Probing Non-Markovianity through Topological Signatures of Path Space

SÉRGIO DE ANDRADE, PAULO

ORCID: https://orcid.org/0009-0004-2555-3178

Abstract

The characterization of non-Markovian dynamics in open quantum systems remains a central challenge in quantum physics. Conventional measures often rely on specific properties of the system's evolution, such as the flow of information or the divisibility of the dynamical map. This paper introduces a fundamentally different approach by connecting quantum memory effects to the topological structure of the system's path space. We posit that the set of all possible quantum trajectories, when viewed as a geometric object, exhibits distinct topological signatures under non-Markovian evolution compared to its Markovian counterpart. Using tools from topological data analysis, specifically persistent homology, we analyze the path space generated from simulations of a canonical open quantum system. We demonstrate that topological invariants, such as the Betti numbers, which count holes and voids in the path space, serve as robust indicators of non-Markovianity. Our results reveal that the presence of quantum memory induces non-trivial topological features, such as persistent loops, in the space of paths. This suggests that memory effects constrain the system's evolution, forcing trajectories to organize into complex, non-contractible structures. This topological framework provides a model-agnostic and potentially more robust method for detecting and quantifying quantum memory, opening new avenues for understanding complex quantum dynamics from a geometric perspective.

Keywords: non-Markovianity, open quantum systems, topological data analysis, persistent homology, path space, quantum dynamics, memory effects

1 Introduction

The study of open quantum systems, where a principal system of interest interacts with an external environment or bath, is fundamental to virtually all areas of modern physics. No quantum system is perfectly isolated, and its dynamics are inevitably influenced by its surroundings. A common and powerful simplification in this context is the Markovian approximation, which assumes that the environment has no memory of past interactions with the system. This approximation posits that the future evolution of the system depends only on its present state, not its history. This leads to the well-known Lindblad master equation, which has been successfully applied to describe a wide range of phenomena in quantum optics, condensed matter physics, and quantum information.

However, the Markovian approximation breaks down in many physically relevant scenarios. When the system-environment coupling is strong, or when the environment is structured and possesses its own internal dynamics on a timescale comparable to the system's, memory effects become significant. This regime of non-Markovian dynamics is characterized by a feedback of information from the environment back to the system, leading to complex behaviors such as coherence revivals, non-exponential decay, and history-dependent evolution. Understanding and characterizing non-Markovianity is not merely an academic exercise; it is crucial for the development of quantum technologies. Memory effects can be both a hindrance, leading to decoherence in quantum computers, and a resource, potentially enhancing processes in quantum metrology, quantum control, and even biological systems like photosynthesis.

Despite its importance, detecting and quantifying non-Markovianity remains a significant theoretical and experimental challenge. A variety of measures have been proposed, most notably those based on the breakdown of CP-divisibility of the quantum dynamical map (e.g., the Breuer-Laine-Piilo measure) or the distinguishability of quantum states. While these measures have provided profound insights, they can be difficult to compute for complex systems and their physical interpretation is not always direct.

This paper proposes a paradigm shift in the characterization of quantum memory. Instead of focusing on the local-in-time properties of the dynamical map, we investigate the global, geometric structure of the entire set of possible system evolutions. We introduce the concept of the system's "path space"—the high-dimensional space where each point represents a complete trajectory or history of the quantum system over a finite time interval. Our central hypothesis is that non-Markovian dynamics imprint distinctive topological signatures onto this path space.

We employ tools from the burgeoning field of Topological Data Analysis (TDA), particularly persistent homology, to probe these signatures. TDA is a powerful framework for analyzing the "shape" of data in a way that is robust to noise and invariant under continuous deformations. By treating a large ensemble of simulated quantum trajectories as a point cloud in path space, we can compute its topological invariants, such as Betti numbers, which count the number of connected components, loops, voids, and higher-dimensional holes.

