.

- In addition to generalizations of the specific model considered in Chapter 5, our results more broadly demonstrate that topological observables can indeed form the basis for new classes of nonlinear wave equations with properties distinct from conventional models. An interesting direction for future work will be to extend this approach to other kinds of topological observables beyond those encountered in persistent homology, and also to the case of systems defined on lattices (or arbitrary graphs) rather than the continuum Schrödinger equation we considered.

Another natural extension of the work presented in this thesis is the experimental realisation of the methods developed across all Chapters:

- The pipeline laid out in Chapter 3 involves the construction of dynamical attractors using the quantum trajectory treatment of open dissipative systems. These quantum trajectories correspond to the conditional pure-state evolution of the system under a particular continuous measurement record of the environment, and can be accessed in experiment via weak continuous measurement of superconducting qubits coupled to microwave cavities [113, 114, 248, 249]. As we have shown in Chapter 3, such an experimental scheme involving observables of quantum trajectories offers an advantage over other existing methods for quantifying chaos in open quantum systems, which can be experimentally difficult to realize or require measurements of the full wavefunction of sets of eigenstates [105, 106, 129–131].

- In Chapter 4, we consider Rydberg atom arrays and lattice spin chains exhibiting quantum many-body scarring behaviour, and use easily accessible computational basis measurement outcomes to device a topological method for detecting non-thermalizing states. As a result, this method translates directly to any quantum simulation platform on which single-shot bit-string samples are routinely accessible, without additional experimental overhead. This broad accessibility is particularly timely given that quantum many-body scarring has now been experimentally observed across a diverse range of hardware, from Rydberg atom arrays [57, 155] and bosonic quantum gases [178] to superconducting qubit processors [179] and trapped ion systems [250].

  In several works following ours, quantum convolutional neural networks and variational quantum algorithms have been used to tackle the same problem, achieving high classification accuracy but at the cost of substantial data requirements and the practical difficulties inherent to variational optimisation [246,247]. By contrast, the unsupervised, topology-based approach introduced in this thesis requires no training data and avoids the trainability challenges associated with variational methods.

- The proposed nonlinear Schrödinger-type model in Chapter 5 can be potentially realized using feed-forward optoelectronic nonlinearities in coupled fiber loops, which are emerging as a powerful platform for nonlinear photonics and quantum simulation [218–220]. Details of this experimental implementation are covered in Section 5.2 of Chapter 5.

  Recent works also propose the realization of density dependent gauge fields and other exotic forms of nonlinearities through Floquet modulation based periodic modulation of the lattice potential and nonlinearity strength [187, 188, 210].

Finally, we would like to comment on the placement of our work within the context of quantum TDA, which has attracted considerable attention in recent years. There is a growing body of work exploring deep connections between some of the objects studied in TDA (e.g. homology groups of simplicial complexes) and certain quantum many-body systems [79, 82, 251]. These connections lie at the heart of various quantum TDA algorithms with the potential to exhibit useful speedups compared to the best classical algorithms, when applied to high-dimensional problems. Examples where high-dimensional cycles naturally emerge (and what their applications might be) remain unknown, and it is now a pressing problem to identify potential use-cases. This thesis establishes baselines for low-dimensional TDA applied to various quantum many-body systems. It will be interesting to compare quantum TDA algorithms and high-dimensional topological features for tackling the problems studied here, once sufficiently powerful quantum processors become available.

References

[1] M. Nakahara. *Geometry, Topology and Physics.* Institute of Physics Publishing, Bristol, UK, 2nd edition, 2003.

[2] R. L. Ricca. An introduction to the geometry and topology of fluid flows. *NATO Sci. Ser. II, Math. Phys. Chem.*, 47, 2001.

[3] P. G. Tait. On Knots. *Trans. R. Soc. Edinburgh*, 24:213–279, 1867.

[4] H. K. Moffatt. The degree of knottedness of tangled vortex lines. *Journal of Fluid Mechanics*, 35(1):117–129, 1969.

[5] M. Kastner. Phase transitions and configuration space topology. *Rev. Mod. Phys.*, 80:167–178, 2008.

[6] M. Z. Hasan and C. L. Kane. Colloquium: Topological insulators. *Rev. Mod. Phys.*, 82:3045–3067, 2010.

[7] T. Ozawa, H. M. Price, A. Amo, N. Goldman, M. Hafezi, L. Lu, M. C. Rechtsman, D. Schuster, J. Simon, O. Zilberberg, and I. Carusotto. Topological photonics. *Rev. Mod. Phys.*, 91:015006, 2019.

[8] J. M. Kosterlitz and D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *Journal of Physics C: Solid State Physics*, 6(7):1181, 1973.

[9] B. A. Bernevig and T. L. Hughes. *Topological Insulators and Topological Superconductors*. Princeton University Press, 2013.

[10] L. Lu, J. Joannopoulos, and M. Soljačić. Topological photonics. *Nature Photon*, 8:821–829, 2014.

[11] E. Witten. Quantum field theory and the Jones polynomial. *Comm. Math. Phys.*, 121:351–399, 1989.

[12] K. von Klitzing, T. Chakraborty, P. Kim, V. Madhavan, X. Dai, J. McIver, Y. Tokura, L. Savary, D. Smirnova, A. M. Rey, C. Felser, J. Gooth, and X. L. Qi. 40 years of the quantum Hall effect. *Nature Reviews Physics*, 2(8):397–401, 2020.

[13] K. von Klitzing, G. Dorda, and M. Pepper. New method for high-accuracy determination of the fine-structure constant based on quantized Hall resistance. *Phys. Rev. Lett.*, 45:494–497, 1980.

[14] R. E. Prange and S. M. Girvin. *The Quantum Hall Effect*. Springer, New York, 1990.

[15] D. W. Sumners. Untangling DNA. *Math. Intell.*, 12:71–80, 1990.

[16] Y. Wang et al. What can topology bring to chemistry? *CCS Chemistry*, 6:2084–2109, 2024.

[17] E. Flapan. *When topology meets chemistry: a topological look at molecular chirality*. Cambridge University Press, 2010.

[18] G. Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.

[19] M. J. Hartmann and G. Carleo. Neural-network approach to dissipative quantum many-body dynamics. *Physical Review Letters*, 122(25):250502, 2019.

[20] C. Molnar. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable.* Independently published, Munich, Germany, 3rd edition, 2025. https://christophm.github.io/interpretable-ml-book Accessed: 2026-05-21.