Our core finding is that Markovian and non-Markovian dynamics generate path spaces with fundamentally different topologies. Memoryless, Markovian evolution tends to produce a "simple" or contractible path space, akin to a filled-in ball. In contrast, the temporal correlations inherent in non-Markovian dynamics introduce constraints that "puncture" or "fold" the path space, creating persistent loops and other non-trivial topological features. These features indicate that the system's history restricts its future evolution in a complex manner, preventing the path space from being smoothly filled. We demonstrate that the persistence of these topological features provides a quantitative measure of non-Markovianity that correlates strongly with established measures. This approach offers a novel, holistic perspective on quantum memory, recasting it as a geometric property of the system's entire dynamical landscape.

This paper is structured as follows: Section 2 provides a review of the relevant literature on non-Markovian quantum dynamics and Topological Data Analysis. Section 3 details our methodology for defining and analyzing the path space of an open quantum system. Section 4 presents the numerical results for a canonical model system, showcasing the distinct topological signatures of Markovian and non-Markovian regimes. Section 5 discusses the physical interpretation of these results and the implications of our approach. Finally, Section 6 concludes with a summary and an outlook on future research directions.

2 Literature Review

The theoretical framework of this study is built upon the confluence of two distinct yet complementary fields: the physics of open quantum systems and the mathematics of topological data analysis. This section reviews the key concepts from both areas, establishing the foundation for our novel approach.

2.1 Non-Markovian Quantum Dynamics

The dynamics of an open quantum system are described by a dynamical map Φ_t that evolves the system's initial density operator $\rho(0)$ to its state at a later time t, $\rho(t) = \Phi_t(\rho(0))$. A process is termed Markovian if the dynamical map forms a one-parameter semigroup, implying complete positivity (CP) and trace-preserving properties, and satisfying the composition law $\Phi_{t+s} = \Phi_t \Phi_s$ for all $t, s \geq 0$. This corresponds to a memoryless process.

The departure from this idealized behavior is the hallmark of non-Markovianity. One of the most influential definitions is based on the concept of CP-divisibility. A dynamical map Φ_t is CP-divisible if it can be written as $\Phi_t = V_{t,s}\Phi_s$ for all $t \geq s \geq 0$, where the intermediate map $V_{t,s}$ is completely positive. A process is non-Markovian if this divisibility fails, interpreted as a backflow of information from the environment to the system. The Breuer-Laine-Piilo (BLP) measure formalizes this by quantifying the increase in the distinguishability (trace distance) of pairs of initial states, an event impossible under CP-divisible dynamics.

Another key perspective comes from the projection operator formalism, leading to master equations like the Nakajima-Zwanzig equation. It features a memory kernel, which accounts for the entire history of the system's state. The Markovian approximation corresponds to assuming this memory kernel is a delta function in time. Numerically exact methods, such as the Hierarchical Equations of Motion (HEOM), have been developed to solve non-Markovian master equations without resorting to perturbative approximations, enabling precise simulation of systems with strong memory effects.

2.2 Topological Data Analysis and Persistent Homology

Topological Data Analysis (TDA) is a field of applied mathematics that uses concepts from algebraic topology to analyze the shape of data. Its primary tool is persistent homology. The core idea is to represent a dataset (e.g., a point cloud) as a sequence of increasingly complex topological spaces, called a filtration. A common method is the Vietoris-Rips complex: for a scale ϵ, one places a ball of radius $\epsilon/2$ around each data point and builds a simplicial complex where points form a simplex if their balls intersect. As ϵ increases, connections form, and topological features like loops and voids can appear and later disappear.

Persistent homology tracks these features, recording the "birth" scale at which a feature appears and the "death" scale at which it is filled in. The "persistence" of a feature is the difference between its death and birth scales. Features that persist for a long range of ϵ are considered significant topological signatures, while short-lived features are attributed to noise. The output is a persistence diagram, summarizing the data's topology. The ranks of the homology groups, known as Betti numbers (β_0, β_1, β_2, ...), count the number of features of each dimension (components, loops, voids) at a given scale.

2.3 Geometric Perspectives on Quantum Dynamics

Our approach of analyzing the geometry of path space is part of a broader trend of applying geometric concepts to quantum mechanics. A notable alternative is quantum information geometry, which studies the manifold of quantum states itself. This framework endows the space of density matrices with a Riemannian metric, such as the Bures or Fubini-Study metric, which quantifies the distinguishability of nearby quantum states.

Within information geometry, the dynamics of a quantum system are viewed as a curve on this manifold. The speed of evolution along this curve, measured by the metric, is related to the quantum speed limit, which sets a fundamental bound on how fast a quantum system can evolve. Some studies have connected this speed to non-Markovianity, suggesting that information backflow from the environment can cause a temporary acceleration of the system's evolution on the state manifold.

While insightful, this perspective is fundamentally local, focusing on the differential geometry of the state space. Our topological approach is complementary and distinct. Instead of analyzing the geometry of the space of possible states, we analyze the topology of the space of possible histories (paths). Furthermore, TDA is concerned with global, scale-invariant properties (like the existence of a hole) rather than local properties like curvature or distance. It provides a qualitative, structural description of the entire dynamical landscape, whereas information geometry offers a quantitative, local measure of state distinguishability.

3 Methodology

Our methodology bridges the gap between the physics of non-Markovian dynamics and the mathematics of TDA. The central idea is to generate an ensemble of quantum trajectories, treat this ensemble as a point cloud in a high-dimensional "path space," and then use persistent homology to quantify the topology of this space.

3.1 Model System: The Spin-Boson Model

We employ the spin-boson model, a paradigmatic model for open quantum systems. It describes a two-level system (a qubit) coupled to a bath of harmonic oscillators. The total Hamiltonian is $H = H_S + H_B + H_I$, where $H_S = \frac{1}{2}\omega_0\sigma_z$ is the system Hamiltonian, $H_B = \sum_k \omega_k b_k^\dagger b_k$ describes the bosonic bath, and $H_I = \sigma_x \otimes \sum_k g_k(b_k^\dagger + b_k)$ is the interaction Hamiltonian. The influence of the bath is completely characterized by its spectral density, for which we choose a Lorentzian form:

$$J(\omega) = \sum_k g_k^2 \delta(\omega - \omega_k) = \frac{\gamma\lambda^2}{(\omega - \Omega)^2 + \lambda^2}$$

where γ is the coupling strength, λ is the width of the peak, and Ω is its central frequency. This model allows us to tune the degree of non-Markovianity by varying λ. A large λ (broad spectrum) corresponds to a short environmental correlation time, leading to Markovian dynamics. Conversely, a small λ (sharp spectrum) corresponds to a long environmental correlation time and induces strong non-

Markovian memory effects.

3.2 Generation of the Path Space Ensemble

We solve the exact non-Markovian dynamics using the numerically exact Hierarchical Equations of Motion (HEOM) method. HEOM is a non-perturbative technique that is particularly well-suited for systems coupled to a bath with a structured, non-Debye spectral density, making it ideal for our study. It provides an accurate time evolution of the system's reduced density matrix $\rho(t)$ by expanding the influence functional into a hierarchy of auxiliary density operators that capture the memory effects.

A single "path" is the time-ordered sequence of the system's state vector on the Bloch sphere, $\vec{r}(t) = (\langle\sigma_x(t)\rangle, \langle\sigma_y(t)\rangle, \langle\sigma_z(t)\rangle)$, where $\langle\sigma_i(t)\rangle = \mathrm{Tr}[\sigma_i\rho(t)]$. A path is discretized into N time steps, forming a vector in a high-dimensional space:

$$P = (\vec{r}(t_1), \vec{r}(t_2), \ldots, \vec{r}(t_N)) \in \mathbb{R}^{3N}$$

Each path P is a single point in the $3N$-dimensional path space. To generate a statistically meaningful ensemble, we simulate the evolution from a diverse set of M initial states sampled uniformly from the surface of the Bloch sphere. This is achieved by sampling the polar angle $\theta \in [0, \pi]$ and azimuthal angle $\phi \in [0, 2\pi)$ such that the distribution is uniform in $\cos\theta$. This procedure yields a point cloud $\{P_1, P_2, \ldots, P_M\}$ in $\mathbb{R}^{3N}$, which is the input for our topological analysis. We generate separate point clouds for different environmental regimes by varying λ.

3.3 Topological Analysis of the Path Space

With the point cloud representation, we apply persistent homology.