[21] C. Zeng, Z. Khan, and N. L. Post. Data-efficient and interpretable inverse materials design using a disentangled variational autoencoder. *AI & Materials*, 2025(1), 2025.

[22] G. Carlsson. Topology and data. *Bull. Amer. Math. Soc.*, 46(4):459–512, 2009.

[23] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction.* American Mathematical Society, 2010.

[24] A. Bukkuri, N. Andor, and I. K. Darcy. Applications of topological data analysis in oncology. *Front. Artif. Intell.*, 4:659037, 2021.

[25] R. Rabadán, Y. Mohamedi, U. Rubin, T. Chu, A. N. Alghalith, O. Elliott, L. Arnés, S. Cal, Á. J. Obaya, A. J. Levine, and P. G. Cámara. Identification of relevant genetic alterations in cancer using topological data analysis. *Nat. Commun.*, 11:3808, 2020.

[26] P. G. Cámara. Topological methods for genomics: present and future directions. *Curr. Opin. Syst. Biol.*, 1:95–101, 2017.

[27] T. Songdechakraiwut and M. K. Chung. Topological learning for brain networks. *Ann. Appl. Stat.*, 17(1):403–433, 2023.

[28] D. Leykam and D. G. Angelakis. Topological data analysis and machine learning. *Advances in Physics: X*, 8(1):2202331, 2023.

[29] J. Murugan and D. Robertson. An introduction to topological data analysis for physicists: From LGM to FRBs. *arXiv preprint*, 2019. arXiv:1904.11044.

[30] Y. Hiraoka, T. Nakamura, A. Hirata, E. G. Escolar, K. Matsue, and Y. Nishiura. Hierarchical structures of amorphous solids characterized by persistent homology. *Proceedings of the National Academy of Sciences*, 113(18):7035–7040, 2016.

[31] S. S. Sørensen, C. A. N. Biscio, M. Bauchy, L. Fajstrup, and M. M. Smedskjaer. Revealing hidden medium-range order in amorphous materials using topological data analysis. *Sci. Adv.*, 6(37), 2020.

[32] T. Nakamura, Y. Hiraoka, A. Hirata, E. G. Escolar, and Y. Nishiura. Persistent homology and manybody atomic structure for medium-range order in the glass. *Nanotechnology*, 26(30):304001, 2015.

[33] R. Mengoni, A. Di Pierro, L. Memarzadeh, and S. Mancini. Persistent homology analysis of multiqubit entanglement. *Quantum Inf. Comput.*, 20(5–6):375–399, 2020.

[34] B. Olsthoorn. Persistent homology of quantum entanglement. *Phys. Rev. B*, 107:115174, 2023.

[35] I. Donato, M. Gori, M. Pettini, G. Petri, S. De Nigris, R. Franzosi, and F. Vaccarino. Persistent homology analysis of phase transitions. *Phys. Rev. E*, 93:052138, 2016.

[36] W. H. Shah, R. Jaimes-Reátegui, G. Huerta-Cuellar, J. H. García-López, and A. N. Pisarchik. Persistent homology approach for uncovering transitions to

chaos. *Chaos, Solitons & Fractals*, 192:116054, 2025.

[37] A. Myers. *Dynamical systems analysis using topological signal processing.* PhD thesis, Michigan State University, 2021.

[38] S. Lloyd, S. Garnerone, and P. Zanardi. Quantum algorithms for topological and geometric analysis of data. *Nat. Commun.*, 7:10138, 2016.

[39] R. Hayakawa. Quantum algorithm for persistent Betti numbers and topological data analysis. *Quantum*, 6:873, 2022.

[40] C. M. G. Leditto, A. Southwell, B. Tonekaboni, G. A. L. White, M. Usman, and K. Modi. Topological signal processing on quantum computers for higher-order network analysis. *Phys. Rev. Appl.*, 23:054054, 2025.

[41] C. Gyurik, C. Cade, and V. Dunjko. Towards quantum advantage via topological data analysis. *Quantum*, 6:855, 2022.

[42] D. W. Berry, Y. Su, C. Gyurik, R. King, J. Basso, A. D. T. Barba, A. Rajput, N. Wiebe, V. Dunjko, and R. Babbush. Analyzing prospects for quantum advantage in topological data analysis. *PRX Quantum*, 5:010319, 2024.

[43] A. Schmidhuber and S. Lloyd. Complexity-theoretic limitations on quantum algorithms for topological data analysis. *PRX Quantum*, 4:040349, 2023.

[44] J. Cai, J. Luo, S. Wang, and S. Yang. Feature selection in machine learning: a new perspective. *Neurocomputing*, 300:70–79, 2018.

[45] Y. Feng, H. You, Z. Zhang, R. Ji, and Y. Gao. Hypergraph neural networks. *Proc. AAAI Conf. Artif. Intell.*, 33:3558–3565, 2019.

[46] J. Zhou et al. Graph neural networks: a review of methods and applications. *AI Open*, 1:57–81, 2020.

[47] F. Noé, A. Tkatchenko, K.-R. Müller, and C. Clementi. Machine learning-accelerated quantum mechanics-based atomistic simulations. *Annu. Rev. Phys. Chem.*, 71:361–390, 2020.

[48] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91:045002, 2019.

[49] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. F. Wang, and L. Yang. Physics-informed machine learning. *Nature Rev. Phys.*, 3:422–440, 2021.

[50] F. Hensel, M. Moor, and B. Rieck. A survey of topological machine learning methods. *Front. Artif. Intell.*, 4:681108, 2021.

[51] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington. A roadmap for the computation of persistent homology. *EPJ Data Sci.*, 6:17, 2017.

[52] P. Skraba and K. Turner. Wasserstein stability for persistence diagrams. *arXiv preprint*, 2025. arXiv:2006.16824.

[53] H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta, and L. Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017.

[54] P. Bubenik. Statistical topological data analysis using persistence landscapes. *J. Mach. Learn. Res.*, 16(3):77–102, 2015.

[55] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33:249–274, 2005.

[56] A. R. Fritsch, Mingwu Lu, G. H. Reid, A. M. Piñeiro, and I. B. Spielman. Creating solitons with controllable and near-zero velocity in Bose-Einstein condensates. *Phys. Rev. A*, 101:053629, 2020.

[57] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 2017.

[58] S. Choi, Y. Salamin, C. Roques-Carmes, J. Sloan, M. Horodynski, and M. Soljacic. Observing the dynamics of quantum states generated inside nonlinear optical cavities. *arXiv preprint*, 2024. arXiv:2412.01772.

[59] M. Cramer, M. B. Plenio, S. T. Flammia, D. Gross, S. D. Bartlett, R. Somma, O. Landon-Cardinal, Y.-K. Liu, and D. Poulin. Efficient quantum state tomography. *Nature Communications*, 1:149, 2010.

[60] V. Robins, N. Rooney, and E. Bradley. Topology-based signal separation. *Chaos*, 14(2):305–313, 2004.

[61] V. Robins. Betti number signatures of homogeneous Poisson point processes. *Physical Review E*, 74(6):061107, 2006.

[62] D. Horak, S. Maletić, and M. Rajković. Persistent homology of complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(03):P03034, 2009.

[63] K. Krishan, H. Kurtuldu, M. F. Schatz, M. Gameiro, K. Mischaikow, and S. Madruga. Homology and symmetry breaking in Rayleigh-Bénard convection: Experiments and simulations. *Physics of Fluids*, 19(11):117105, 2007.

[64] H. Kurtuldu, K. Mischaikow, and M. F. Schatz. Extensive scaling from computational homology and Karhunen-Loève decomposition analysis of Rayleigh-Bénard convection experiments. *Physical Review Letters*, 107(3):034503, 2011.

[65] S. Hong and D. Kim. Medium-range order in amorphous ices revealed by persistent homology. *Journal of Physics: Condensed Matter*, 31(45):455403, 2019.

[66] S. Ardanza-Trevijano, I. Zuriguel, R. Arévalo, and D. Maza. Topological analysis of tapped granular media using persistent homology. *Physical Review E*, 89(5):052212, 2014.

[67] M. Saadatfar, H. Takeuchi, V. Robins, N. Francois, and Y. Hiraoka. Pore configuration landscape of granular crystallization. *Nature Communications*, 8:15082, 2017.

[68] B. Olsthoorn, J. Hellsvik, and A. V. Balatsky. Finding hidden order in spin models with persistent homology. *Phys. Rev. Res.*, 2:043308, 2020.

[69] D. Sehayek and R. G. Melko. Persistent homology of Z_2 gauge theories. *Physical Review B*, 106(8):085111, 2022.

[70] D. Ormrod Morley, P. S. Salmon, and M. Wilson. Persistent homology in two-dimensional atomic networks. *The Journal of Chemical Physics*, 154(12):124109, 2021.

[71] S. Park, Y. Hwang, and B.-J. Yang. Unsupervised learning of topological phase diagram using topological data analysis. *Physical Review B*, 105(19):195115, 2022.

[72] D. Leykam and D. G. Angelakis. Photonic band structure design using persistent homology. *APL Photonics*, 6(3):030802, 2021.

[73] G. Hamilton and F. Leditzky. Probing multipartite entanglement through persistent homology. *Commun. Math. Phys.*, 405, 2024.

[74] Z. Su, X. Liu, L. B. Hamdan, V. Maroulas, J. Wu, G. Carlsson, and G. W. Wei. Topological data analysis and topological deep learning beyond persistent homology: a review. *Artificial Intelligence Review*, 59(2):58, 2026.

[75] T. Madeleine, N. Podoliak, O. Buchnev, I. Membrillo Solis, T. Orlova, M. van Rossem, M. Kaczmarek, G. D'Alessandro, and J. Brodzki. Topological learning for the classification of disorder: An application to the design of metasurfaces. *ACS Nano*, 18(1), 2024.

[76] X. Crean, J. Giansiracusa, and B. Lucini. Topological data analysis of monopole current networks in $U(1)$ lattice gauge theory. *SciPost Phys.*, 17:100, 2024.

[77] D. Spitz, J. M. Urban, and J. M. Pawlowski. Topological data analysis of the deconfinement transition in SU(3) lattice gauge theory. *Phys. Rev. D*, 111:114519, 2025.

[78] V. Noel and D. Spitz. Detecting defect dynamics in relativistic field theories far from equilibrium using topological data analysis. *Phys. Rev. D*, 109:056011, 2024.

[79] M. Crichigno and T. Kohler. Clique homology is QMA-hard. *Nature Communications*, 15(1):9846, 2024.

[80] S. Apers, S. Gribling, S. Sen, and D. Szabó. A (simple) classical algorithm for estimating Betti numbers. *Quantum*, 7:1202, December 2023.

[81] C. Gyurik, A. Schmidhuber, R. King, V. Dunjko, and R. Hayakawa. Provable quantum speedups for computing persistence in topological data analysis. *PRX Quantum*, 7:020361, 2026.

[82] S. Scali, C. Umeano, and O. Kyriienko. The topology of data hides in quantum thermal states. *APL Quantum*, 1(3):036106, 2024.

[83] S. Scali, C. Umeano, and O. Kyriienko. Quantum topological data analysis via the estimation of the density of states. *Phys. Rev. A*, 110:042616, 2024.

[84] C. M. G. Leditto, A. Southwell, B. Tonekaboni, M. Usman, and K. Modi. Quantum HodgeRank: Topology-based rank aggregation on quantum computers. *Phys. Rev. Appl.*, 23:L061005, 2025.

[85] K. Mittal and S. Gupta. Topological characterization and early detection of bifurcations and chaos in complex systems using persistent homology. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(5):051102, 2017.

[86] A. Myers, E. Munch, and F. A. Khasawneh. Persistent homology of complex networks for dynamic state detection. *Phys. Rev. E*, 100:022314, 2019.

[87] J. R. Tempelman and F. A. Khasawneh. A look into chaos detection through topological data analysis. *Physica D: Nonlinear Phenomena*, 406:132446, 2020.

[88] L. Polterovich and E. Shelukhin. Autonomous Hamiltonian flows, Hofer's geometry and persistence modules. *Selecta Math.*, 22:227–296, 2016.

[89] H. Cao, D. Leykam, and D. G. Angelakis. Unravelling quantum chaos using persistent homology. *Phys. Rev. E*, 107:044204, 2023.

[90] E. Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 2 edition, 2002.

[91] M. Berry. Quantum chaology, not quantum chaos. *Physica Scripta*, 40(3):335–336, 1989.

[92] M. V. Berry and M. Tabor. Level Clustering in the Regular Spectrum. *Proc. R. Soc. London A*, 356(1686):375–394, 1977.

[93] F. Haake, M. Kuś, and R. Scharf. Classical and quantum chaos for a kicked top. *Z. Physik B - Condensed Matter*, 65(3):381–395, September 1987.