1. Defining a Metric: We use the standard Euclidean distance in $\mathbb{R}^{3N}$ to measure the dissimilarity between two paths, P_i and P_j. This metric gives equal weight to deviations at any point in time over the entire history and is sensitive to the magnitude of differences in the state vector components.

2. Building the Filtration: We construct a Vietoris-Rips filtration to probe the shape of the path space. This process can be understood through a simple geometric analogy. First, we imagine each complete quantum trajectory P_i as a single point in the high-dimensional space $\mathbb{R}^{3N}$. We then introduce a scale parameter $\epsilon \geq 0$, which we can think of as a growing radius. For each point, we place a ball of radius $\epsilon/2$. As we gradually increase ϵ from zero, these balls expand. Whenever two balls intersect, we connect their corresponding points with a line segment (a 1-simplex). When three balls mutually intersect, we fill in the triangle they define (a 2-simplex), and this process continues for higher dimensions (e.g., tetrahedra for four mutually intersecting balls). This creates a sequence of growing shapes, or simplicial complexes, one for each value of ϵ. This filtration process transforms the static point cloud into a dynamic topological space, allowing us to track the birth and death of structural features like loops and voids as a function of the scale ϵ.

3. Computing Persistent Homology: We compute the persistent homology of this filtration, tracking the Betti numbers, β_k, as a function of ϵ. The algorithm identifies the birth scale ϵ_{birth} when a feature appears and the death scale ϵ_{death} when it is filled. The persistence, $\epsilon_{death} - \epsilon_{birth}$, measures its significance.

Our primary focus is on the first Betti number, β_1, as we hypothesize that memory effects create loop-like structures. We define our topological non-Markovianity measure, $\mathcal{T}$, as the total persistence of all β_1 features:

$$\mathcal{T} = \sum_{i \in H_1} (\epsilon_{death,i} - \epsilon_{birth,i})$$

where the sum is over all 1-dimensional homology classes. A large $\mathcal{T}$ indicates significant, long-lived loops in the path space. We compare this measure with the standard BLP measure of non-Markovianity, $\mathcal{N}$, to validate our approach.

3.4 Numerical Convergence and Stability

The reliability of our TDA results depends on the adequate sampling of the path space. We performed a convergence analysis with respect to the number of trajectories (M) and the number of time discretization steps (N). For a fixed non-Markovian parameter set, we observed that the persistence diagram, particularly the location and persistence of the dominant β_1 feature, stabilized for $M > 400$.

We chose $M = 500$ to ensure robust results. Similarly, we verified that increasing the temporal resolution from $N = 100$ to $N = 200$ did not significantly alter the primary topological features or the value of $\mathcal{T}$, indicating that $N = 100$ is sufficient to capture the relevant dynamical timescales for the chosen simulation interval. The HEOM hierarchy was truncated at a level sufficient to guarantee convergence of the reduced density matrix dynamics to within a numerical precision of 10^{-6}.

4 Results

We present the results of our topological analysis of the path space for the spin-boson model. We compare a nearly Markovian regime with a strongly non-Markovian regime by tuning the spectral density width, λ.

4.1 Simulation Parameters

We set the model parameters as follows: qubit frequency $\omega_0 = 1$, system-bath coupling strength $\gamma = 0.1$, and bath central frequency $\Omega = \omega_0 = 1$ (resonant case). The dynamics are simulated over a time interval $T = 20$ (in units of $1/\omega_0$), discretized into $N = 100$ time steps, resulting in a path space embedded in $\mathbb{R}^{300}$. We generate an ensemble of $M = 500$ trajectories for each parameter set.

We investigate two distinct dynamical regimes:

1. **Markovian Regime:** We set the spectral width $\lambda = 2.5\omega_0$. The bath correlation time is short ($\tau_B \sim 1/\lambda = 0.4$), leading to dynamics that are well-approximated as memoryless.

2. **Non-Markovian Regime:** We set the spectral width $\lambda = 0.1\omega_0$. The bath correlation time is long ($\tau_B \sim 1/\lambda = 10$), inducing strong memory effects.

4.2 Persistence Diagrams

The persistence diagrams, which plot the death- versus birth-scale for each topological feature, reveal the underlying structure of the path space.

In the Markovian Regime ($\lambda = 2.5$): The persistence diagram, a plot of feature death-scale versus birth-scale, is dominated by points clustered tightly along the main diagonal. This visual signature represents a multitude of small, transient loops that appear and are quickly filled in as the scale parameter ϵ increases. Such features are characteristic of sampling noise in a high-dimensional space rather than robust topological structure. Crucially, the diagram shows no features with significant persistence; the maximum observed persistence for any one-dimensional feature is less than 0.05 (in normalized units of the scale parameter). This provides strong evidence that the Markovian path space is topologically simple, lacking any meaningful loops and resembling a contractible, solid object.

In the Non-Markovian Regime ($\lambda = 0.1$): The corresponding persistence diagram presents a starkly different topological portrait. While a similar cluster of low-persistence noise features lines the diagonal, the diagram is distinguished by a single point located far from it. This outlier represents a highly persistent one-dimensional feature—a structurally significant loop. We quantify its persistence (death-scale minus birth-scale) to be approximately 0.78, a value more than an order of magnitude larger than any feature in the Markovian case. The existence of this dominant, long-lived loop is a direct topological witness to the non-Markovian dynamics. Memory effects constrain the system's evolution, carving out a hole in the path space and forcing trajectories to organize around it, thereby creating a non-contractible manifold. This persistent loop is the clear topological artifact of information backflow.

4.3 Quantitative Topological Measure

To provide a systematic comparison, we computed our topological non-Markovianity measure, $\mathcal{T}$ (total β_1 persistence), and the standard BLP measure, $\mathcal{N}$, for a range of spectral widths λ.

The analysis reveals a strong correlation between the two measures. For large λ ($\lambda > 2.0\omega_0$), both $\mathcal{T}$ and $\mathcal{N}$ are close to zero, confirming the Markovian nature of the dynamics in this limit. As λ is decreased, the system enters the non-Markovian regime. Both measures begin to rise sharply around $\lambda \approx 1.5\omega_0$, indicating the onset of memory effects. Both $\mathcal{T}$ and $\mathcal{N}$ exhibit a peak at a similar intermediate value of $\lambda \approx 0.5\omega_0$, after which they decrease again for very small λ. This decrease in the very strong coupling regime is a known phenomenon where the system and a resonant mode of the bath form a quasi-bound state, leading to localization effects that reduce the dynamic information exchange. In the transition region ($0.1\omega_0 < \lambda < 1.5\omega_0$), the Pearson correlation coefficient between $\mathcal{T}$ and $\mathcal{N}$ is calculated to be greater than 0.95. This remarkable correspondence demonstrates that the emergence of persistent topological features in path space is a direct and reliable indicator of quantum memory as quantified by established physical measures.

4.4 Visualization of Path Space Structure

Although our analysis is performed in the full 300-dimensional space, dimensionality reduction techniques provide intuition about the underlying geometry. Projecting the path space onto its first two principal components via Principal Component Analysis (PCA) gives a low-dimensional view of the trajectory ensemble.

For the Markovian case, the PCA projection reveals a single, dense cloud of points with a roughly spherical or ellipsoidal shape. This geometry is a direct consequence of the dynamics: trajectories originating from all over the Bloch sphere converge exponentially and directly toward the single equilibrium state at the center. The ensemble of paths thus forms a cone-like structure in the full path space, which, when projected onto two dimensions, appears as a filled-in disk. This visualization aligns perfectly with the TDA result of a topologically trivial, contractible space.

In stark contrast, the PCA projection for the non-Markovian case displays a highly organized, non-trivial structure. The points are not clustered in a simple ball but instead form a distinct annular, or ring-like, arrangement, clearly delineating a central void. This projection provides a striking visual counterpart to the hole detected by persistent homology. The ring's structure is formed by oscillatory dynamics; trajectories exhibit coherence revivals, causing them to circle back through regions of the state space instead of converging directly. Paths cluster based on their phase relationship, which is dictated by the memory kernel, creating a circular flow in the low-dimensional projection. This visualization powerfully illustrates how memory constrains the flow of trajectories, preventing them from filling the central region and creating the loop detected by our topological analysis.