[94] A. R. Kolovsky and A. Buchleitner. Quantum chaos in the Bose-Hubbard model. *EPL*, 68(5):632, 2004.

[95] F. Haake, H. Wiedemann, and K. Zyczkowski. Lyapunov exponents from quantum dynamics. *Annalen der Physik*, 504(7):531–539, 1992.

[96] V. I. Man'ko and R. Vilela Mendes. Lyapunov exponent in quantum mechanics. A phase-space approach. *Physica D: Nonlinear Phenomena*, 145(3):330–348, 2000.

[97] T. Bhattacharya, S. Habib, and K. Jacobs. Continuous quantum measurement and the emergence of classical chaos. *Phys. Rev. Lett.*, 85:4852–4855, 2000.

[98] E. B. Rozenbaum, S. Ganeshan, and V. Galitski. Lyapunov exponent and out-of-time-ordered correlator's growth rate in a chaotic system. *Phys. Rev. Lett.*, 118:086801, 2017.

[99] J. Z. Wang, G. Benenti, G. Casati, and W. G. Wang. Quantum chaos and the correspondence principle. *Phys. Rev. E*, 103(3):L030201, Mar 2021.

[100] A. Sahu, S. PG, and V. Madhok. Effect of chaos on information gain in quantum tomography. *Phys. Rev. E*, 106(2):024209, 2022.

[101] F. Borgonovi, F. M. Izrailev, L. F. Santos, and V. G. Zelevinsky. Quantum chaos and thermalization in isolated systems of interacting particles. *Physics Reports*, 626:1–58, 2016.

[102] L. K. Joshi, A. Elben, A. Vikram, B. Vermersch, V. Galitski, and P. Zoller. Probing many-body quantum chaos with quantum simulators. *Phys. Rev. X*, 12:011018, 2022.

[103] K. M. Frahm, R. Fleckinger, and D. L. Shepelyansky. Quantum chaos and random matrix theory for fidelity decay in quantum computations with static imperfections. *Eur. Phys. J. D*, 29:139–155, 2004.

[104] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices. *Nature Phys.*, 14:595–600, 2018.

[105] B. Swingle, G. Bentsen, M. Schleier-Smith, and P. Hayden. Measuring the scrambling of quantum information. *Phys. Rev. A*, 94(4), 2016.

[106] J. Li, R. H. Fan, H. Y. Wang, B. T. Ye, B. Zeng, H. Zhai, X. H. Peng, and J. F. Du. Measuring out-of-time-order correlators on a nuclear magnetic resonance quantum simulator. *Phys. Rev. X*, 7(3):031011, 2017.

[107] S. Maletić, Y. Zhao, and M. Rajkovic. Persistent topological features of dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26:053105, 2015.

[108] Q. H. Tran and Y. Hasegawa. Topological time-series analysis with delay-variant embedding. *Phys. Rev. E*, 99:032209, 2019.

[109] A. Wolf, J. Swift, H. Swinney, and J. Vastano. Determining Lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena*, 16:285–317, 1985.

[110] A. Margaris, N. Kofidis, and M. Roumeliotis. A detailed study of the Wolf's algorithm. *International Journal of Computer Mathematics*, 86(7):1135–1148, 2009.

[111] H. P. Breuer and F. Petruccione. *The theory of open quantum systems*. Oxford University Press, 2002.

[112] K. Mølmer, Y. Castin, and J. Dalibard. Monte Carlo wave-function method in quantum optics. *J. Opt. Soc. Am. B*, 10(3):524–538, 1993.

[113] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi. Observing single quantum trajectories of a superconducting quantum bit. *Nature*, 502(7470):211–214, 2013.

[114] S. J. Weber, K. W. Murch, M. E. Schwartz, N. Roch, and I. Siddiqi. Quantum trajectories of superconducting qubits. *Computus Rendus Phys.*, 17(7):766–777, 2016.

[115] F. Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980*, pages 366–381, Berlin, Heidelberg, 1981. Springer Berlin Heidelberg.

[116] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Phys. Rev. A*, 33:1134–1140, 1986.

[117] M. B. Kennel, R. Brown, and H. D. I. Abarbanel. Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Phys. Rev. A*, 45:3403–3411, 1992.

[118] P. Szlachetka, K. Grygiel, and J. Bajer. Chaos and order in a kicked anharmonic oscillator: Classical and quantum analysis. *Phys. Rev. E*, 48:101–108, 1993.

[119] T. P. Spiller and J. F. Ralph. The emergence of chaos in an open quantum system. *Phys. Lett. A*, 194(4):235–240, 1994.

[120] I. I. Yusipov, O. S. Vershinina, S. V. Denisov, and M. V. Ivanchenko. Photon waiting-time distributions: A keyhole into dissipative quantum chaos. *Chaos*, 30(2):023107, 2020.

[121] M. B. Plenio and P. L. Knight. The quantum-jump approach to dissipative dynamics in quantum optics. *Rev. Mod. Phys.*, 70:101–144, 1998.

[122] P. S. Linsay. Period doubling and chaotic behavior in a driven anharmonic oscillator. *Phys. Rev. Lett.*, 47:1349–1352, 1981.

[123] T. A. Brun, I. C. Percival, and R. Schack. Quantum chaos in open systems: a quantum state diffusion analysis. *J. Phys. A: Mathematical and General*, 29(9):2077–2090, 1996.

[124] Y. Ota and I. Ohba. Crossover from classical to quantum behavior of the Duffing oscillator through a pseudo-Lyapunov-exponent. *Phys. Rev. E*, 71:015201(R), 2005.

[125] J. R. Johansson, P. D. Nation, and F. Nori. QuTiP: An open-source Python framework for the dynamics of open quantum systems. *Comp. Phys. Comm.*, 183:1760–1772, 2012.

[126] I. Güzel and A. Kaygun. Classification of stochastic processes with topological data analysis. *arXiv preprint*, 2022. arXiv:2206.03973.

[127] I. I. Yusipov, O. S. Vershinina, S. Denisov, S. P. Kuznetsov, and M. V. Ivanchenko. Quantum Lyapunov exponents beyond continuous measurements. *Chaos*, 29(6):063130, 2019.

[128] B. Pokharel, M. Misplon, and W. Lynn. Chaos and dynamical complexity in the quantum to classical transition. *Scientific Reports*, 8:2108, 2018.