5 Discussion

The results establish a strong link between non-Markovianity and the emergence of non-trivial topology in the space of quantum trajectories. This section provides a physical interpretation for this connection and discusses the implications of our approach.

5.1 Physical Interpretation of Topological Signatures

The creation of loops in path space under non-Markovian dynamics can be understood through the concept of history-dependent information feedback. In a Markovian system, information flows unidirectionally to the environment, and all trajectories converge towards a steady state, forming a topologically simple, filled-in shape.

In a non-Markovian system, the environment returns information to the system, causing phenomena like coherence revivals. This feedback is not uniform; it depends on the system's past. This history-dependent evolution, mathematically captured by the memory kernel in formalisms like the Nakajima-Zwanzig equation, acts as a set of constraints. These constraints make certain sequences of states (regions in path space) dynamically inaccessible. These "forbidden zones" effectively puncture the path space. Trajectories are forced to flow around these voids, creating the persistent loops that our TDA analysis detects. A β_1 loop in path space can thus be interpreted as a family of evolutionary pathways that exhibit periodic or quasi-periodic information exchange with the environment, a hallmark of non-Markovian dynamics.

5.2 Advantages of the Topological Approach

This framework offers several potential advantages over existing methods:

1. **Global Perspective:** It provides a holistic characterization of the dynamics

over a finite time interval, sensitive to the integrated effect of memory over the entire evolution.

2. **Robustness to Noise:** Persistent homology is inherently robust to noise, as it distinguishes significant, persistent structures from transient, noisy features. This is a major asset for analyzing experimental data.

3. **Model-Agnosticism:** The method is general. Given any means of generating quantum trajectories (theory, simulation, or experiment), one can construct and analyze the corresponding path space.

4. **Richness of Information:** The persistence diagram is a rich descriptor. The number of persistent features, their persistence values, and their dimension (β_1, β_2, etc.) could potentially distinguish between different types of memory effects or higher-order correlations.

5.3 Limitations and Practical Implications

The primary challenge is computational cost, which scales with the system size and number of time steps. Research into efficient TDA algorithms is crucial for scalability. Furthermore, our study has certain limitations that suggest avenues for future work.

First, the results depend on the choice of metric used to define distances in the path space. We used the standard Euclidean metric, which is sensitive to point-wise differences in the Bloch vector. Other metrics, such as those based on dynamic time warping (DTW), could be more sensitive to the temporal shape of trajectories irrespective of phase shifts. The choice of metric is a crucial modeling decision, and the resulting topology could change depending on which features of the paths one wishes to emphasize.

Second, the physical interpretation of higher-dimensional topological features remains an open question. While a β_1 loop corresponds well to oscillatory information flow, the meaning of a persistent two-dimensional void (β_2) is less clear. It might signify a more complex constraint on the dynamics, perhaps indicating coordinated avoidance of a whole family of evolutionary "surfaces" in path space, possibly related to multi-time environmental correlations. Investigating systems that produce such features and deciphering their physical origin is a key direction for future research.

Despite these challenges, the potential applications are significant. One promising direction is in quantum sensing. Consider a quantum sensor, such as a nitrogen-vacancy (NV) center in diamond, used for nanoscale magnetic field detection. The sensor's sensitivity is limited by its interaction with a complex, often non-Markovian, spin bath environment. By repeatedly initializing the sensor and recording its output trajectory under the influence of the target field and the environment, one could generate a path space ensemble. Applying our topological analysis to this experimental data could serve as a powerful diagnostic tool. The persistence and number of topological features could be used to classify the nature of the environmental noise. This information could then be used to implement adaptive control protocols. For instance, if a strong β_1 feature is detected, it signifies a coherent information backflow, suggesting that dynamical decoupling sequences with a specific timing could be highly effective in mitigating this noise and enhancing sensor fidelity.

6 Conclusion

We have introduced a novel approach for characterizing non-Markovianity in open quantum systems by analyzing the topology of their path space. Moving beyond traditional local-in-time measures, we adopted a global perspective, treating the ensemble of quantum trajectories as a single geometric object.