[129] T. Macrì, A. Smerzi, and L. Pezzè. Loschmidt echo for quantum metrology. *Phys. Rev. A*, 94:010102(R), 2016.

[130] P. Zanardi and N. Anand. Information scrambling and chaos in open quantum systems. *Phys. Rev. A*, 103(6):062214, Jun 2021.

[131] D. Huber, O. V. Marchukov, H.-W. Hammer, and A. G. Volosniev. Morphology of three-body quantum states from machine learning. *New Journal of Physics*, 23(6):065009, 2021.

[132] J. Carrasquilla and R. G. Melko. Machine learning phases of matter. *Nature Communications*, 8:180, 2017.

[133] L. Huijse, R. G. Melko, and J. Carrasquilla. Unsupervised learning of phase transitions: from principal component analysis to variational autoencoders. *Phys. Rev. B*, 98:045129, 2018.

[134] H. Cao, D. G. Angelakis, and D. Leykam. Unsupervised learning of quantum many-body scars using intrinsic dimension. *Machine Learning: Science and Technology*, 2024.

[135] J. M. Deutsch. Eigenstate thermalization hypothesis. *Reports on Progress in Physics*, 81(8):082001, 2018.

[136] C. J. Turner, A. A. Michailidis, and D. A. Abanin. Weak ergodicity breaking from quantum many-body scars. *Nature Phys.*, 14:745–749, 2018.

[137] Michael Schecter and Thomas Iadecola. Weak ergodicity breaking and quantum many-body scars in spin-1 XY magnets. *Physical Review Letters*, 123(14), October 2019.

[138] A. Chandran, T. Iadecola, V. Khemani, and R. Moessner. Quantum many-body scars: A quasiparticle perspective. *Annual Review of Condensed Matter Physics*, 14(1):443–469, 2023.

[139] K. M. Lee, R. Melendrez, A. Pal, and H. J. Changlani. Exact three-colored quantum scars from geometric frustration. *Phys. Rev. B*, 101:241111, 2020.

[140] S. Moudgalya, N. Regnault, and B. A. Bernevig. Entanglement of exact excited states of Affleck-Kennedy-Lieb-Tasaki models: Exact results, many-body scars, and violation of the strong eigenstate thermalization hypothesis. *Phys. Rev. B*, 98:235156, 2018.

[141] D. K. Mark, C. J. Lin, and O. I. Motrunich. Exact eigenstates in the Lesanovsky model, proximity to integrability and the PXP model, and approximate scar states. *Phys. Rev. B*, 101:094308, 2020.

[142] J.-Y. Desaules, D. Banerjee, A. Hudomal, Z. Papić, A. Sen, and J. C. Halimeh. Weak ergodicity breaking in the Schwinger model. *Phys. Rev. B*, 107:L201105, 2023.

[143] J.-Y. Desaules, A. Hudomal, D. Banerjee, A. Sen, Z. Papić, and J. C. Halimeh. Prominent quantum many-body scars in a truncated Schwinger model. *Phys. Rev. B*, 107:205112, 2023.

[144] T. Iadecola and M. Schecter. Quantum many-body scar states with emergent kinetic constraints and finite-entanglement revivals. *Physical Review B*, 101(2), 2020.

[145] S. Scherg, T. Kohlert, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger. Observing non-ergodicity due to kinetic constraints in tilted Fermi-Hubbard chains. *Nature Communications*, 12(1), 2021.

[146] A. Hudomal, J.-Y. Desaules, B. Mukherjee, G.-X. Su, J. C. Halimeh, and Z. Papić. Driving quantum many-body scars in the PXP model. *Physical Review B*, 106(10), 2022.

[147] G. X. Su, H. Sun, A. Hudomal, J.-Y. Desaules, Z. Y. Zhou, B. Yang, J. C. Halimeh, Z. S. Yuan, Z. Papić, and J. W. Pan. Observation of many-body scarring in a Bose-Hubbard quantum simulator. *Phys. Rev. Res.*, 5(2), 2023.

[148] J. Carrasquilla. Machine learning for quantum matter. *Advances in Physics: X*, 5(1):1797528, 2020.

[149] H. Atanasova, L. Bernheimer, and G. Cohen. Stochastic representation of many-body quantum states. *Nat. Commun.*, 14:3601, 2023.

[150] L. Wang. Discovering phase transitions with unsupervised learning. *Physical Review B*, 94(19), 2016.

[151] W. J. Hu, R. R. P. Singh, and R. T. Scalettar. Discovering phases, phase transitions, and crossovers through unsupervised machine learning: A critical examination. *Phys. Rev. E*, 95:062122, 2017.

[152] A. Lidiak and Z. X. Gong. Unsupervised machine learning of quantum phase transitions using diffusion maps. *Phys. Rev. Lett.*, 125:225701, 2020.

[153] J. F. Rodriguez-Nieva and M. S. Scheurer. Identifying topological order through unsupervised machine learning. *Nat. Phys.*, 15:790–795, 2019.

[154] T. Mendes-Santos, A. Angelone, Alex Rodriguez, R. Fazio, and M. Dalmonte. Intrinsic dimension of path integrals: Data-mining quantum criticality and emergent simplicity. *PRX Quantum*, 2:030332, 2021.

[155] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, S. Choi, M. Serbyn, M. Greiner, V. Vuletić, and M. D. Lukin. Controlling quantum many-body dynamics in driven Rydberg atom arrays. *Science*, 371(6536):1355–1359, March 2021.

[156] T. Szołdra, P. Sierant, M. Lewenstein, and J. Zakrzewski. Unsupervised detection of decoupled subspaces: Many-body scars and beyond. *Phys. Rev. B*, 105:224205, 2022.

[157] C. D. Han, C. Z. Wang, and Y. C. Lai. Meta-machine-learning-based quantum scar detector. *Phys. Rev. Appl.*, 19:064042, 2023.

[158] K. Omiya and M. Müller. Quantum many-body scars in bipartite Rydberg arrays originating from hidden projector embedding. *Phys. Rev. A*, 107:023318, 2023.

[159] H. R. Wang, D. Yuan, S. Y. Zhang, Z. Wang, D. L. Deng, and L. M. Duan. Embedding quantum many-body scars into decoherence-free subspaces. *Phys. Rev. Lett.*, 132:150401, 2024.

[160] I. Borg and P. Groenen. *Modern multidimensional scaling: Theory and applications.* Springer Series in Statistics, 2005.

[161] C. J. Cao, S. M. Carroll, and S. Michalakis. Space from hilbert space: Recovering geometry from bulk entanglement. *Phys. Rev. D*, 95:024031, 2017.

[162] D. J. Skinner, B. Song, H. Jeckel, E. Jelli, K. Drescher, and J. Dunkel. Topological metric detects hidden order in disordered media. *Phys. Rev. Lett.*, 126:048101, 2021.

[163] J. B. Kruskal. Nonmetric multidimensional scaling: A numerical method. *Psychometrika*, 29:115–129, 1964.

[164] G. De Palma, M. Marvian, D. Trevisan, and S. Lloyd. The quantum Wasserstein distance of order 1. *IEEE Transactions on Information Theory*, 67(10):6627–6643, 2021.

[165] M. Iversen and A. E. B. Nielsen. Tower of quantum scars in a partially many-body localized system. *Phys. Rev. B*, 107:205140, 2023.

[166] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White. Measurement of qubits. *Physical Review A*, 64(5), 2001.

[167] P. Campadelli, E. Casiraghi, C. Ceruti, and A. Rozza. Intrinsic dimension estimation: Relevant techniques and a benchmark framework. *Mathematical Problems in Engineering*, 2015:21, 2015.

[168] D. Granata and V. Carnevale. Accurate estimation of the intrinsic dimension using graph distances: Unraveling the geometric complexity of datasets. *Sci Rep*, 6:31377, 2016.

[169] T. Mendes-Santos, X. Turkeshi, M. Dalmonte, and A. Rodriguez. Unsupervised learning universal critical behavior via the intrinsic dimension. *Phys. Rev. X*, 11:011040, 2021.

[170] R. K. Panda, R. Verdel, A. Rodriguez, H. Sun, G. Bianconi, and M. Dalmonte. Non-parametric learning critical behavior in Ising partition functions: PCA entropy and intrinsic dimension. *SciPost Phys. Core*, 6:086, 2023.

[171] R. Verdel, V. Vitale, R. K. Panda, E. D. Donkor, A. Rodriguez, S. Lannig, Y. Deller, H. Strobel, M. K. Oberthaler, and M. Dalmonte. Data-driven discovery of relevant information in quantum simulators. *arXiv preprint*, 2023. arXiv:2307.10040.

[172] F. Camastra and A. Staiano. Intrinsic dimension estimation: Advances and open problems. *Information Sciences*, 328:26–41, 2016.

[173] E. Facco, M. d'Errico, and A. Rodriguez. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Sci. Rep.*, 7:12140, 2017.

[174] I. Macocco, A. Glielmo, J. Grilli, and A. Laio. Intrinsic dimension estimation for discrete metrics. *Phys. Rev. Lett.*, 130:067401, 2023.

[175] IonQ. IonQ Aria: Practical performance, 2023. https://www.ionq.com/resources/ionq-aria-practical-performance. Accessed: 2023-12-12.

[176] S. Dooley. Robust quantum sensing in strongly interacting systems with many-body scars. *PRX Quantum*, 2(2), 2021.

[177] J.-Y. Desaules, F. Pietracaprina, Z. Papić, J. Goold, and S. Pappalardi. Extensive multipartite entanglement from SU(2) quantum many-body scars. *Physical Review Letters*, 129(2), 2022.

[178] W. Kao, K. Y. Li, K. Y. Lin, S. Gopalakrishnan, and B. L. Lev. Topological pumping of a 1D dipolar gas into strongly correlated prethermal states. *Science*, 371(6526):296–300, 2021.

[179] P. F. Zhang et al. Many-body Hilbert space scarring on a superconducting processor. *Nature Physics*, 19(1), 2022.

[180] B. Bhattacharjee, S. Sur, and P. Nandy. Probing quantum scars and weak ergodicity breaking through quantum complexity. *Physical Review B*, 106(20), 2022.

[181] V. Vitale, T. Mendes-Santos, A. Rodriguez, and M. Dalmonte. Topological Kolmogorov complexity and the Berezinskii-Kosterlitz-Thouless mechanism. *arXiv preprint*, 2023. arXiv:2305.05396.

[182] T. Mendes-Santos, M. Schmitt, A. Angelone, A. Rodriguez, P. Scholl, H. J. Williams, D. Barredo, T. Lahaye, A. Browaeys, M. Heyl, and M. Dalmonte. Wave function network description and Kolmogorov complexity of quantum many-body systems. *arXiv preprint*, 2023. arXiv:2301.13216.

[183] S. Nandy, B. Mukherjee, A. Bhattacharyya, and A. Banerjee. Quantum state complexity meets many-body scars. *arXiv preprint*, 2023. arXiv:2305.13322.

[184] W. W. Ho, S. Choi, H. Pichler, and M. D. Lukin. Periodic orbits, entanglement, and quantum many-body scars in constrained models: Matrix product state approach. *Phys. Rev. Lett.*, 122:040603, 2019.

[185] V. Khemani, C. R. Laumann, and A. Chandran. Signatures of integrability in the dynamics of Rydberg-blockaded chains. *Phys. Rev. B*, 99:161101, 2019.

[186] K. Bull, J.-Y. Desaules, and Z. Papić. Quantum scars as embeddings of weakly broken Lie algebra representations. *Phys. Rev. B*, 101:165139, 2020.

[187] L. W. Clark, B. M. Anderson, L. Feng, A. Gaj, K. Levin, and C. Chin. Observation of density-dependent gauge fields in a Bose-Einstein condensate

based on micromotion control in a shaken two-dimensional lattice. *Phys. Rev. Lett.*, 121:030402, 2018.

[188] F. Görg, K. Sandholzer, and J. Minguzzi. Realization of density-dependent Peierls phases to engineer quantized gauge fields coupled to ultracold matter. *Nat. Phys.*, 15:1161–1167, 2019.

[189] H. Cao and D. Leykam. Cusp solitons mediated by a topological nonlinearity. *Chaos*, 36(3):033104, 2026.

[190] W. Królikowski and O. Bang. Solitons in nonlocal nonlinear media: Exact solutions. *Phys. Rev. E*, 63:016610, 2000.

[191] W. Krolikowski, O. Bang, J. J. Rasmussen, and J. Wyller. Modulational instability in nonlocal nonlinear kerr media. *Phys. Rev. E*, 64:016612, 2001.

[192] W. Królikowski, O. Bang, N. I. Nikolov, D. Neshev, J. Wyller, J. J. Rasmussen, and D. Edmundson. Modulational instability, solitons and beam propagation in spatially nonlocal nonlinear media. *Journal of Optics B: Quantum and Semiclassical Optics*, 6(5):S288, 2004.