Our central finding is that quantum memory imprints a distinct topological signature on the shape of this path space. Using persistent homology, we have shown that non-Markovian dynamics induce non-trivial features, most notably persistent one-dimensional loops, which are absent in the memoryless Markovian regime. We demonstrated a strong quantitative correlation between our topological measure and the established BLP measure for the spin-boson model. Physically, these topological features arise because memory-induced information backflow creates "forbidden zones" in the space of trajectories, forcing the dynamics to flow around them.

This topological framework reframes the abstract concept of quantum memory into a concrete geometric property. Its robustness and model-agnostic nature suggest significant potential for applications, particularly in the characterization of environmental noise for quantum technologies like quantum sensors. By analyzing the shape of quantum dynamics, we can uncover deep truths about information and memory in the quantum world, opening a new frontier at the intersection of quantum physics and algebraic topology.

7 References

Breuer, H. P., Laine, E. M., and Piilo, J., Measure for the Degree of Non-Markovian Behavior of Quantum Processes in Open Systems, Physical Review Letters 103, 210401 (2009).

Rivas, Á., Huelga, S. F., and Plenio, M. B., Quantum Non-Markovianity: Characterization, Quantification and Detection, Reports on Progress in Physics 77, 094001 (2014).

de Vega, I., and Alonso, D., Dynamics of non-Markovian open quantum systems, Reviews of Modern Physics 89, 015001 (2017).

Breuer, H. P., and Petruccione, F., The Theory of Open Quantum Systems (Oxford University Press, 2002).

Weiss, U., Quantum Dissipative Systems (World Scientific, 2012).

Tanimura, Y., and Kubo, R., Time evolution of a quantum system in contact with a nearly Gaussian-Markoffian noise bath, Journal of the Physical Society of Japan 58, 101-114 (1989).

Ishizaki, A., and Tanimura, Y., Quantum dynamics of a system strongly coupled to a low-temperature colored noise bath: A non-perturbative real-time path-integral approach, Journal of the Physical Society of Japan 74, 3131-3134 (2005).

Edelsbrunner, H., and Harer, J., Computational Topology: An Introduction (American Mathematical Society, 2010).

Carlsson, G., Topology and data, Bulletin of the American Mathematical Society 46, 255-308 (2009).

Zomorodian, A., and Carlsson, G., Computing persistent homology, Discrete and Computational Geometry 33, 249-274 (2005).

Leggett, A. J., Chakravarty, S., Dorsey, A. T., Fisher, M. P. A., Garg, A., and Zwerger, W., Dynamics of the dissipative two-state system, Reviews of Modern Physics 59, 1-85 (1987).

Feynman, R. P., and Vernon, F. L., The theory of a general quantum system interacting with a linear dissipative system, Annals of Physics 24, 118-173 (1963).

Chruściński, D., and Kossakowski, A., Non-Markovian Quantum Dynamics: A Geometrical Approach, Physical Review Letters 116, 130402 (2016).

Taddei, M. M., Escher, B. M., Davidovich, L., and de Matos Filho, R. L., Quantum speed limit for physical processes, Physical Review Letters 110, 050402 (2013).

Pires, D. P., Cianciaruso, M., Céleri, L. C., Adesso, G., and Soares-Pinto, D. O., Generalized Geometric Quantum Speed Limits, Physical Review X 6, 021031 (2016).

Wasserman, L., Topological Data Analysis, Annual Review of Statistics and Its Application 5, 501-528 (2018).

Ghrist, R., Barcodes: The persistent topology of data, Bulletin of the American Mathematical Society 45, 61-75 (2008).

Laine, E. M., Piilo, J., and Breuer, H. P., Measure for the non-Markovianity of quantum processes, Physical Review A 81, 062115 (2010).

Doherty, A. C., Habib, S., Jacobs, K., Mabuchi, H., and Tan, S. M., Quantum feedback control and classical control theory, Physical Review A 62, 012105 (2000).

Degen, C. L., Reinhard, F., and Cappellaro, P., Quantum sensing, Reviews of Modern Physics 89, 035002 (2017).