[193] S. K. Turitsyn and A. V. Mikhailov. Chapter 6.2.5 - Applications of Solitons. In *Scattering*, pages 1741–1753. Academic Press, London, 2002.

[194] K. Porsezian and V. C. Kuriakose. *Optical Solitons: Theoretical and Experimental Challenges*, volume 613. Springer Science & Business Media, 2003.

[195] N. S. Ginsberg, J. Brand, and L. V. Hau. Observation of hybrid soliton vortex-ring structures in Bose-Einstein condensates. *Phys. Rev. Lett.*, 94:040403, 2005.

[196] M. Porkolab and M. V. Goldman. Upper-hybrid solitons and oscillating-two-stream instabilities. *Phys. Fluids*, 19(6), 1976.

[197] C. Montes. Nonlinear kinetic equation in plasma physics leading to soliton structures. In *Nonlinear Problems in Theoretical Physics*, pages 205–216, Berlin, Heidelberg, 1979. Springer Berlin Heidelberg.

[198] A. Ankiewicz, Y. Wang, S. Wabnitz, and N. Akhmediev. Extended nonlinear Schrödinger equation with higher-order odd and even terms and its rogue wave solutions. *Phys. Rev. E*, 89:012907, 2014.

[199] L. Li, Z. D. Li, B. A. Malomed, D. Mihalache, and W. M. Liu. Exact soliton solutions and nonlinear modulation instability in spinor Bose-Einstein condensates. *Phys. Rev. A*, 72:033611, 2005.

[200] O. Bang, W. Krolikowski, J. Wyller, and J. J. Rasmussen. Collapse arrest and soliton stabilization in nonlocal nonlinear media. *Phys. Rev. E*, 66:046619, 2002.

[201] F. W. Ye, Y. V. Kartashov, and L. Torner. Stabilization of dipole solitons in nonlocal nonlinear media. *Phys. Rev. A*, 77:043821, 2008.

[202] Z. Y. Xu, Y. V. Kartashov, and L. Torner. Stabilization of vector soliton complexes in nonlocal nonlinear media. *Phys. Rev. E*, 73:055601, 2006.

[203] W. B. Rabie, H. M. Ahmed, I. Samir, and M. Alnahhass. Optical solitons and stability analysis for NLSE with nonlocal nonlinearity, nonlinear chromatic dispersion and Kudryashov's generalized quintuple-power nonlinearity. *Results in Physics*, 59:107589, 2024.

[204] M. Edmonds, M. Valiente, G. Juzeliūnas, L. Santos, and P. Ohberg. Simulating an interacting gauge theory with ultracold Bose gases. *Physical Review*

Letters, 110:085301, 2013.

[205] J. H. Zheng, B. Xiong, G. Juzeliūnas, and D. W. Wang. Topological condensate in an interaction-induced gauge potential. *Phys. Rev. A*, 92:013604, 2015.

[206] A. Biswas. Topological 1-soliton solution of the nonlinear Schrodinger's equation with Kerr law nonlinearity in 1+2 dimensions. *Commun. Nonlinear Sci. Numer. Simul.*, 14(7):2845–2847, 2009.

[207] W. N. Faugno, M. Salerno, and T. Ozawa. Density dependent gauge field inducing emergent Su-Schrieffer-Heeger physics, solitons, and condensates in a discrete nonlinear Schrödinger equation. *Phys. Rev. Lett.*, 132:023401, 2024.

[208] I. A. Bhat, S. Sivaprakasam, and B. A. Malomed. Modulational instability and soliton generation in chiral Bose-Einstein condensates with zero-energy nonlinearity. *Phys. Rev. E*, 103:032206, 2021.

[209] K. X. Yao, Z. D. Zhang, and C. Chin. Domain-wall dynamics in Bose–Einstein condensates with synthetic gauge fields. *Nature*, 602:68–72, 2022.

[210] E. Martinez, C. Muschik, and P. Schindler. Real-time dynamics of lattice gauge theories with a few-qubit quantum computer. *Nature*, 534:516–519, 2016.

[211] R. Gordon, J. T. Blakely, and D. Sinton. Particle-optical self-trapping. *Phys. Rev. A*, 75:055801, 2007.

[212] Y. Lamhot, A. Barak, O. Peleg, and M. Segev. Self-trapping of optical beams through thermophoresis. *Phys. Rev. Lett.*, 105:163906, 2010.

[213] R. Gautam, Y. X. Xiang, J. Lamstein, Y. Liang, A. Bezryadina, G. Liang, T. Hansson, B. Wetzel, D. Preece, A. White, M. Silverman, S. Kazarian, J. J. Xu, R. Morandotti, and Z. G. Chen. Optical force-induced nonlinearity and self-guiding of light in human red blood cell suspensions. *Light: Science & Applications*, 8(1):31, 2019.

[214] J. P. Gordon, R. C. C. Leite, R. S. Moore, S. P. S. Porto, and J. R. Whinnery. Long-transient effects in lasers with inserted liquid samples. *Journal of Applied Physics*, 36(1):3–8, 1965.

[215] M. Horowitz, R. Daisy, O. Werner, and B. Fischer. Large thermal nonlinearities and spatial self-phase modulation in $Sr_xBa_{1-x}Nb_2O_6$ and $BaTiO_3$ crystals. *Opt. Lett.*, 17(7):475–477, 1992.

[216] Y. S. Kivshar and G. P. Agrawal. *Optical Solitons: From Fibers to Photonic Crystals*. Academic Press, 2003.

[217] S. Abe and A. Ogura. Solitary waves and their critical behavior in a nonlinear nonlocal medium with power-law response. *Phys. Rev. E*, 57:6066–6070, 1998.

[218] L. Q. Yuan, Q. Lin, M. Xiao, and S. H. Fan. Synthetic dimension in photonics. *Optica*, 5(11):1396–1405, 2018.

[219] D. Leykam and L. Q. Yuan. Topological phases in ring resonators: recent progress and future prospects. *Nanophotonics*, 9(15):4473–4487, 2020.

[220] S. L. Wang, B. Wang, C. Y. Liu, C. Z. Qin, L. G. Zhao, W. W. Liu, S. Longhi, and P. X. Lu. Nonlinear non-Hermitian skin effect and skin solitons in temporal photonic feedforward lattices. *Phys. Rev. Lett.*, 134:243805, 2025.

[221] D. D. Weller, M. Hefenbrock, M. Beigl, J. Aghassi-Hagmann, and M. B. Tahoori. Realization and training of an inverter-based printed neuromorphic

computing system. *Scientific Reports*, 11(1):9554, 2021.

[222] C. Pappas, S. Kovaios, M. Moralis-Pegios, A. Tsakyridis, G. Giamougiannis, M. Kirtas, J. Van Kerrebrouck, G. Coudyzer, X. Yin, N. Passalis, A. Tefas, and N. Pleros. Programmable Tanh-, ELU-, Sigmoid-, and Sin-based nonlinear activation functions for neuromorphic photonics. *IEEE Journal of Selected Topics in Quantum Electronics*, 29(6):1–10, 2023.

[223] D. Z. Zheng, S. Y. Xiang, N. Q. Li, Y. H. Zhang, X. X. Guo, L. Y. Zhao, and Y. Hao. Neuromorphic network with photonic weighting and photoelectronic nonlinear activation based on SOA and APD. *ACS Photonics*, 11(10):4193–4199, 2024.

[224] F. Chazal, L. J. Guibas, S. Y. Oudot, and P. Skraba. Persistence-based clustering in Riemannian manifolds. *J. ACM*, 60(6), 2013.

[225] J. Harer D. Cohen-Steiner, H. Edelsbrunner. Lipschitz functions have L_p-stable persistence. *Found. Comput. Math*, 10:127–139, 2010.

[226] G. Hamilton, T. Dore, and C. Plumberg. Applications of persistent homology in nuclear collisions. *Phys. Rev. C*, 106:064912, 2022.

[227] A. Cole, G. J. Loges, and G. Shiu. Quantitative and interpretable order parameters for phase transitions from persistent homology. *Phys. Rev. B*, 104:104426, 2021.

[228] M. Carriere, F. Chazal, M. Glisse, Y. Ike, H. Kannan, and Y. Umeda. Optimizing persistent homology based functions. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proc. Mach. Learn. Res.*, pages 1294–1303, 2021.

[229] V. E. Zakharov and L. A. Ostrovsky. Modulation instability: The beginning. *Physica D: Nonlinear Phenomena*, 238(5):540–548, 2009.

[230] A. Hasegawa and F. Tappert. Transmission of stationary nonlinear optical pulses in dispersive dielectric fibers. I. Anomalous dispersion. *Applied Physics Letters*, 23(3):142–144, 1973.

[231] T. R. Taha and M. I. Ablowitz. Analytical and numerical aspects of certain nonlinear evolution equations. II. Numerical, nonlinear Schrödinger equation. *Journal of Computational Physics*, 55(2):203–230, 1984.

[232] M. Gell-Mann and F. Low. Bound states in quantum field theory. *Phys. Rev.*, 84:350–354, 1951.

[233] L. Lehtovaara, J. Toivanen, and J. Eloranta. Solution of time-independent Schrödinger equation by the imaginary time propagation method. *Journal of Computational Physics*, 221(1):148–157, 2007.

[234] M. J. Ma and Z. Huang. Bright soliton solution of a Gross–Pitaevskii equation. *Applied Mathematics Letters*, 26(7):718–724, 2013.

[235] N. Karjanto. Bright soliton solution of the nonlinear Schrödinger equation: Fourier spectrum and fundamental characteristics. *Mathematics*, 10(23), 2022.

[236] G.-D. Lin, Wei Zhang, and L.-M. Duan. Characteristics of Bose-Einstein condensation in an optical lattice. *Phys. Rev. A*, 77:043626, 2008.

[237] M. Tajiri and Y. Watanabe. Breather solutions to the focusing nonlinear Schrödinger equation. *Phys. Rev. E*, 57:3510–3519, 1998.

[238] W. Chen, J. Mao, K. Chen, X. Li, Y. Tan, M. Li, and L. Zhang. Generation of high uniformity flat-top beams by reconstructing the amplitude distribution at descending edges. *Optics & Laser Technology*, 186:112776, 2025.

[239] Y. P. Zheng, W. J. Tan, D. D. Han, K. L. Ren, Y. W. Wang, F. Zhao, and J. H. Si. Propagation of high-power optical flat-topped beams in strongly nonlinear media. *Results in Physics*, 31:105016, 2021.

[240] J. Larson. Integrability versus quantum thermalization. *J. Phys. B: At. Mol. Opt. Phys.*, 46:224016, 2013.

[241] J. K. Eastman, S. S. Szigeti, J. J. Hope, and A. R. R. Carvalho. Controlling chaos in the quantum regime using adaptive measurements. *Phys. Rev. A*, 99:012111, 2019.

[242] B. Batistić and M. Robnik. Dynamical localization of chaotic eigenstates in the mixed-type systems: spectral statistics in a billiard system after separation of regular and chaotic eigenstates. *J. Phys. A: Math. Theor.*, 46(315102), 2013.

[243] G. Carlsson and V. de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.

[244] S. Tymochko, E. Munch, and F. A. Khasawneh. Using zigzag persistent homology to detect Hopf bifurcations in dynamical systems. *Algorithms*, 13(11):278, 2020.

[245] A. Piga, M. Lewenstein, and J. Q. Quach. Quantum chaos and entanglement in ergodic and nonergodic systems. *Phys. Rev. E*, 99:032213, 2019.

[246] J. J. Feng, B. Zhang, Z.-C. Yang, and Q. T. Zhuang. Uncovering quantum many-body scars with quantum machine learning. *npj Quantum Information*,

11:42, 2025.

[247] J. Ren, A. Hallam, L. Ying, and Z. Papić. Scarfinder: A detector of optimal scar trajectories in quantum many-body dynamics. *PRX Quantum*, 6:040332, 2025.

[248] N. Roch, M. E. Schwartz, F. Motzoi, C. Macklin, R. Vijay, A. W. Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi. Observation of measurement-induced entanglement and quantum trajectories of remote superconducting qubits. *Phys. Rev. Lett.*, 112:170501, 2014.

[249] S. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi. Mapping the optimal route between two quantum states. *Nature*, 511:570–573, 2014.

[250] L. Logarić, J. Goold, and S. Dooley. Dynamical signatures of conventional and asymptotic quantum many-body scars on a trapped ion simulator. *arXiv preprint*, 2026. arXiv:2604.12296.

[251] X. G. Wen and Z. H. Wang. Volume and topological invariants of quantum many-body systems. *Phys. Rev. Res.*, 2:033030, 2020.